

# The Cage and the Mirror:

Compression, Selection, and Institutional Self-Deception

Jeremy McEntire

Cage & Mirror Press

[jmc@cageandmirror.com](mailto:jmc@cageandmirror.com)

April 2026

## Executive Summary

Every organization that has destroyed itself already knew what would save it. Boeing’s engineers flagged the problem. NASA’s data showed the risk. Theranos’s scientists knew the device did not work. The information existed. The organizations failed anyway.

This paper explains why. The answer is not stupidity, greed, or groupthink. It is architecture. When organizations grow, they build systems to coordinate decisions: metrics, approval chains, reporting structures. These systems compress messy reality into clean dashboards. The compression is necessary—no one can process everything—but it creates blind spots. What the system does not measure becomes invisible. And the system cannot see its own blind spots, because the only tool available for looking is the system itself.

Three forces make this worse. First, people shade their reports to match what their bosses want to hear. Second, producing a careful, honest assessment is expensive, while producing a reassuring summary is cheap—so reassuring summaries flood the system. Third, once a belief takes hold, contradicting it is career suicide. These three forces compound. By the time information travels from the people who know the truth to the people who make decisions, the truth has been filtered out.

The same mechanism that produces organizational blindness also produces creative breakthroughs. Whether compression generates dysfunction or innovation depends entirely on what the organization rewards. This is the central finding: the problem and the opportunity share a single cause.

The solution is not to eliminate the system—coordination requires it. The solution is to build structures that make the system’s blind spots visible from within. We call this *meta-compliance*: instead of pretending the governance system captures everything, the organization formally documents what it knows it cannot see, tracks the assumptions it relies on, and protects channels for information that contradicts the official story. A game-theoretic model shows this approach dominates across 70–82% of operating conditions. But a trap emerges: in turbulent

industries, legal pressure rises alongside uncertainty, pushing organizations toward rigid compliance precisely when flexibility matters most.

Practitioners should read Sections I, V–VII, and VIII. Theorists should read the full paper.

### Abstract

Every organization that has destroyed itself possessed the information that would have saved it. Boeing’s engineers flagged the MCAS system as a single-point-of-failure risk. Theranos’s laboratory directors knew the Edison device could not meet clinical accuracy standards. NASA’s O-ring data showed temperature dependence before the Challenger launch. The information existed. The organizations failed anyway.

This paper identifies the structural mechanism, develops the formal theory, provides a measurement architecture, and demonstrates the mechanism operating in legal systems. The problem is formalization itself: governance frames that enable coordination and satisfy fiduciary duty also compress decision variance, creating blind spots invisible from within. We test this against 75 SEC filings from 25 companies, finding universal variance compression (5.8–33%) following IPO, with magnitude proportional to legal exposure ( $R^2 = 0.80$ ). The mechanism is information-theoretic: three compound dynamics—strategic communication degradation (Crawford–Sobel), adverse selection in idea markets (Akerlof), and transmission bias (Boyd–Richerson)—produce *dysmemic pressure*, a selection force that favors fit over truth. An agent-based simulation sweeping 126 configurations confirms monotonic information destruction through hierarchical depth. The theory is substrate-independent: a generative lossy channel framework proves five sufficient conditions for net-beneficial noise (Theorem 1, validated computationally with zero counterexamples across 500 Monte Carlo configurations), establishes dual valence of compression artifacts (creative emergence vs. pathological drift depend on selection criteria, not the channel), and formalizes the compression ratchet that traps organizations in self-reinforcing equilibria. A formal foundation—the strategic rate-distortion-perception function unifying Crawford–Sobel and Blau–Michaeli—quantifies the cost of conformity: 22–72% of total distortion across representative parameter regimes. A stigmergic mesh architecture detects organizational dysfunction from work artifacts (TPR 0.82 for coordination failures, 22,500 Monte Carlo runs; 30-day field deployment with 23 confirmed findings). Mandatory arbitration combined with platform dominance creates a six-filter multiplicative pipeline suppressing 99.998% of meritorious claims; single-filter reform is structurally insufficient. The solution is architectural, not cultural: *meta-compliance*—formal governance structures that satisfy fiduciary duty by documenting awareness of incompleteness rather than claiming completeness. A game-theoretic model shows meta-compliance dominates 70–82% of the parameter space, but a *turbulence paradox* emerges when legal exposure rises with environmental turbulence, trapping organizations in variance elimination precisely when adaptation is most needed. Testable propositions with measurement protocols and falsification criteria are derived throughout.

**Keywords:** organizational incompleteness, dysmemic pressure, generative lossy channel, meta-compliance, fiduciary duty, variance compression, stochastic resonance, information degradation, institutional design, rate-distortion-perception, stigmergic mesh, structural immunity, mandatory arbitration

# Contents

<b>1</b>	<b>Introduction</b>	<b>6</b>
1.1	Positioning in the Literature . . . . .	7
<b>2</b>	<b>The Problem: Formalization Creates Blind Spots</b>	<b>8</b>
2.1	When Compression Is Beneficial . . . . .	8
2.2	Frame Dependence and Incompleteness-Like Properties . . . . .	9
2.3	The Legal Amplifier . . . . .	9
2.4	Propositions . . . . .	10
2.5	The Geometry of Defensibility . . . . .	10
2.6	Empirical Evidence: Variance Compression in SEC Filings . . . . .	11
2.6.1	Methodology . . . . .	11
2.6.2	Results: Universal Compression . . . . .	11
2.6.3	The Legal Amplifier Effect . . . . .	12
2.6.4	Founder Insulation . . . . .	12
2.6.5	Born Caged Confirmed . . . . .	13
2.6.6	Regression Analysis . . . . .	13
2.6.7	Genre Confound Controls . . . . .	14
2.6.8	Summary of Empirical Findings . . . . .	14
<b>3</b>	<b>The Mechanism: How Information Dies in Hierarchies</b>	<b>14</b>
3.1	Dynamic 1: Strategic Communication Degradation . . . . .	15
3.1.1	Formal Derivation . . . . .	15
3.2	Dynamic 2: Adverse Selection in Idea Markets . . . . .	16
3.3	Dynamic 3: Transmission Bias . . . . .	17
3.4	The Compound: Dysmemic Pressure . . . . .	17
3.5	Agent-Based Simulation of Hierarchical Information Degradation . . . . .	18
3.5.1	Model Architecture . . . . .	18
3.5.2	Parameter Sweep . . . . .	18
3.5.3	Key Findings . . . . .	18
3.5.4	Operationalizing Dysmemic Pressure . . . . .	19
3.6	Boeing as Running Example . . . . .	19
3.7	Connection to Forecasting . . . . .	20
<b>4</b>	<b>The Theory: When Noise Helps, When It Kills</b>	<b>20</b>
4.1	The Generative Lossy Channel . . . . .	21
4.1.1	Scope and Limits of the Organizational Mapping . . . . .	22
4.2	Level-Crossing: Noise at $N$ Becomes Information at $N+1$ . . . . .	22
4.3	When Does Noise Help? The Forbidden Interval Theorem . . . . .	23
4.4	Theorem 1: Five Sufficient Conditions for Net-Beneficial Noise . . . . .	24

4.5	Corollaries . . . . .	26
4.6	The Tolerance Location Principle . . . . .	26
4.7	Dual Valence . . . . .	27
4.8	The Compression Ratchet . . . . .	27
4.9	Computational Validation . . . . .	29
4.9.1	Strategy . . . . .	29
4.9.2	Stochastic Resonance Class . . . . .	29
4.9.3	Non-SR Mechanisms . . . . .	29
4.9.4	Monte Carlo Sensitivity Analysis . . . . .	29
4.10	Cross-Domain Evidence . . . . .	30
4.10.1	Organizational Disasters: Normalized Deviance as the Ratchet . . . . .	30
4.10.2	Jazz Improvisation: Creative Emergence Through Controlled Channel Noise . . . . .	31
4.10.3	AI Alignment: Sycophancy as Equilibrium Compression Artifact . . . . .	31
4.10.4	Cross-Domain Summary . . . . .	31
<b>5</b>	<b>The Formal Foundation: Strategic Rate-Distortion-Perception</b>	<b>32</b>
5.1	The Strategic RDP Function . . . . .	32
5.2	Theorem A: Generalized RDP Tradeoff . . . . .	33
5.3	Theorem B: The Gaussian Closed Form and the Generative Residual . . . . .	33
5.4	Numerical Illustration: Three Regimes . . . . .	34
5.5	Organizational Sufficient Conditions . . . . .	34
5.6	A Motivating Conjecture: RLHF Sycophancy . . . . .	35
<b>6</b>	<b>Measurement: Detecting Organizational Dysfunction</b>	<b>35</b>
6.1	Architecture: A Self-Organizing Mesh . . . . .	35
6.2	Three Detection Signatures for Normalized Deviance . . . . .	36
6.3	Controlled Evaluation: 22,500 Monte Carlo Runs . . . . .	37
6.4	Field Deployment . . . . .	37
<b>7</b>	<b>Application: Structural Immunity in Legal Systems</b>	<b>38</b>
7.1	The Six-Filter Multiplicative Pipeline . . . . .	38
7.2	The Compound Suppression Result . . . . .	39
7.3	Sensitivity Analysis . . . . .	39
7.4	The Independence Assumption and the FKG Defense . . . . .	40
7.5	Cross-Sector Variation . . . . .	40
7.6	Policy Conclusion . . . . .	40
<b>8</b>	<b>The Solution: Meta-Compliance</b>	<b>40</b>
8.1	Meta-Compliance and Sanctioned Paradoxes . . . . .	41
8.2	A Game-Theoretic Model of Governance Mode Selection . . . . .	42
8.2.1	Calibration . . . . .	43

8.2.2	Sensitivity . . . . .	43
8.3	The Turbulence Paradox . . . . .	44
8.4	Empirical Illustrations . . . . .	44
8.4.1	NASA Post-Columbia . . . . .	44
8.4.2	Pixar Braintrust . . . . .	45
8.4.3	Bridgewater Associates . . . . .	45
8.5	Consolidated Propositions . . . . .	45
8.6	Failure Modes: When Meta-Compliance Becomes Compliance Theater . . . . .	46
<b>9</b>	<b>Discussion and Conclusion</b>	<b>47</b>
9.1	The Full Causal Chain . . . . .	47
9.2	Policy Implications . . . . .	48
9.3	Open Problems . . . . .	48
9.4	Testable Predictions Summary . . . . .	49
9.5	Conclusion . . . . .	49

# 1 Introduction

In January 1986, engineers at Morton Thiokol told NASA managers that the Space Shuttle Challenger’s O-ring seals had never been tested at the launch-day temperature. The managers overruled the engineers. Seven crew members died seventy-three seconds after liftoff. In October 2018, a Boeing 737 MAX dove into the Java Sea, killing 189 people. Boeing’s own engineers had flagged the MCAS flight-control system as a single-point-of-failure risk years earlier. Five months later, a second MAX crashed in Ethiopia, killing 157 more. Between 2003 and 2018, Theranos claimed revolutionary blood-testing technology while its laboratory directors knew the Edison device could not meet clinical accuracy standards. Employees who raised concerns were fired or threatened with legal action.

These are not stories about bad people. The engineers at Boeing were competent. The O-ring data at NASA was available to anyone who looked. Theranos employed qualified laboratory scientists who understood the chemistry. In each case, the organization possessed the information that would have saved it. The information existed in documents, databases, and the minds of employees throughout the hierarchy. The organization failed anyway.

The conventional explanations invoke psychology: cognitive bias, groupthink, willful blindness. These explanations describe individual failures. They do not explain why the pattern recurs across industries, decades, and organizational forms. They do not explain why organizations staffed by people who privately know the truth collectively act on falsehood. The mechanism is not psychological. It is structural.

This paper develops that structural mechanism in seven parts. Section 2 establishes that formalization—the governance frames organizations adopt to coordinate at scale and satisfy legal requirements—compresses decision variance, creating blind spots invisible from within. This is tested empirically against 75 SEC filings. Section 3 identifies the information-theoretic mechanism: three compound dynamics produce a selection force we call *dysmemic pressure*, formalized through a hierarchical Crawford–Sobel model and validated with agent-based simulation. Section 4 proves the theory is substrate-independent: a generative lossy channel framework establishes when compression artifacts are pathological and when they are beneficial, with the valence determined by selection criteria rather than the channel itself. Section 5 provides the formal information-theoretic foundation: a strategic rate-distortion-perception function that unifies Crawford–Sobel strategic communication with Blau–Michaeli perception constraints, quantifying how much distortion is attributable to conformity pressure versus strategic misalignment. Section 6 develops a measurement architecture—a stigmergic mesh that detects organizational dysfunction from work artifacts without requiring anyone to know what to look for. Section 7 demonstrates the mechanism operating in legal systems: mandatory arbitration combined with platform dominance creates a multiplicative suppression pipeline that renders legal rights unexercisable in practice. Section 8 develops the architectural solution: *meta-compliance*, governance structures that satisfy fiduciary duty by documenting awareness of incompleteness rather than claiming completeness.

The seven parts form a causal chain, then extend it. Formalization creates frames. Frames

compress information. Compression creates selection environments. Selection environments produce systematic self-deception. The mechanism is substrate-independent—it appears wherever compression and selection interact. The formal foundation (Section 5) proves this mathematically: conformity pressure accounts for 22–72% of total distortion in strategic communication channels. The measurement section (Section 6) shows the mechanism can be detected in practice. The legal application (Section 7) shows it operating in legal architecture, where six multiplicative filters suppress 99.998% of meritorious claims. Crucially, the same mechanism that produces organizational blindness also produces creative emergence. Whether compression generates dysfunction or innovation depends on the selection pressure applied to the output. This dual-valence is not a caveat—it is the central theoretical finding. And the solution is not escape from the cage but conscious placement of mirrors within it: structures that make the cage’s boundaries visible to those operating inside.

## 1.1 Positioning in the Literature

The organizational literatures have documented every piece of this puzzle without assembling it. March (1991) formalizes the drift from exploration to exploitation. Weber describes rationalization becoming self-perpetuating (Weber, 1958). Christensen (1997) documents incumbent failure during disruption. Agency theory explains incentive misalignment (Jensen and Meckling, 1976). Organizational learning theory distinguishes single-loop from double-loop but does not explain why organizations systematically select against double-loop even when leaders understand its value (Argyris and Schön, 1978). DiMaggio and Powell (1983) identify isomorphic pressures without explaining why isomorphism produces rigidity, not just similarity. Meyer and Rowan (1977) show that organizations adopt formal structures as “myths and ceremonies” for legitimacy. Scott (2013) synthesizes institutional theory around regulative, normative, and cultural-cognitive pillars.

In information economics, Crawford and Sobel (1982) prove that strategic communication under preference divergence is endogenously lossy. Akerlof (1970) shows adverse selection in markets with quality uncertainty. In cultural evolution, Boyd and Richerson (1985) document transmission biases independent of truth value. In information theory, Blau and Michaeli (2018; 2019) prove that lossy compression under a perceptual quality constraint forces the decoder to generate. Atlan (1979) establishes that noise at one level can produce information at a higher level.

In high-reliability organizations, Weick and Sutcliffe (2007) identify five characteristics of organizations that sustain safety in high-hazard environments. Edmondson (1999; 2019) demonstrates that psychological safety predicts learning behavior. In military doctrine, mission command formalizes bounded autonomy within hierarchical structures (Shamir, 2011). In legal theory, the business judgment rule protects directors who demonstrate informed process (*Smith v. Van Gorkom*, 1985). Tetlock (2005) shows that fox-style forecasters dramatically outperform hedgehogs, and superforecasting teams implement structural features that map directly onto anti-dysmemic architecture (Tetlock and Gardner, 2015).

Each of these literatures has identified a fragment. This paper identifies the common mechanism

that unifies them: compression creates information gaps; selection operating in those gaps determines whether the system self-deceives or self-corrects. The contribution is the integration—and the formal machinery that makes the integration precise.

## 2 The Problem: Formalization Creates Blind Spots

Organizations grow. Growth demands coordination. Coordination demands formalization: policies, procedures, metrics, approval chains, reporting structures. Public ownership adds a second pressure: fiduciary duty requires that strategic decisions be *demonstrably* sound—not merely correct, but documentably so. The two pressures converge. Both push toward governance frames that compress rich realities into legible dashboards.

The compression is not neutral. What survives compression is what the frame measures. What the frame does not measure becomes invisible—not contested, not debated, but absent from the decision space entirely. Continuous improvement makes the organization better at optimizing within its metrics while worsening what those metrics miss. And the frame’s limitations cannot be stated in the frame’s own terms: challenges to the frame are evaluated by the frame’s criteria and found wanting. The frame defends itself.

### 2.1 When Compression Is Beneficial

An important boundary condition: compression is not inherently pathological. In stable, routine domains—manufacturing, military operations, franchise systems—compression is essential and beneficial. McDonald’s succeeds precisely because it has compressed the variability out of food preparation; the environment is engineered to be stable, and the frame matches reality closely enough that what it excludes does not matter. Military standard operating procedures save lives by eliminating decision latency in well-understood scenarios. Franchise models scale by compressing local variation into reproducible templates.

The pathology arises specifically when two conditions hold simultaneously: (a) the environment changes faster than governance frames update, and (b) the organization faces novel or uncertain challenges that require information the frame excludes. Boeing did not fail because it had procedures. It failed because aerospace safety requires detecting novel failure modes—precisely the kind of information that compression removes. The MCAS system was a new risk in an environment that the existing certification frame was not designed to evaluate. McDonald’s compression works because hamburgers do not present novel failure modes. Boeing’s compression killed people because aircraft systems do.

The theory developed in this paper applies to the second case: organizations operating in environments where the gap between the governance frame and reality is large, growing, or consequential. In stable environments with well-matched frames, compression is a feature, not a bug.

## 2.2 Frame Dependence and Incompleteness-Like Properties

Organizations exhibit incompleteness-like properties when four conditions hold:

1. **Frame dependence.** Decisions are justified using internal metrics. Success is defined by the frame; the frame is validated by success.
2. **Systematic blind spots.** Compression removes dimensions. What is lost becomes invisible—not merely unknown but structurally absent from the decision space.
3. **Optimization deepens blindness.** Continuous improvement along measured dimensions makes the organization more efficient at its own metrics while worsening what those metrics miss.
4. **Frame invisibility and self-defense.** The frame’s limitations cannot be stated in the frame’s terms. Challenges read as irrational, uninformed, or disloyal from within.

These four properties are operational, not metaphorical. They generate testable predictions about what happens when organizations formalize—predictions tested in Section 2.6.

## 2.3 The Legal Amplifier

The business judgment rule protects directors from liability when they demonstrate that their decision was informed and made in good faith. Over decades of litigation, “informed” evolved from a substantive standard to an evidentiary one: can they *show* they followed reasonable procedures?

The transformation crystallized in *Smith v. Van Gorkom* (*Smith v. Van Gorkom*, 1985), where Delaware’s Supreme Court held directors personally liable despite acting without self-interest and producing a premium for shareholders. The board approved a merger after a two-hour meeting with limited documentation. The outcome was favorable. The process was deficient. Proof of prudence matters more than correctness of outcome.

Directors learn the lesson. Documentation quality predicts litigation survival better than decision quality. A mediocre choice supported by thick documentation survives challenge. A superior choice lacking formal justification invites liability. The incentive is unambiguous: when outcomes are uncertain, optimize for defensibility over effectiveness.

The effect scales with legal exposure. Private companies with concentrated ownership face limited derivative suit risk. Public companies with dispersed shareholders face continuous litigation risk; every significant decision must be defensible to strangers with adversarial incentives. The Sarbanes-Oxley Act (*Sarbanes-Oxley*, 2002) and Dodd-Frank Act (*Dodd-Frank*, 2010) formalized this trend, mandating additional controls, certifications, and audit trails.

Consider two decision paths. Path A: follow documented metrics, cite industry precedent, secure formal approvals. If it fails, the director is protected. Path B: contradict metrics based on external perspective, deviate from industry norms. If it fails, the director is exposed. Rational directors converge on Path A regardless of which might produce better outcomes. The system produces the problem and forecloses the solution.

## 2.4 Propositions

From this mechanism follow three testable propositions.

**Proposition 2.1** (Variance Compression). *Following formalization events, the variance of strategic decision justifications decreases measurably. Let  $V(t)$  denote variance in strategic decision justifications at time  $t$ , operationalized as linguistic entropy in strategy documents, number of formally documented alternatives, and dispersion in approval votes. Let  $F(t)$  denote a formalization event (IPO, major compliance mandate, significant litigation). For a post-event window  $\Delta \in [6, 24]$  months:*

$$V(t + \Delta) < V(t)$$

*for decisions of comparable magnitude.*

**Boundary conditions.** Strongest in high-liability industries with dispersed shareholders; weakest in founder-controlled firms with fast feedback loops.

**Falsification.** If variance does not compress post-event when controlling for decision complexity and environment, the formalization–compression link is not supported.

**Corollary 2.2** (Born Caged). *In mature regulatory environments, the formalization event moves upstream of the IPO itself. S-1 registration statements drafted in post-SOX, high-litigation environments exhibit lower initial linguistic variance than S-1s from earlier regulatory eras, holding industry and business-model novelty constant.*

**Proposition 2.3** (Independence–Novelty). *Units that maintain independent evaluation criteria produce more category-creating innovations than units evaluated by parent metrics. Let  $N$  be the count of category-creating innovations over five years. Let  $I \in [0, 1]$  be an independence score. Then:*

$$\mathbb{E}[N \mid I > 0.7] > \mathbb{E}[N \mid I < 0.3].$$

**Proposition 2.4** (Authority–Impact). *Dissent mechanisms with authority to raise concerns without prior approval change more strategic decisions than mechanisms that require approval to speak. Let  $C$  be the count of significant strategic changes attributed to a dissent mechanism over two years. Let authority score  $A \in \{0, 1, 2, 3\}$  count the presence of (i) direct escalation rights, (ii) formal delay authority, (iii) gate-free issuance of reports. Then:*

$$\mathbb{E}[C \mid A = 3] > \mathbb{E}[C \mid A = 0].$$

## 2.5 The Geometry of Defensibility

A space defined by formalization density (horizontal) and legal exposure (vertical) produces four diagnostic regions:

- **Lower-left: “Coin Flip Zone.”** Low formalization and exposure; internal learning suffices.
- **Lower-right: “Efficient Private Formalization.”** High formalization, low exposure; efficiency dominates.
- **Upper-left: “Growth Gauntlet.”** Rising exposure with loose structures; transient, pressure to formalize.
- **Upper-right: “Fiduciary Trap.”** High formalization and exposure; defensibility dominates; external perspective most necessary and least permissible.

Typical trajectories move from lower-left toward upper-right (growth → IPO → formalization). The upper-right corner marks the strongest predicted co-occurrence of all three proposition effects.

## 2.6 Empirical Evidence: Variance Compression in SEC Filings

### 2.6.1 Methodology

For each of 25 target companies, we identified the foundational S-1 (or F-1/S-1/A) registration statement and the first two subsequent annual reports (10-K or 20-F/40-F). The complete text from two key strategic sections—“Item 1: Business” and “Item 7: Management’s Discussion and Analysis (MD&A)” —was extracted from each filing. Two metrics were computed:

- **Lexical Diversity (LD):** The ratio of unique words to total words. Lower values indicate more repetitive vocabulary.
- **Shannon Entropy (SE):** Information uncertainty (base 2). Lower values indicate more predictable text.

Companies were organized into five analytical cohorts: (1) Pre-SOX Tech (IPO 1997–1999), (2) Post-SOX/Web 2.0 (IPO 2002–2012), (3) Modern Cloud/Social (IPO 2012–2020), (4) High-Liability/Regulated, and (5) Founder Control. Two mature public companies (Adobe, Intuit) served as controls.

### 2.6.2 Results: Universal Compression

Across all 25 companies and all three temporal cohorts, the transition from S-1 to Y1 10-K produces a decline in both lexical diversity and Shannon entropy. The compression is universal in direction and substantial in magnitude, ranging from 5.8% (Meta, founder-controlled) to 33% (Coinbase, high-liability). The tandem movement of both metrics confirms that post-IPO language is not merely narrower but structurally simpler.

The control cases sharpen the finding. Adobe’s 10-K metrics from 1995–1997 hover at  $LD = 0.1305 \pm 0.0004$  and  $SE = 10.44 \pm 0.02$ . Intuit shows comparable stasis. Absent a formalization event, the linguistic profile of a mature public company remains remarkably fixed.

Three patterns within the baseline confirm the theoretical predictions:

We emphasize that this evidence is correlational. The compression we measure in language may reflect appropriate professionalization rather than harmful information loss. We do not claim that lexical diversity directly measures decision quality. We claim that the pattern—universal, moderated by legal exposure, absent in low-formalization controls—is consistent with the mechanism we formalize in Section 3 and inconsistent with genre-convention explanations alone.

**Category creators compress most.** Firms that must invent strategic language—Amazon describing e-commerce in 1997, Google defining search-as-business-model in 2004, Salesforce explaining SaaS in 2004—show the largest absolute drops (14.6–17.6% in Y1). The S-1’s task is to explain a novel concept; the 10-K’s task is to defend financial results. The shift from evangelism to defense produces the sharpest compression.

**Market followers start lower but still compress.** Workday (S-1 LD: 0.1504) and ServiceNow (S-1 LD: 0.1533) entered the SaaS market as challengers, not creators. The compression is more moderate (9.2–11.8%) but equally consistent.

**Compression continues but decelerates.** In nearly all cases, the Y1-to-Y2 decline is smaller than the S-1-to-Y1 decline. The sharpest formalization occurs at the IPO boundary; subsequent years show continued but diminishing compression as the language approaches its formalized steady state.

### 2.6.3 The Legal Amplifier Effect

The Group 4 firms test the legal amplifier prediction directly. Coinbase’s S-1 attempted to evangelize the “cryptoeconomy” (LD: 0.1944, SE: 12.48). Its Y1 and Y2 10-Ks were filed while the company was under active SEC investigation. Lexical diversity collapsed by 33% to 0.1302 in Y1 and continued falling to 0.1226 in Y2—the lowest endpoint in the dataset. Robinhood presents a mirror image: “democratizing finance” compressed 30% in Y1 following the GameStop controversy and congressional hearings. Moderna’s S-1 described mRNA as a “new class of medicines” (LD: 0.1910); the 10-Ks show a 16.2% Y1 drop. Blackstone, already “born caged” by financial industry standards (S-1 LD: 0.1412, the lowest initial variance in the dataset), still compressed a further 7.5%.

The pattern is consistent: when legal threat becomes existential, variance compression is not gradual drift but rapid, forced capitulation.

### 2.6.4 Founder Insulation

Meta’s S-1 contained a high-variance founder letter (LD: 0.1709). The Y1 drop was only 5.8%—substantially milder than the 15.3% drop at LinkedIn, a non-founder-controlled peer from the same era. Snap’s S-1 famously defined Snap as a “camera company” (LD: 0.1730); despite intense

criticism and a collapsing stock price, the Y1 drop was only 7.2%. The dual-class structure allowed the founder to resist.

Founder control acts as an insulator, not a shield. The metrics still decline. But the rate and magnitude are significantly attenuated relative to non-founder-controlled peers.

### 2.6.5 Born Caged Confirmed

Snowflake’s 2020 S-1 (LD: 0.1495) and Okta’s 2017 S-1 (LD: 0.1519) registered substantially lower initial variance than category creators of earlier eras: Amazon 1997 (LD: 0.1742), Google 2004 (LD: 0.1825). The legal and banking ecosystem has internalized the requirements of demonstrable soundness so thoroughly that the cage begins to operate before the firm is formally public.

### 2.6.6 Regression Analysis

We estimate three nested OLS models (HC1 robust standard errors) with lexical diversity as the dependent variable across 60 filing observations (20 IPO companies  $\times$  3 filings each; mature controls excluded).

Table 1: Regression Results: Lexical Diversity

	(1) Base	(2) Interactions	(3) Company FE
Intercept	0.1689*** (0.0033)	0.1667*** (0.0040)	0.1668*** (0.0049)
years_since_ipo	-0.0141*** (0.0018)	-0.0120*** (0.0022)	-0.0120*** (0.0018)
is_high_liability	-0.0067** (0.0030)	-0.0016 (0.0058)	
is_founder_controlled	0.0185*** (0.0038)	0.0170*** (0.0053)	
is_born_caged	0.0107* (0.0061)	0.0115 (0.0085)	
years $\times$ high_liability		-0.0102 (0.0065)	-0.0102 (0.0064)
years $\times$ founder		0.0015 (0.0037)	0.0015 (0.0034)
years $\times$ born_caged		-0.0008 (0.0034)	-0.0008 (0.0032)
Company FE	No	No	Yes
$R^2$	0.599	0.644	0.798
Adj. $R^2$	0.554	0.580	0.669
$N$	60	60	60

Standard errors (HC1) in parentheses. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ . Cohort dummies included in Models (1)–(2) but suppressed for space. Company FE in Model (3) absorb all time-invariant firm characteristics.

Each additional year post-IPO reduces lexical diversity by 0.012 ( $p < 0.001$ ), confirming that variance compression is not a one-time genre effect but a continuing process. The interaction term

for high-liability firms is negative ( $-0.010$ ), indicating that legally exposed firms compress at nearly double the baseline rate—an 85% amplification effect. The company fixed-effects model explains 80% of the variance in lexical diversity.

### 2.6.7 Genre Confound Controls

A natural objection: the S-1 to 10-K transition reflects genre differences, not compression. Three controls address this:

Table 2: Genre Confound Controls

Transition	$N$	Mean $\Delta$ LD	Paired $t$	$p$ -value	Cohen’s $d$
S-1 $\rightarrow$ Y1	20	-0.02452	-7.772	<0.0001	-1.738
Y1 $\rightarrow$ Y2	20	-0.00373	-9.370	<0.0001	-2.095
S-1 $\rightarrow$ Y2	20	-0.02825	-8.271	<0.0001	-1.850
<i>Mature Controls (Adobe, Intuit)</i>					
Adobe (3-yr range)	3	0.0008	—	—	—
Intuit (3-yr range)	3	0.0005	—	—	—

Compression continues within the 10-K genre itself: Y1 to Y2 shows a statistically significant decline ( $t = -9.37$ ,  $p < 0.0001$ , Cohen’s  $d = -2.10$ ). Same document type, same company, one year apart—yet variance still compresses. Mature public companies show effectively zero drift over three years (LD range  $< 0.001$ ). All three transitions produce changes significantly different from zero with large effect sizes ( $|d| > 1.7$ ).

### 2.6.8 Summary of Empirical Findings

1. **Universal compression.** All 25 companies show LD and SE declining from S-1 to Y1, with continued decline to Y2.
2. **Legal amplifier confirmed.** High-liability firms show the most extreme compression (16–33%).
3. **Founder insulation confirmed.** Dual-class firms show attenuated compression (5.8–7.2%).
4. **Born Caged confirmed.** Modern S-1s exhibit lower initial variance than earlier-era S-1s.
5. **Control cases confirm attribution.** Adobe and Intuit show linguistic stasis ( $\pm 0.3\%$ ).
6. **Strategic discontinuity bounds the model.** Netflix shows that radical pivots can temporarily overcome compression.

## 3 The Mechanism: How Information Dies in Hierarchies

Section 2 established that formalization compresses variance. This section identifies the information-theoretic mechanism. Three dynamics compound to produce what we call *dysmemic pressure*: the

selection force that favors fit over truth in environments shaped by compressed representations.

The compound formulation is the contribution. Each dynamic has been documented separately. Their interaction explains what the components alone cannot: why intelligent organizations deceive themselves despite incentives and intentions pointing the other way.

Consider a pharmaceutical company whose scientists discover that a lead drug candidate has a safety signal. The discovery must travel from bench scientist to project lead to division head to chief medical officer to CEO to board. At each interface, three forces operate simultaneously. First, the scientist’s interests diverge from the manager’s—Crawford–Sobel strategic communication degrades the signal. Second, producing a rigorous safety assessment is expensive while producing a reassuring summary is cheap—Akerlof adverse selection floods the internal information market with optimistic signals. Third, the CEO’s public commitment to the drug candidate makes skepticism low-status—Boyd–Richerson transmission bias ensures the optimistic narrative propagates while the safety concern dies. By the time “information” about the drug reaches the board, it bears no meaningful relationship to the bench data.

That is dysmemic pressure. The compound of three well-understood dynamics, producing a result none of them predicts alone.

### 3.1 Dynamic 1: Strategic Communication Degradation

When someone with private information sends a message to someone who must act on it, how much truth gets transmitted? The answer depends on how aligned their interests are.

Crawford and Sobel (1982) worked out the mathematics. Perfect alignment produces full transmission. As interests diverge, the sender lumps states into coarser categories, smoothing away details that would hurt the sender. At sufficient divergence, messages carry no information at all—the *babbling equilibrium*, where the sender’s messages become statistically independent of reality.

Organizations stack these interfaces. The engineer’s interests diverge from the manager’s. The manager’s from the director’s. The director’s from the executive’s. Each interface is a degradation point. Each responds rationally to local incentives. The aggregate is organizational self-deception built from individually rational choices.

#### 3.1.1 Formal Derivation

Consider the canonical Crawford–Sobel setup with uniform-quadratic preferences:

- The state  $\theta$  is drawn uniformly from  $[0, 1]$ .
- The Sender observes  $\theta$  and transmits a costless message  $m$ .
- The Receiver observes  $m$  and chooses an action  $a$ .
- Sender utility:  $U_S(a, \theta; b) = -(a - \theta - b)^2$ , so the Sender prefers  $a = \theta + b$ .

- Receiver utility:  $U_R(a, \theta) = -(a - \theta)^2$ , so the Receiver prefers  $a = \theta$ .
- The bias parameter  $b \in [0, 1/2)$  measures preference divergence.

**Proposition 3.1** (Partition count). *In the most-informative equilibrium, the Sender partitions  $[0, 1]$  into  $N^*(b)$  intervals, where*

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

The partition boundaries are

$$a_i = \frac{i}{N} + 2b \cdot i(i - N), \quad i = 0, 1, \dots, N, \quad (2)$$

satisfying  $a_0 = 0$  and  $a_N = 1$ . The Receiver's equilibrium action in interval  $i$  is the midpoint  $(a_{i-1} + a_i)/2$ .

*Proof.* The boundary  $a_i$  is determined by the Sender's indifference condition: the Sender at state  $\theta = a_i$  is indifferent between inducing the Receiver's action for interval  $i$  and that for interval  $i + 1$ . With quadratic loss and uniform priors, the indifference condition yields the recurrence  $a_{i+1} - a_i = a_i - a_{i-1} + 4b$ , with boundary conditions  $a_0 = 0$ ,  $a_N = 1$ . Solving this linear recurrence gives (2). The constraint that all intervals have positive length ( $a_i < a_{i+1}$  for all  $i$ ) yields the maximum  $N$  satisfying  $N(N - 1) < 1/(2b)$ , which gives (1).  $\square$

**Proposition 3.2** (Babbling boundary). *As  $b \rightarrow 1/4$ , the equilibrium partition count  $N^*(b) \rightarrow 1$  and the mutual information  $I(\theta; m) \rightarrow 0$ . The channel degenerates to babbling: the Sender's message becomes statistically independent of the true state.*

*Proof.* At  $b = 1/4$ , equation (1) gives  $N^*(1/4) = \lfloor (-1 + \sqrt{1 + 8})/2 \rfloor = \lfloor 1 \rfloor = 1$ . With  $N = 1$ , the Sender's message is constant regardless of  $\theta$ , so  $I(\theta; m) = 0$ . For  $b$  slightly below  $1/4$ ,  $N^* = 2$  and the partition is  $\{[0, a_1], [a_1, 1]\}$  with  $a_1 = 1/2 - 2b$ , which approaches 0 as  $b \rightarrow 1/4$ , so the mutual information approaches 0 continuously.  $\square$

**Definition 3.3** (Channel mutual information). Since  $\theta \sim \text{Uniform}[0, 1]$  and the message  $m$  identifies which of  $N$  intervals  $\theta$  falls in, with interval lengths  $L_i = a_i - a_{i-1}$ :

$$I(\theta; m) = h(\theta) - h(\theta | m) = - \sum_{i=1}^N L_i \log_2 L_i. \quad (3)$$

This equals the entropy of the discrete distribution over partition intervals.

## 3.2 Dynamic 2: Adverse Selection in Idea Markets

Akerlof (1970) showed what happens when buyers cannot assess quality before purchase. The same dynamic operates in organizational information markets. Producing accurate assessments is

expensive: it requires gathering data, doing analysis, acknowledging uncertainty, and delivering conclusions people would rather not hear. Producing optimistic assessments is cheap: it requires confidence and alignment with what receivers prefer. At the moment of presentation, accurate and optimistic look identical.

When receivers cannot verify quality at the moment of consumption, cheap signals flood the market. Producers of accurate signals face a problem: they bear higher costs for signals indistinguishable from cheap ones. Rationally, some reduce their investment in accuracy. Some stop producing altogether. The market settles at noise.

### 3.3 Dynamic 3: Transmission Bias

Boyd and Richerson (1985) documented how ideas spread independent of their truth value. Selection operates on transmissibility.

*Content bias*: simple ideas spread faster than complex ones. “We just need to execute better” is lighter cognitive load than “Our architecture has accumulated technical debt requiring a multi-quarter remediation effort with uncertain ROI.”

*Prestige bias*: ideas associated with successful people spread faster than identical ideas from unknown sources. If the CEO believes the competitor is irrelevant, that belief cascades downward regardless of evidence (Henrich and McElreath, 2003).

*Conformity bias*: once a belief reaches critical mass, deviation becomes costly. If everyone reports green, reporting red marks you as the problem (Bikhchandani et al., 1992).

### 3.4 The Compound: Dymemic Pressure

The three dynamics compound. Strategic degradation means accurate signals face friction proportional to preference divergence. Adverse selection means accuracy is costly and underrewarded. Transmission bias means whatever survives spreads based on spreadability, not accuracy.

In plain terms, dymemic pressure is the tendency of organizations to forget inconvenient truths. It operates through three reinforcing dynamics: people shade their reports to please their bosses, good ideas get driven out by safe ones, and each retelling loses nuance.

**Definition 3.4** (Dymemic Pressure). *Dymemic pressure* is the compound selective force favoring fit over truth in environments shaped by compressed representations. It intensifies with compression ratio, preference divergence, verification cost, and stakes. A *dymeme* is a signal optimized for survival in the gap between representation and reality. Its survival depends on alignment with receiver preferences, ease of transmission, and conformity with the acceptability distribution. Truth is not the selection criterion.

The pressure creates a ratchet. Each dymeme that establishes itself tilts the landscape. The next dymeme becomes easier to establish. The next accurate signal becomes harder to transmit.

## 3.5 Agent-Based Simulation of Hierarchical Information Degradation

### 3.5.1 Model Architecture

We model an organizational hierarchy as a cascade of Crawford–Sobel channels. Ground truth  $\theta \sim \text{Uniform}[0, 1]$  enters at Layer 1. Each successive layer  $k$  observes the previous layer’s action  $a_{k-1}$  and communicates through its own biased channel with parameter  $b_k$ .

**Definition 3.5** (Hierarchical channel). A  $K$ -layer hierarchical channel is a Markov chain  $\theta \rightarrow a_1 \rightarrow a_2 \rightarrow \dots \rightarrow a_K$ , where each transition  $a_{k-1} \rightarrow a_k$  is a Crawford–Sobel channel with bias  $b_k$  and optional additive Gaussian noise  $\epsilon_k \sim \mathcal{N}(0, \sigma^2)$ .

By the *data processing inequality* (Cover and Thomas, 2006), information can only be lost through the hierarchy:

$$I(\theta; a_K) \leq I(\theta; a_{K-1}) \leq \dots \leq I(\theta; a_1) \leq H(\theta). \quad (4)$$

This is the information-theoretic expression of dysmemic pressure: each layer of strategic communication can only destroy information about ground-level reality, never create it.

### 3.5.2 Parameter Sweep

The simulation performs a full parameter sweep:

- **Hierarchy depths:**  $K \in \{3, 4, 5, 6, 7, 8, 10\}$  layers.
- **Per-layer bias:**  $b \in \{0.01, 0.05, 0.10, 0.15, 0.20, 0.24\}$ .
- **Observation noise:**  $\sigma^2 \in \{0, 0.01, 0.05\}$ .

This yields  $7 \times 6 \times 3 = 126$  configurations, each evaluated over 50 Monte Carlo iterations with 1,000 samples per iteration.

### 3.5.3 Key Findings

**Finding 1: Without noise, partition structure is preserved.** When  $\sigma^2 = 0$ , the effective partition count remains stable through all hierarchy depths. At  $b = 0.01$ , the channel preserves  $N^* = 7$  partitions through 10 layers; at  $b = 0.1$ , it preserves  $N^* = 2$  partitions throughout. The noiseless Crawford–Sobel channel maps each input interval to a single output action deterministically: once quantized, re-quantization does not further degrade.

**Finding 2: With noise, mutual information degrades monotonically.** When  $\sigma^2 > 0$ , noise perturbs the output of each layer, causing the input to the next layer to fall in different partition intervals than it would noiseless. This breaks the partition-preservation property and causes monotonic MI decay consistent with the data processing inequality. At  $b = 0.1$  with  $\sigma^2 = 0.01$ ,

MI drops from 1.68 bits at Layer 1 to 1.58 bits at Layer 3 in a 3-layer hierarchy. With  $\sigma^2 = 0.05$ , the drop is more severe.

**Finding 3: Phase diagram reveals informative vs. blind regimes.** The phase diagram shows mutual information at the top layer as a function of per-layer bias and hierarchy depth. Two regimes are clearly visible:

- **Informative regime** ( $b \lesssim 0.10$ , shallow hierarchies): top-layer MI remains substantial ( $> 0.5$  bits).
- **Blind regime** ( $b \gtrsim 0.20$  or deep hierarchies with noise): top-layer MI approaches zero.

**Finding 4: Bias near the babbling boundary is catastrophic.** At  $b = 0.24$  (just below the babbling boundary), even the first-layer channel transmits only  $\approx 0.14$  bits. Through a 10-layer hierarchy, the top receives virtually no information about ground truth regardless of noise level. The organization is effectively blind.

### 3.5.4 Operationalizing Dysmemic Pressure

**Definition 3.6** (Dysmemic Pressure Index). The *Dysmemic Pressure Index* at layer  $k$  of an organizational hierarchy is:

$$DP(k) = 1 - \frac{I(\text{report}_k, \theta)}{I(\text{report}_0, \theta)}, \quad (5)$$

where  $\theta$  is the ground truth,  $\text{report}_k$  is the information available at layer  $k$ , and  $\text{report}_0$  is direct observation (the maximally informative baseline).

The index has clean boundary behavior:  $DP(k) = 0$  means no information loss (transparent hierarchy);  $DP(k) = 1$  means complete information destruction (blind hierarchy). By the data processing inequality (4),  $DP(k)$  is monotonically non-decreasing in  $k$ : dysmemic pressure can only accumulate, never reverse.

*Remark 3.7* (Practical measurement). Computing  $DP(k)$  requires access to ground truth  $\theta$ , which is often unavailable in real organizations. Proxy approaches include: (1) comparing internal reports against subsequently revealed outcomes; (2) comparing reports at different hierarchical levels against external benchmarks; (3) using prediction markets or proper scoring rules to elicit beliefs that can be compared across layers.

## 3.6 Boeing as Running Example

The Boeing 737 MAX case illustrates all three dynamics operating simultaneously. After the 1997 merger with McDonnell Douglas, Boeing’s corporate culture shifted from engineering-led to finance-led (Useem, 2019). Engineers’ preferences (design a safe aircraft, even at additional cost) diverged sharply from management’s preferences (certify quickly, minimize pilot training differences).

In Crawford–Sobel terms, the bias  $b$  between engineering assessment and management reception was large. Engineers who reported MCAS risks accurately faced schedule pressure and career friction. Management reframed engineering concerns as schedule issues: the partition coarsened from nuanced risk assessment to binary—certified or not certified.

Adverse selection operated through the Organization Designation Authorization program, which allowed Boeing employees to conduct certifications on the FAA’s behalf. Accurate risk signals and schedule-compliant signals were indistinguishable to the FAA at the moment of certification. The cheap signal (“MCAS is within acceptable parameters”) flooded out the expensive signal (“MCAS creates a single-point-of-failure requiring redesign”).

Prestige bias operated through Boeing’s leadership hierarchy: the CEO’s focus on stock price set the prestige frame, and assessments aligned with those priorities propagated downward. Conformity bias locked in the launch-oriented culture. Content bias favored “evolutionary upgrade to a proven design” over “novel flight-control system compensating for aerodynamic changes introduces new failure modes.”

Ground truth: the MCAS design was unsafe without adequate safeguards. The dysmemic pressure index was near 1: virtually complete information destruction between the engineering assessment and the board-level representation. Three hundred forty-six people died.

### 3.7 Connection to Forecasting

Tetlock (2005) demonstrated that expert political judgment is, on average, only slightly better than chance—with enormous variance. “Foxes” (who draw on multiple frameworks and update frequently) significantly outperformed “hedgehogs” (organized around a single grand theory). In the dysmemic pressure framework, hedgehog organizations exhibit stronger pressure because a single narrative becomes the selection criterion for internal signals.

Superforecasting teams (Tetlock and Gardner, 2015) implement structural features that map directly onto anti-dysmemic architecture: independent assessment before discussion (bypassing hierarchical compression), aggregation without social compression, proper scoring rules (reducing  $b$  toward zero), frequent updating (preventing conformity lock-in), and cognitive diversity (reducing prestige amplification of a single perspective). These are not exhortations. They are engineering specifications.

## 4 The Theory: When Noise Helps, When It Kills

The previous sections established a problem (formalization compresses) and a mechanism (compression selects for fit over truth). This section asks the deeper question: is compression *always* pathological? The answer turns out to be no—and the conditions under which compression artifacts become beneficial rather than destructive reveal the full structure of the theory.

The key insight is substrate-independence. The mechanism producing systematic self-deception in organizations and AI systems is the *same mechanism*, operating on the same formal objects (lossy

channels, conformity constraints, selection environments), and formalized by the same mathematical framework. Human psychology cannot explain the pattern in AI systems lacking human psychology. Machine learning training procedures cannot explain the pattern in organizations lacking gradient descent. The pattern appears in both because the cause is what they share: compression creates information gaps, and selection operating in those gaps drives systematic drift from accuracy toward fit.

#### 4.1 The Generative Lossy Channel

**Definition 4.1** (Generative Lossy Channel). A *generative lossy channel* is a communication system  $(S, C, R, Q)$  where:

- $S$  is a source producing states  $\theta$  from a distribution  $p_\theta$  over state space  $\Theta$ .
- $C$  is a lossy channel with rate  $R_C < H(\theta)$ , producing a compressed representation  $m$  of  $\theta$ .
- $R$  is a receiver that maps  $m$  to an output  $\hat{\theta}$ .
- $Q$  is an *acceptability distribution* over outputs: the receiver’s output must satisfy  $d(p_{\hat{\theta}}, Q) \leq P$  for some divergence  $d$  and tolerance  $P$ .

The channel is *generative* when  $P$  is finite and  $d(p_{\hat{\theta}^*}, Q) > 0$ , where  $\hat{\theta}^*$  is the distortion-minimizing reconstruction absent the acceptability constraint. Under these conditions, the receiver must produce outputs that diverge from the source beyond what compression alone requires.

**Definition 4.2** (The Generative Residual). For a generative lossy channel, the *generative residual* is:

$$\Delta_{\text{gen}} = \mathbb{E}[\Delta(\hat{\theta}, \theta) \mid d(p_{\hat{\theta}}, Q) \leq P] - \mathbb{E}[\Delta(\hat{\theta}^*, \theta)] \tag{6}$$

where  $\Delta$  is the distortion measure, the first term is the minimum achievable distortion under the acceptability constraint, and the second is the unconstrained minimum. When  $\Delta_{\text{gen}} > 0$ , the receiver is forced to generate.

The generative residual arises from the interaction of three constraints, each established independently:

1. **Endogenous lossy compression** (Crawford–Sobel). When sender and receiver preferences diverge, the channel quantizes reality into at most  $N^*$  discrete categories, where  $N^* = \lfloor -\frac{1}{2} + \sqrt{\frac{1}{4} + \frac{1}{2b}} \rfloor$  and  $b$  is the bias parameter. The channel is lossy for game-theoretic reasons, not bandwidth limitations.
2. **Forced generation under conformity** (Blau–Michaeli). The rate-distortion-perception tradeoff (Blau and Michaeli, 2019) proves that when the reconstruction must satisfy a distributional constraint ( $d(p_{\hat{X}}, Q) \leq P$ ), the minimum achievable rate is strictly elevated. The decoder must inject structure—generate—to close the gap.

3. **Selection on the residual.** The generative residual is a divergence from the source, not a directed error. Its direction is determined by the selection criterion operating on the receiver’s output: convergent selection (rewarding fit with existing frames) produces pathological drift; divergent selection (rewarding functional novelty) produces creative emergence.

#### 4.1.1 Scope and Limits of the Organizational Mapping

The Blau–Michaeli result is a theorem about signal compression under a distributional constraint on the reconstruction. Applying it to organizational communication requires identifying the corresponding constraint. We make the mapping explicit and bound its status.

**Where the mapping tightens.** The structural consequence of the tradeoff—that compression combined with a distributional output constraint forces divergence from the specific source—transfers directly. The organizational receiver cannot simultaneously (1) accept a heavily compressed input (coarse partition), (2) produce an output that matches the true state (low distortion), and (3) produce an output that satisfies institutional acceptability (distributional constraint). The three-way tradeoff’s *structure* applies: any two, but not all three.

**Where the mapping loosens.** The Blau–Michaeli theorem specifies the constraint as distributional match between output and source ( $P_{\hat{X}} \approx P_X$ ). Organizational acceptability is a constraint on  $P_{\hat{X}}$  alone—the output must be drawn from an acceptable distribution, but that distribution need not match the source distribution. This is a weaker constraint than Blau–Michaeli’s, which means the organizational case is *not* a direct application of the theorem but a *motivated structural analogy*: the tradeoff’s shape (compression + output constraint  $\rightarrow$  source divergence) is preserved; the specific divergence bound is not. The formal chain from Crawford–Sobel to Theorem 4.6 does not depend on Blau–Michaeli holding in its exact form—Theorem 4.6’s proof relies on the Kosko forbidden interval theorem (C2a path), Jensen’s inequality (C2b path), and first-order calculus (C2c path), none of which require the perception-distortion tradeoff. Blau–Michaeli provides the *interpretive* bridge explaining *why* the reconstruction diverges from the source, not the *formal* mechanism by which noise produces benefit.

## 4.2 Level-Crossing: Noise at $N$ Becomes Information at $N+1$

Henri Atlan’s “complexity from noise” principle (Atlan, 1979), building on von Foerster’s “order from noise” (von Foerster, 1960): in a hierarchically organized system, perturbation (noise) at organizational level  $N$  can become functional information at level  $N+1$ , provided the higher level has sufficient integrative capacity to exploit the variation.

**Definition 4.3** (Level-Crossing System). A system exhibits *generative level-crossing* when:

1. There exist at least two organizational levels  $L_N$  and  $L_{N+1}$  with distinct state spaces.
2. Level  $L_N$  produces variation  $V_N$  (noise relative to  $L_N$ ’s functional requirements).

3. Level  $L_{N+1}$  has an integrative mapping  $\phi : V_N \rightarrow S_{N+1}$  that transforms  $L_N$ -noise into  $L_{N+1}$ -signal.
4. The mapping  $\phi$  satisfies:  $I(V_N; S_{N+1}) > 0$ —the variation carries nonzero mutual information with the higher-level signal.

Four instances illustrate the principle’s reach:

- A genetic mutation is noise within an organism’s developmental program but information within the population’s adaptive capacity.
- A misheard chord is noise within a musical performance but information within the evolution of harmonic language.
- A mistranslated text is noise within the author’s intended meaning but information within the intellectual tradition that grows from the mistranslation.
- A Crawford–Sobel residual is noise within the sender’s intended communication but information within the organization’s capacity for novel synthesis.

The open question—under what conditions on the system’s nonlinearity profile does  $\phi$  exist and produce net benefit?—connects to the Kosko forbidden interval theorem (Section 4.3) and is resolved by Theorem 4.6.

### 4.3 When Does Noise Help? The Forbidden Interval Theorem

In nonlinear systems with a threshold, the addition of an optimal level of noise *enhances* signal detection. The signal-to-noise ratio follows an inverted-U curve, peaking at nonzero noise. This is the stochastic resonance phenomenon, first described by Benzi et al. (1981) in the context of paleoclimate oscillations and surveyed comprehensively by Gammaitoni et al. (1998).

Kosko and Mitaim (2003) prove necessary and sufficient conditions for stochastic resonance in threshold detectors:

**Theorem 4.4** (Forbidden Interval—Kosko et al.). *A threshold signal detector exhibits a noise benefit (stochastic resonance) if and only if the noise mean  $\mu$  does not lie in the “forbidden interval”  $(\theta - s_1, \theta - s_0)$ , where  $\theta$  is the detection threshold and  $s_0, s_1$  are the subthreshold signal levels.*

Three properties make this result foundational for the present framework. First, the condition is both necessary *and* sufficient—not merely a sufficient condition for benefit but a complete characterization. Second, the result holds for all noise probability density functions with finite variance. Third, it extends to the entire uncountably infinite class of  $\alpha$ -stable distributions (heavy-tailed noise), meaning the theorem applies far beyond Gaussian assumptions.

The forbidden interval theorem provides the formal condition that Atlan’s level-crossing principle lacks. Atlan-type level crossing—where noise at level  $N$  becomes information at level  $N+1$ —occurs

when: (1) the higher-level integration function  $\phi$  has a nonlinear threshold, (2) the variation  $V_N$  has a mean that falls outside the forbidden interval of  $\phi$ 's threshold structure, and (3) the variation magnitude is moderate. This yields the inverted-U prediction: there exists an optimal noise level  $\sigma^*$  that maximizes the mutual information  $I(V_N; S_{N+1})$ , and this optimum is nonzero. These conditions ground the C2a path of Theorem 4.6 below.

#### 4.4 Theorem 1: Five Sufficient Conditions for Net-Beneficial Noise

**Definition 4.5** (Two-Level System with Noise). A *two-level system* consists of:

- Level- $N$  performance  $P_N(\sigma)$ , degrading with noise amplitude  $\sigma$ .
- Level- $(N+1)$  performance  $P_{N+1}(\sigma)$ , representing emergent or higher-order outcomes.
- System performance  $P_{\text{sys}}(\sigma) = f(P_N(\sigma), P_{N+1}(\sigma))$  for some aggregation function  $f$ .

The system exhibits *net-beneficial noise* if there exists  $\sigma^* > 0$  such that  $P_{\text{sys}}(\sigma^*) > P_{\text{sys}}(0)$ .

**Theorem 4.6** (Five Sufficient Conditions for Net-Beneficial Noise). *Let  $(P_N, P_{N+1}, P_{\text{sys}})$  be a two-level system. The following five conditions are jointly sufficient for the existence of  $\sigma^* > 0$  with  $P_{\text{sys}}(\sigma^*) > P_{\text{sys}}(0)$ :*

**C1 (Suboptimality).** *The noiseless system is not at the global optimum of  $P_{N+1}$ : there exists a state achieving higher  $P_{N+1}$  that is inaccessible at  $\sigma = 0$ . Formally: the system operates below its level- $(N+1)$  capacity.*

**C2 (Nonlinear integration).** *The function mapping noise to level- $(N+1)$  benefit is nonlinear, through one of three mechanisms:*

C2a: **Threshold stochastic resonance.** *A subthreshold signal is boosted by noise through a nonlinear detector.*

C2b: **Jensen-gap rectification.** *A convex function rectifies symmetric noise into a net positive:  $\mathbb{E}[g(x + \xi)] > g(x)$  when  $g'' > 0$  in the operating region.*

C2c: **Opposing monotones.** *The product of an increasing function (benefit given success) and a decreasing function (probability of success) has an interior maximum.*

**C3 (Accessibility).** *The noise distribution has support in the improvement region: the perturbation can reach states with higher  $P_{N+1}$ .*

**C4 (Asymmetric cost–benefit).** *The level- $(N+1)$  gain exceeds the level- $N$  cost over some interval:*

$$\exists \sigma^* > 0 : \quad \beta \int_0^{\sigma^*} \Delta_{N+1}(\sigma) d\sigma > \alpha \int_0^{\sigma^*} |\Delta_N(\sigma)| d\sigma \quad (7)$$

where  $\Delta_{N+1}(\sigma) = P_{N+1}(\sigma) - P_{N+1}(0)$ ,  $\Delta_N(\sigma) = P_N(\sigma) - P_N(0)$ , and  $\alpha, \beta > 0$  are system-dependent weights.

**C5 (Graceful degradation).** *Level- $N$  performance degrades gradually: there is no catastrophic collapse at small  $\sigma$ . The system has sufficient redundancy or slack to absorb moderate noise without structural failure.*

*Under these conditions,  $P_{\text{sys}}(\sigma)$  exhibits an inverted-U: there exists  $0 < \sigma^* < \bar{\sigma} < \infty$  such that  $P_{\text{sys}}(\sigma^*) > P_{\text{sys}}(0)$  and  $P_{\text{sys}}(\sigma) < P_{\text{sys}}(0)$  for  $\sigma > \bar{\sigma}$ . The optimal noise level  $\sigma^*$  is bounded; there is always too much noise.*

The proof proceeds through three independent paths, one for each C2 sub-mechanism.

**Proof sketch for C2a (Threshold SR).** Under C1, the system has a subthreshold signal: a state with higher  $P_{N+1}$  that the deterministic system cannot reach. Under C2a, a nonlinear threshold creates a detection probability  $P_{N+1}(\sigma) = \Phi\left(\frac{s-\theta}{\sigma}\right)$  that increases from zero as  $\sigma$  increases. Under C3, the noise distribution reaches the threshold. Under C4, the integral gain exceeds the integral cost. Under C5, the cost function has no discontinuity. The existence of the interior maximum follows from continuity:  $P_{N+1}(0) = 0$  (subthreshold),  $P_{N+1}(\sigma)$  increases for moderate  $\sigma$ , and  $P_N(\sigma)$  degrades continuously. The aggregation  $P_{\text{sys}}(\sigma) = f(P_N, P_{N+1})$  inherits the inverted-U from C4's integral condition.  $\square$

**Proof sketch for C2b (Jensen gap).** Under C2b,  $g$  is convex in the operating region, so  $\mathbb{E}[g(x + \xi)] - g(x) = J(\sigma) \geq 0$  by Jensen's inequality. The Jensen gap  $J(\sigma)$  is the level- $(N+1)$  gain. Under C4,  $\beta J(\sigma)$  exceeds  $\alpha |\Delta_N(\sigma)|$  for moderate  $\sigma$ . Under C5, the cost  $|\Delta_N|$  grows continuously from zero. The inverted-U follows from  $J(\sigma) \sim \frac{1}{2}g''(x)\sigma^2$  for small  $\sigma$  (quadratic onset) versus linear or faster cost growth for large  $\sigma$ .  $\square$

**Proof sketch for C2c (Opposing monotones).** Under C2c, benefit given success  $g(\sigma)$  is increasing and probability of success  $h(\sigma)$  is decreasing. The product  $g(\sigma) \cdot h(\sigma)$  has an interior maximum by first-order calculus:  $\frac{d}{d\sigma}[g \cdot h] = g'h + gh'$ , which changes sign because  $g' > 0$  and  $h' < 0$ . The maximum  $\sigma^*$  satisfies  $g'(\sigma^*)/g(\sigma^*) = -h'(\sigma^*)/h(\sigma^*)$ . Under C4 and C5, the net system performance inherits this interior maximum.  $\square$

**Lemma 4.7** (Integral vs. Marginal C4). *For threshold systems satisfying C2a, the marginal derivative condition*

$$\beta \cdot \left. \frac{\partial P_{N+1}}{\partial \sigma} \right|_{\sigma \rightarrow 0^+} > \alpha \cdot \left| \left. \frac{\partial P_N}{\partial \sigma} \right|_{\sigma \rightarrow 0^+} \right| \quad (8)$$

*is neither necessary nor sufficient for net benefit. The correct condition is the integral formulation in C4.*

*Proof.* For a threshold detector with signal  $s < \theta$  and Gaussian noise,  $P_{N+1}(\sigma) = \Phi\left(\frac{s-\theta}{\sigma}\right) \rightarrow 0$  exponentially fast as  $\sigma \rightarrow 0^+$ , so the marginal benefit is exponentially small. Meanwhile, level- $N$  cost begins immediately:  $|\partial P_N / \partial \sigma|_{\sigma \rightarrow 0^+} = L_N > 0$ . The marginal condition is never satisfied, yet the integral condition can be met at finite  $\sigma$  where detection probability becomes substantial.  $\square$

*Remark 4.8* (The diagnostic role of C4). C4 is the hardest of the five conditions to verify a priori, because checking it requires either running the noise sweep or estimating the integral from the system’s response profile. The theorem’s practical value lies in the other four conditions: C1, C2, C3, and C5 are *easy to check* from the system’s structure without running any experiment. They serve as a screening tool. If any of C1–C3 or C5 fails, the system cannot benefit from noise regardless of the weighting, and C4 need not be checked. If all four pass, then the question reduces to C4: does the gain–cost ratio favor the higher level at some finite  $\sigma$ ? The practical workflow is: screen with C1–C3, C5 (cheap)  $\rightarrow$  if all pass, invest in checking C4 (expensive)  $\rightarrow$  if C4 holds, benefit is guaranteed.

## 4.5 Corollaries

**Corollary 4.9** (The Brittleness Trap). *A system that minimizes redundancy at level  $N$  violates C5 and cannot exhibit generative level-crossing, regardless of how well C1–C4 are satisfied. Efficiency kills generativity when it eliminates the buffer that absorbs noise.*

**Corollary 4.10** (The Alignment Trap). *A system with perfectly aligned incentives ( $b = 0$ ) has a nearly lossless channel, producing minimal variation. Perfect alignment can satisfy level- $(N+1)$  needs directly, eliminating the need for level-crossing—but also eliminating the capacity for it when the environment shifts.*

**Corollary 4.11** (The Over-Noise Catastrophe). *For any system satisfying C1–C5, there exists  $\bar{\sigma}$  beyond which  $P_{sys}(\sigma) < P_{sys}(0)$ . The inverted- $U$  is bounded. There is always too much noise.*

## 4.6 The Tolerance Location Principle

Synthesizing Crawford–Sobel, Blau–Michaeli, the Atlan level-crossing principle (Definition 4.3), the forbidden interval theorem (Theorem 4.4), and the five sufficient conditions of Theorem 4.6:

**Theorem 4.12** (Tolerance Location Principle). *A system’s resilience and generative capacity are maximized when:*

1. ***Distinguishable core and interface.*** *The system has a structurally identifiable boundary between its internal processing (core) and its interaction with the environment (interface). The core maintains rigidity (invariant structures, stable representations). The interface maintains tolerance (absorptive capacity for environmental variance).*
2. ***Sufficient absorptive capacity at the interface.*** *The interface must satisfy Ashby’s Law of Requisite Variety (Ashby, 1956): its response repertoire must be at least as rich as the variance it faces. Formally:  $H(\text{interface responses}) \geq H(\text{environmental perturbations relevant to the system})$ .*
3. ***Task-relevant information preservation across the boundary.*** *The transformation from environmental noise to internal signal must satisfy Tishby’s Information Bottleneck*

criterion (Tishby et al., 1999): minimize  $I(X;T)$  (compression) subject to maintaining  $I(T;Y)$  (relevance).

**Corollary 4.13** (Inverted Architecture Failure). *When tolerance is placed at the core and rigidity at the interface—flexible internals, rigid boundaries—the system loses both resilience and generative capacity.*

**Corollary 4.14** (The Sterility of Perfection). *A system with zero tolerance at the interface ( $H(\text{interface responses}) = 0$ ) has zero generative capacity. Perfect, lossless transmission eliminates the generative residual, producing a system that is deterministic and sterile.*

The third natural corollary—the inseparability of dysfunction and creativity—is stated as Corollary 4.16 in Section 4.7 below, where the dual-valence framework provides the context for it.

## 4.7 Dual Valence

**Proposition 4.15** (Dual Valence of the Generative Residual). *The generative residual  $\Delta_{\text{gen}}$  (Definition 4.2) has dual valence:*

- Under convergent selection (selection rewards fit with the existing frame),  $\Delta_{\text{gen}}$  produces systematic drift toward internally-fit outputs. This is the mechanism underlying compliance theater, impression management, and sycophancy.
- Under divergent selection (selection rewards functional novelty or accuracy to external reality),  $\Delta_{\text{gen}}$  produces reconstruction that occasionally outperforms the source. This is the mechanism underlying creative emergence.

The dual valence follows from the fact that  $\Delta_{\text{gen}}$  is a divergence from the source, not a directed error. The direction is determined by the selection criterion, not by the compression mechanism.

**Corollary 4.16** (The Inseparability of Dysfunction and Creativity). *Since both dysfunctional drift and creative novelty arise from the same generative residual under the same conformity constraint, no system can eliminate dysfunction potential without simultaneously eliminating creative potential. The design problem is not elimination but channeling.*

## 4.8 The Compression Ratchet

**Proposition 4.17** (The Compression Ratchet). *Let  $\mathcal{S}$  be a generative lossy channel (Definition 4.1) satisfying the following conditions:*

- (R1) **Endogenous acceptability.** *The acceptability distribution  $Q_t$  at time  $t$  is a function of prior outputs:  $Q_{t+1} = \Gamma(Q_t, \{\hat{\theta}_\tau\}_{\tau \leq t})$ , where  $\Gamma$  is a distribution update rule that moves  $Q$  toward the empirical distribution of recent outputs.*
- (R2) **Positive rate with preference divergence.** *The channel has bias  $0 < b < 1/4$ , so the channel is endogenously lossy with a positive generative residual.*

(R3) **Convergent selection.** *The selection criterion operating on outputs rewards conformity with  $Q_t$ : outputs closer to  $Q_t$  receive higher fitness.*

(R4) **No external correction.** *The system receives no input from sources outside the compression-selection loop: no external audit, no independent verification, no exogenous signal forcing  $Q_t$  toward the source distribution  $p_\theta$ .*

Then the system has a self-reinforcing equilibrium: a fixed point  $Q^*$  of the dynamics  $Q_{t+1} = \Gamma(Q_t, \cdot)$  such that:

- (i)  $d(Q^*, p_\theta) > d(Q_0, p_\theta)$ : *the equilibrium acceptability distribution is farther from reality than the initial one.*
- (ii) *The basin of attraction is expanding: once  $d(Q_t, p_\theta) > \delta$  for a threshold  $\delta$  determined by the system parameters, the dynamics are monotonically drift-increasing.*
- (iii) *Recovery requires violation of (R4): external intervention that introduces signals not subject to the internal selection environment.*

*Proof.* Under (R2), the channel is endogenously lossy with generative residual  $\Delta_{\text{gen}} > 0$ . The receiver's outputs therefore deviate from the source:  $d(p_{\hat{\theta}}, p_\theta) > 0$ .

Under (R1), these deviant outputs update the acceptability distribution:  $Q_{t+1}$  moves toward  $p_{\hat{\theta}}$ . Since  $p_{\hat{\theta}}$  deviates from  $p_\theta$ , the updated  $Q_{t+1}$  also deviates from  $p_\theta$ .

Under (R3), the next period's outputs are selected for conformity with  $Q_{t+1}$ , which is now farther from  $p_\theta$  than  $Q_t$  was. The generative residual at  $t + 1$  produces outputs deviating from  $p_\theta$  in the same direction, because the conformity constraint now pulls toward a  $Q$  that has already drifted.

The monotonicity follows:  $d(Q_{t+1}, p_\theta) \geq d(Q_t, p_\theta)$  whenever the generative residual reinforces the drift direction, which is guaranteed under convergent selection (R3). Each dysmeme that establishes itself tilts the landscape further. The next dysmeme becomes easier to establish. The next accurate signal becomes harder to transmit.

The equilibrium  $Q^*$  is the fixed point where the drift rate equals zero: the acceptability distribution has stabilized at a point where the generative residual exactly reproduces the current  $Q^*$ . The constructed reality generates signals that reinforce the selection criteria that generated the constructed reality.

Under (R4), no mechanism exists to correct the drift. Reform proposals are themselves signals subject to the internal selection environment. A reform that would displace the equilibrium threatens agents whose positions depend on the current arrangement and is selected against. A reform that changes vocabulary while preserving fitness criteria is selected for, producing surface change without equilibrium displacement. Recovery therefore requires violating (R4): introducing observation, verification, or authority from outside the compression-selection loop.  $\square$

*Remark 4.18* (Connection to high-reliability organizations). Organizations that resist the ratchet—aircraft carriers, nuclear plants, air traffic control—do so by institutionalizing violations of condition (R4): redundant verification from independent channels, licensed dissent roles that create protected paths for signals the internal selection environment would otherwise filter, and deference to expertise that migrates decision authority to whoever has the least compressed information about the current state (Weick and Sutcliffe, 2007).

## 4.9 Computational Validation

### 4.9.1 Strategy

Theorem 4.6 asserts five sufficient conditions for net generative level-crossing. This section subjects that claim to adversarial computational testing across six mechanistically distinct domains: threshold detection, sigmoid detection, polynomial detection (null model), simulated annealing, ensemble diversity, and multi-armed bandit exploration.

### 4.9.2 Stochastic Resonance Class

A two-level system with level- $N$  degradation under three models (linear, exponential, catastrophic) and level- $(N+1)$  detection under four detector types. Each configuration swept across 100 noise levels with 10,000 trials per level and 200 bootstrap resamples.

Configuration	Net Benefit	Optimal $\sigma^*$	Inverted-U	Conditions Met
Threshold, all met	<b>0.170</b>	0.45	Yes	All 5
Sigmoid, all met	<b>0.091</b>	0.33	Yes	All 5
Polynomial, all met	0.000	—	No	C2 marginal
Linear (null)	0.000	—	No	C2 violated
C1 violated	0.000	—	No	C1 violated
C5 violated	0.000	—	No	C5 violated

### 4.9.3 Non-SR Mechanisms

**Simulated annealing** (optimization landscape escape). Net benefit 0.254 when conditions met; zero when C1 or C4 violated.

**Ensemble diversity** (error decorrelation). Twenty-one predictors on a hard XOR problem. Net benefit 0.068 when conditions met.

**Multi-armed bandit** (exploration–exploitation). Ten arms with unknown rewards. Net benefit 0.345 when conditions met; zero when C1 violated.

### 4.9.4 Monte Carlo Sensitivity Analysis

500 random parameter configurations. Conditions checked dynamically.

Category	Count	Percentage
All conditions met AND benefit observed	75	15.0%
All conditions met AND no benefit	0	0.0%
Some condition violated AND benefit observed	38	7.6%
Some condition violated AND no benefit	387	77.4%
<b>Total</b>	<b>500</b>	

**Sufficiency rate: 100%.** Zero counterexamples. The conditions are sufficient but not necessary (7.6% of violations still showed benefit, indicating the conditions define a conservative boundary).

## 4.10 Cross-Domain Evidence

A theory of substrate-independent generative lossy channels predicts that the same formal structure—compression forces reconstruction, reconstruction diverges from source, selection determines valence—should appear in substrates that share no surface features. Three domains are mapped onto the C1–C5 framework using published external evidence. These instances verify qualitative satisfaction of conditions C1–C5; full formal verification within each domain is future work.

### 4.10.1 Organizational Disasters: Normalized Deviance as the Ratchet

Vaughan’s (1996) ethnographic study of the Challenger disaster provides the most thoroughly documented case of organizational self-deception through compression-selection dynamics. Her *normalized deviance* maps directly onto the ratchet mechanism (Proposition 4.17):

Cond.	Mapping to Challenger normalized deviance
C1	<b>Met.</b> The system was suboptimal: O-ring erosion was a known failure mode, and the engineering data showed temperature dependence.
C2	<b>Met (inverted).</b> The acceptability threshold created a step-function response: erosion below threshold produced no organizational response. Under convergent selection, this nonlinearity amplified drift rather than recovery.
C3	<b>Met (inverted).</b> Most flights showed acceptable erosion, reinforcing the compressed signal that erosion was manageable.
C4	<b>Met for dysfunction.</b> Schedule pressure ( $\alpha$ ) dominated engineering caution ( $\beta$ ). Each successful flight with anomalies increased the weight on “it worked before.”
C5	<b>Met.</b> Organizational capacity for processing bad news degraded gradually through budget cuts and schedule pressure.

All ratchet conditions (R1–R4) of Proposition 4.17 are satisfied in Vaughan’s account. The Columbia Accident Investigation Board found seventeen years later that the same failures had reasserted themselves—the ratchet had resumed after post-Challenger reforms decayed.

#### 4.10.2 Jazz Improvisation: Creative Emergence Through Controlled Channel Noise

Berliner’s (1994) ethnographic study documents ensemble improvisation as a generative lossy channel under divergent selection. Each musician transmits intent through an inherently lossy channel; the residual is filled by the listening musician’s priors. All five conditions are met, but under divergent selection (aesthetic novelty rewarded), the same mechanism that produces normalized deviance produces creative emergence.

Duke Ellington’s creative method inverted the standard big-band model. Where other bandleaders wrote parts for instrument sections, Ellington wrote for individual musicians’ tonal signatures. He described hearing each player’s unique sound and composing to exploit it—not the trumpet part, but the “Cootie Williams part,” crafted to the specific way Williams bent notes and shaped phrases (Ellington, 1973; Tucker, 1995). Ellington calculated how each musician would tend to deviate from the written page, designing passages where the likely deviation would produce the desired effect. The written score was intentionally lossy: it omitted details the musician’s individual style would fill in. The gap between composition and performance was the generative material. When ASCAP licensing disputes in 1941 forced the band to stop performing Ellington’s existing compositions on radio, Billy Strayhorn composed “Take the ‘A’ Train” in a single night. The constraint forced reconstruction from a different musician’s interpretive frame, producing what became the band’s signature piece. The lossy channel between compositional intent and each musician’s execution was the source of the orchestra’s distinctive sound—the same channel structure as organizational communication, under divergent selection (aesthetic novelty) rather than convergent selection (procedural compliance). This is Atlan level-crossing (Definition 4.3) by design: the “error” at level  $N$  (deviation from the written score) became information at level  $N+1$  (the tonal signature Ellington composed *for*), with C5 satisfied by each musician’s idiosyncratic mastery—the redundancy that allowed the gap to become signal rather than noise.

The dual-valence claim (Proposition 4.15) is directly testable here: same channel structure, same compression, same forced reconstruction. Different selection criterion. Different valence.

#### 4.10.3 AI Alignment: Sycophancy as Equilibrium Compression Artifact

RLHF training creates a strategic channel between the model and human evaluators. The model’s output must conform to the distribution of “helpful-sounding” responses. Perez et al. (2022) developed evaluations for sycophantic behavior. Sharma et al. (2023) directly measured that models shift stated views to match user opinions—precisely the pattern predicted by compression-selection under convergent selection.

The substrate-independence claim is strongest here: the same mechanism produces sycophancy in silicon with no human psychology involved.

#### 4.10.4 Cross-Domain Summary

Domain	Substrate	Channel	Selection	Valence
Challenger/Columbia	Org. hierarchy	Pref. divergence	Schedule press.	Conv. → deviance
Jazz improvisation	Musical interact.	Acoustic/cognitive	Aesthetic novelty	Div. → emergence
RLHF sycophancy	AI training	Reward compression	Reward signal	Conv. → sycophancy

No two domains share a substrate. The formal structure is identical across all three.

## 5 The Formal Foundation: Strategic Rate-Distortion-Perception

When a middle manager rewrites an engineer’s safety assessment before passing it to the executive team, three constraints operate simultaneously. The rewrite compresses: it reduces a detailed technical report to a summary, incurring distortion. The manager’s incentives diverge from the engineer’s: the compression is strategic, not bandwidth-limited. And the summary must *conform*: it must look like what executive summaries look like—professional, measured, within the range of acceptable organizational communication. These three constraints—compression, strategic misalignment, and distributional conformity—interact in ways that the previous sections described qualitatively. This section provides the formal information-theoretic foundation.

The generative lossy channel framework of Section 4 established that compression under a conformity constraint forces generation. The formal foundation developed here—the *strategic rate-distortion-perception function*—makes this precise: it defines the minimum information rate required to achieve a given distortion when encoder and decoder have misaligned objectives and the decoder’s output must conform to a target distribution. The key quantity is the *generative residual*: the excess distortion attributable to conformity pressure beyond what strategic misalignment alone requires.

### 5.1 The Strategic RDP Function

Three foundational results, each established independently, converge here. Shannon’s rate-distortion theory (Cover and Thomas, 2006) characterizes the minimum rate for a given distortion. Crawford and Sobel (1982) characterize strategic communication equilibria with endogenous information loss. Blau and Michaeli (2018; 2019) prove that lossy compression under a perceptual quality constraint—where the reconstruction must resemble a target distribution—forces the decoder to generate, strictly elevating the required rate.

No existing result combines strategic misalignment with a distributional conformity constraint. The *strategic rate-distortion-perception function* fills this gap.

**Definition 5.1** (Strategic Rate-Distortion-Perception Function). For source  $X \sim p_X$ , reconstruction  $\hat{X}$ , distortion measure  $\Delta$ , divergence  $d$ , and arbitrary target distribution  $Q$ :

$$R(D, P, Q) = \min_{p(\hat{x}|x)} I(X; \hat{X}) \quad \text{subject to:} \quad \mathbb{E}[\Delta(X, \hat{X})] \leq D, \quad d(p_{\hat{X}}, Q) \leq P \quad (9)$$

When  $Q = p_X$ , this reduces to the standard Blau–Michaeli function  $R(D, P)$ .

The generalization from  $Q = p_X$  to arbitrary  $Q$  is essential for organizational applications: the distribution of “acceptable” outputs ( $Q$ ) is not the distribution of reality ( $p_X$ ). Executive summaries must look like executive summaries, not like the messy ground truth they summarize.

## 5.2 Theorem A: Generalized RDP Tradeoff

**Theorem 5.2** (Generalized RDP Tradeoff). *For source  $X \sim p_X$ , distortion measure  $\Delta$  (not constant), divergence  $d$  convex in its first argument, and target distribution  $Q$  with non-empty feasibility ( $D \geq D_{\min}(P, Q)$ ):*

- (i)  $R(D, P, Q)$  is non-increasing in  $D$  and  $P$ .
- (ii)  $R(D, P, Q)$  is convex in  $D$ .
- (iii)  $R(D, P, Q) \geq R(D)$  for all finite  $P$ , with equality only when the unconstrained optimizer already satisfies  $d(p_{\hat{X}^*}, Q) \leq P$ .

At  $P = 0$  (strict distributional match), the minimum achievable distortion equals the optimal transport cost between  $p_X$  and  $Q$  under  $\Delta$ .

Property (iii) is the key result for organizational theory: whenever the conformity constraint binds—whenever “looking right” differs from “being right”—the required communication rate is strictly elevated. The organization must transmit more to achieve the same accuracy, or accept more distortion at the same rate. In practice, organizations accept the distortion.

## 5.3 Theorem B: The Gaussian Closed Form and the Generative Residual

For a Gaussian source  $\theta \sim \mathcal{N}(0, \sigma_\theta^2)$  with quadratic distortion and KL-divergence perception constraint, the strategic RDP equilibrium admits a closed form.

**Definition 5.3** (The Generative Residual). For a strategic RDP game with perception-constrained equilibrium distortion  $D_R(P)$  and unconstrained equilibrium distortion  $D_R(\infty)$ :

$$\Delta_{\text{gen}} = D_R(P) - D_R(\infty) \tag{10}$$

When  $\Delta_{\text{gen}} > 0$ , the receiver produces outputs that diverge from the source beyond what strategic misalignment alone requires. The excess is the cost of conformity.

**Theorem 5.4** (Gaussian Strategic RDP Equilibrium). *For Gaussian source, quadratic utilities with bias  $b > 0$ , and perception constraint  $D_{\text{KL}}(p_\alpha \| Q) \leq P$  where  $Q = \mathcal{N}(\mu_Q, \sigma_Q^2)$ , the receiver’s MSE decomposes as:*

$$D_R = \underbrace{(1 - a)^2 \sigma_\theta^2 + a^2 D_R^0}_{\text{strategic misalignment}} + \underbrace{c^2 + \sigma_Z^2}_{\text{conformity cost}} \tag{11}$$

where  $a$  is a scaling parameter,  $c$  is a mean shift toward  $\mu_Q$ ,  $\sigma_Z^2$  is injected noise, and  $D_R^0$  is the Crawford–Sobel baseline distortion. Whenever the perception constraint binds:  $\Delta_{\text{gen}} = D_R(P) - D_R^0 > 0$ .

Three structural properties emerge. First, a *mean-variance tradeoff*: the receiver allocates distortion budget between shifting toward  $Q$ ’s mean and matching  $Q$ ’s variance—a tradeoff absent from standard RDP where  $Q = p_X$ . Second, *variance injection*: when the target is more variable than the unconstrained reconstruction, the receiver must add noise. Third, *variance suppression*: when the target is narrow, the receiver shrinks toward the prior mean—conformity to a narrow norm forces conservatism. This last property formalizes the variance compression documented empirically in Section 2.6.

#### 5.4 Numerical Illustration: Three Regimes

Three worked examples span qualitatively different regimes, all with source variance  $\sigma_\theta^2 = 1$ .

Regime	$b$	$Q$	$P$	$N^*$	$\Delta_{\text{gen}}$	Conformity %
Moderate bias, shifted target	0.10	$\mathcal{N}(0.2, 1.69)$	0.50	2	0.274	22.2%
Low bias, tight constraint	0.05	$\mathcal{N}(1.0, 0.25)$	0.01	4	0.938	71.9%
High bias (babbling), wide target	0.20	$\mathcal{N}(0, 4.0)$	0.50	1	0.634	39.5%

When the target is close to the source and the constraint is moderate, about one-fifth of the receiver’s total error comes from conformity. When the target is far from the source and the constraint is tight, conformity dominates—nearly three-quarters of all distortion is attributable to the requirement that outputs “look right.” Even in the babbling regime, where the communication channel transmits zero information, conformity pressure creates substantial additional distortion: the receiver must inflate the variance of outputs to match the target, adding noise beyond what the uninformative channel already imposes. Conformity pressure accounts for 22–72% of total distortion across these regimes, with a sharp phase transition at the Crawford–Sobel babbling boundary ( $b = 1/4$ ).

#### 5.5 Organizational Sufficient Conditions

Four checkable conditions are jointly sufficient for the strategic RDP tradeoff to hold in any organizational communication channel: (O1) the channel carries information ( $b < 1/4$ ); (O2) sender and receiver have non-identical objectives ( $b > 0$ ); (O3) there exists an enforceable distribution of acceptable outputs; (O4) the acceptable distribution differs from the unconstrained reconstruction. When all four hold, the generative residual is strictly positive, and the dual valence of Proposition 4.15 applies: convergent selection produces systematic drift toward internally-fit outputs; divergent selection can produce creative emergence.

## 5.6 A Motivating Conjecture: RLHF Sycophancy

The strategic RDP framework suggests a structural interpretation of sycophancy in AI systems trained via reinforcement learning from human feedback (RLHF). The mapping is: the language model is the sender observing the user’s actual need; the reward model defines the target distribution  $Q$  of “good” outputs; the KL penalty coefficient  $\beta$  constrains how far the fine-tuned distribution can deviate from the reference, corresponding to the perception constraint  $P$ . Sycophancy—the tendency to produce agreeable rather than truthful responses—would then correspond to the generative residual  $\Delta_{\text{gen}}$ : excess distortion from the model conforming to reward-maximizing outputs rather than truthful ones.

This mapping is a *conjecture*, not an established application. Three caveats bound its status. First, the Gaussian closed form assumes continuous distributions; RLHF operates on discrete token sequences, and the gap between continuous and discrete settings is non-trivial. Second, the Nash equilibrium framework assumes simultaneous strategic interaction; RLHF training is a sequential optimization process. Third, the Crawford–Sobel partition structure assumes the sender partitions a continuous state space; language model generation does not partition in this sense. The structural predictions—monotonicity (more RLHF training increases sycophancy), shape dependence (biased reward models produce more sycophancy), and irreducibility ( $\Delta_{\text{gen}} > 0$  for any finite  $P$ )—are qualitatively consistent with empirical findings (Sharma et al., 2023), but quantitative correspondence requires empirical calibration that has not been performed.

## 6 Measurement: Detecting Organizational Dysfunction

The previous sections established that organizations systematically compress information and that the compression has formal structure. A natural question follows: can the dysfunction be detected before it produces catastrophe? This section describes a measurement architecture—a stigmergic mesh—that detects organizational dysfunction from work artifacts without requiring anyone to know what to look for.

The core difficulty is structural, not informational. Philip Armour’s taxonomy distinguishes 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—all require a query. The query assumption is precisely what second-order ignorance violates.

### 6.1 Architecture: A Self-Organizing Mesh

The Ambient Structure Discovery architecture (McEntire, 2026a) operates on three principles. First, *continuous ingestion* of sematectonic traces—work artifacts (code commits, deployment logs,

issue tracker activity) that carry structural information regardless of whether anyone intended to communicate it. These are costly signals in the economic sense: producing them requires the work to actually have been done. Second, *self-organizing topological specialization*: signals route through a competitive mesh of processing nodes grounded in Adaptive Resonance Theory (Carpenter and Grossberg, 1987), which provides the only self-organizing neural architecture with formal stability-plasticity guarantees. Nodes develop affinity for signals they have successfully processed, producing an emergent topology that is itself a structural model of the organization’s actual work patterns. Third, *geometric pattern detection*: two signals that activate similar patterns across mesh nodes are structurally related regardless of whether they share vocabulary. Detection operates on topology, not semantics, which is why it can surface relationships that keyword search, semantic search, and human inspection would miss.

A multiplicative scoring function prioritizes findings by structural importance:

$$S(f) = d_{\text{bridge}}(f) \cdot c_{\text{entity}}(f) \cdot r_{\text{coupling}}(f) \cdot \gamma(f) \tag{12}$$

where bridge distance measures how far outside the organization’s known subspace the finding lives, entity confidence measures identity resolution quality, risk coupling measures proximity to control surfaces, and communicability  $\gamma$  uses the Crawford–Sobel partition count to gate whether the finding can propagate through the communication channel to the human recipient. All four components are topological—none measure satisfaction, sentiment, or self-report.

## 6.2 Three Detection Signatures for Normalized Deviance

Normalized deviance—the gradual redefinition of unacceptable risk as routine through repeated exposure without consequence—caused the *Challenger* and *Columbia* disasters. It is the organizational pathology most relevant to the dysmemic pressure framework of Section 3. Three measurable signatures formalize it as a detection target.

**Signature 1: Eigenvalue collapse.** Under healthy conditions, different teams maintain distinct perspectives on a shared risk. The correlation matrix across mesh workers processing signals about that risk has full rank. As normalized deviance progresses, framing converges: teams that previously discussed the risk with distinct vocabulary begin using the same phrases. The matrix loses rank. Eigenvalues merge. The effective dimensionality of the organization’s perception decreases even as the risk itself persists.

**Signature 2: Linguistic compression.** The transition from active risk management to normalized deviance has a measurable linguistic signature that *precedes* operational failure. Hedging density increases (“will fail” becomes “may experience degradation”). Passive voice ratio rises. Nominalization frequency increases (“the system malfunctions” becomes “malfunction events have been observed”). Commitment specificity decreases. These features compose a compression index

that tracks the same phenomenon documented in the SEC filing analysis of Section 2.6—but in real time, before the annual report is filed.

**Signature 3: Action decoupling.** Discussion frequency decouples from resolution activity. The ratio of discussion signals to action signals (commits, deployments, ticket closures) drops below 0.1 for a given risk topic. The organizational response becomes the organizational substitute for action. Status meetings review the risk. Dashboards display 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.

### 6.3 Controlled Evaluation: 22,500 Monte Carlo Runs

A controlled Monte Carlo evaluation—22,500 runs across 225 configurations (3 anomaly types  $\times$  5 severity levels  $\times$  3 mesh sizes  $\times$  5 noise levels  $\times$  100 iterations)—establishes detection performance.

Anomaly Type	Average TPR	High-Severity TPR	FPR
Coordination failure	0.661	0.82	< 0.15
Knowledge silo	0.140	0.108	< 0.10
Normalized deviance	0.199	0.228	< 0.10

Two findings merit emphasis. First, normalized deviance is detected at much lower rates than coordination failures (TPR 0.199 vs. 0.661), confirming that gradual distributional drift resists per-signal scoring. The mesh misses approximately 80% of normalized deviance signals. This is an honest limitation, not a parameter tuning problem: normalized deviance produces no individual anomalous signals, only gradual distributional shift.

Second, the evaluation reveals a *severity inversion* in knowledge silo detection: deeper silos produce *lower* per-signal detection rates (TPR decreases from 0.166 at severity 0.1 to 0.108 at severity 0.9). The mechanism: at low severity, the silo is leaky—some knowledge terms appear in the struggling team’s signals, creating detectable hybrid patterns. At high severity, the silo is complete: the isolation mechanism that constitutes the silo structurally prevents the cross-boundary signal patterns that detection requires. Any correlation-based detection method faces this limit.

Detection is robust across mesh sizes: maximum TPR spread < 0.07 for all anomaly types.

### 6.4 Field Deployment

A complementary 30-day field deployment at a technology company (3,520 signals from GitHub and Linear, 23 findings surfaced from 3,146 candidates—an 11% survival rate from candidate to promoted finding) validates the architecture’s self-organizing properties. The mesh self-organized from 3 initial workers to 47, with emergent specializations corresponding to the organization’s functional domains without being given that structure. Processing cost was \$0.012 per signal.

Four illustrative findings: a systemic coordination failure connecting six independently-worked tickets across three vendors; a  $G_D/G_C$  divergence where a code solution existed but the organizational

impasse persisted; a policy communication failure where one team was unaware of a change other teams had been executing for months; and a P1 ticket unassigned for 26 days that was being actively cited to explain related failures but never discussed as a work item—the ticket’s existence substituting for resolution.

Cross-source validation using internal communication records (a data source the mesh does not ingest) confirmed all five investigated findings and extended three. The mesh, reading costly signals, detected patterns that existed in the organization’s cheap-signal channels but could not propagate through them.

## 7 Application: Structural Immunity in Legal Systems

When a bank charges you an illegal five-dollar fee, you cannot practically challenge it even if you are legally right. Mandatory arbitration eliminates class actions. Individual arbitration costs more than the fee. The filing process is opaque. The outcome is private. Each barrier alone might seem manageable; together they create *structural immunity*—a condition where legal rights exist on paper but cannot be exercised in practice.

Structural immunity is the dysmemic pressure mechanism operating at the interface between corporations and the legal system. Where dysmemic pressure filters information before it reaches organizational decision-makers, arbitration clauses filter disputes before they reach courts. The mechanisms differ but the function is identical: ensuring that inconvenient challenges never aggregate into actionable patterns. This section formalizes the suppression architecture as a multiplicative pipeline and shows that single-filter reform is structurally insufficient (McEntire, 2026b).

### 7.1 The Six-Filter Multiplicative Pipeline

**Definition 7.1** (Suppression Pipeline). Let  $N$  denote the number of consumers experiencing actionable harm. The suppression pipeline consists of  $k$  sequential filters  $F_1, \dots, F_k$ , each with suppression rate  $s_i \in [0, 1]$ . The compound survival rate is:

$$P(\text{survival}) = \prod_{i=1}^k (1 - s_i) \tag{13}$$

The six filters are: (1) *Awareness failure*—many consumers never recognize they have an actionable claim; (2) *Rational apathy*—pursuit costs exceed expected recovery (“Only a lunatic or a fanatic sues for \$30”—Justice Kagan, *Italian Colors* dissent); (3) *Information asymmetry*—confidential proceedings prevent pattern recognition across cases while corporations accumulate institutional knowledge; (4) *Mandatory arbitration friction*—procedural barriers including truncated discovery and abbreviated timelines; (5) *Class action waiver*—elimination of the only economically rational mechanism for addressing small-dollar harms; (6) *Regulatory capture and operational edge closure*—legislative failure and customer-facing nodes structurally incapable of issuing remedies.

The filters are multiplicative, not additive. A consumer who overcomes rational apathy still faces information asymmetry, mandatory arbitration friction, and class action waivers. A regulator who identifies a pattern still cannot aggregate affected consumers.

## 7.2 The Compound Suppression Result

Calibrating against CFPB data—600 observed filings from 80,000 expected claims—yields the following per-filter rates for consumer finance:

Filter	Suppression	Pass-through
Awareness failure	85.0%	15.0%
Rational apathy	99.1%	0.87%
Information asymmetry	40.0%	60.0%
Mandatory arbitration	60.0%	40.0%
Class action waiver	85.0%	15.0%
Regulatory capture / edge closure	50.0%	50.0%
<b>Compound survival</b>		<b>0.002%</b>

The compound survival rate of 0.002% means approximately 2 claims from 80,000 harmed consumers. Rational apathy alone—at 99.1% suppression—eliminates nearly all claims before any other filter operates.

Removing the single most impactful filter (rational apathy) produces a 108× improvement, but surviving claims increase only from 2 to 216 out of 80,000. The remaining five filters still block 99.7% of claims. No other single filter removal exceeds a 6.5× improvement. Achieving even 10% survival requires dismantling at least three filters simultaneously.

## 7.3 Sensitivity Analysis

The calibration depends on the assumed harm rate. Table 3 tests whether the central conclusion holds across the full plausible range.

Table 3: Robustness of structural immunity to harm rate assumptions

Harm rate	Expected claims	Compound suppression	Per-filter suppression (equal model)	Still blocked after removing strongest filter
0.01%	8,000	92.5%	35.1%	88.5%
0.05%	40,000	98.5%	50.3%	97.0%
0.10%	80,000	99.25%	55.8%	98.3%
0.50%	400,000	99.85%	66.2%	99.6%
1.00%	800,000	99.925%	69.9%	99.8%

Across two orders of magnitude in assumed harm rate, removing the single strongest filter never reduces blockage below 88%. Even at the most conservative assumption (0.01% harm rate, where compound suppression is “only” 92.5%), the remaining five filters still block nearly nine out of ten

claims after the strongest filter is completely eliminated. The paper’s central policy conclusion—that piecemeal reform targeting one filter at a time cannot restore meaningful access to adjudication—is robust to the harm rate assumption.

#### 7.4 The Independence Assumption and the FKG Defense

The pipeline model assumes filters operate independently. In practice, consumers who overcome one barrier may be more likely to overcome subsequent barriers (positive correlation). If filters are positively correlated, then by the FKG inequality,  $P(\text{survive all}) \geq \prod_i P(\text{survive}_i)$ : the independence assumption is *conservative*—it understates actual survival, making the estimated suppression a lower bound. The policy conclusion—compound suppression is robust to single-filter reform—holds *a fortiori* under positive correlation, because the remaining filters’ compound effect would also be stronger than the independence model suggests.

#### 7.5 Cross-Sector Variation

Cross-sector extension to gig economy, social media, and healthcare reveals that the dominant filter varies by industry: rational apathy in consumer finance and social media (where per-unit harm is low), class action waivers in the gig economy (where harm is high but claims benefit from aggregation), information asymmetry in healthcare (where harm is high and claims are individually viable but procedurally opaque). This variation means that sector-specific intervention strategies are necessary. One-size-fits-all reform is mathematically suboptimal under the pipeline model.

#### 7.6 Policy Conclusion

The multiplicative structure of the pipeline explains why decades of piecemeal reform have failed to restore access to adjudication and why structural immunity is robust to incremental policy change. The design constraint any effective reform must satisfy: target the dominant filter and at least two others simultaneously. For consumer finance, this means addressing rational apathy (filing subsidies, fee-shifting), restoring aggregation (limiting class action waivers below a dollar threshold), and increasing awareness (mandatory disclosure of arbitration outcomes). Removing these three filters would achieve approximately 12% survival—still imperfect but an improvement of over 5,000× from the current 0.002%.

### 8 The Solution: Meta-Compliance

The cage cannot be eliminated because it is a consequence of coordination requirements and legal obligations. The question is not how to escape but how to see clearly from within. This section develops the architectural answer: *meta-compliance*, governance structures that satisfy fiduciary duty by documenting awareness of incompleteness rather than claiming completeness.

The core insight is that the business judgment rule does not require elimination of uncertainty. It requires evidence of informed, reasonable process. An organization that documents its awareness of where its formal systems will prove insufficient, establishes governance structures to manage that insufficiency, and maintains oversight of those structures has demonstrated informed process—arguably more convincingly than one that claims its formal systems are adequate.

## 8.1 Meta-Compliance and Sanctioned Paradoxes

Meta-compliance means documenting what you know you don't know. Instead of pretending the formal system captures everything, organizations explicitly record the blind spots, track the assumptions, and create protected channels for information that contradicts the official narrative.

**Definition 8.1** (Meta-compliance). An organization practices *meta-compliance* when it satisfies the legal requirement for demonstrable soundness by demonstrating awareness that soundness cannot be fully demonstrated—proving prudence not by claiming its processes are sufficient but by showing it recognizes where processes will fail and has designed structures to manage that recognition.

An obvious objection: meta-compliance is itself a formal system, subject to the same compression dynamics it addresses. This regress is real but bounded. The purpose of meta-compliance is not to achieve perfect awareness—that is impossible—but to shift the legal and institutional incentive structure so that acknowledging uncertainty is rewarded rather than punished. The regress terminates at the point where documenting “we don't know what we don't know” is sufficient to satisfy fiduciary duty. This is a legal standard, not an epistemic one.

Meta-compliance creates space for *sanctioned paradoxes*: governance structures that formally require what formal logic cannot validate.

**Definition 8.2** (Sanctioned paradox). A *sanctioned paradox* is a governance structure that the organization formally requires and documents as necessary, whose function depends on operating outside or against the organization's primary governance frame. Its value is destroyed if it is absorbed into the frame it exists to challenge.

Four types emerge:

**1. Apprenticeship as formalized judgment transmission.** The organization formally requires learning that formal evaluation cannot fully measure.

**2. Red teams as required contradiction.** Dissent mechanisms receive explicit authority to challenge strategic decisions. The charter acknowledges that normal governance may miss critical factors.

**3. Operational roles as mandated improvisation.** Boundary-spanning roles receive authority to violate standard procedures when conditions require.

**4. Innovation spaces as protected variance.** Units receive charters separating them from parent metrics, acknowledging pursuit of opportunities that cannot be justified using current evaluation criteria.

For sanctioned paradoxes to survive legal scrutiny, they must meet three conditions simultaneously: explicit documentation of necessity, clear boundaries and supervision, and structural protection from absorption. These three conditions constitute *trust architectures*.

## 8.2 A Game-Theoretic Model of Governance Mode Selection

**Definition 8.3** (Two modes of demonstrable soundness). *Mode A* proves soundness by eliminating variance: documenting that procedures were followed, metrics were met, standards were satisfied. *Mode B* proves soundness by documenting bounded variance: recording that uncertainty exists, that procedures will be insufficient, and that judgment under supervision is the appropriate response.

**Definition 8.4** (Environment parameters). Let  $\tau \in [0, 1]$  denote *environmental turbulence*—the rate at which conditions deviate from governance frame assumptions. Let  $\lambda \in [0, 1]$  denote *legal exposure*—the degree to which governance decisions face fiduciary scrutiny and litigation risk.

Each governance mode has three structural parameters. Mode A:  $d_A$  (defensibility in stable conditions, high),  $r_A$  (rigidity cost when environments shift),  $v_A$  (variance liability, low). Mode B:  $d_B$  (defensibility, lower than  $d_A$ ),  $a_B$  (adaptation benefit when environments shift),  $v_B$  (variance liability, higher than  $v_A$ ).

The payoff functions are:

$$U_A(\tau, \lambda) = (1 - \tau) d_A - \tau r_A - \lambda v_A \quad (14)$$

$$U_B(\tau, \lambda) = (1 - \tau) d_B + \tau a_B - \lambda v_B \quad (15)$$

In stable environments (low  $\tau$ ), Mode A captures nearly all its defensibility while Mode B's lower defensibility is a drag. As turbulence rises, Mode A pays increasing rigidity cost while Mode B gains from adaptation. Legal exposure imposes a drag on both modes proportional to variance liability, but heavier on Mode B.

**Theorem 8.5** (Crossover turbulence). *For given legal exposure  $\lambda$ , the crossover turbulence is:*

$$\tau^*(\lambda) = \frac{(d_A - d_B) - \lambda(v_A - v_B)}{(d_A - d_B) + r_A + a_B} \quad (16)$$

*For  $\tau < \tau^*$ , Mode A dominates. For  $\tau > \tau^*$ , Mode B dominates.*

*Proof.* Expand  $U_A = U_B$ :

$$(1 - \tau) d_A - \tau r_A - \lambda v_A = (1 - \tau) d_B + \tau a_B - \lambda v_B$$

Rearranging:

$$d_A - d_B - \lambda(v_A - v_B) = \tau[(d_A - d_B) + r_A + a_B]$$

Since  $v_A < v_B$  (so  $v_A - v_B < 0$ ), increasing  $\lambda$  increases the numerator, pushing  $\tau^*$  upward. Dividing by the positive denominator yields (16).  $\square$

*Remark 8.6.* The sign of  $\partial\tau^*/\partial\lambda$  is  $-(v_A - v_B)/(d_A - d_B + r_A + a_B)$ . Since  $v_A < v_B$ , this derivative is positive: higher legal exposure raises  $\tau^*$ , contracting the region where Mode B dominates. This is the formal expression of how fiduciary pressure pushes organizations toward Mode A.

### 8.2.1 Calibration

Table 4: Baseline parameter calibration.

Parameter	Interpretation	Value	Rationale
$d_A$	Mode A defensibility	0.9	Conformity is easy to document
$d_B$	Mode B defensibility	0.6	Bounded variance harder to justify
$r_A$	Mode A rigidity cost	0.6	Rigid systems fail under disruption
$a_B$	Mode B adaptation benefit	0.8	Adaptive capacity is highly valuable
$v_A$	Mode A variance liability	0.1	Low variance, low litigation surface
$v_B$	Mode B variance liability	0.3	Managed variance invites scrutiny

Under this calibration:

- At  $\lambda = 0.0$ :  $\tau^* = 0.176$  (Mode B dominates for 82.4% of turbulence range)
- At  $\lambda = 0.50$ :  $\tau^* = 0.235$  (Mode B dominates for 76.5%)
- At  $\lambda = 1.0$ :  $\tau^* = 0.294$  (Mode B dominates for 70.6%)

**Key result:** Across the full range of legal exposure, Mode B dominates 70–82% of the turbulence parameter space.

### 8.2.2 Sensitivity

The equilibrium is most sensitive to: closing the defensibility gap ( $d_B \rightarrow d_A$ ) drops  $\tau^*$  and expands Mode B’s region; higher Mode A rigidity cost ( $r_A$ ) pushes  $\tau^*$  downward; higher Mode B variance liability ( $v_B$ ) pushes  $\tau^*$  upward, contracting Mode B’s region.

**Corollary 8.7** (Van Gorkom Shift). *A legal precedent that increases  $\lambda$  from  $\lambda_0$  to  $\lambda_1$  shifts the crossover turbulence by:*

$$\Delta\tau^* = \frac{(\lambda_1 - \lambda_0)(v_B - v_A)}{(d_A - d_B) + r_A + a_B} \quad (17)$$

*This is always positive when  $v_B > v_A$ , meaning increased legal exposure always contracts the Mode B region.*

Under baseline parameters, a shift from  $\lambda = 0.3$  to  $\lambda = 0.7$  raises  $\tau^*$  by +0.047—organizations that were rationally choosing meta-compliance are pushed toward variance elimination.

### 8.3 The Turbulence Paradox

The most consequential finding emerges when legal exposure rises with turbulence. In practice, turbulent environments attract regulatory attention, shareholder litigation, and media scrutiny. We model this coupling as  $\lambda = c \cdot \tau$ .

**Theorem 8.8** (Turbulence Paradox). *When legal exposure is coupled to turbulence ( $\lambda = c\tau$ ), the crossover turbulence under coupling is:*

$$\tau_c^* = \frac{d_A - d_B}{(d_A - d_B) + r_A + a_B - c(v_B - v_A)} \quad (18)$$

*There exists a critical coupling  $c^*$  at which  $\tau_c^* = 1$  and Mode A dominates for all turbulence levels:*

$$c^* = \frac{r_A + a_B}{v_B - v_A} \quad (19)$$

*Under baseline parameters,  $c^* = 7.0$ .*

*Proof.* Substituting  $\lambda = c\tau$  into  $U_A = U_B$  and solving for  $\tau$ :

$$\begin{aligned} (1 - \tau)d_A - \tau r_A - c\tau v_A &= (1 - \tau)d_B + \tau a_B - c\tau v_B \\ d_A - d_B &= \tau[(d_A - d_B) + r_A + a_B - c(v_B - v_A)] \end{aligned}$$

yielding (18). Setting  $\tau_c^* = 1$  and solving for  $c$  gives (19). □

The paradox: as industries become more turbulent, Mode B becomes more valuable because adaptive capacity matters more. But if legal exposure rises with turbulence, the legal penalty on variance increases simultaneously. At sufficiently high coupling, Mode A dominates universally—organizations are trapped in variance elimination precisely when variance management is most needed.

*Remark 8.9.* The turbulence paradox is not merely theoretical. Industries experiencing rapid technological disruption often simultaneously face increased regulatory scrutiny—consider financial technology, autonomous vehicles, or artificial intelligence governance. The model predicts these industries face the strongest structural pressure toward Mode A governance despite having the greatest need for Mode B adaptation.

## 8.4 Empirical Illustrations

### 8.4.1 NASA Post-Columbia

The Columbia Accident Investigation Board concluded in 2003 that NASA’s organizational culture was “as much a cause of the accident as the foam that struck the left wing” (CAIB, 2003). Despite formal safety processes established after Challenger, the organization had reverted to Mode A

governance: safety reviews existed as documented procedures but had lost independence and authority.

Post-Columbia reforms implemented sanctioned paradoxes: Independent Technical Authority (engineering assessments separated from program management, with authority to halt operations); Safety and Mission Assurance (reporting directly to the NASA Administrator); and dissent documentation (formal processes requiring that dissenting opinions be recorded and addressed before proceeding). This was a deliberate shift from Mode A to Mode B—from documenting conformity to documenting awareness of insufficiency. The history of post-Challenger reforms reverting to Mode A within two decades provides a natural experiment in decay dynamics.

### 8.4.2 Pixar Braintrust

Pixar’s Braintrust meets every few months to assess films under development. The director presents rough cuts. The Braintrust—comprising other directors, writers, and heads of story—provides candid feedback. The structural innovation is the separation of authority from assessment: the Braintrust has no authority to override the director. The director retains sole responsibility for all decisions. This is a sanctioned paradox: a formally constituted evaluative body whose power is explicitly limited to challenge without override. Ed Catmull, Pixar’s co-founder, described the critical design choice as ensuring that the group “should have no authority, so the person responsible for the production would never have to come into the room in a defensive posture knowing the group could undermine him” (Catmull and Wallace, 2014). The result is a feedback channel structured to reward divergent signal (what is wrong with this film?) rather than convergent signal (what will make the director happy?). This is meta-compliance in creative production: the governance structure formally documents that the current version is incomplete and structurally protects the channel for information that says so.

### 8.4.3 Bridgewater Associates

Bridgewater operates under “radical transparency” (Dalio, 2017). All meetings are recorded. Employees rate each other via dot-voting. The “believability-weighted” decision process evaluates reasoning quality independently of hierarchy. Bridgewater’s “Principles” explicitly state that models and processes will be wrong. This is meta-compliance in its most literal form: compliance documentation consisting substantially of evidence that models are incomplete.

## 8.5 Consolidated Propositions

The Mirror framework generates three consolidated testable predictions.

**Proposition 8.10** (Variance Preservation). *Organizations with documented trust architectures show higher variance preservation following formalization events than organizations without such mechanisms, measured by sustained diversity in decision justifications, strategic approaches, and resource allocations.*

**Measurement protocol.** Classify trust architecture quality  $T \in [0, 1]$  before the formalization event. Measure variance  $V(t)$  at  $t_0 - 12$  months through  $t_0 + 24$  months. Compute variance preservation ratio  $\rho = V(t_0 + 24)/V(t_0)$ . Test whether  $\rho$  is significantly higher for  $T > 0.7$  vs.  $T < 0.3$ .

**Falsification.** High- $T$  organizations show compression equal to or greater than low- $T$  organizations.

**Proposition 8.11** (Equilibrium Threshold). *Organizations operating in environments with turbulence above  $\tau^*(\lambda)$  that adopt Mode A governance will show declining performance relative to Mode B adopters, with the performance gap widening as  $\tau$  increases beyond  $\tau^*$ .*

**Falsification.** Mode B organizations do not outperform Mode A organizations in high-turbulence environments.

**Proposition 8.12** (Turbulence Paradox). *In industries where legal exposure rises with environmental turbulence (coupling  $c > 0$ ), organizations will exhibit higher rates of Mode A adoption than predicted by turbulence alone, and this “excess Mode A adoption” will correlate with increased organizational failure rates during disruption.*

**Falsification.** Coupling  $c$  does not predict excess Mode A adoption, or excess Mode A adoption does not correlate with higher failure rates during disruption.

## 8.6 Failure Modes: When Meta-Compliance Becomes Compliance Theater

Meta-compliance is not self-sustaining. Four failure modes merit analysis.

**Evaluation by the frame it transcends.** A red team whose recommendations are assessed by the strategic planning group it was chartered to critique has lost its function. The evaluation criterion recaptures the mechanism.

**Metric optimization of meta-compliance.** If “number of red team challenges issued” becomes the metric, teams optimize for challenge quantity rather than quality. Goodhart’s Law applied to meta-compliance (Goodhart, 1975).

**The recursive cage.** If regulators require evidence of meta-compliance—documented proof that organizations acknowledge incompleteness—then meta-compliance itself becomes a compliance requirement. Organizations optimize documentation of incompleteness-recognition rather than actually recognizing incompleteness. The mirror becomes another wall.

**The turbulence paradox as trap.** Firms in turbulent, litigious environments rationally choose Mode A even though Mode B would produce better outcomes absent the legal penalty. They *know* they need adaptive governance but cannot justify it under fiduciary scrutiny. The failure is systemic.

Three diagnostic criteria distinguish genuine meta-compliance from theater: *consequence asymmetry* (the mechanism can halt operations), *discomfort generation* (decisions are delayed, leaders are challenged), and *independence verification* (the mechanism has contradicted the dominant frame and been sustained).

## 9 Discussion and Conclusion

### 9.1 The Full Causal Chain

The seven sections of this paper trace a single causal chain, then extend it from theory through measurement to application:

1. **Formalization creates frames** (Section 2). Growth and fiduciary duty push organizations toward governance frames that compress reality into legible metrics. Empirically confirmed: 75 SEC filings show universal variance compression (5.8–33%) following IPO, with magnitude proportional to legal exposure ( $R^2 = 0.80$ ).
2. **Frames create selection environments** (Section 3). Three compound dynamics—strategic communication degradation, adverse selection, and transmission bias—produce dysmemic pressure: a selection force favoring fit over truth. The hierarchical Crawford–Sobel simulation confirms monotonic information destruction through organizational depth across 126 configurations.
3. **Selection determines valence** (Section 4). The generative lossy channel framework proves that compression artifacts can be either pathological (under convergent selection) or beneficial (under divergent selection). Five sufficient conditions for net-beneficial noise are validated with zero counterexamples across 500 Monte Carlo configurations. The compression ratchet formalizes why recovery requires external intervention.
4. **The formal foundation quantifies the cost of conformity** (Section 5). The strategic rate-distortion-perception function proves that conformity pressure accounts for 22–72% of total distortion in strategic communication channels, with a phase transition at the babbling boundary. The generative residual—the excess distortion from distributional conformity—is strictly positive whenever the perception constraint binds.
5. **The dysfunction can be measured** (Section 6). A stigmergic mesh architecture detects organizational dysfunction from work artifacts with TPR 0.82 for coordination failures and TPR 0.199 for normalized deviance (22,500 Monte Carlo runs). The 80% miss rate on normalized deviance is an honest limitation: gradual distributional drift resists per-signal

scoring. Field deployment validates the architecture in production (30 days, 23 findings, cross-source confirmation).

6. **The mechanism operates in legal systems** (Section 7). Mandatory arbitration combined with platform dominance creates a six-filter multiplicative pipeline that suppresses 99.998% of meritorious claims. Single-filter reform is structurally insufficient: removing the strongest filter still leaves 99.7% of claims blocked. The dominant filter varies by industry, requiring sector-specific intervention.
7. **Architecture determines outcome** (Section 8). Meta-compliance provides the architectural response: governance structures that satisfy fiduciary duty by documenting awareness of incompleteness. The game-theoretic model shows Mode B dominates 70–82% of the parameter space, but the turbulence paradox reveals structural traps when legal exposure rises with environmental turbulence.

## 9.2 Policy Implications

The turbulence paradox has immediate policy relevance. If legal regimes that increase fiduciary scrutiny in turbulent industries inadvertently trap organizations in Mode A governance, then the legal structure designed to protect stakeholders may increase systemic fragility. This does not argue against fiduciary duty but suggests that legal standards might explicitly accommodate Mode B governance—recognizing documented variance management as evidence of prudence, not recklessness.

The Van Gorkom standard already accommodates this. The court faulted not uncertainty but failure to demonstrate deliberation about uncertainty. *In re Caremark* (*In re Caremark*, 1996) held that good-faith effort to establish information systems satisfies oversight duty even if systems prove insufficient. Trust architectures are information systems: formalized structures for accessing perspectives that standard reporting misses.

## 9.3 Open Problems

**O1: Multi-level cascades.** Theorem 4.6 covers two levels. Does it compose across  $N$  levels?

**O2: Dynamic environments.** Static conditions assumed. In non-stationary environments, the optimal noise level must track shifting thresholds.

**O3: Quantitative  $\sigma^*$  prediction.** The theorem guarantees existence but not a closed form for the optimal noise level.

**O4: Endogenous acceptability dynamics.** The ratchet proposition treats acceptability distribution dynamics qualitatively. A full characterization of the fixed-point  $Q^*$  as a function of system parameters is needed.

**O5: Ratchet speed measurement.** Operationalizing ratchet speed as a function of compression ratio and preference divergence would yield quantitative organizational diagnostics.

**O6: Empirical validation of the turbulence paradox.** Longitudinal study tracking governance mode, turbulence, legal exposure, and outcomes across industries with varying legal-turbulence coupling.

**O7: Non-Gaussian strategic RDP.** Extension of the closed-form results to discrete distributions and non-quadratic loss functions would strengthen the RLHF mapping from conjecture to formal application.

**O8: Longitudinal normalized deviance detection.** The 80% miss rate on normalized deviance (Section 6) identifies the critical research direction: developing detectors that track distributional properties of organizational language over time windows longer than the current rolling baseline.

**O9: Dynamic filter interaction.** The structural immunity pipeline model (Section 7) assumes filter independence. Modeling dynamic interactions—where reform in one filter changes others through feedback—could make multi-filter intervention more effective than the static model predicts.

## 9.4 Testable Predictions Summary

#	Prediction	Falsification
1	Variance compresses post-formalization	No compression in controlled sample
2	Compression proportional to legal exposure	No exposure gradient
3	Founder control attenuates compression	No founder-control effect
4	Information degrades monotonically through hierarchy	MI increases at any layer
5	Trust architectures preserve variance	No preservation differential
6	Mode B outperforms Mode A above $\tau^*$	No performance differential
7	Turbulence paradox traps organizations	No excess Mode A in high- $c$ industries
8	Recovery requires external intervention	Internal reform displaces ratchet equilibrium
9	Conformity accounts for 22–72% of distortion	Generative residual negligible in calibrated channels
10	Normalized deviance detectable via linguistic compression	No linguistic signature precedes operational failure
11	Single-filter reform insufficient for structural immunity	Single-filter removal achieves >5% survival

Predictions 1–3 have empirical support from the SEC filing analysis (Section 2.6). Prediction 4 has computational support from the hierarchical simulation (Section 3.5). Prediction 9 has numerical support from the Gaussian strategic RDP computation (Section 5). Prediction 10 has computational support from the Monte Carlo evaluation (Section 6). Prediction 11 has analytical support from the pipeline model (Section 7). Predictions 5–8 await empirical testing.

## 9.5 Conclusion

Organizations, like formal systems, cannot be both complete and consistent. They can only choose which incompleteness they will tolerate. Those optimizing for demonstrable soundness accept

frame-dependence and slow adaptation. Those maintaining external perspective accept coordination costs and legal risk. Most organizations choose demonstrable soundness because the costs of external perspective are immediate while the costs of rigidity are deferred.

This paper has traced the full arc: the problem exists (Section 2—universal variance compression in SEC filings), the mechanism explains it (Section 3—dysmemic pressure from three compound dynamics), the theory formalizes it (Section 4—generative lossy channels with dual valence), the formal foundation quantifies it (Section 5—22–72% of distortion from conformity), it can be measured (Section 6—stigmergic mesh with 82% TPR on coordination failures and honest 20% on normalized deviance), it appears in legal systems (Section 7—six-filter pipeline suppressing 99.998% of claims), and meta-compliance addresses it (Section 8—governance structures dominating 70–82% of the parameter space).

The mechanism is embedded in governance architecture, not in the character of individuals who occupy roles within it. This is why reforms repeatedly fail. Each generation discovers the problem and implements solutions. The solutions work temporarily. Then they decay—not because leaders lose commitment but because the legal and operational pressures that created formalization remain unchanged. The structural immunity analysis makes this concrete: decades of piecemeal arbitration reform have failed not because the reforms were poorly designed but because the multiplicative architecture is robust to single-filter intervention.

The trap has no villain. It has a law.

But the law admits a strategy. Meta-compliance does not promise freedom. It promises awareness. And awareness, sustained across time and leadership changes through governance structures that institutionalize paradox rather than resolve it, is the only sustainable adaptation to structural incompleteness.

Trust architectures do not escape the cage. They place mirrors within it.

## References

- Akerlof, G. A. (1970). The market for “lemons”: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3):488–500.
- Adams, R. B. and Ferreira, D. (2009). Women in the boardroom and their impact on governance and performance. *Journal of Financial Economics*, 94(2):291–309.
- Argyris, C. and Schön, D. A. (1978). *Organizational Learning: A Theory of Action Perspective*. Addison-Wesley, Reading, MA.
- Army, U.S. (2019). *ADP 6-0: Mission Command*. Department of the Army, Washington, DC.
- Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall, London.
- Atlan, H. (1979). *Entre le cristal et la fumée*. Éditions du Seuil, Paris.

- Beer, M. and Nohria, N. (2000). Cracking the code of change. *Harvard Business Review*, 78(3):133–141.
- Benzi, R., Sutera, A., and Vulpiani, A. (1981). The mechanism of stochastic resonance. *Journal of Physics A: Mathematical and General*, 14(11):L453–L457.
- Berliner, P. F. (1994). *Thinking in Jazz: The Infinite Art of Improvisation*. University of Chicago Press, Chicago.
- Bernstein, S. (2015). Does going public affect innovation? *Journal of Finance*, 70(4):1365–1403.
- Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026.
- Blau, Y. and Michaeli, T. (2018). The perception-distortion tradeoff. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6228–6237.
- Blau, Y. and Michaeli, T. (2019). Rethinking lossy compression: The rate-distortion-perception tradeoff. In *Proceedings of the 36th International Conference on Machine Learning*, pages 675–685.
- Boyd, R. and Richerson, P. J. (1985). *Culture and the Evolutionary Process*. University of Chicago Press, Chicago.
- Brown, T. (2008). Design thinking. *Harvard Business Review*, 86(6):84–92.
- Buchanan, J. M. and Tullock, G. (1962). *The Calculus of Consent: Logical Foundations of Constitutional Democracy*. University of Michigan Press, Ann Arbor.
- Columbia Accident Investigation Board (2003). *Report of the Columbia Accident Investigation Board*, Volume 1. NASA, Washington, DC.
- Carreyrou, J. (2018). *Bad Blood: Secrets and Lies in a Silicon Valley Startup*. Alfred A. Knopf, New York.
- Catmull, E. and Wallace, A. (2014). *Creativity, Inc.: Overcoming the Unseen Forces That Stand in the Way of True Inspiration*. Random House, New York.
- Chapeau-Blondeau, F. (1997). Stochastic resonance in the Hopfield model driven by non-Gaussian noise. *Physical Review E*, 55(2):2016–2019.
- Christensen, C. M. (1997). *The Innovator’s Dilemma: When New Technologies Cause Great Firms to Fail*. Harvard Business School Press, Boston.
- Cover, T. M. and Thomas, J. A. (2006). *Elements of Information Theory*. Wiley-Interscience, 2nd edition.

- Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50(6):1431–1451.
- Dalio, R. (2017). *Principles: Life and Work*. Simon & Schuster, New York.
- DiMaggio, P. J. and Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, 48(2):147–160.
- Dodd-Frank Wall Street Reform and Consumer Protection Act (2010). Pub. L. 111–203, 124 Stat. 1376.
- Dreyfus, H. L. and Dreyfus, S. E. (1986). *Mind Over Machine: The Power of Human Intuition and Expertise in the Era of the Computer*. Free Press, New York.
- Edmondson, A. C. (1999). Psychological safety and learning behavior in work teams. *Administrative Science Quarterly*, 44(2):350–383.
- Edmondson, A. C. (2019). *The Fearless Organization: Creating Psychological Safety in the Workplace for Learning, Innovation, and Growth*. Wiley, Hoboken, NJ.
- Ellington, D. (1973). *Music Is My Mistress*. Doubleday, Garden City, NY.
- Gammaitoni, L., Hänggi, P., Jung, P., and Marchesoni, F. (1998). Stochastic resonance. *Reviews of Modern Physics*, 70(1):223–287.
- Goodhart, C. A. E. (1975). Problems of monetary management: The U.K. experience. In *Papers in Monetary Economics*, Reserve Bank of Australia, Sydney.
- Hannan, M. T. and Freeman, J. (1984). Structural inertia and organizational change. *American Sociological Review*, 49(2):149–164.
- Hanson, R. (2003). Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119.
- Henrich, J. and McElreath, R. (2003). The evolution of cultural evolution. *Evolutionary Anthropology*, 12(3):123–135.
- House Committee on Transportation and Infrastructure (2020). *The Design, Development & Certification of the Boeing 737 MAX: Final Committee Report*. U.S. House of Representatives, Washington, DC.
- In re Caremark International Inc. Derivative Litigation* (1996). 698 A.2d 959 (Del. Ch. 1996).
- Jensen, M. C. and Meckling, W. H. (1976). Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4):305–360.
- Johnson-Laird, P. N. (2002). How jazz musicians improvise. *Music Perception*, 19(3):415–442.

- Kosko, B. and Mitaim, S. (2003). Stochastic resonance in noisy threshold neurons. *Neural Networks*, 16(5–6):755–761.
- La Porte, T. R. and Consolini, P. M. (1991). Working in practice but not in theory: Theoretical challenges of “high-reliability organizations.” *Journal of Public Administration Research and Theory*, 1(1):19–48.
- Lave, J. and Wenger, E. (1991). *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press, Cambridge.
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1):71–87.
- McCord, P. (2014). How Netflix reinvented HR. *Harvard Business Review*, 92(1):71–76.
- McCrum, D. (2022). *Money Men: A Hot Startup, a Billion Dollar Fraud, a Fight for the Truth*. Scribner, New York.
- Meyer, J. W. and Rowan, B. (1977). Institutionalized organizations: Formal structure as myth and ceremony. *American Journal of Sociology*, 83(2):340–363.
- Nembhard, I. M. and Edmondson, A. C. (2006). Making it safe: The effects of leader inclusiveness and professional status on psychological safety and improvement efforts in health care teams. *Journal of Organizational Behavior*, 27(7):941–966.
- Nemeth, C. J. (1986). Differential contributions of majority and minority influence. *Psychological Review*, 93(1):23–32.
- Orr, J. E. (1996). *Talking About Machines: An Ethnography of a Modern Job*. Cornell University Press, Ithaca, NY.
- Perez, E., Ringer, S., Lukošiuūtė, K., Nguyen, K., Chen, E., Heiner, S., Pettit, C., Olsson, C., Kundu, S., Kadavath, S., et al. (2022). Discovering language model behaviors with model-written evaluations. *arXiv preprint arXiv:2212.09251*.
- Pigeau, R. and McCann, C. (2002). Re-conceptualizing command and control. *Canadian Military Journal*, 3(1):53–64.
- Prendergast, C. (1993). A theory of “yes men.” *American Economic Review*, 83(4):757–770.
- Pressing, J. (1988). Improvisation: Methods and models. In Sloboda, J. A., editor, *Generative Processes in Music*, pages 129–178. Oxford University Press.
- Roberts, K. H. (1990). Some characteristics of one type of high reliability organization. *Organization Science*, 1(2):160–176.

- Robinson, P. (2019). Boeing’s 737 MAX: A failure of management, not just technology. *Stanford Business Insights*, November 2019.
- Sarbanes-Oxley Act (2002). Pub. L. 107–204, 116 Stat. 745.
- Schulz-Hardt, S., Brodbeck, F. C., Mojzisch, A., Kerschreiter, R., and Frey, D. (2006). Group decision making in hidden profile situations: Dissent as a facilitator for decision quality. *Journal of Personality and Social Psychology*, 91(6):1080–1093.
- Scott, W. R. (2013). *Institutions and Organizations: Ideas, Interests, and Identities*. Sage Publications, 4th edition.
- Shamir, E. (2011). *Transforming Command: The Pursuit of Mission Command in the U.S., British, and Israeli Armies*. Stanford University Press.
- Sharma, M., Tong, M., Korbak, T., Duvenaud, D., Askell, A., Bowman, S. R., Cheng, N., Durmus, E., Hatfield-Dodds, Z., Johnston, S. R., et al. (2023). Towards understanding sycophancy in language models. *arXiv preprint arXiv:2310.13548*.
- Smith v. Van Gorkom* (1985). 488 A.2d 858 (Del. 1985).
- Sturdy, A. (2011). Consultancy’s consequences? A critical assessment of management consultancy’s impact on management. *British Journal of Management*, 22(3):517–530.
- Teece, D. J. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7):509–533.
- Tetlock, P. E. (2005). *Expert Political Judgment: How Good Is It? How Can We Know?* Princeton University Press.
- Tetlock, P. E. and Gardner, D. (2015). *Superforecasting: The Art and Science of Prediction*. Crown Publishers, New York.
- Tishby, N., Pereira, F. C., and Bialek, W. (1999). The information bottleneck method. In *Proceedings of the 37th Allerton Conference on Communication, Control, and Computing*, pages 368–377.
- Tucker, M. (1995). *Ellington: The Early Years*. University of Illinois Press, Urbana.
- Tushman, M. L. and O’Reilly, C. A. (1996). Ambidextrous organizations: Managing evolutionary and revolutionary change. *California Management Review*, 38(4):8–29.
- Useem, J. (2019). The long-forgotten flight that sent Boeing off course. *The Atlantic*, November 2019.
- Vaughan, D. (1996). *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. University of Chicago Press, Chicago.

- von Foerster, H. (1960). On self-organizing systems and their environments. In Yovits, M. C. and Cameron, S., editors, *Self-Organizing Systems*, pages 31–50. Pergamon Press, New York.
- Weber, M. (1958). Bureaucracy. In Gerth, H. H. and Mills, C. W., editors, *From Max Weber: Essays in Sociology*, pages 196–244. Oxford University Press, New York. (Originally published 1922.)
- Weick, K. E. and Sutcliffe, K. M. (2007). *Managing the Unexpected: Resilient Performance in an Age of Uncertainty*. Jossey-Bass, 2nd edition.
- Werr, A. and Styhre, A. (2003). Management consultants—friend or foe? *International Studies of Management & Organization*, 32(4):43–66.
- Carpenter, G. A. and 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.
- McEntire, J. (2026). Ambient structure discovery via stigmergic mesh: Detecting normalized deviance through distributed organizational signal processing. Working paper, Cage & Mirror Press.
- McEntire, J. (2026). Structural immunity: How platform dominance and mandatory arbitration create corporate impunity. Working paper, Cage & Mirror Press.
- Consumer Financial Protection Bureau (2015). *Arbitration Study: Report to Congress, Pursuant to Dodd-Frank Wall Street Reform and Consumer Protection Act § 1028(a)*. CFPB, Washington, DC.