Dysmemic Pressure: Selection Dynamics in Organizational Information Environments

An Agent-Based Simulation of Information Degradation Through Hierarchical Strategic Communication

Jeremy McEntire
Cage & Mirror Press
jmc@kindex.tools

March 2026

Abstract

Organizations staffed by intelligent people converge on collective delusion with predictable regularity. This paper identifies a structural mechanism—*dysmemic pressure*—the compound selection force that emerges when organizations compress information to coordinate at scale. Three dynamics drive the pressure: strategic communication degradation (Crawford & Sobel, 1982), adverse selection in idea markets (Akerlof, 1970), and transmission bias (Boyd & Richerson, 1985). We formalize the first dynamic through a hierarchical Crawford–Sobel model and validate it with agent-based Monte Carlo simulation. In the single-channel case, we derive that as preference divergence b approaches 1/4, the equilibrium partition count $N^*(b) \to 1$ and mutual information $I(\theta; y) \to 0$: the channel degenerates to babbling. In the hierarchical case, a parameter sweep over 126 configurations (7 layer counts $\times$ 6 bias levels $\times$ 3 noise levels, 50 Monte Carlo iterations each) reveals that without noise, partition structure is preserved through the hierarchy, but with noise, mutual information degrades monotonically as predicted by the data processing inequality. We operationalize dysmemic pressure as $\text{DP}(k) = 1 - I(\text{report}_k, \theta)/I(\text{report}_0, \theta)$, providing a measurable index of information destruction at each hierarchical layer. Three case studies—Boeing 737 MAX, Theranos, and Wirecard—illustrate the mechanism across industries. The analysis reframes organizational dysfunction from moral failure to selection outcome, explaining why standard interventions fail and what structural countermeasures might succeed.

Keywords: organizational behavior, information economics, cultural evolution, strategic communication, cheap talk, agency theory, agent-based simulation, organizational failure, epistemic dysfunction

Contents

1 Introduction

The modern organization is supposed to be an information-processing machine. Distributed knowledge flows upward through reporting structures, gets aggregated by managers, and emerges as coordinated action. The hierarchy exists, in theory, to make the whole smarter than any individual part. This premise underlies nearly a century of organizational theory, from Weber's bureaucratic rationality through Simon's bounded rationality to contemporary work on organizational learning and knowledge management.

The premise does not survive contact with observation.

Boeing's engineers flagged the MCAS flight-control system as a single-point-of-failure risk years before the 737 MAX crashes killed 346 people (Robinson and Shay, 2019; House Committee on Transportation and Infrastructure, 2020). Theranos engineers knew the Edison blood-testing device could not meet clinical accuracy standards; those who raised concerns were fired or silenced (Carreyrou, 2018). Wirecard's internal auditors and external journalists documented a €1.9 billion hole in the balance sheet; regulators and auditors accepted management's denials for years (McCrum, 2022). In each case, the organization possessed the information that would have saved it. The information existed in documents, in databases, in 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 structural. Organizations compress information to coordinate at scale. Compression creates gaps between representation and reality. Those gaps become selection environments. What survives in those environments is what fits—regardless of whether it maps reality. And fitness, in organizational contexts, is often negatively correlated with truth.

I call this selection force *dysmemic pressure*. The term names the compound dynamic that emerges when compression and selection interact. Three component forces drive it: strategic communication degradation, adverse selection in idea markets, and transmission bias. Each force is documented in separate literatures. The synthesis explains what the components alone cannot: why intelligent organizations deceive themselves despite incentives and intentions pointing the other way.

Contributions of this paper. Beyond the theoretical synthesis, this paper makes three advances over the earlier version (McEntire, 2025):

1. A **formal derivation** of partition coarsening and information loss in the Crawford–Sobel channel, establishing the babbling boundary at $b = 1/4$ and computing mutual information as a function of bias (Section 3).

2. An **agent-based Monte Carlo simulation** of hierarchical strategic communication, sweeping 126 parameter configurations to characterize information degradation through organizational depth (Section 4).

3. An **operationalization** of dysmemic pressure as a measurable quantity—the Dysmemic Pressure Index DP(k)—enabling empirical assessment of information destruction at each layer of an organization (Section 5).

The paper proceeds as follows. Section 2 reviews the three component dynamics. Section 3 develops the formal Crawford–Sobel partition coarsening derivation. Section 4 presents the agent-based simulation results. Section 5 operationalizes dysmemic pressure as a measurable index. Section 6 presents three case studies. Section 7 connects to the forecasting literature. Section 8 describes the compound pressure and the Cage. Section 9 derives testable propositions. Section 10 evaluates structural countermeasures. Section 11 discusses limitations. Section 12 concludes.

2 The Three Dynamics

2.1 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, transmission degrades: the sender lumps states into coarser categories, smoothing away details that would hurt the sender. At sufficient divergence, messages carry no information at all about the true state—what Crawford and Sobel call a *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 diverge from the director's. The director's diverge from the executive's. Each interface is a potential degradation point. The engineer softens bad news to avoid scrutiny. The manager filters for relevance to avoid difficult conversations. The director packages for palatability to maintain relationships. The executive presents optimistically to the board to protect their position. Each responds rationally to local incentives. The aggregate is organizational self-deception built from individually rational choices.

Prendergast (1993) demonstrated that when organizations rely on subjective performance evaluations, subordinates face powerful incentives to conform to the principal's prior beliefs. The Yes Man emerges from this structure as the equilibrium output of subjective evaluation, not a character flaw in particular individuals.

2.2 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.

2.3 Transmission Bias

Boyd and Richerson (1985) documented how ideas spread independent of their truth value. Selection operates on transmissibility. Whether an idea represents reality is a separate question.

Content bias: simple ideas spread faster than complex ones because they are easier to remember and repeat. "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. The pressure to conform locks in whatever belief reached threshold first (Bikhchandani et al., 1992).

These biases operate largely below conscious awareness. The transmission properties of an idea determine its spread. Truth is along for the ride, or not.

3 Crawford–Sobel Partition Coarsening: Formal Derivation

We now formalize the first dynamic—strategic communication degradation—to establish the quantitative foundation for the simulation that follows.

3.1 The Model

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

- The state θ is drawn uniformly from $[0, 1]$.
- The Sender observes θ 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.

3.2 Equilibrium Partition Structure

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

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

The partition boundaries are

$$a_i = \frac{i}{N} + 2b \cdot i(i - N), \qquad i = 0, 1, \ldots, N, \tag{2}$$

which satisfy $a_0 = 0$ *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). □

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

Proof. At $b = 1/4$, equation (1) gives $N^*(1/4) = \lfloor(-1+\sqrt{1+8})/2\rfloor = \lfloor(-1+3)/2\rfloor = \lfloor 1 \rfloor = 1$. With $N = 1$, the Sender's message is constant regardless of θ, so $I(\theta; m) = 0$. For b slightly below $1/4$, $N^* = 2$ and the partition is $\{[0, a_1], [a_1, 1]\}$ with $a_1 = 1/2 + 2b(1-2) = 1/2 - 2b$, which approaches 0 as $b \to 1/4$, so the mutual information approaches 0 continuously. □

3.3 Mutual Information

Since $\theta \sim \text{Uniform}[0,1]$ and the message m identifies which of N intervals θ falls in, with interval lengths $L_i = a_i - a_{i-1}$:

Definition 3.3 (Channel mutual information)**.** The mutual information of the Crawford–Sobel channel is the reduction in differential entropy:

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

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

At $b = 0.1$, the most-informative equilibrium has $N^* = 3$ partitions. The theoretical mutual information is $I(\theta; m) = 0.8813$ bits. Our Monte Carlo simulation (10,000 samples) recovers $\hat{I} = 0.8778$ bits, validating the implementation (see Figure 1).

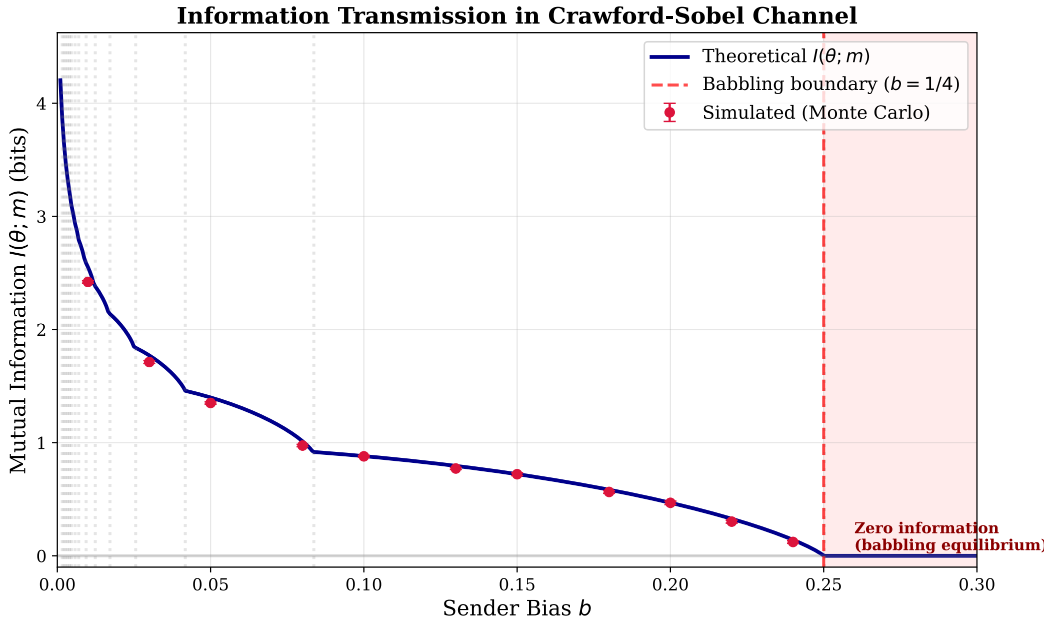

Figure 1: Mutual information $I(\theta; m)$ in the Crawford–Sobel channel as a function of sender bias b. The solid curve shows the theoretical value from equation (3); red circles show Monte Carlo estimates (50 iterations, 5000 samples each). The babbling boundary at $b = 1/4$ is marked. Gray vertical lines indicate partition-count transitions. At $b = 0.1$: theoretical $I = 0.8813$ bits, simulated $\hat{I} = 0.8778$ bits.

4 Agent-Based Simulation of Hierarchical Information Degradation

4.1 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} (not the ground truth) and communicates through its own biased channel with parameter b_k.

Definition 4.1 (Hierarchical channel)**.** A K-layer hierarchical channel is a Markov chain $\theta \to a_1 \to a_2 \to \cdots \to a_K$, where each transition $a_{k-1} \to a_k$ 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:

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

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.

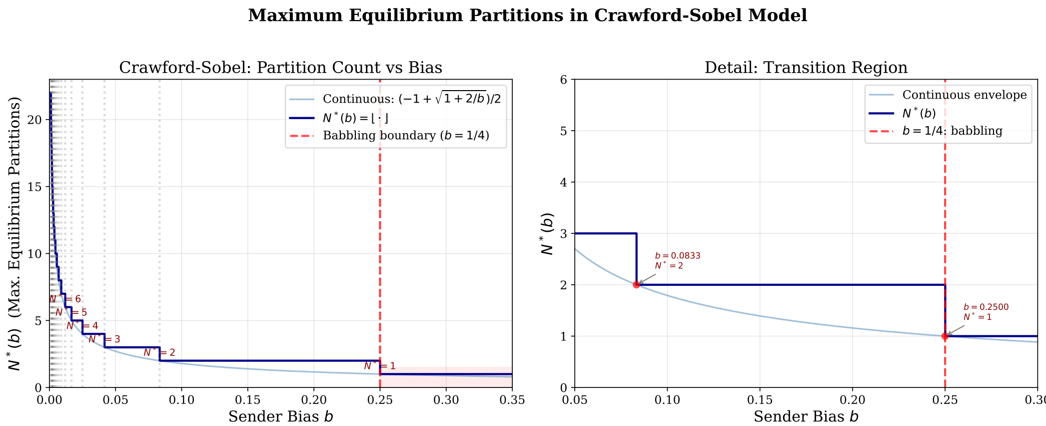

Figure 2: Maximum equilibrium partition count $N^*(b)$ as a function of sender bias b. Left: full range showing the step-function decrease with the continuous envelope $(-1+\sqrt{1+2/b})/2$. Right: detail near the babbling boundary. As b increases, the state space is partitioned into progressively fewer intervals until only one remains ($N^* = 1$, babbling) at $b \geq 1/4$.

4.2 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.

For each configuration, we measure:

1. Mutual information between ground truth and each layer's report: $I(\theta; a_k)$.
2. Effective partition count at each layer (number of distinct actions observed).
3. Information retention ratio: $I(\theta; a_k)/I(\theta; a_1)$.
4. Expected quadratic loss: $E[(a_k - \theta)^2]$.

4.3 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 each bias level (Figure 3). At $b = 0.01$, the channel preserves $N^* = 7$ partitions through 10 layers; at $b = 0.1$, it preserves $N^* = 2$ partitions throughout. This occurs because the noiseless Crawford–Sobel channel maps each input interval to a single output action deterministically: once the state has been quantized, re-quantization through the same partition structure does not further degrade it.

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 (Figure 4). 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: from 1.77 to 1.60 bits.

Finding 3: Phase diagram reveals informative vs. blind regimes. The phase diagram (Figure 5) 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.

The transition between regimes sharpens with noise and deepens with hierarchy depth.

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 ≈ 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.

5 Operationalizing Dysmemic Pressure

The simulation enables a precise operationalization of dysmemic pressure as a measurable quantity.

Definition 5.1 (Dysmemic Pressure Index)**.** The *Dysmemic Pressure Index* at layer k of an organizational hierarchy is:

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

where θ is the ground truth, report_k is the information available at layer k, and report_0 is direct observation (the maximally informative baseline).

The index has clean boundary behavior:

- $\mathrm{DP}(k) = 0$: no information loss at layer k relative to direct observation. The hierarchy is transparent.
- $\mathrm{DP}(k) = 1$: complete information destruction. The report at layer k is independent of ground truth. The hierarchy is blind.
- Intermediate values measure the fractional information loss.

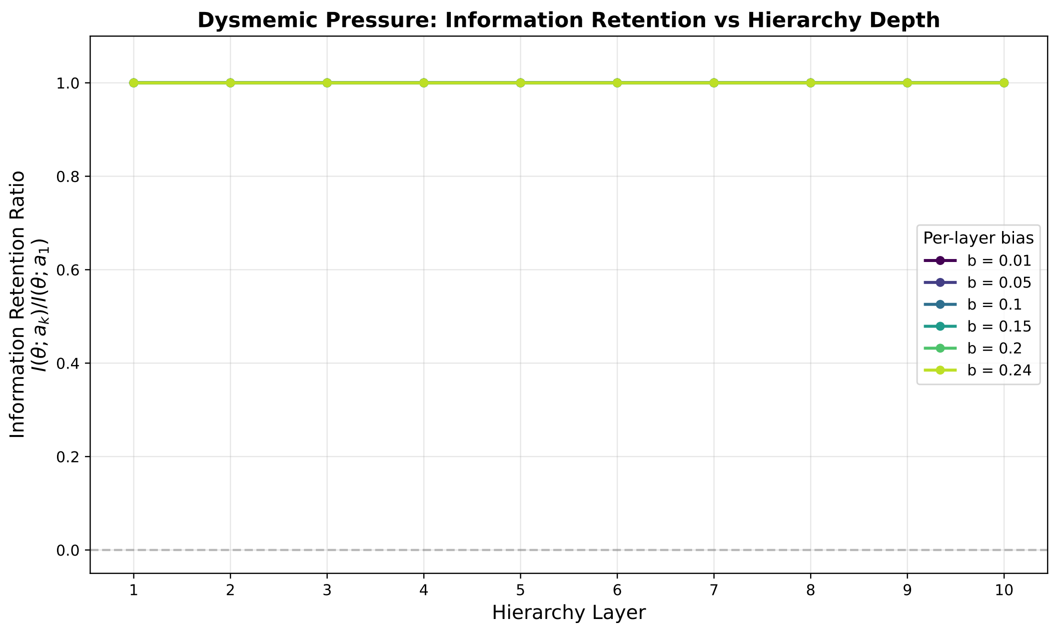

Figure 3: Information retention ratio $I(\theta; a_k)/I(\theta; a_1)$ vs. hierarchy layer for the 10-layer hierarchy with zero noise ($\sigma^2 = 0$). Each curve represents a different per-layer bias b. Without noise, partition structure is preserved and retention remains at 1.0 across all depths. Error bands show ± 1 standard deviation over 50 Monte Carlo iterations.

By the data processing inequality (4), DP(k) is monotonically non-decreasing in k: dysmemic pressure can only accumulate through a hierarchy, never reverse.

Remark 5.2 (Practical measurement)**.** Computing DP(k) requires access to ground truth θ, 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. The index is most directly applicable in domains where outcomes eventually become observable (product launches, financial projections, risk assessments) and less applicable where ground truth remains contested.

From the simulation data, we can compute DP(k) for each configuration. For example, in a 5-layer hierarchy with $b = 0.1$ and $\sigma^2 = 0.01$:

- DP(1) $= 0$ (baseline layer).
- DP(2) ≈ 0.035 (3.5% information loss).
- DP(5) ≈ 0.060 (6.0% information loss at the top).

At $b = 0.24$ with $\sigma^2 = 0.05$, the top-layer index reaches DP(5) ≈ 0.10, indicating substantial information destruction even in a moderately deep hierarchy.

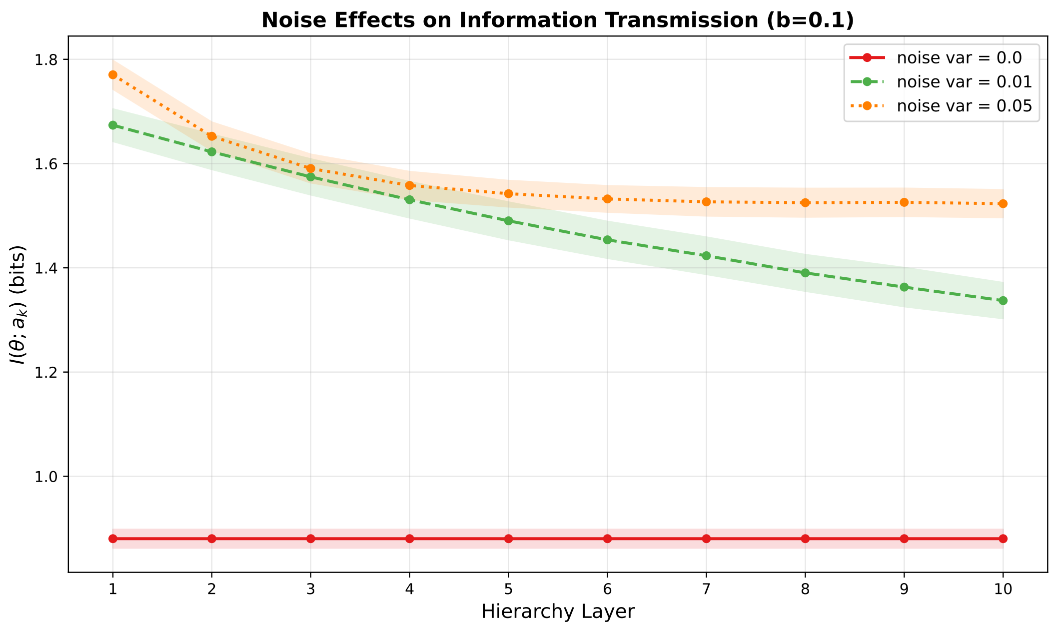

Figure 4: Effect of observation noise on information transmission ($b = 0.1$, 10-layer hierarchy). Without noise, MI remains constant. With noise ($\sigma^2 = 0.01$ or 0.05), MI degrades monotonically through the hierarchy, consistent with the data processing inequality. Each additional layer irreversibly destroys information.

6 Case Studies

Three cases illustrate dysmemic pressure operating across different industries, failure modes, and organizational contexts. Each demonstrates the mechanism's key features: preference divergence generating communication degradation, adverse selection in idea markets, transmission biases locking in dysfunction, and the self-reinforcing nature of the resulting equilibrium.

6.1 Boeing 737 MAX: Engineering Truth vs. Certification Fitness

The Boeing 737 MAX crashes—Lion Air Flight 610 (October 2018) and Ethiopian Airlines Flight 302 (March 2019)—killed 346 people. The proximate cause was the Maneuvering Characteristics Augmentation System (MCAS), a flight-control system that repeatedly pushed the aircraft's nose down based on a single angle-of-attack sensor, without adequate pilot awareness or override capability (Robinson and Shay, 2019; House Committee on Transportation and Infrastructure, 2020).

The information existed within Boeing. Engineers had identified MCAS as a single-point-of-failure risk. Test pilots had experienced uncommanded nose-down events during simulation. Internal communications, revealed during congressional investigation, show engineers expressing alarm: one wrote that he had "basically lied to the regulators (unknowingly)" about MCAS (House Committee on Transportation and Infrastructure, 2020). The House Transportation Committee's final report documented systematic suppression of safety

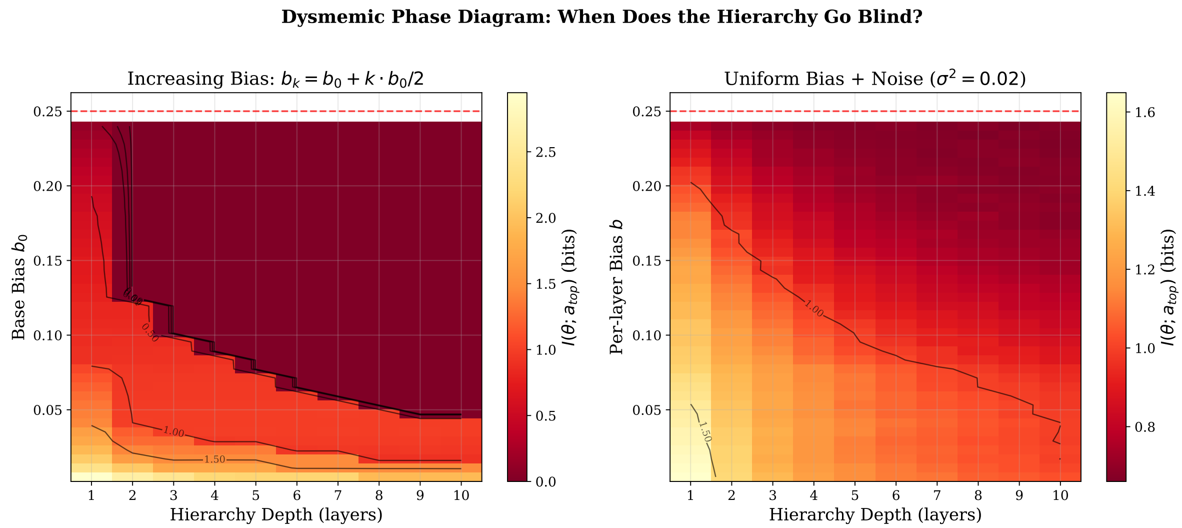

Figure 5: Phase diagram of hierarchical information transmission. Color indicates $I(\theta; a_{\text{top}})$ in bits. Left: increasing bias model ($b_k = b_0 + k \cdot b_0/2$). Right: uniform bias with noise ($\sigma^2 = 0.02$). The red dashed line marks the babbling boundary ($b = 0.25$). Contour lines show iso-information curves. Deep hierarchies with moderate bias can be as informationally impoverished as shallow hierarchies near babbling.

concerns throughout the development process.

Strategic communication degradation. The preference divergence at Boeing was structural. 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 if it requires additional development time and cost) diverged sharply from management's preferences (certify the aircraft quickly, minimize pilot training differences from the 737 NG to preserve the competitive position against the Airbus A320neo).

In Crawford–Sobel terms, the bias b between engineering assessment and management reception was large. Engineers who reported MCAS risks accurately faced schedule pressure, career friction, and the implication that they were obstructing the program. Management reframed engineering concerns as schedule issues: the question shifted from "Is this safe?" to "Can we certify on time?" The partition coarsened from a nuanced risk assessment to a binary: certified or not certified.

Adverse selection. The information asymmetry between Boeing and the FAA was exploited through the Organization Designation Authorization (ODA) program, which allowed Boeing employees to conduct certifications on the FAA's behalf. The employees performing certification faced the same internal selection pressures as other Boeing staff—their careers depended on the same management that needed certification to proceed. 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").

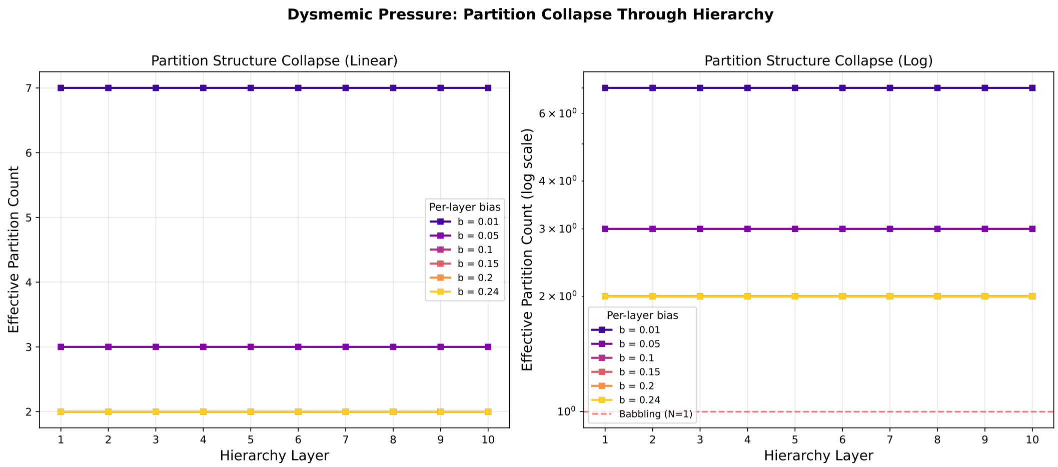

Figure 6: Effective partition count vs. hierarchy depth for the 10-layer hierarchy (zero noise). Left: linear scale. Right: log scale with the babbling threshold ($N = 1$) marked. Low-bias channels ($b = 0.01$) maintain high partition counts through all layers; high-bias channels ($b = 0.24$) collapse to 2 effective partitions immediately.

Transmission bias. Prestige bias operated through Boeing's leadership hierarchy: the CEO and board's focus on stock price and delivery schedule set the prestige frame, and assessments aligned with those priorities propagated downward. Conformity bias locked in the launch-oriented culture—once the 737 MAX program was committed, deviation from the schedule narrative marked the deviator. Content bias favored the simple narrative ("evolutionary upgrade to a proven design") over the complex truth ("novel flight-control system compensating for aerodynamic changes introduces new failure modes requiring extensive testing").

Mapping to the model. Ground truth: the MCAS design was unsafe without adequate safeguards. This information was compressed through the organizational hierarchy—engineer to lead, lead to program manager, program manager to division head, division head to CEO, CEO to board—with each layer's bias parameter oriented toward schedule and certification. The report that reached the top was "certified safe." The dysmemic pressure index was near 1: virtually complete information destruction between the engineering assessment and the board-level representation.

6.2 Theranos: Selection for Silence

Between 2003 and 2018, Theranos claimed to have developed technology for comprehensive blood testing from a single finger prick. The company reached a peak valuation of approximately $9 billion. The technology did not work (Carreyrou, 2018).

Strategic communication degradation. Elizabeth Holmes created an information environment with extreme preference divergence. Engineers who raised concerns about test accuracy—the core product's fundamental capability—were fired, threatened with legal ac-

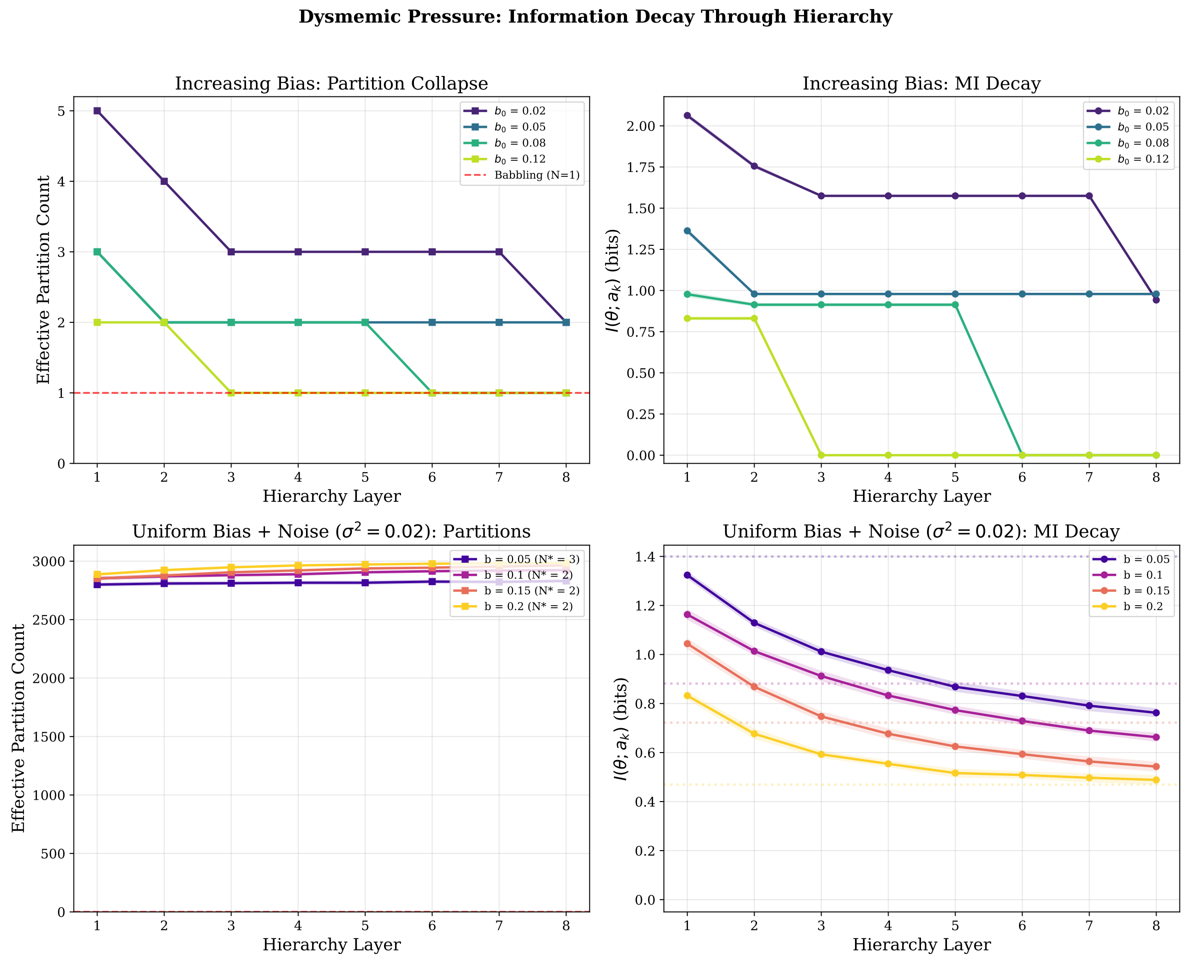

Figure 7: Simulated partition counts and MI decay through 8-layer hierarchies. Top row: increasing bias schedule ($b_k = b_0 + k \cdot 0.4 b_0$) without noise. Bottom row: uniform bias with observation noise ($\sigma^2 = 0.02$). In both cases, information is progressively destroyed as it moves up the hierarchy, with the rate of destruction depending on the bias-to-noise ratio.

tion, or marginalized (Carreyrou, 2018). The Crawford–Sobel bias parameter b was pushed to its maximum: the cost of truthful reporting was job loss, and the benefit of optimistic reporting was continued employment and stock option value.

The communication channel between laboratory reality and executive representation was in permanent babbling equilibrium. Laboratory directors who reported that the Edison device produced inaccurate results were replaced. Their replacements learned the lesson. The rational strategy for any laboratory employee was to report what Holmes wanted to hear, regardless of what the data showed.

Adverse selection. The idea market within Theranos was comprehensively captured. Accurate signals about test performance were costly to produce (requiring rigorous validation) and career-destroying to transmit. Optimistic signals were cheap (requiring only compliance with the official narrative) and career-sustaining. Over time, employees capable of producing accurate signals left or were expelled. The remaining workforce was selected for willingness to produce signals aligned with the official narrative—a textbook adverse selection outcome (Akerlof, 1970).

Transmission bias. Holmes cultivated extraordinary prestige bias. Her board included Henry Kissinger, George Shultz, James Mattis, and other luminaries whose endorsement cascaded through the investment and media ecosystem (Carreyrou, 2018). The prestige of Theranos's directors made skeptical signals about the technology's viability low-status relative to enthusiastic signals. Conformity bias operated through the investor community: once major names invested, questioning the technology meant questioning the judgment of respected figures.

Mapping to the model. Ground truth: the Edison blood-testing device could not meet clinical accuracy standards. Selection dynamics eliminated accurate signals from the internal information environment. Engineers who knew the truth were systematically removed. The remaining information population consisted entirely of dysmemes—signals optimized for internal fitness (continued employment, funding) rather than correspondence with external reality (test accuracy). The dysmemic pressure index reached 1: complete information destruction. Only external intervention—investigative journalism by Carreyrou (2018), followed by regulatory and legal action—penetrated the information environment.

6.3 Wirecard: When External Signals Cannot Penetrate

Wirecard AG, a German payment processor and financial services company, collapsed in June 2020 when auditors could not confirm the existence of €1.9 billion in cash balances purportedly held in Philippine bank accounts. The fraud had been ongoing for years, and multiple external signals had pointed to it (McCrum, 2022).

Strategic communication degradation. Within Wirecard, the preference divergence was between the executives perpetrating the fraud (who needed information environments that confirmed the company's legitimacy) and employees, auditors, and regulators (who nominally needed accurate financial information). Wirecard's management created internal communication channels that systematically suppressed disconfirming information. Employees who questioned the legitimacy of certain business units or partner relationships were reassigned or dismissed.

Adverse selection. The adverse selection dynamic at Wirecard extended beyond the firm to its external ecosystem. Ernst & Young (EY), the external auditor, faced an information asymmetry: verifying Wirecard's claims required independent confirmation from third-party banks in Southeast Asia, which was costly and logistically difficult. For over a decade, EY relied on confirmations provided by Wirecard itself rather than independently verifying them (McCrum, 2022). The cheap signal (Wirecard's self-produced documentation) crowded out the expensive signal (independent verification).

External signal failure. What makes the Wirecard case distinctive is that external signals *existed* but could not penetrate the dysmemic equilibrium. Short sellers (including Muddy Waters and individual analysts) publicly documented evidence of fraud beginning in 2008. Dan McCrum of the *Financial Times* published a series of investigative articles beginning in 2015 detailing accounting irregularities (McCrum, 2022). Rather than investigating these external signals, the German financial regulator BaFin filed a criminal complaint against the *Financial Times* journalists and short sellers for alleged market manipulation.

This is dysmemic pressure operating at the ecosystem level. The internal selection environment—where Wirecard's management controlled the information—extended outward through regulatory capture and legal intimidation. Accurate signals produced outside the organization were actively suppressed by institutions whose own selection environments favored accepting Wirecard's narrative (BaFin avoiding the embarrassment of a major German company being fraudulent; EY avoiding the cost and conflict of aggressive auditing).

Mapping to the model. Ground truth: €1.9 billion in reported assets did not exist. The internal organizational hierarchy operated near complete babbling—management controlled the information environment absolutely. But uniquely, the external information environment was also captured: regulators, auditors, and investors faced their own dysmemic pressures that aligned with accepting the internal narrative. Only persistent external signals from actors with different fitness landscapes—short sellers (who profited from truth) and investigative journalists (whose professional fitness rewarded exposure)—eventually forced the constructed reality to confront external reality. This case illustrates that dysmemic pressure is not confined to single organizations but can cascade through institutional ecosystems.

7 Connection to Forecasting Literature

The dysmemic pressure framework connects directly to research on expert judgment and forecasting, particularly the work of Philip Tetlock and collaborators.

7.1 Foxes, Hedgehogs, and Organizational Architecture

Tetlock (2005) demonstrated that expert political judgment is, on average, only slightly better than chance—but with enormous variance. He identified two cognitive styles among forecasters: "hedgehogs" (who organize their thinking around a single grand theory) and "foxes" (who draw on multiple frameworks, update frequently, and are comfortable with ambiguity). Foxes significantly outperformed hedgehogs.

In the dysmemic pressure framework, hedgehog organizations exhibit stronger dysmemic pressure. A hedgehog organization commits to a single strategic narrative ("we are a smartphone company," "our technology works," "our business units are legitimate"). This narrative becomes the selection criterion for internal signals: information that confirms the narrative is fit; information that contradicts it is unfit. The partition structure coarsens around the narrative boundary—nuanced assessments of risk, performance, or competitive position are compressed into binary compatibility with the master theory.

Fox organizations—those that maintain multiple frameworks, tolerate ambiguity, and update beliefs incrementally—resist dysmemic pressure through structural pluralism. When no single narrative dominates, the selection environment does not select as strongly against any particular class of signals.

7.2 Superforecasting as Anti-Dysmemic Architecture

Tetlock and Gardner (2015) documented how superforecasting teams achieve remarkable predictive accuracy. The structural features that distinguish these teams map directly onto anti-dysmemic architecture:

- **Independent assessment before discussion.** Team members generate forecasts independently before sharing, preventing prestige and conformity bias from shaping initial estimates. In our model, this is equivalent to bypassing the hierarchical channel entirely for the initial signal.
- **Aggregation without compression.** Superforecasting teams aggregate diverse perspectives through structured processes (e.g., extremizing algorithms) rather than compressing them into consensus through social dynamics. This preserves partition structure rather than coarsening it.
- **Scoring and accountability.** Brier scores and similar proper scoring rules create incentive alignment between forecasters and accuracy. This reduces the Crawford–Sobel bias parameter b toward zero, increasing the equilibrium partition count and information transmission.
- **Frequent updating.** Superforecasters revise estimates frequently in response to new evidence, preventing the conformity lock-in that stabilizes dysmemic equilibria.
- **Cognitive diversity.** Teams include members with different backgrounds and analytical frameworks, reducing the probability that prestige bias amplifies a single perspective.

In the hierarchical model, superforecasting teams implement what amounts to a parallel architecture with low bias—multiple independent channels feeding into a well-designed aggregator—rather than a serial architecture with cascading bias. The data processing inequality (4) tells us that serial processing can only lose information; the superforecasting architecture avoids serial processing precisely where it would be most destructive.

7.3 Prediction Markets and Evaluation Decoupling

Hanson (2003) proposed prediction markets as mechanisms for eliciting private beliefs. In the dysmemic framework, prediction markets work because they decouple evaluation from information. A trader in a prediction market is rewarded for accuracy, not for agreeing with their manager. The Crawford–Sobel bias b is structurally reduced because the payoff function rewards correspondence with reality rather than correspondence with the receiver's preferences.

The forecasting literature thus provides empirical evidence for a key claim of the dysmemic pressure framework: organizations can be designed to resist information degradation, but only through structural changes to the selection environment, not through exhortation or culture change.

8 The Compound Pressure and the Cage

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.

Definition 8.1 (Dysmemic pressure)**.** Dysmemic pressure is the compound selection force in an organization that favors cultural variants (ideas, signals, practices) whose internal fitness is negatively correlated with external accuracy. Its intensity is a function of preference divergence b, verification cost c, and transmission ease τ:

$$\text{DP} \propto f(b, c, \tau). \tag{6}$$

This is a heuristic summarizing the mechanism's structure: pressure intensifies as any component increases.

A *dysmeme* is a signal optimized for the gap between representation and reality. Its survival depends on alignment with receiver preferences, ease of transmission, association with prestige, conformity with apparent consensus, and defensibility under scrutiny. It may be true, false, or somewhere between. Truth is not the selection criterion.

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

8.1 The Cage

When compression and selection interact long enough, the organization reaches a stable state where constructed reality, signal population, and selection criteria become mutually consistent. The constructed reality generates signals that reinforce the selection criteria that generated the constructed reality.

This equilibrium is the *Cage*. It resists displacement because:

- People have adapted to the current landscape (careers, skills, relationships, identities depend on it).
- Constructed reality confirms itself (the frame filters evidence about the frame).
- Coordination has locked in (unilateral deviation is punished).

The Cage is not a conspiracy. No one designed it. It emerges from the interaction of compression and selection operating over time.

8.2 Why Reform Fails

Reform is a signal. Signals face selection. A reform proposal must survive the environment it aims to change. Decision-makers evaluate it using the current frame. Reform that would actually displace the equilibrium threatens everyone invested in it and faces maximum resistance. Reform that changes vocabulary while leaving fitness criteria intact faces minimal resistance. Selection favors the latter.

A frequently cited estimate puts organizational change failure rates between 60 and 80 percent (Beer and Nohria, 2000). The rate makes sense: the initiatives target behavior without targeting the selection environment that produces the behavior.

9 Testable Propositions

The mechanism generates specific, testable predictions.

Proposition 9.1 (Preference divergence)**.** *Communication accuracy should correlate with incentive alignment. Layers with aligned incentives should transmit information more accurately than layers with divergent incentives. The information retention ratio* $I(\theta; a_k)/I(\theta; a_1)$ *should decrease as the bias parameter* b_k *increases.*

This follows directly from the Crawford–Sobel model and is confirmed by the simulation: at $b = 0.01$, the 10-layer hierarchy retains 100% of Layer 1 information (without noise); at $b = 0.24$, the single-layer channel transmits only 0.14 bits.

Proposition 9.2 (Evaluation coupling)**.** *When the recipient of information is also responsible for evaluating the sender, information quality decreases. Organizations that structurally separate "who needs to know" from "who controls your career" will exhibit less dysmemic pressure.*

Proposition 9.3 (Process capture)**.** *Organizational processes whose outputs are used to evaluate participants will, over time, optimize for evaluation success rather than process purpose. Dysfunction should cluster in unmeasured dimensions.*

Proposition 9.4 (Intervention decay)**.** *Interventions that change expressed norms without changing payoff structures will exhibit initial improvement followed by regression to the pre-intervention equilibrium. The rate of regression should correlate with the strength of unchanged selection pressure.*

Proposition 9.5 (External correction)**.** *Organizations under strong dysmemic pressure can only be corrected by external shock—information or consequences from outside the selection environment. Internal reform attempts will be absorbed into the dysmemic equilibrium.*

The Wirecard case provides striking support for Proposition 9.5: internal information was entirely captured, and only external actors with different fitness landscapes (short sellers, journalists) forced correction.

10 Structural Countermeasures

If dysmemic pressure is structural, effective countermeasures must alter the selection environment rather than exhort different behavior within it.

Evaluation decoupling. Separating the recipient of information from the evaluator of its source reduces bias in the Crawford–Sobel sense. Independent audit functions that report to boards rather than management, ombudsman offices with protected status, and anonymous reporting channels with credible confidentiality provide this separation. Boeing's ODA program violated this principle by having Boeing employees certify Boeing's own aircraft—evaluation and information production were maximally coupled.

External validation. Requiring external review of key assessments introduces information from outside the internal selection environment. The Wirecard case demonstrates both the necessity and difficulty of external validation: EY's audits failed because the auditor accepted internally produced documentation rather than independently verifying it.

Prediction tracking and proper scoring rules. Formal tracking of predictions against outcomes makes accuracy visible over time. When forecasters must own their predictions, adverse selection pressure on idea markets decreases. Proper scoring rules reward accurate probability assessments regardless of what the predictor wanted to be true, directly reducing the bias parameter b.

Anti-dysmemic architecture. Drawing on the superforecasting literature (Section 7), organizations can implement:

- Independent assessment before aggregation (reducing prestige and conformity bias).
- Parallel information channels rather than serial hierarchies (avoiding cascading information loss).
- Explicit scoring and accountability for prediction accuracy.
- Protected channels for dissent that bypass hierarchical compression.

The maintenance problem. Any structure that counterweights dysmemic pressure faces continuous pressure toward absorption. The red team that becomes too influential will be defunded. The independent audit that creates too much friction will see its mandate narrowed. Durable countermeasures require governance independence (reporting lines that bypass those being assessed), resource independence (budgets protected from retaliation), and evaluation independence (career consequences not controlled by those under scrutiny).

11 Limitations

The framework has important limitations that bound its applicability. We state them explicitly.

1. **Crawford–Sobel assumptions.** The CS model assumes rational agents with a known, fixed bias parameter b. Real organizations have unknown, heterogeneous, and time-varying biases. Agents may not know their own bias, and the bias may depend on the state being communicated (e.g., bad news triggers higher bias than good news). The uniform-quadratic specification is analytically convenient but not the only possible preference structure.

2. **Simplified hierarchy.** The agent-based model uses a strict serial hierarchy: information flows upward through a linear chain. Real organizations have matrix structures, informal networks, skip-level relationships, and lateral information flows that can partially bypass hierarchical compression. The model captures the worst case (pure serial degradation) but may overstate information loss in organizations with redundant channels.

3. **Case studies are illustrative, not controlled.** The Boeing, Theranos, and Wirecard cases are selected to illustrate the mechanism, not to test it. They are subject to hindsight bias, narrative selection, and the absence of counterfactual comparison. Rigorous testing requires controlled or quasi-experimental designs that are difficult to implement in organizational settings.

4. **Ground truth access.** The operationalization of DP(k) requires access to ground truth θ, which is often unavailable in real organizations until after the fact (when outcomes are observed). The index is most useful retrospectively or in domains where outcomes materialize on known timescales.

5. **Quantification of transmission bias.** The second and third dynamics (adverse selection and transmission bias) are characterized qualitatively rather than formally modeled in the simulation. Integrating these dynamics into the agent-based framework—e.g., modeling heterogeneous agent types who differ in signal production cost, or modeling network effects on belief propagation—remains future work.

6. **Scale and generalizability.** The mechanism operates most powerfully at scale. Small organizations with direct observation and tight feedback loops may not develop strong dysmemic pressure. The framework applies primarily to organizations large enough that information must flow through multiple nodes and slow enough that consequences are temporally distant from actions.

12 Conclusion

Organizations fail because they select against the information that would save them. The selection is structural. Compression creates gaps between representation and reality. Gaps become selection environments. What survives is what fits—regardless of correspondence with reality.

Three dynamics drive the pressure. Strategic communication degrades as preferences diverge—formalized here through the Crawford–Sobel partition theorem and validated by agent-based simulation showing that information can only be lost, never gained, as it propagates through an organizational hierarchy. Adverse selection floods idea markets with cheap signals. Transmission biases favor spread over accuracy.

The simulation demonstrates that the informational consequences are severe and quantifiable. A parameter sweep over 126 configurations reveals a sharp transition between informative and blind regimes in the hierarchy's phase space. Without noise, the partition structure is preserved—a theoretical limit that no real organization achieves. With even modest noise (representing the real-world imperfections of communication, interpretation, and memory), mutual information degrades monotonically through each hierarchical layer.

The Dysmemic Pressure Index $\mathrm{DP}(k) = 1 - I(\mathrm{report}_k, \theta)/I(\mathrm{report}_0, \theta)$ provides a measurable operationalization: a number between 0 and 1 indicating the fraction of ground-truth information destroyed at each layer. This transforms dysmemic pressure from a metaphor into a quantity.

The case studies—Boeing 737 MAX, Theranos, Wirecard—illustrate the mechanism operating in engineering, biotech, and financial services. In each case, the organization possessed the information that would have saved it. In each case, the selection environment destroyed that information before it could reach decision-makers in usable form.

The connection to forecasting research provides a constructive complement to the diagnosis. Superforecasting teams implement structural features—independent assessment, aggregation without compression, proper scoring rules, frequent updating—that map directly onto anti-dysmemic architecture. These are not exhortations. They are engineering specifications.

Standard interventions fail because they treat equilibrium outputs as behavioral problems. Effective intervention requires altering the selection environment itself: decoupling evaluation from information, introducing external validation, tracking predictions against outcomes, protecting channels that bypass fitness pressure.

The question for any organization is not whether dysmemic pressure exists—it does, always, at scale. The question is whether anything counterweights it. Understanding the mechanism makes engineering possible. Building structures that offset it makes accuracy achievable. Maintaining those structures against erosion is the ongoing cost of clarity.

References

Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3):488–500.

Beer, M. and Nohria, N. (2000). Cracking the code of change. *Harvard Business Review*, 78(3):133–141.

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.

Boyd, R. and Richerson, P. J. (1985). *Culture and the Evolutionary Process*. University of Chicago Press.

Carreyrou, J. (2018). *Bad Blood: Secrets and Lies in a Silicon Valley Startup*. Alfred A. Knopf.

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.

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). Final committee report: The design, development and certification of the Boeing 737 MAX. Technical report, U.S. House of Representatives. 116th Congress.

McCrum, D. (2022). *Money Men: A Hot Startup, A Billion Dollar Fraud, A Fight for the Truth*. Bantam Press.

McEntire, J. (2025). Dysmemic pressure: Selection dynamics in organizational information environments. *arXiv preprint arXiv:2512.14716*. Also available at SSRN: `https://ssrn.com/abstract=6015814`.

Prendergast, C. (1993). A theory of "yes men". *American Economic Review*, 83(4):757–770.

Robinson, P. and Shay, G. (2019). What led to Boeing's 737 Max crisis. *Bloomberg Businessweek*.

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.

Useem, J. (2019). The long-forgotten flight that sent Boeing off course. *The Atlantic*.