The Cage: How Fiduciary Duty Creates Organizational Incompleteness

Jeremy McEntire
Cage & Mirror Press

March 2026

Abstract

Formalization enables scale but creates frames that compress reality into legible metrics and procedures. Under fiduciary law, boards must *demonstrate* prudence, pushing organizations toward decisions that are defensible inside those frames. The result is a predictable, regime-dependent tendency toward rigidity: variance in justifications compresses after formalization events; mechanisms that could supply external perspective are absorbed; and the pressure intensifies with legal exposure. This paper develops the mechanism, connects it to Delaware doctrine as an evidentiary "legal amplifier," derives falsifiable propositions, and tests them empirically against 75 SEC filings from 25 companies spanning three decades. The empirical analysis confirms that lexical diversity compresses 5.8–33% following IPO, with compression magnitude proportional to legal exposure and inversely proportional to founder control. We clarify the scope of the Gödel analogy (structural, not mathematical), separate universal tension from regime-dependent intensification, and address competing explanations. Boundary conditions specify when the tendency becomes a trap (high formalization, high exposure, long lock-in) and when it weakens (founder control, fast feedback). A companion paper, *The Mirror*, shows how organizations can remain legally defensible while preserving adaptation through meta-compliance. Our concluding claim is modest and testable: where demonstrable soundness dominates evaluation, organizations predictably trade adaptation capacity for defensibility unless countervailing structures are continuously protected.

1 Introduction

For more than fifty years, organizational scholars have documented a persistent pattern. Organizations converge on dominant practices and struggle to deviate. Successful companies fail when their industries face disruption. Bureaucracy intensifies with age and size. Strategies that worked become constraints that bind. The phenomenon appears across contexts—manufacturing and services, technology and healthcare, startups and conglomerates—and persists despite waves of reform.

The proposed solutions show equally consistent convergence. Create structural separation between exploration and exploitation [Tushman and O'Reilly, 1996]. Establish skunkworks isolated

from mainstream operations [Christensen, 1997]. Institute red teams to challenge prevailing assumptions [Nemeth, 1986, Schulz-Hardt et al., 2006]. Hire external consultants to provide fresh perspective [Werr and Styhre, 2003, Sturdy, 2011]. Diversify boards to bring different viewpoints [Adams and Ferreira, 2009]. Train leaders in double-loop learning to question frames rather than merely optimize within them [Argyris and Schön, 1978]. The recommendations recur because the problem recurs.

Yet the solutions prove temporary. Skunkworks that initially produce breakthrough innovations gradually adopt parent company metrics and revert to incremental output. Red teams that begin with authority to challenge decisions find their scope narrowed through "coordination requirements." External consultants who arrive with independent frameworks learn which recommendations clients will accept and which create tension. Board diversity measured by demographics shows weak or null effects on performance because demographic difference does not guarantee cognitive difference. Double-loop learning remains rare in practice despite decades of advocacy and training.

The pattern suggests that rigidity is not an accident—not the result of poor leadership, weak culture, or misaligned incentives—but a structural consequence of something deeper. Organizations do not become rigid because they fail to implement solutions. They become rigid, implement solutions, watch the solutions decay, and face the same rigidity again. The cycle repeats across generations, industries, and governance models. This persistence indicates mechanism, not mistake.

1.1 The Theoretical Gap

The literature has identified the pattern without explaining its structural necessity. March [1991] formalizes the drift from exploration to exploitation but does not explain why formalization makes this drift inevitable rather than merely common. Weber [1958] describes rationalization becoming self-perpetuating but offers only charismatic disruption as escape. Christensen [1997] documents incumbent failure during disruption but treats skunkworks as best practice rather than structural requirement emerging from incompleteness. Agency theory [Jensen and Meckling, 1976] explains incentive misalignment but not why the misalignment produces frame-dependent blind spots. Organizational learning literature [Argyris and Schön, 1978] distinguishes single-loop from double-loop but does not account for why organizations under formalization systematically select against double-loop even when leaders understand its value.

Three puzzles remain unresolved. First, why do measurement problems recur across literatures attempting to study external perspective, ambidexterity, or cognitive diversity? Each literature confronts variants of the same challenge: capturing perspective-taking from within a research frame encounters the same difficulties organizations face when validating insights from outside their operational frames.

Second, why do solutions based on structural separation work when they work, yet fail so predictably? Some skunkworks produce category-creating innovations while others become expensive distractions. The solutions are structurally similar—all create external perspective through separation—yet outcomes diverge. What determines when separation enables adaptation versus

when it becomes performance theater?

Third, why does the problem intensify with organizational scale, age, and success? The very organizations with the most resources to invest in adaptation are the most vulnerable to rigidity. Every attempt to maintain agility at scale fails in similar ways, yet we lack theory explaining why scale should structurally produce rigidity independent of leadership quality, cultural strength, or competitive pressure.

1.2 The Contribution

This paper argues that these three puzzles derive from a single mechanism: organizations under formalization exhibit incompleteness-like properties analogous to Gödel's Incompleteness Theorem. Just as formal logical systems powerful enough to express arithmetic contain true statements unprovable within the system's rules, formalized organizations operating under frame-dependent metrics create strategic truths that cannot be validated using the frame's internal logic. The analogy is structural, not mathematical.

This framework unifies findings across economics, sociology, strategy, and organizational behavior that previous syntheses have treated as separate phenomena. March's exploration-exploitation trade-off, Weber's Iron Cage, Christensen's Innovator's Dilemma, and organizational ambidexterity research are not independent observations. They are manifestations of the same structure: formalization creates frames, frames create blind spots, and those blind spots require external perspective for validation.

The framework further shows that this mechanism intensifies at scale by legal design. Corporate law, through fiduciary duty and the business judgment rule, requires that strategic decisions be demonstrably sound. Demonstrable soundness demands formalization. But formalization is precisely what creates incompleteness. The legal structure creates the problem and constrains the solution. This is the fiduciary trap.

We develop three testable propositions that avoid the circularity plaguing prior research. Rather than measuring whether external perspective produces better outcomes, we measure structural properties. Proposition 1 predicts that formalization events create measurable discontinuities in decision variance. Proposition 2 predicts that units with independent evaluation criteria produce more category-creating innovations. Proposition 3 predicts that dissent mechanisms with authority to raise concerns without approval change more strategic decisions. Each proposition includes falsification criteria and boundary conditions. Section 6 tests Proposition 1 against 75 SEC filings from 25 companies, confirming variance compression across all cases studied.

1.3 Scope and Limitations

This framework applies to organizations operating under conditions of high formalization and high legal exposure. It predicts strong effects for public companies in regulated industries with dispersed shareholders and long decision lock-in periods. It predicts weak or null effects for small private

companies with concentrated ownership, fast feedback loops, and low liability exposure. The model is regime-dependent by design.

We do not claim that external perspective guarantees success or that formalization guarantees failure. The framework describes a forced trade-off, not a normative recommendation. Most organizations rationally choose formalization because its costs are deferred while the costs of maintaining external perspective—legal risk, coordination friction, board discomfort—are immediate. Understanding the trade-off enables conscious choice rather than unconscious drift.

1.4 Roadmap

Section 2 reviews the convergent findings across organizational literatures. Section 3 introduces the structural analogy to Gödel's Incompleteness Theorem and defines what "incompleteness-like properties" mean operationally. Section 4 derives three testable propositions and specifies the Geometry of Defensibility. Section 5 explains how corporate law amplifies the problem at scale. Section 6 presents empirical evidence from SEC filings testing Proposition 1. Section 7 discusses boundary conditions, moderators, solution decay, alternative explanations, and falsification criteria. Section 8 describes regression analysis in progress. Section 9 considers implications. Section 10 concludes.

2 Literature Review

2.1 The Convergence Pattern

Across domains, the empirical record shows that organizations converge on dominant frames. March [1991] formalizes the drift toward exploitation. Weber's analysis of rationalization describes how rule-based coordination becomes self-perpetuating [Weber, 1958]. Argyris and Schön [1978] distinguish single-loop from double-loop learning and observe the rarity of the latter. Hannan and Freeman [1984] document increasing structural inertia with age and size. Together these streams describe patterned rigidity but do not explain why convergence is structurally inevitable under formalization.

2.2 The Separation Solution

When adaptation does occur, scholars repeatedly find that separation helps. Tushman and O'Reilly [1996] show that ambidexterity requires structural separation. Christensen [1997] argues that skunkworks or separate units are often necessary. Adversarial processes improve decision quality by institutionalizing dissent [Nemeth, 1986, Schulz-Hardt et al., 2006]. Dynamic capabilities research highlights sensing and transforming as distinct from execution [Teece et al., 1997]. These literatures converge on separation as effective but treat it as best practice rather than structural necessity.

2.3 The Measurement Problem

Many streams report recurring empirical difficulties. Board diversity studies use demographic proxies that show small or mixed effects [Adams and Ferreira, 2009]. Consultant effectiveness evidence is mixed because "externality" collapses when consultants adopt client frames [Sturdy, 2011]. Ambidexterity research struggles to measure exploration without back-coding on innovation outcomes. Cognitive complexity studies risk circularity when language used to describe complexity is itself the dependent measure [Barr et al., 1992]. These challenges share a core feature: attempts to measure perspective-taking from within a perspective.

2.4 Scale and Legal Connection

Agency theory analyzes principal-agent misalignment [Jensen and Meckling, 1976]. Public choice extends self-interest models to political actors [Buchanan and Tullock, 1962]. Corporate law's business-judgment rule protects directors who can demonstrate informed, reasonable process [Smi, 1985]. Goodhart's Law warns that measures become targets [Goodhart, 1984, Strathern, 1997]. These accounts explain why formalization grows with scale and exposure but do not articulate the mechanism by which formalization produces frame-dependent blind spots.

2.5 Institutional Theory and the Cage

The institutional tradition provides the closest antecedent to the Cage framework, but the relationship is one of subsumption rather than overlap. DiMaggio and Powell [1983] identify three pressures—coercive (regulatory), mimetic (copying peers), and normative (professional standards)—that drive organizations toward isomorphism. Their "iron cage" metaphor describes the observable pattern: organizations in the same field come to resemble one another. The Cage framework identifies the *mechanism* producing this pattern. Formalization creates shared frames; shared frames create convergent blind spots; legal amplification makes frame-dependent reasoning binding. Isomorphism is the symptom. Frame-dependent optimization under fiduciary duty is the cause.

Meyer and Rowan [1977] show that organizations adopt formal structures as "myths and ceremonies"—rational-looking practices adopted for legitimacy rather than technical efficiency. This ceremonial conformity is a special case of frame-dependent optimization: organizations adopt the frame's visible markers (documented processes, standard metrics, industry-benchmarked governance) because those markers constitute demonstrable soundness under fiduciary review. The Cage specifies *why* ceremonial conformity is rational—it satisfies evidentiary standards—and *when* it becomes binding—when legal exposure makes non-conformity costly.

Scott [2013] synthesizes institutional theory around three pillars: regulative (laws and enforcement), normative (professional obligations and standards), and cultural-cognitive (shared frameworks of meaning). These three pillars map onto the three dimensions of the Cage. The regulative pillar corresponds to the legal amplifier (Section 5). The normative pillar corresponds to the formalization pressure that creates shared metrics and procedures. The cultural-cognitive pillar cor-

responds to the frame-dependence that makes blind spots invisible from within. Scott's framework describes institutional structure; the Cage explains why that structure creates incompleteness-like properties and why it resists reform.

2.6 Synthesis

The literature documents (i) convergence, (ii) separation as effective, (iii) persistent measurement problems, (iv) the intensifying role of scale and law, and (v) institutional pressures toward isomorphism and ceremonial conformity. What is missing is a unifying mechanism that renders these patterns necessary under formalization and explains why institutional isomorphism produces not just homogeneity but structural rigidity. Section 3 develops that mechanism as a structural analogy to Gödel's incompleteness.

3 Theoretical Framework: Gödel Analogy

3.1 Gödel's Incompleteness Theorem

Gödel (1931) proved that any formal system powerful enough to express basic arithmetic contains true statements that cannot be proven using only the system's rules [Gödel, 1931]. Formal systems have three properties: (1) explicit rules, (2) mechanical application, and (3) exhaustive specification. For systems meeting these criteria, consistency and completeness cannot both be achieved.

An accessible analogue is a judge confined to codified statutes. Some cases present truths—"this outcome is just"—that cannot be reached by statute alone. Principles external to the code make the answer evident; the code cannot derive it.

Modern organizations, especially at scale and under legal exposure, approximate formal systems: policies supply explicit rules; process compliance supplies mechanical application; dashboards and KPIs seek exhaustive specification. As formalization advances, organizations inherit incompleteness-like properties.

3.2 Explicit Caveat

This is a structural analogy, not a claim of mathematical identity. Organizations lack complete axiomatization and decidable propositions. The claim is narrower: formalization produces incompleteness-like properties—frame dependence, blind spots, self-defense—that the analogy predicts. Its value is predictive utility [Bunge, 1977, Hesse, 1966].

3.3 Organizations as Formalized Systems

Table 1 maps formal system components to organizational equivalents.

3.4 Incompleteness-Like Properties (Operational)

Organizations exhibit incompleteness-like properties when four conditions hold:

Table 1: Organizational Formalization as Structural Analogue to Formal Systems

Formal system	**Organizational equivalent**	**Mechanism**	**Example**
Explicit rules (axioms)	Policies, procedures	Written constraints	Employee handbook, approval matrices
Inference rules	Decision protocols	Defined paths	Investment criteria, hiring rubrics
Mechanical application	Process compliance	Executable by any trained operator	SOPs, checklists, workflow software
Exhaustive specification	KPIs, dashboards	Everything measured or out-of-scope	OKRs, "what gets measured gets managed"
Consistency requirement	Demonstrable soundness	All decisions justifiable post hoc	Business-judgment rule, audit requirements
Inferred conclusions	Strategic choices	Derived from rules and data	"Enter market X because metrics Y"
Unprovable truths	Insights outside the frame	Value exists but cannot be validated	"This segment will matter" when KPIs do not show it
External observer	Different axioms	Operates outside the dominant frame	Consultant with independent charter; red team

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.

3. **Optimization deepens blindness.** Continuous improvement 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 from within.

3.5 Why the Analogy is Productive

The analogy explains four persistent puzzles: measurement problems (research frames mirror organizational frames), why separation works (different axioms permit validation), scale effects (more formalization, stronger frames), and solution decay (external mechanisms lose value when absorbed). It also yields falsifiable propositions. Section 4 specifies three.

4 Mechanism and Propositions

4.1 The Causal Chain

Growth raises coordination demands; public ownership raises legal exposure. Organizations respond by formalizing. Formalization constructs frames that compress rich realities into legible dashboards. Compression creates blind spots. Optimization deepens those blind spots. Frame-dependent reasoning self-defends. External perspective becomes necessary but legally risky. The system creates the problem and limits the solution.

From this mechanism follow three testable propositions.

4.2 Proposition 1: Variance Compression Thesis

Statement. Following formalization events, the variance of strategic decision justifications decreases measurably.

Derivation. If demonstrable soundness requires justification within formal frames, the range of acceptable justifications narrows. Pre-formalization, decisions invoke intuition, analogy, and non-standard evidence; post-formalization, they cite metrics, precedent, and documented analysis.

Formal statement. Let $V(t)$ denote variance in strategic decision justifications at time t, operationalized as (i) linguistic entropy in strategy documents, (ii) number of formally documented alternatives, and (iii) 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, large and dispersed shareholder bases, post-crisis regulatory environments; weakest in founder-controlled firms, small private firms, fast-feedback contexts, and cultures with documented psychological safety.

Measurement protocol. Sample strategic decisions pre- and post-IPO across a matched set of firms. Code linguistic entropy (lexical diversity and reasoning patterns), alternatives considered, and voting dispersion. Control for decision stakes, industry volatility, and macro conditions.

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

Empirical support. Evidence supporting this proposition appears in Section 6: lexical diversity compresses 5.8–33% across 25 companies following IPO, with the magnitude of compression proportional to legal exposure and inversely proportional to founder control.

4.2.1 Corollary 1.1: Born Caged

Statement. 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.

Derivation. If the legal and banking ecosystem internalizes the requirements of demonstrable soundness, the "cage" begins to operate before the firm is formally public. Legal counsel, underwriters, and auditors—themselves products of formalized institutions—impose frame-constrained language on the S-1, producing pre-compression through mimetic isomorphism [DiMaggio and Powell, 1983]. The formalization event is no longer a discrete boundary at IPO; it diffuses into the pre-IPO preparation process.

Formal statement. Let $V_{\text{S-1}}(c)$ denote the linguistic variance of S-1 filings for cohort c (defined by IPO era). For cohorts c_1 (pre-SOX, 1997–2002) and c_2 (post-2015):

$$\mathbb{E}[V_{\text{S-1}}(c_2)] < \mathbb{E}[V_{\text{S-1}}(c_1)]$$

controlling for industry and business-model novelty.

Empirical support. Section 6 confirms this pattern. Snowflake's 2020 S-1 (LD: 0.1495) and Okta's 2017 S-1 (LD: 0.1519) exhibit substantially lower initial variance than Amazon's 1997 S-1 (LD: 0.1742) or Google's 2004 S-1 (LD: 0.1825), despite comparable business-model novelty. The cage is already partially constructed before IPO in modern regulatory regimes.

4.3 Proposition 2: Independence–Novelty Thesis

Statement. Units that maintain independent evaluation criteria produce more category-creating innovations than units evaluated by parent metrics.

Derivation. Frames exclude opportunities that do not register on existing metrics. Independence—separate P&L, distinct KPI set, autonomous decision rights—constitutes different axioms, allowing validation of strategies the parent frame cannot justify.

Formal statement. Let N be the count of category-creating innovations over five years. Let $I \in [0, 1]$ be an independence score. Then:

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

Boundary conditions. Strongest in high-disruption industries and exploration-chartered units; weakest in low-disruption industries and exploitation-focused units.

Falsification. If high-independence units do not show higher category creation in high-disruption contexts, after controls, the frame-dependence mechanism is not supported.

4.4 Proposition 3: Authority–Impact Thesis

Statement. Dissent mechanisms with authority to raise concerns without prior approval change more strategic decisions than mechanisms that require approval to speak.

Derivation. If frames self-defend, dissent embedded in the same approval chains is filtered by the frame's criteria. Veto-free authority breaks the circularity.

Formal statement. 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:

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

Boundary conditions. Strongest for high-stakes, long-lock-in, high-complexity decisions; weakest for routine, reversible, short-feedback decisions.

Falsification. If veto-free mechanisms do not change more high-stakes decisions, after controls, the frame-defense account is not supported.

4.5 Interaction Structure

If the mechanism is common, organizations high in formalization and legal exposure should jointly exhibit variance compression, stronger independence–novelty effects, and larger authority–impact differentials. Lack of co-movement bounds the theory.

4.6 Geometry of Defensibility

A space defined by formalization density (horizontal) and legal exposure (vertical) produces four 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 effects.

5 The Legal Amplifier

5.1 Delaware's Evidentiary Logic

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* [Smi, 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 that 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.

5.2 From Effectiveness to Defensibility

The legal logic produces a systematic transformation. 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. The asymmetry is sharp. Rational directors converge on Path A regardless of which might produce better outcomes.

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 [SOX, 2002] and Dodd-Frank Act [Dod, 2010] formalized this trend, mandating additional controls, certifications, and audit trails.

5.3 Why This Is Not Failure

Corporate law's evolution toward demonstrable soundness is the rational consequence of applying fiduciary principles to complex decisions under uncertainty. Courts cannot know whether a decision was correct; they assess whether it was reasonable given what was knowable at the time. The side effect—systematic bias toward formalized, defensible decisions over potentially superior but less documentable choices—is the price we pay. It is trade-off, not failure.

5.4 Empirical Signatures

The legal amplification mechanism should leave measurable traces. Three patterns follow:

First, the **IPO discontinuity**: variance compression should exhibit a discontinuity at the IPO event, controlling for size and market conditions.

Second, **industry gradients**: high-exposure industries should show stronger variance compression than low-exposure industries.

Third, **temporal shocks**: major legislative events (SOX, Dodd-Frank) should create measurable shifts in decision variance.

These patterns are jointly diagnostic. Section 6 tests the first two predictions directly.

5.5 The Closing Trap

The sequence forms a closed loop. Scale and legal exposure mandate formalization. Formalization creates frames. Frames create blind spots. Blind spots require external perspective. Legal requirements prevent external perspective from being formalized. The system produces the problem and forecloses the solution.

6 Empirical Evidence: Variance Compression in SEC Filings

The theoretical framework predicts that formalization events produce measurable compression in the variance of organizational language. This section tests Proposition 1 against a dataset of 75 SEC filings from 25 companies, measuring linguistic variance across the S-1 registration statement and the first two annual reports (10-K) following IPO.

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 of Financial Condition and Results of Operations (MD&A)"—was extracted from each filing. This combined text corpus was processed using two metrics:

- **Lexical Diversity (LD):** The ratio of unique words to total words ($Unique\ Tokens / Total\ Tokens$). Lower values indicate more repetitive, homogeneous vocabulary.

- **Shannon Entropy (SE):** A measure of information uncertainty (base 2). Lower values indicate more predictable text, with a few terms dominating discourse.

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 control cases to establish a baseline for linguistic stasis absent a recent formalization event. Table 2 presents the full results.

	10-K Y1 (2000)	0.1511	11.14
	10-K Y2 (2001)	0.1495	11.09
Group 2: Post-SOX / Web 2.0 (IPO 2002–2012)			
Google (GOOGL)	S-1 (2004)	0.1825	12.05
	10-K Y1 (2005)	0.1503	11.24
	10-K Y2 (2006)	0.1466	11.10
Salesforce (CRM)	S-1 (2004)	0.1776	11.84
	10-K Y1 (2005)	0.1492	11.06
	10-K Y2 (2006)	0.1450	10.97
Netflix (NFLX)	S-1 (2002)	0.1650	11.33
	10-K Y1 (2003)	0.1421	10.78
	10-K Y2 (2004)	0.1399	10.71
LinkedIn (LNKD)	S-1 (2011)	0.1688	11.57
	10-K Y1 (2012)	0.1430	10.90
	10-K Y2 (2013)	0.1394	10.82
Workday (WDAY)	S-1 (2012)	0.1504	11.12
	10-K Y1 (2013)	0.1365	10.63
	10-K Y2 (2014)	0.1340	10.55
ServiceNow (NOW)	S-1 (2012)	0.1533	11.20
	10-K Y1 (2013)	0.1381	10.70
	10-K Y2 (2014)	0.1352	10.61
Group 3: Modern Cloud/Social (IPO 2012–2020)			
Meta (META)	S-1 (2012)	0.1709	11.81
	10-K Y1 (2013)	0.1610	11.50
	10-K Y2 (2014)	0.1575	11.42
Shopify (SHOP)	F-1 (2015)	0.1670	11.45
	40-F Y1 (2016)	0.1481	10.98
	40-F Y2 (2017)	0.1442	10.91
Snowflake (SNOW)	S-1 (2020)	0.1495	10.99
	10-K Y1 (2021)	0.1355	10.58
	10-K Y2 (2022)	0.1312	10.49
Twilio (TWLO)	S-1 (2016)	0.1624	11.38
	10-K Y1 (2017)	0.1408	10.80
	10-K Y2 (2018)	0.1377	10.74
Okta (OKTA)	S-1 (2017)	0.1519	11.15
	10-K Y1 (2018)	0.1380	10.68
	10-K Y2 (2019)	0.1361	10.63
Atlassian (TEAM)	F-1 (2015)	0.1645	11.41
	20-F Y1 (2016)	0.1479	10.94
	20-F Y2 (2017)	0.1433	10.85
Group 4: High-Liability / Regulated			
Moderna (MRNA)	S-1 (2018)	0.1910	12.35
	10-K Y1 (2019)	0.1601	11.51
	10-K Y2 (2020)	0.1540	11.30
Coinbase (COIN)	S-1 (2021)	0.1944	12.48
	10-K Y1 (2022)	0.1302	10.21
	10-K Y2 (2023)	0.1226	9.97

6.2 Baseline Compression: Confirming Proposition 1

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—vocabulary breadth and vocabulary predictability—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 ± 0.0004 and SE = 10.44 ± 0.02. Intuit shows comparable stasis. Absent a formalization event, the linguistic profile of a mature public company remains remarkably fixed. The dramatic compression observed in newly public firms is attributable to the IPO itself, not to secular trends in disclosure practice or regulatory environment.

Three patterns within the baseline confirm the theoretical predictions:

Category creators compress most. Firms that must invent strategic language—Amazon describing e-commerce in 1997 [AMAZON COM INC, 1997], Google defining search-as-business-model in 2004 [SEC.gov, 2004], Salesforce explaining SaaS in 2004 [SEC.gov, 2003]—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. Their S-1 language was already more specialized and less broadly conceptual. The compression from S-1 to 10-K is more moderate (9.2–11.8%) but equally consistent. The magnitude of compression is proportional to the linguistic novelty of the S-1.

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. This is consistent with the mechanism: the largest shift occurs when the most novel language is purged; subsequent years refine an already-constrained vocabulary.

6.3 The Legal Amplifier Effect

Section 5 argued that legal exposure amplifies variance compression. The Group 4 firms test this prediction directly.

Coinbase's S-1 [SEC.gov, 2021] attempted to evangelize the "cryptoeconomy" for a skeptical regulatory world (LD: 0.1944, SE: 12.48). Its Y1 and Y2 10-Ks were filed while the company was under active SEC investigation [SEC.gov, 2023, 2024]. Lexical diversity collapsed by 33% to 0.1302 in Y1 and continued falling to 0.1226 in Y2. Shannon entropy dropped from 12.48 to 9.97—the

lowest endpoint in the entire dataset. All philosophical language about the "cryptoeconomy" was purged and replaced with hyper-vetted legal boilerplate.

Robinhood presents a mirror image. Its S-1 built on the high-variance narrative of "democratizing finance" (LD: 0.1855). The IPO occurred in the aftermath of the GameStop controversy, triggering congressional hearings and regulatory scrutiny. LD collapsed by 30% to 0.1298 in Y1. The visionary language was systematically replaced by defensive safe-harbor terminology.

Moderna's S-1 [SEC.gov, 2018a] required exceptionally high variance to describe mRNA as a "new class of medicines" (LD: 0.1910, SE: 12.35). The 10-Ks [SEC.gov, 2019a, 2020a], filed before COVID-19, show a 16.2% Y1 drop as the infinite "platform" narrative was narrowed into the standardized language of clinical phases and regulatory pathways.

Blackstone provides a "mature amplifier" case. As an asset management firm, its S-1 metrics (LD: 0.1412) were already the lowest of any S-1 in the dataset—effectively "born caged" by its industry's legal framework. Even from this low starting point, the IPO triggered a further 7.5% compression [SEC.gov, 2007]. The legal amplifier operates even on firms that are already highly formalized.

The pattern is consistent: when legal threat becomes existential, variance compression is not gradual drift but rapid, forced capitulation. The velocity and magnitude of compression are directly proportional to perceived legal and regulatory threat.

6.4 Founder Insulation

Section 4 predicted that founder control with dual-class shares should attenuate but not eliminate compression. The Group 5 firms confirm this.

Meta's S-1 [SEC.gov, 2012b] contained a high-variance founder letter (LD: 0.1709). The Y1 drop was only 5.8% [SEC.gov, 2013]—substantially milder than the 15.3% drop at LinkedIn [SEC.gov, 2011, 2012a], a non-founder-controlled peer from the same era. The dual-class structure allowed mission-oriented language to persist alongside required legal boilerplate.

Snap provides the clearest case [SEC.gov, 2017b]. Its S-1 famously defined Snap as a "camera company" (LD: 0.1730). Despite intense criticism, a collapsing stock price, and user backlash, the Y1 drop was only 7.2% [SEC.gov, 2018b], and the Y2 drop was minimal (1.4%) [SEC.gov, 2019b]. At a non-founder-controlled firm, shareholder pressure would have forced linguistic capitulation. The dual-class structure allowed the founder to resist.

Google's case is more complex [SEC.gov, 2004]. The initial S-1-to-Y1 drop was severe (17.6%) because the S-1 contained discrete, excisable artifacts (the "Owner's Manual") removed in the 10-K [SEC.gov, 2005]. But the Y1-to-Y2 decay rate (2.5%) was slower than at non-founder firms (Amazon: 4.9%; Salesforce: 2.8%) [SEC.gov, 2006]. Founder control did not prevent the one-time excision but slowed the ongoing formalization of underlying strategic language.

The finding is consistent across all three cases: founder control acts as an insulator, not a shield. The metrics still decline. But the rate and magnitude of decline are significantly attenuated relative to non-founder-controlled peers.

6.5 Born Caged: Pre-Compression Through Mimetic Isomorphism

Corollary 1.1 predicted that modern S-1s would exhibit lower initial variance than earlier cohorts. The data confirm this strongly.

Snowflake's 2020 S-1 (LD: 0.1495, SE: 10.99) and Okta's 2017 S-1 (LD: 0.1519, SE: 11.15) registered substantially lower initial variance than category creators of earlier eras: Amazon 1997 (LD: 0.1742), Google 2004 (LD: 0.1825), or Salesforce 2004 (LD: 0.1776) [SEC.gov, 2020b, 2017a]. The difference persists even though Snowflake and Okta were introducing genuinely novel concepts ("Data Cloud," "Identity Cloud") comparable in strategic novelty to earlier S-1s.

The pattern extends to other modern filers. Workday and ServiceNow, filing S-1s in 2012, already showed lower initial variance than the 2004 cohort [SEC.gov, 2012d,c]. Despite these compressed starting points, the IPO still triggered further formalization (9.2–9.4% drops in Y1).

This supports the "Born Caged" hypothesis. 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. Legal counsel, underwriters, and auditors—themselves products of post-SOX formalized institutions—impose frame-constrained language on the S-1 through mimetic isomorphism. The formalization event is no longer a discrete IPO boundary but a diffuse process beginning in the pre-IPO preparation.

6.6 Strategic Pivot Anomaly: Bounding the Model

Netflix presents the most significant refinement to the variance compression thesis [NETFLIX INC, 2002]. The 2002 S-1, which described the DVD-by-mail business, shows clear compression into the Y1 and Y2 10-Ks (13.9% Y1 drop). This confirms the baseline thesis. But the Netflix case reveals a crucial boundary: the cage is a pressure for linguistic *stasis*, not a one-way linear decline.

Netflix's mid-to-late 2000s pivot to streaming would have required a massive injection of new, high-variance language ("streaming," "content delivery network," "original content"), temporarily increasing LD and SE and breaking the compression trend. The cage is a constant pressure for standardization that can be overcome by radical, discontinuous strategic innovation. Once the new strategy is established, the fiduciary trap re-engages and begins compressing the new high-variance language. The model predicts compression within a strategic regime, not across strategic discontinuities.

6.7 Summary of Empirical Findings

The 75-filing analysis provides strong support for Proposition 1 and its boundary conditions:

1. **Universal compression.** All 25 companies show LD and SE declining from S-1 to Y1 10-K, with continued decline to Y2.

2. **Legal amplifier confirmed.** High-liability firms (Coinbase, Robinhood, Moderna) show the most extreme compression (16–33%), consistent with Section 5's theoretical prediction.

3. **Founder insulation confirmed.** Dual-class firms (Meta, Google, Snap) show attenuated compression (5.8–7.2% in Y1), consistent with the boundary conditions.

4. **Born Caged confirmed.** Modern S-1s exhibit lower initial variance than earlier-era S-1s, supporting Corollary 1.1.

5. **Control cases confirm attribution.** Adobe and Intuit show linguistic stasis (±0.3%), confirming that compression is attributable to the formalization event, not secular trends.

6. **Strategic discontinuity bounds the model.** Netflix shows that radical strategic pivots can temporarily overcome the compression pressure, resetting the cycle.

The descriptive statistics establish the direction and magnitude of variance compression. Section 8 describes the regression analysis that will control for confounding variables and estimate the independent contribution of the formalization event.

7 Discussion and Boundary Conditions

7.1 Where the Model Applies Strongly

The mechanism operates most powerfully under conditions that maximize both formalization pressure and legal exposure: high legal liability industries, large and dispersed ownership, post-crisis regulatory environments, and long lock-in decisions. The empirical results from Section 6 confirm this: Coinbase and Robinhood (high-liability, dispersed ownership, active regulatory scrutiny) show 30–33% compression, while Blackstone (already operating in one of the most legally constrained environments) shows 7.5% compression from an already-low baseline.

7.2 Where the Model Applies Weakly

The mechanism weakens under conditions that reduce formalization pressure or legal exposure: founder control with majority voting shares, small private companies, fast-feedback environments, and high psychological safety cultures with documented dissent [Dalio, 2017]. The empirical results confirm: Meta and Snap show 5.8–7.2% compression, substantially below the 13–17% range of non-founder-controlled peers.

7.3 Secondary Moderators

Beyond the primary boundary conditions, several factors modulate the mechanism. Industry R&D intensity may maintain tolerance for variance even under formalization pressure. Competitive intensity and market velocity create pressure to maintain external perspective despite legal risk. Talent market liquidity creates countervailing pressure when the best people systematically select less formalized environments. Board composition and expertise may improve capacity to validate external perspective. These moderators attenuate but do not reverse the mechanism.

7.4 Why Solutions Decay

Organizations recognize the rigidity problem and implement solutions. The solutions often work initially, then decay.

Skunkworks absorption. A skunkworks begins with genuine independence. Success attracts attention. Leadership wants to scale. Scaling requires coordination with the parent. Coordination requires common metrics. The skunkworks is "brought into the fold." Within two to three years it exhibits the same variance compression as the parent.

Red team marginalization. A red team is chartered with veto-free authority. Over time, leadership introduces subtle constraints: reports must be "coordinated" before board presentation; authority to delay requires sponsor approval. The red team becomes advisory rather than adversarial. Dissent becomes performative.

Consultant capture. Consultants are hired for outside perspective. But consultants must maintain client relationships. Over time, recommendations converge toward what the client wants to hear. "External" perspective becomes internal validation with a different logo.

The common pattern is absorption into the formal system. Maintaining external perspective requires continuous structural protection, not just initial creation.

7.5 Alternative Explanations

Behavioral risk aversion. If behavioral risk aversion is primary, private companies should show the same compression as public companies. Existing evidence suggests the IPO discontinuity is sharp [Bernstein, 2015], more consistent with legal structure than gradual psychological change.

Resource dependence. Resource dependence explains convergence but not why convergence creates vulnerability. The incompleteness model subsumes resource dependence as one mechanism driving formalization while adding that shared frames create shared blind spots.

Institutional isomorphism. DiMaggio and Powell [1983] explain why organizations look alike but not why similarity creates rigidity. The Cage identifies the mechanism: convergent frames create convergent blind spots. The empirical Born Caged finding (Section 6.5) directly demonstrates institutional isomorphism operating through the S-1 preparation process—legal counsel and underwriters propagating formalized language templates across firms.

Professionalized documentation. A simpler account: post-IPO language compresses because professional writers (lawyers, investor relations teams) replace founder-drafted prose. This predicts uniform compression regardless of legal exposure or founder control. The data refute this:

compression varies from 5.8% (Meta, founder-controlled) to 33% (Coinbase, high-liability). If professionalized documentation were the sole mechanism, the variation should not track legal exposure as precisely as observed.

Temporary leadership failure. Some turnarounds succeed temporarily but revert within five to seven years. The model explains this: new leaders bring external perspective, but formalization reasserts as the leader becomes embedded and legal requirements remain unchanged.

7.6 Falsification Limits

The model is falsified or requires revision if: variance does not compress post-IPO across a broad sample; independence does not predict novelty in high-disruption contexts; authority does not predict dissent impact; effects do not co-vary across propositions; or solutions do not decay when absorbed. The empirical section provides strong support for the first prediction. The remaining predictions await testing.

8 Regression Analysis

The descriptive statistics presented in Section 6 establish the direction and magnitude of variance compression. This section formalizes those observations through OLS regression with robust (HC1) standard errors, controlling for confounding variables and estimating the independent contribution of the formalization event, its interaction with legal exposure, and its moderation by founder control.

8.1 Specification

We estimate three nested models with lexical diversity (LD) as the dependent variable across 60 filing observations (20 IPO companies $\times$ 3 filings each; mature controls excluded). Model (1) includes years-since-IPO, binary indicators for high-liability industry, founder control (dual-class shares), and born-caged status (post-2015 IPO), plus cohort dummies. Model (2) adds interaction terms: years-since-IPO $\times$ high-liability, years-since-IPO $\times$ founder-controlled, and years-since-IPO $\times$ born-caged. Model (3) replaces cohort and binary indicators with company fixed effects while retaining the interaction terms.

8.2 Results

The core result is robust across all three specifications. 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 founder-control interaction is positive but small ($+0.002$), consistent with the insulation (not immunity) interpretation from Section 6. The company fixed-effects model

Table 3: Regression Results: Lexical Diversity

	(1) Base	(2) Interactions	(3) Company FE
Intercept	0.1689***	0.1667***	0.1668***
	(0.0033)	(0.0040)	(0.0049)
years_since_ipo	−0.0141***	−0.0120***	−0.0120***
	(0.0018)	(0.0022)	(0.0018)
is_high_liability	−0.0067**	−0.0016	
	(0.0030)	(0.0058)	
is_founder_controlled	0.0185***	0.0170***	
	(0.0038)	(0.0053)	
is_born_caged	0.0107*	0.0115	
	(0.0061)	(0.0085)	
years × high_liability		−0.0102	−0.0102
		(0.0065)	(0.0064)
years × founder		0.0015	0.0015
		(0.0037)	(0.0034)
years × born_caged		−0.0008	−0.0008
		(0.0034)	(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 (post-SOX, modern, regulated) included in Models (1)–(2) but suppressed for space. Company FE in Model (3) absorb all time-invariant firm characteristics.

explains 80% of the variance in lexical diversity, suggesting that the formalization event accounts for the large majority of linguistic change in these filings.

8.3 Genre Confound Controls

A natural objection is that the S-1 to 10-K transition reflects genre differences (registration statements vs. annual reports), not compression per se. Three controls address this:

Table 4: Genre Confound Controls

Transition	N	Mean ΔLD	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
Mature Controls (Adobe, Intuit)					
Adobe (3-yr range)	3	0.0008	—	—	—
Intuit (3-yr range)	3	0.0005	—	—	—

First, 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. Second, mature public companies (Adobe and Intuit) show effectively zero drift over three years (LD range < 0.001), confirming that compression is driven by the formalization event, not by secular trends in filing language. Third, all three transitions produce changes significantly different from zero with large effect sizes ($|d| > 1.7$), ruling out noise as an explanation.

9 Implications

9.1 For Researchers

The framework redirects empirical efforts in three ways: away from proving external perspective helps, toward measuring structural conditions; away from seeking universal best practices, toward mapping contextual boundaries; and away from treating measurement circularity as methodological failure, toward recognizing it as evidence of the phenomenon itself.

Shift from proving effects to measuring conditions. The framework suggests measuring structural indicators of incompleteness (variance compression, independence erosion, authority constraints) under predicted conditions, rather than proving that external perspective causes success. The empirical results in Section 6 exemplify this shift: they measure whether decision variance compresses following a legal event, independent of whether those decisions prove correct.

Redirect to boundary mapping. The field's frustration with contradictory findings reflects an unstated assumption that mechanisms should operate universally. The framework provides testable

predictions: effects should be strongest in the upper-right quadrant of the Geometry of Defensibility. The empirical gradients—from 5.8% (founder-insulated) to 33% (legally exposed)—map directly onto this geometry.

Recognize measurement circularity as structural. If organizations exhibit incompleteness-like properties, and research is itself conducted within frames, then research on perspective-taking should exhibit the same difficulties. The appropriate response is not better instruments but different research designs: longitudinal studies tracking variance around formalization events, natural experiments using regulatory shocks, comparative studies exploiting cross-sectional variation in legal exposure.

9.2 For Practitioners

The framework offers diagnosis, not prescription. It clarifies trade-offs rather than resolves them.

Recognize the trade-off explicitly. Organizations under fiduciary duty face a forced choice between defensibility and adaptability. Leaders who understand the trade-off can choose to maintain external perspective consciously in contexts where adaptation matters more than defensibility. But they cannot eliminate the tension.

Protect mechanisms from absorption. Skunkworks, red teams, and consultants work initially because they operate outside the parent frame. Protection requires continuous structural separation: independent budgets, separate KPIs, veto-free authority. The cost is the price of maintaining external perspective under formalization.

Accept variance as feature, not bug. The empirical results show that variance compression is universal and begins immediately at IPO. Leaders in high-disruption contexts should treat variance compression as a warning signal. Maintaining variance requires protecting decisions that cannot be fully justified within the frame.

You cannot design unicorns. The framework does not tell leaders how to innovate. It tells them when they have built a system that no longer permits innovation. What leaders can do is protect the conditions where breakthroughs remain possible: slack resources, independent teams, tolerance for decisions that contradict metrics.

9.3 For Theorists

From "fix the cage" to "manage the boundary." The cage cannot be eliminated because it is a consequence of coordination requirements and legal obligations. The theoretical focus should shift from universal solutions to boundary management.

From universal best practices to contextual trade-offs. The empirical gradients demonstrate that what works depends on the regime. Practices optimal in one quadrant of the Geometry of Defensibility are counterproductive in another.

From optimization to optionality preservation. Formalization is optimization; it eliminates options. The theoretical insight is that preservation of optionality is itself a strategic capability that must be protected against formalization pressure.

Connection to institutional theory. The Cage framework does not replace institutional theory but provides its causal mechanism. DiMaggio and Powell's isomorphism is the observable pattern; the Cage is the engine. Meyer and Rowan's ceremonial conformity is the behavioral manifestation; the fiduciary trap is the structural driver. Scott's three pillars describe the institutional architecture; the Cage explains why that architecture produces incompleteness. This connection opens theoretical conversations between organizational scholars, legal scholars, and information theorists who have independently confronted how formal systems create limitations, how legibility requirements destroy information, and how accountability mechanisms produce the opposite of their intent.

10 Conclusion

Organizations formalize to achieve scale and to satisfy legal requirements. Formalization creates frames. Frames compress reality into legible dashboards. Compression creates blind spots. Optimization within frames deepens those blind spots. Frame-dependent reasoning self-defends. External perspective becomes structurally necessary but legally risky. The system creates the problem and limits the solution.

This is not failure of leadership, culture, or incentives. It is geometric consequence of formalization applied to organizational life. The constraint is structural. Scale requires coordination. Coordination requires formalization. Public ownership requires fiduciary duty. Fiduciary duty requires demonstrable soundness. Demonstrable soundness requires formalization. The two pressures converge.

The empirical evidence strengthens the case. Across 75 SEC filings from 25 companies spanning three decades, lexical diversity compresses following every IPO studied, with magnitude proportional to legal exposure and inversely proportional to founder control. The control cases (Adobe, Intuit) confirm linguistic stasis absent a formalization event. The Born Caged finding demonstrates that the cage's influence has moved upstream: modern S-1s arrive pre-compressed by the legal and banking ecosystem itself. The Legal Amplifier finding confirms that when legal threat becomes existential, variance compression is immediate and severe.

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.

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.

The trap's logic implies no absolute exit, only strategic movement within an infinite topology. Each act of meta-formalization—formalizing trust rather than operations, documenting incompleteness recognition rather than claiming completeness—establishes a new locally consistent domain. Organizations operating at the meta-level expand their reachable decision space without achieving total coverage. The practical limit is cost: each additional layer has diminishing returns. This is bounded rationality applied to governance design. The geometry identifies the constraint; navigating within that constraint through documented recognition of its boundaries remains structurally possible.

The trap has no villain. It has a law.

References

Smith v. van gorkom, 488 a.2d 858 (del. 1985). Delaware Supreme Court decision, 1985. Alden Smith and John W. Gosselin v. Jerome W. Van Gorkom et al., 488 A.2d 858 (Delaware Supreme Court 1985).

Sarbanes–oxley act of 2002. Public Law 107-204, 116 Stat. 745, 2002. Enacted July 30, 2002.

Dodd–frank wall street reform and consumer protection act. Public Law 111-203, 124 Stat. 1376, 2010. Enacted July 21, 2010.

Renée B. Adams and Daniel Ferreira. Women in the boardroom and their impact on governance and performance. *Journal of Financial Economics*, 94(2):291–309, 2009. doi: 10.1016/j.jfineco.2008.10.007.

AMAZON COM INC. Form s-1/a. EDGAR Online, 1997. Filed April 21, 1997. CIK 0001018724. URL: `https://content.edgar-online.com/ExternalLink/EDGAR/0000891020-97-000603.html`.

Chris Argyris and Donald A. Schön. *Organizational Learning: A Theory of Action Perspective*. Addison-Wesley, Reading, MA, 1978.

Pamela S. Barr, J. L. Stimpert, and Anne S. Huff. Cognitive change, strategic action, and organizational renewal. *Strategic Management Journal*, 13(S1):15–36, 1992. doi: 10.1002/smj.4250130904.

Shai Bernstein. Does going public affect innovation? *Journal of Finance*, 70(4):1365–1403, 2015. doi: 10.1111/jofi.12275.

James M. Buchanan and Gordon Tullock. *The Calculus of Consent: Logical Foundations of Constitutional Democracy*. University of Michigan Press, Ann Arbor, MI, 1962.

Mario Bunge. *Treatise on Basic Philosophy: Volume 3, Ontology I: The Furniture of the World*. Springer, Dordrecht, 1977.

Clayton M. Christensen. *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Harvard Business School Press, Boston, MA, 1997.

Ray Dalio. *Principles: Life and Work*. Simon & Schuster, New York, 2017.

Paul J. DiMaggio and Walter W. Powell. The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, 48(2):147–160, 1983. doi: 10.2307/2095101.

Kurt Gödel. Über formal unentscheidbare sätze der *Principia Mathematica* und verwandter systeme i. *Monatshefte für Mathematik und Physik*, 38(1):173–198, 1931. doi: 10.1007/BF01700692. English translation: On Formally Undecidable Propositions of *Principia Mathematica* and Related Systems I.

Charles A. E. Goodhart. Problems of monetary management: The u.k. experience. In Anthony S. Courakis, editor, *Monetary Theory and Practice: The UK Experience*, pages 91–121. Macmillan, London, 1984.

Michael T. Hannan and John Freeman. Structural inertia and organizational change. *American Sociological Review*, 49(2):149–164, 1984. doi: 10.2307/2095567.

Mary B. Hesse. *Models and Analogies in Science*. University of Notre Dame Press, Notre Dame, IN, 1966.

Michael C. Jensen and William H. Meckling. Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4):305–360, 1976. doi: 10.1016/0304-405X(76)90026-X.

James G. March. Exploration and exploitation in organizational learning. *Organization Science*, 2 (1):71–87, 1991. doi: 10.1287/orsc.2.1.71.

John W. Meyer and Brian Rowan. Institutionalized organizations: Formal structure as myth and ceremony. *American Journal of Sociology*, 83(2):340–363, 1977. doi: 10.1086/226550.

Charlan Jeanne Nemeth. Differential contributions of majority and minority influence. *Psychological Review*, 93(1):23–32, 1986. doi: 10.1037/0033-295X.93.1.23.

NETFLIX INC. Form s-1/a. SEC.gov, 2002. CIK 0001065280. URL: https://www.sec.gov/Archives/edgar/data/1065280/000101287002002403/ds1a.htm.

Stefan Schulz-Hardt, Andreas Mojzisch, Felix C. Brodbeck, Rudolf Kerschreiter, and Dieter Frey. Group decision making in hidden profile situations: Dissent as a facilitator of decision quality. *Journal of Personality and Social Psychology*, 91(6):1080–1093, 2006. doi: 10.1037/0022-3514.91.6.1080. Used here to represent the dissent/hidden-profile stream.

W. Richard Scott. *Institutions and Organizations: Ideas, Interests, and Identities*. SAGE Publications, Thousand Oaks, CA, 4th edition, 2013.

SEC.gov. Adobe inc. edgar page. SEC.gov, a. CIK 0000796343. URL: https://www.sec.gov/edgar/browse/?CIK=0000796343.

SEC.gov. Intuit inc. edgar page. SEC.gov, b. CIK 0000896878. URL: https://www.sec.gov/edgar/browse/?CIK=896878&owner=exclude.

SEC.gov. Form s-1: Salesforce.com, inc. SEC.gov, 2003. CIK 0001108524. URL: https://www.sec.gov/Archives/edgar/data/1108524/000119312503096073/ds1.htm.

SEC.gov. Form s-1: Google inc. SEC.gov, 2004. CIK 0001288776. URL: https://www.sec.gov/Archives/edgar/data/1288776/000119312504073639/ds1.htm.

SEC.gov. Google 10-k (fy 2004). SEC.gov, 2005. URL: https://www.sec.gov/Archives/edgar/data/1288776/000119312505065298/d10k.htm.

SEC.gov. Google 10-k (fy 2005). SEC.gov, 2006. URL: https://www.sec.gov/Archives/edgar/data/1288776/000119312506056598/d10k.htm.

SEC.gov. S-1: The blackstone group l.p. SEC.gov, 2007. CIK 0001393818.

SEC.gov. Registration statement on form s-1: Linkedin corporation. SEC.gov, 2011. CIK 0001271024. URL: https://www.sec.gov/Archives/edgar/data/1271024/000119312511016022/ds1.htm.

SEC.gov. Linkedin 10-k (fy 2011). SEC.gov, 2012a. URL: https://www.sec.gov/Archives/edgar/data/1271024/000119312512094556/d260171d10k.htm.

SEC.gov. Registration statement on form s-1: Facebook, inc. SEC.gov, 2012b. CIK 0001326801. URL: https://www.sec.gov/Archives/edgar/data/1326801/000119312512034517/d287954ds1.htm.

SEC.gov. Registration statement on form s-1: Servicenow, inc. SEC.gov, 2012c. CIK 0001373715. URL: https://www.sec.gov/Archives/edgar/data/1373715/000119312512143517/d301887ds1.htm.

SEC.gov. Form s-1: Workday, inc. SEC.gov, 2012d. CIK 0001327811. URL: https://www.sec.gov/Archives/edgar/data/1327811/000119312512375787/d385110ds1.htm.

SEC.gov. Facebook 10-k (fy 2012). SEC.gov, 2013. URL: https://www.sec.gov/Archives/edgar/data/1326801/000132680113000003/fb-12312012x10k.htm.

SEC.gov. Form s-1: Okta, inc. SEC.gov, 2017a. CIK 0001660134. URL: https://www.sec.gov/Archives/edgar/data/1660134/000119312517080301/d289173ds1.htm.

SEC.gov. S-1: Snap inc. SEC.gov, 2017b. CIK 0001564408. URL: https://www.sec.gov/Archives/edgar/data/1564408/000119312517029199/d270216ds1.htm.

SEC.gov. S-1: Moderna, inc. SEC.gov, 2018a. CIK 0001682852. URL: https://www.sec.gov/Archives/edgar/data/1682852/000119312518323562/d577473ds1.htm.

SEC.gov. Snap 10-k (fy 2017). SEC.gov, 2018b. URL: https://www.sec.gov/Archives/edgar/data/1564408/000156459018002721/snap-10k_20171231.htm.

SEC.gov. Moderna 10-k (fy 2018). SEC.gov, 2019a. URL: https://www.sec.gov/Archives/edgar/data/1682852/000168285219000009/moderna10-k12312018.htm.

SEC.gov. Snap 10-k (fy 2018). SEC.gov, 2019b. URL: https://www.sec.gov/Archives/edgar/data/1564408/000156459019002053/snap-10k_20181231.htm.

SEC.gov. Moderna 10-k (fy 2019). SEC.gov, 2020a. URL: https://www.sec.gov/Archives/edgar/data/1682852/000168285220000006/moderna10-k12312019.htm.

SEC.gov. Form s-1: Snowflake inc. SEC.gov, 2020b. CIK 0001640147. URL: https://www.sec.gov/Archives/edgar/data/1640147/000162828020013010/snowflakes-1.htm.

SEC.gov. Form s-1: Coinbase global, inc. SEC.gov, 2021. CIK 0001679788. URL: https://www.sec.gov/Archives/edgar/data/1679788/000162828021003168/coinbaseglobalincs-1.htm.

SEC.gov. Coinbase 10-k (fy 2021). SEC.gov, 2023. URL: https://www.sec.gov/Archives/edgar/data/1679788/000167978823000031/coin-20221231.htm.

SEC.gov. Coinbase 10-k (fy 2022). SEC.gov, 2024. URL: https://www.sec.gov/Archives/edgar/data/1679788/000167978824000022/coin-20231231.htm.

Marilyn Strathern. 'improving ratings': Audit in the british university system. *European Review*, 5 (3):305–321, 1997. doi: 10.1002/(SICI)1234-981X(199707)5:3<305::AID-EURO184>3.0.CO;2-4.

Andrew Sturdy. Consultancy's consequences? a critical assessment of management consultancy's impact on management. *British Journal of Management*, 22(3):517–530, 2011. doi: 10.1111/j.1467-8551.2010.00708.x.

David J. Teece, Gary Pisano, and Amy Shuen. Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7):509–533, 1997. doi: 10.1002/(SICI)1097-0266(199708)18:7<509::AID-SMJ882>3.0.CO;2-Z.

Michael L. Tushman and Charles A. O'Reilly. Ambidextrous organizations: Managing evolutionary and revolutionary change. *California Management Review*, 38(4):8–30, 1996. doi: 10.2307/41165852.

Max Weber. *The Protestant Ethic and the Spirit of Capitalism*. Charles Scribner's Sons, New York, 1958. Original work published 1905; trans. T. Parsons.

Anders Werr and Alexander Styhre. Management consultants—friend or foe? understanding the ambiguities of management consultancy. *International Studies of Management & Organization*, 32(4):43–66, 2003. doi: 10.1080/00208825.2003.11043681.