Righting the Ship: The Succession Paradox

Jeremy McEntire

Abstract

Organizations facing decline reach for the same lever: replace the leader. Reported failure rates for such successions vary widely—studies cite figures from thirty to seventy percent depending on definition and timeframe—but the modal outcome is disappointing regardless of the precise estimate. This consistency demands structural explanation. This paper argues that succession failure is not primarily selection error but selection outcome: the same dynamics that produce organizational dysfunction also govern the response to dysfunction. New leaders inherit the information architecture that trapped their predecessors. The filters remain. The fitness landscape persists. The faces change; the Cage does not.

This analysis builds on prior work on dysmemic pressure (McEntire, 2025) and organizational compression dynamics (McEntire, 2026), integrating these mechanisms with established literatures on upper echelons theory, institutional isomorphism, and high-reliability organizations. Two intervention patterns historically succeed: comprehensive disruption of the information environment, or architectural reform by a leader who understands the mechanism. The modal approach—new leader, preserved architecture—fails because it treats personnel as the problem when the selection environment is the problem. Case evidence from IBM, Ford, Best Buy, and Home Depot illustrates the dynamics. The framework extends to authoritarian regimes and institutional decline more broadly. Testable propositions and falsification criteria are specified, including the meta-prediction that the theory will be adopted at lower rates than simpler alternatives regardless of accuracy.

Keywords: organizational succession, information architecture, institutional failure, corporate turnaround, leadership selection, dysmemic pressure

1. Introduction

Organizations facing decline have a favorite move: replace the leader. The board convenes. Heads shake gravely. A search commences. A new CEO arrives with fresh eyes and a mandate for change. Eighteen months later, the pattern repeats.

Estimates of succession failure rates vary considerably. Meta-analyses of CEO succession find that post-succession performance depends heavily on context, timeframe, and how failure is defined (Schepker, Kim, Patel, Thatcher, & Campion, 2017). Some studies focus on early CEO departure; others on stock performance; still others on strategic change implementation. The widely cited claim that seventy percent of change initiatives fail has been persuasively critiqued as lacking empirical foundation, traceable to an acknowledged unscientific estimate that calcified into received wisdom (Hughes, 2011). What the evidence

does support is that succession outcomes cluster toward disappointment more often than success, and that this pattern holds across industries, economic conditions, and whether the predecessor was fired or retired. The precise failure rate matters less than its stability: something structural is happening. The question is what.

The conventional explanations fall into two categories. The first blames selection: the board chose wrong. The second blames execution: the board chose right, but the new leader lacked sufficient support, faced unexpected obstacles, or made implementation errors. Both explanations assume the strategy is sound. Neither asks whether changing the leader is itself the problem—whether succession is what organizations do instead of changing themselves.

This paper argues that succession failure is better understood as selection outcome than selection error. The same dynamics that produce organizational dysfunction also govern the organizational response to dysfunction. The insight is not entirely new. Scholars have long recognized that leadership change can function as ritual scapegoating—a symbolic act to appease stakeholders independent of performance outcomes (Gamson & Scotch, 1964). The romance of leadership literature has documented systematic overattribution of organizational outcomes to individual executives (Meindl, Ehrlich, & Dukerich, 1985). Institutional theory explains how organizations adopt practices for legitimacy rather than effectiveness, copying peers in ways that may have little connection to actual performance improvement (DiMaggio & Powell, 1983).

What these literatures leave unexplained is the mechanism by which succession fails even when boards genuinely attempt to address dysfunction. A board engaging in ritual scapegoating knows, at some level, that the gesture is symbolic. But most boards believe they are solving a real problem. They hire thoughtfully. They invest in onboarding. They provide resources and authority. The intervention still fails. Why?

The framework proposed here provides an answer grounded in information dynamics. Drawing on prior work synthesizing strategic communication theory, agency theory, and cultural evolution (McEntire, 2025), and applying the analysis of organizational compression and self-deception (McEntire, 2026), this paper demonstrates how succession fails structurally rather than incidentally. The new leader enters an information environment that has been shaped by years of selection pressure. The signals reaching the top have been filtered, shaded, optimized for fit rather than accuracy. The organization's sensemaking apparatus is compromised. The new leader plugs into this apparatus and receives its outputs, which are statistically independent of reality in precisely the dimensions that matter for diagnosing dysfunction.

The argument proceeds as follows. Section 2 briefly integrates the theoretical framework, establishing mechanisms without rehearsing arguments made elsewhere. Section 3 applies the framework to succession, explaining why the modal approach fails and what alternatives exist. Section 4 presents case evidence from four organizations that illustrate success and failure modes with unusual clarity: IBM under Gerstner, Ford under Mulally, Best Buy under Joly, and Home Depot under Nardelli. Section 5 extends the analysis to authoritarian regimes

and institutional decline, demonstrating substrate independence. Section 6 states testable propositions and falsification criteria. Section 7 concludes.

A note on scope. This paper does not claim that all succession failures result from preserved information architecture, nor that personnel quality is irrelevant. Leadership matters. The claim is narrower: that a significant and underappreciated fraction of succession failures result from structural dynamics, and that recognizing these dynamics changes what interventions make sense. The framework specifies when architectural intervention is necessary—but necessity is not sufficiency. Changing the information architecture in the wrong direction, as Home Depot discovered, produces failure as surely as preserving dysfunction.

2. Theoretical Framework

2.1 Related Explanations

Before presenting the dysmemic framework, it is worth situating it against existing explanations for succession failure and organizational dysfunction.

Upper echelons theory (Hambrick & Mason, 1984) argues that organizational outcomes reflect the characteristics of top management—their experiences, values, and cognitive bases. This literature has generated productive research on how executive demographics predict strategic choices. What it does not explain is why executive characteristics should matter less than the information environment those executives inhabit. A CEO's experience matters only if it shapes interpretation of signals; if the signals themselves are corrupted, interpretation quality becomes secondary.

Structural inertia theory (Hannan & Freeman, 1984) proposes that organizations are selected for reliability and accountability, which makes them resistant to change. Inertia explains persistence but not the specific failure mode of succession: why the new leader, hired precisely to overcome inertia, becomes captured by it.

Threat rigidity (Staw, Sandelands, & Dutton, 1981) describes how organizations under threat restrict information processing and centralize control. This explains some dysfunction during crisis but not the persistence of dysfunction across leadership transitions when the threat has been acknowledged and response mandated.

Institutional isomorphism (DiMaggio & Powell, 1983) explains why organizations converge on similar practices—including the practice of CEO replacement—through coercive, mimetic, and normative pressures. This explains adoption of the succession strategy but not its failure. If succession worked, convergence on succession would be adaptive. The puzzle is convergence on something that does not work.

Normalization of deviance (Vaughan, 1996) describes how organizations gradually accept risk as normal through incremental accommodation. Vaughan's analysis of the Challenger

disaster showed how signals of O-ring erosion were reinterpreted rather than acted upon. This mechanism—signals being absorbed into existing frames rather than triggering response—is central to the dysmemic framework, which provides formal structure for understanding why such absorption occurs and when it can be interrupted.

The framework proposed here integrates elements from these literatures while providing a unifying mechanism: selection pressure operating on signals within information architectures. Where upper echelons theory emphasizes executive characteristics, the dysmemic framework emphasizes the selection environment those executives encounter. Where structural inertia emphasizes organizational resistance to change, the dysmemic framework specifies what resists: the fitness landscape for signals, which captures new leaders as readily as it captured old ones. Where institutional theory emphasizes legitimacy-seeking, the dysmemic framework explains why legitimacy-seeking produces dysfunction: because selection favors signals that are legible and defensible over signals that are accurate.

2.2 The Dysmemic Mechanism

The mechanism operates through three linked dynamics, presented here in summary form. Fuller treatment appears in McEntire (2025) and McEntire (2026).

Compression. Organizations coordinate at scale by compressing information. Ten thousand customer interactions become one summary report. A thousand engineering assessments become one dashboard metric. Shannon (1948) proved that compression has cost: push information through a channel smaller than the information itself and something gets lost. Every compression operator has a null space—the set of distinctions it cannot make. Quarterly revenue cannot distinguish happy customers from departing ones. Green-yellow-red status cannot distinguish correlated small risks from catastrophic exposure. The executive sees what the metric shows and cannot see what the metric misses.

Selection. Where compression discards distinctions, gaps open between representation and reality. Signals compete to fill those gaps. What survives is what fits the local environment—regardless of correspondence with the world. Strategic communication degrades as preferences diverge: senders shade messages toward what receivers want to hear, and stacked interfaces compound distortion until signals become statistically independent of states they describe (Crawford & Sobel, 1982). Adverse selection floods idea markets with low-quality signals when verification is costly (Akerlof, 1970). Transmission bias favors signals based on properties unrelated to truth—simplicity, prestige, conformity (Boyd & Richerson, 1985). The compound force is dysmemic pressure: selection favoring signals that increase individual payoff while decreasing collective adaptability.

Equilibrium. The drift stabilizes. Compression creates gaps; selection fills them; surviving signals become the environment to which new signals must adapt. At equilibrium, constructed reality, signal population, and selection criteria become mutually consistent. This is the Cage. The walls are invisible from inside. Reform fails because reform proposals are themselves signals facing selection in the environment they aim to change.

2.3 Information Architecture: Operationalization

For the framework to generate testable predictions, information architecture must be defined with sufficient precision to be measured. The construct comprises five observable components. Channel topology refers to who can speak to whom, what bypass routes exist, whether skip-level access is available or forbidden. Compression operators specify what gets aggregated, into what, at what cadence—the specific metrics, reports, and summaries that reduce high-dimensional reality to manageable representations. Verification mechanisms include independent audit functions, red teams, external board access to operations, prediction markets, and protected dissent channels—structures that create observation outside local selection pressure. Selection criteria determine what is rewarded and punished, formally through compensation and promotion, informally through status, access, and reputation. Protected variance encompasses mechanisms that keep alternative signals alive despite poor fit with the dominant frame—minority reports, devil's advocates, explicit disagreement norms.

Fitness landscape change means alteration to one or more of these components sufficient to shift what signals are selected. A succession that replaces executives without changing any of these components preserves the information architecture. A succession that alters reporting structures, verification mechanisms, or incentive criteria changes it. The direction of change matters: architectural change that selects for accurate signals differs from architectural change that selects for different-but-still-inaccurate signals.

3. The Succession Problem

3.1 The Standard Story

The conventional narrative runs as follows. Organization declines. Board identifies leadership as the problem. Search produces candidates. Board selects. New leader arrives with mandate. Transformation ensues, or does not, depending on execution quality.

The narrative contains an assumption so embedded it becomes invisible: that leadership is in fact the problem. The organization struggled; the leader was present during the struggle; therefore the leader caused the struggle, or failed to prevent it. The inference feels obvious. It is also, in a significant fraction of cases, wrong.

Consider what the new leader inherits. The reporting structures remain. The meeting cadences persist. The metrics that shaped behavior yesterday shape behavior today. The people who learned what signals survive—who learned through repeated iteration what to say and what to suppress—remain in place, their learned responses intact. The information environment is unchanged. The new leader sees what the old leader saw, because the new leader looks through the same apparatus.

The standard story treats the leader as cause. The dysmemic framework suggests the leader is more often effect—the output of a selection process operating on signals, not the origin of the signals themselves.

3.2 Why the Middle Ground Fails

The modal succession strategy is the middle ground: new leader, preserved architecture. The board rejects the extremes. Keeping the old leader signals complacency. Comprehensive restructuring signals panic. The middle ground signals measured response.

The middle ground fails for reasons the framework makes precise. First, the new leader enters an information environment optimized against accuracy. The previous leader did not fail for lack of intelligence or effort. The previous leader failed, in significant part, because signals reaching the top had been selected for fit rather than truth. The new leader plugs into the same apparatus. The same signals arrive. The same blind spots persist.

Second, the preserved management layer has adapted to the preserved fitness landscape. Middle managers learned what gets rewarded and what gets punished. They learned which messages advance careers and which end them. The new leader announces transformation. The middle layer waits. They have seen announcements before. They know leaders come and go while fitness landscapes remain. The rational strategy is to perform compliance while waiting for the landscape to reassert.

Third, the dysmemic ecosystem adapts faster than the leader can. New vocabulary proliferates. The organization learns to speak transformation while practicing continuity. The leader hears new words and mistakes them for new reality.

3.3 Two Paths That Work

Two approaches have historically succeeded. Both work by changing what gets selected for. They differ in what they preserve.

Comprehensive disruption resets enough of the information architecture simultaneously that accumulated gaming strategies lose purchase. Signals optimized for the old landscape find themselves unfit for the new one. Presidential term limits function this way: the incoming administration inherits institutional constraints but not the specific selection environment its predecessor learned to navigate. The risk is obvious. Comprehensive disruption cannot easily distinguish structures that should be destroyed from structures that should be preserved. Institutional knowledge lives in the same systems that house accumulated dysfunction.

Architectural reform from within works when a leader who understands the mechanism dismantles it piece by piece. Not by replacing people, but by changing what the system selects for. The fitness landscape shifts. Signals that were unfit become fit. The same people, facing different selection pressure, produce different signals. This path preserves institutional knowledge while changing the selection environment. It requires a leader who sees the

Cage—who understands that the problem is not the people but what the people have adapted to. It requires sufficient authority to change fitness criteria against resistance. It requires time, and that the leader not be captured before the work is complete.

The two paths share common structure: both alter selection pressure rather than selecting different people to face unchanged pressure.

3.4 A Necessary Clarification: Necessary but Not Sufficient

The framework claims that information architecture change is necessary for successful turnaround. It does not claim sufficiency. One can change information architecture in the wrong direction—impose a fitness landscape that selects against the capabilities that made the organization valuable—and fail catastrophically. The Home Depot case below illustrates this precisely.

Architectural change is necessary because unchanged architecture produces unchanged signals, which produce unchanged dysfunction regardless of who occupies leadership positions. But the direction of architectural change determines outcome. An architecture that selects for accurate customer feedback, genuine risk assessment, and candid communication enables adaptation. An architecture that selects for command compliance, metric optimization, and hierarchy deference may produce different dysfunction than the original—but dysfunction nonetheless.

3.5 Why Boards Choose the Middle Ground Anyway

If the middle ground fails reliably, why do sophisticated boards keep choosing it?

Boards face dysmemic pressure like everyone else. The board's decision is a signal competing for survival in the board's environment. The selection criteria are legibility, defensibility, and stakeholder palatability. Legibility means the narrative must be simple: organization struggled, board acted, new leader hired—stakeholders can follow the story. Defensibility means the decision must generate documentation: search process conducted, candidates evaluated, references checked—if the new leader fails, the board demonstrated prudence. Stakeholder palatability means the middle ground creates minimal disruption to constituencies that judge the board: analysts continue covering a recognizable company, institutional investors see continuity, employees do not panic.

Comprehensive disruption fails all three tests. Architectural reform fails on speed—boards face pressure to act visibly and immediately, and a leader who spends the first year understanding the information environment looks passive. The selection environment governing board decisions rewards the choice that does not work. The boards are not stupid. They are rational actors in an environment that selects against effective action.

There is a deeper reason boards cannot see this. The board's information about the organization arrives through the same architecture that trapped the CEO. Management

presents to the board; management controls what the board sees. Independent directors rely on the CEO as their primary sensor, and that sensor is embedded in the system being evaluated. Consultants and search firms operate in the same legitimacy market—recommending what peers have done, what can be documented, what will be defensible if it fails. The twenty-billion-dollar executive search industry earns fees by placing candidates, not by recommending architectural reform that generates no placement. The board is not outside the Cage looking in. The board is inside.

3.6 Boundary Conditions

The claim that architectural change is necessary applies above a threshold of dysmemic saturation. For mild dysfunction—where feedback loops remain tight, where external signals punish inaccuracy quickly, where the gap between representation and reality has not yet widened into equilibrium—leadership change alone may suffice.

Small organizations with short cycle times form Cages more slowly. Compression is light because coordination does not require heavy summarization. Preferences are more likely to align because everyone can see everyone else. Verification is cheap because the distance between action and outcome is short. A startup with twenty people and weekly iterations faces different structural conditions than a conglomerate with two hundred thousand employees and quarterly reporting cycles.

Highly instrumented environments with hard external truth signals leave less room for selection to operate. A trading desk receives immediate feedback from markets. A manufacturing line receives immediate feedback from defect rates. Where constructed reality confronts external consequences quickly, the Cage cannot fully form. Dysfunction requires soft feedback, long delays, and ambiguous causation.

The theory's predictions are strongest where compression is severe, preference divergence is high, and verification is costly—the conditions under which babbling equilibrium forms. Where these conditions are absent or weak, the middle ground may work. Where they are present and saturated, it cannot. Saturation indicators include universal green reporting in the presence of obvious loss, bad news that consistently surprises leadership, forecasts that are never backtested, incentives that reward local optimization against global outcomes, dissent that is punished or routed through social risk, and boards whose primary sensor is management narrative rather than independent verification.

4. Case Evidence

Four cases illustrate the dynamics. These are mechanism-tracing illustrations, not validation evidence—they show how the theory plays out in specific contexts, not proof that the theory is correct. The proof, if it comes, will require systematic analysis across samples large enough to support inference. What the cases provide is clarity about what architectural change looks like in practice, and what happens when it is present, absent, or misdirected.

Two succeeded through architectural reform that preserved institutional knowledge. One succeeded through architectural reform while keeping existing personnel entirely. One failed by changing architecture in the wrong direction. Together they demonstrate that the predictive variable is not leader quality but whether intervention addressed the selection environment appropriately.

4.1 IBM: The Cage Made Visible

When Lou Gerstner arrived at IBM in April 1993, the company was losing sixteen billion dollars over three years and planning its own dissolution. The prevailing wisdom held that IBM's size had become fatal. Gerstner was hired to preside over breakup.

Gerstner refused. His experience as a customer had taught him something IBM's executives could not see from inside: that customers needed integration, not fragmentation. IBM's scale was not the problem. IBM's inability to deploy scale on behalf of customers was the problem.

What Gerstner found confirmed the diagnosis. IBM had become a collection of fiefdoms. Geographic divisions operated as independent kingdoms. The company's internal transfer pricing consumed more executive attention than customer pricing. Gerstner later wrote of finding a company totally focused on its internal rules and conflicts (Gerstner, 2002).

The information architecture was visibly broken. Gerstner discovered his emails to European employees were being intercepted by the head of IBM Europe, who filtered which communications from the CEO would reach his territory. The executive was not hiding malfeasance; he was protecting his fiefdom. The CEO's messages were signals that might disrupt local equilibrium. The rational response, from within the local fitness landscape, was to block them.

Gerstner's response targeted the fitness landscape rather than personnel. Compensation shifted from divisional performance to corporate performance. Geographic fiefdoms were subordinated to global industry teams. Transfer pricing debates were eliminated by fiat. The people did not change. The selection environment changed. Executives who had optimized for territorial defense found that territorial defense no longer paid.

Upon becoming CEO, Gerstner initially changed only the CFO, the HR chief, and three key line executives (Gerstner, 2002). The mass layoffs IBM undertook were cost reduction, not cultural transformation. Cultural transformation came from changing what the system rewarded. IBM went from losing eight billion dollars in 1993 to earning eight billion in profit by 2001.

4.2 Ford: Architecture Without Personnel Change

Alan Mulally arrived at Ford in September 2006 with the company anticipating a seventeen billion dollar loss. Ford's leadership team was fractured by silos and ego-driven competition. No one wanted to admit what was not working.

Mulally's intervention is often cited by critics of the dysmemic framework as evidence that personnel need not change—as if this contradicts the theory. It does not. Mulally's success perfectly illustrates architectural reform. He kept the personnel and changed the information architecture.

The centerpiece was the Business Plan Review, or BPR: a weekly meeting held on the same day, at the same time, with mandatory attendance by all senior executives. Each leader would present status using color-coded charts—green for on track, yellow for at risk, red for off-plan. The format was rigid. Debate was forbidden in the main meeting. Issues requiring discussion moved to a separate Special Attention Review.

In early BPRs, every chart was green. Mulally knew this could not be true. A company losing seventeen billion dollars does not have everything on track. The breakthrough came when executive Mark Fields presented a red chart—a production failure on a new vehicle launch. The room froze, expecting punishment. Mulally clapped. Who can help Mark? he asked.

This single moment changed the fitness landscape. Truth-telling was rewarded. Problem-hiding was no longer the path to survival. The same executives, facing different selection criteria, produced different signals. Mulally did not fire his way to transformation. He changed what the system selected for. Ford went from seventeen billion in losses to eleven consecutive quarters of profitability, achieving eight billion in annual profit by 2013—without requiring a government bailout when competitors did.

The case confirms the framework: architectural change is necessary. Mulally achieved it without mass personnel change, demonstrating that architecture, not personnel, is the operative variable.

4.3 Best Buy: Listening to Discarded Information

When Hubert Joly became CEO of Best Buy in 2012, the company was expected to follow Circuit City into extinction. Analysts recommended closing stores en masse. The standard playbook said: cut costs, reduce headcount, shrink to survivability.

Joly rejected the playbook. His first action was to spend a week working in a Best Buy store—not touring, but working. He wore the blue shirt, talked to customers, listened to employees. What he learned contradicted what the executive suite believed. The employees knew what was wrong. They knew which processes frustrated customers, which policies made no sense. They had been telling headquarters for years. Headquarters had not been listening. The information was discarded because the fitness landscape did not reward surfacing problems from below.

Joly built a Mirror—structures that create observation outside local selection pressure. He created channels through which discarded information could reach decision-makers without passing through filters that had blocked it. The transformation he announced was not strategy imposed from above. It was permission for the organization to surface what it already knew.

What Joly preserved is as important as what he changed. The blue shirts—front-line employees with customer knowledge—were institutional knowledge carriers. The dysfunction was not in their expertise but in headquarters' inability to access it. A leader following the standard playbook would have cut the workforce, destroying exactly the capability that differentiated Best Buy from Amazon. Best Buy did not close hundreds of stores. It matched Amazon's prices. It used stores as distribution hubs. It invested in employees. Stock price increased more than tenfold from 2012 lows.

4.4 Home Depot: The Wrong Architecture

Bob Nardelli arrived at Home Depot in December 2000 with credentials no one could question. He had run GE's power systems division with distinction. He brought the GE playbook: process discipline, operational metrics, Six Sigma methodology.

Home Depot's founders had built a different kind of company. Store managers had authority to make decisions without seeking approval. Employees were hired for expertise and empowered to help customers with projects. The culture was entrepreneurial, decentralized, and deliberately unstructured.

Nardelli provided discipline. He centralized decision-making. He implemented rigorous metrics. He replaced experienced full-time employees with less expensive part-timers. He shifted resources from retail to a wholesale supply business that could be managed with precision. The metrics improved. Costs fell. Inventory turns increased. By operational measures Nardelli understood, the company was executing well.

The customer experience collapsed. Home Depot had built its reputation on knowledgeable employees who could guide customers through complex projects. Under Nardelli's fitness landscape, they were selected against. Customer satisfaction indices fell. Customers defected to Lowe's. The stock price went nowhere for six years while Lowe's doubled.

Nardelli was not incompetent. He was applying a fitness landscape appropriate to GE's industrial businesses to a retail business where competitive advantage lived in exactly what those methods destroyed. The null space of Nardelli's compression operators contained precisely what made Home Depot valuable.

This case does not contradict the framework—it confirms a crucial element. Architectural change is necessary but not sufficient. Nardelli changed the architecture. He changed it in the wrong direction. He selected against the signals that carried Home Depot's competitive advantage. The failure was not lack of change but change that destroyed capability-bearing structures while building dysfunction-producing ones. His successor, Frank Blake, reversed course—restored investment in employees, rebuilt service culture, returned decision authority to stores. The stock recovered.

4.5 Synthesis

The four cases share structure beneath surface differences. IBM and Ford succeeded because leaders identified information architecture as the problem and changed it. They did not assume new strategy through old structures would produce new results. They changed what the system selected for. Best Buy succeeded because Joly identified discarded information as the solution and built channels to recover it. He changed what could be heard. Home Depot failed because Nardelli changed architecture without understanding what it contained. He got the transformation he sought, and the transformation destroyed value.

Personnel change is neither necessary nor sufficient. Gerstner succeeded with minimal turnover. Mulally succeeded with no turnover. Nardelli failed with extensive turnover. The dimension that matters is information architecture: what gets reported, what gets filtered, what gets rewarded, what cannot be said—and whether changes to these select for accuracy or against it.

5. Beyond Organizations

The mechanism is general. Corporations are one substrate; there are others. This section extends the analysis to authoritarian governance and institutional decline. The purpose is to demonstrate that the dynamics are substrate-independent—that they follow from compression and selection as such, appearing in different material wherever those forces interact.

5.1 The Dictator's Dilemma

Authoritarian regimes represent the limiting case of dysmemic pressure. Preference divergence between ruler and subordinate is maximal. Punishment for deviation is severe. The fitness landscape for signals reaching the top is maximally misaligned with accuracy. The result is predictable. Communication degrades toward babbling equilibrium. The dictator becomes the least informed person in the room.

Consider the American invasion of Iraq in 2003. As coalition forces approached Baghdad, Saddam Hussein apparently believed the war was going well. His generals reported success. The regime was built on lies because everyone knew that Saddam severely punished anybody who told him unpleasant truths (Woods, Lacey, & Murray, 2006). At each interface, the fitness landscape selected for optimism.

The Russian invasion of Ukraine displayed similar dynamics. Putin apparently expected Kyiv to fall within days. Analysts of Russian intelligence culture have documented how Soviet-era institutional patterns created environments ill-suited to assessments that could be construed as critiques of senior officers or political leaders. Officers avoided analysis that might produce unwelcome conclusions.

Stalin had eighty-four warnings of Hitler's invasion available before Operation Barbarossa (Whaley, 1973). He ignored them. The modus operandi of his inner circle was sniff out, suck up, survive (Wirtz, 2023). The fitness landscape for signals reaching Stalin was

unambiguous. Accurate signals about threats were fatal for the messenger. The Wehrmacht arrived regardless.

The dictator's dilemma is that power to punish deviation is also power to destroy information. Democratic institutions function partly as Mirrors—opposition parties, free press, independent judiciary create observation points outside the ruling coalition's fitness landscape.

5.2 Institutional Lifecycles

Complex adaptive systems accumulate gaming strategies over time. Participants learn what the system rewards. They optimize. As optimization spreads, the gap between measured performance and actual function widens. The meta-game displaces the object-level game. Constitutions last, on average, nineteen years (Elkins, Ginsburg, & Melton, 2009). The mechanism predicts this: gaming strategies accumulate until external reality delivers correction that the internal frame cannot absorb.

Presidential term limits function as reset mechanisms. They force periodic disruption of information architecture surrounding the executive. Corporate succession attempts the same reset but typically fails because it replaces personnel while preserving architecture.

High-reliability organizations demonstrate that the ratchet can be interrupted. Aircraft carriers, nuclear plants, and air traffic control systems maintain function for decades by building Mirrors into operating structure—preoccupation with failure, reluctance to simplify, sensitivity to operations (Weick & Sutcliffe, 2007). The interval between Challenger (1986) and Columbia (2003) shows that even organizations that learned can forget.

6. Propositions and Falsification

A theory that cannot be wrong is not a theory. The framework generates predictions that can be tested.

6.1 Propositions

P1: Architectural change predicts turnaround success better than personnel change. Organizations that alter information architecture should outperform organizations that replace executives while preserving reporting structures, holding other factors constant. Operationalizing information architecture change requires coding changes to channel topology, compression operators, verification mechanisms, selection criteria, and protected variance—the five components defined in Section 2.

P2: Middle-ground successions fail at higher rates than alternatives. Successions coded as new leader, preserved architecture should underperform both comprehensive regime change and incumbent-led architectural reform on three-year performance measures.

P3: Mirror mechanisms accelerate diagnosis. Organizations with independent audit, external board access to operations, protected dissent channels, or internal prediction markets should identify dysfunctional structures faster than organizations relying on internal assessment alone.

P4: Vocabulary-only interventions decay. Interventions that change expressed norms without changing selection criteria should exhibit initial improvement followed by regression. Decay rate should correlate with gap between espoused criteria and actual selection criteria.

P5: Board decisions correlate with memetic fitness over empirical effectiveness. Boards should prefer strategies that are simple, legible, precedented, and defensible over strategies that are complex, ambiguous, novel, and difficult to document—regardless of outcome evidence.

6.2 Falsification Criteria

The theory would be weakened by evidence that middle-ground successions succeed at rates comparable to architectural interventions, that personnel characteristics predict turnaround outcomes as well as or better than architectural change, that vocabulary-only interventions produce durable behavioral change, or that high-compression high-divergence systems exhibit low dysfunction.

The theory would be falsified by evidence that dysfunction distributes randomly with respect to what metrics capture versus what metrics miss, or that reform succeeds durably without fitness-landscape change.

The null-space prediction deserves elaboration because it is the framework's most distinctive empirical claim. If dysfunction clusters in what metrics miss rather than distributing randomly, specific patterns should appear. Safety culture failures should be masked by incident metrics that count reportable events but miss near-misses and normalized risk. Customer churn should be masked by revenue metrics when price increases temporarily offset departures. Technical debt should accumulate behind velocity metrics that reward feature delivery without measuring maintenance burden. Fraud should hide behind growth metrics until the gap between reported and actual performance becomes too large to conceal.

In failed successions specifically, new leaders should optimize measured dimensions harder—the dimensions they can see—while unmeasured dimensions worsen. The pattern is testable with publicly available data: earnings calls, risk disclosures, employee reviews, customer satisfaction proxies, incident rates. Perfect measurement is unnecessary. The prediction is directional: dysfunction in the null space, improvement in the measured space, divergence over time.

6.3 The Meta-Prediction

The theory predicts its own competitive disadvantage. Change information architecture while preserving institutional knowledge is complex. Fire the CEO is simple. Transmission bias favors simple over complex.

This prediction is testable. Track adoption of different frameworks. The theory predicts personnel-focused interventions will dominate regardless of evidence, that architectural frameworks will be adopted primarily after external shock forces frame-breaking, and that regression to simpler approaches will follow as shock recedes. If the theory spread rapidly, achieving adoption rates comparable to simpler frameworks, the meta-prediction would be wrong.

7. Conclusion

Organizations change leaders to avoid changing themselves. This paper has argued that the pattern is equilibrium—the predictable output of selection dynamics operating on the response to dysfunction as surely as on the dysfunction itself.

The argument is structural, not psychological. Boards are not stupid. New leaders are not incompetent. The failure rate holds because the modal intervention cannot address a problem located in architecture rather than personnel. The information environment that trapped the predecessor traps the successor.

Two paths work. Comprehensive disruption clears accumulated gaming strategies at cost of institutional knowledge. Architectural reform preserves institutional knowledge at cost of time and risk of capture. Both succeed by changing what the system selects for. Both can fail if the direction of change is wrong—if new architecture selects against accuracy as reliably as old architecture did.

The middle ground fails because it changes neither. It satisfies the board's need to act without requiring the organization to change.

The mechanism extends beyond corporations. Dictators believe their own propaganda because the fitness landscape for signals reaching them selects against accuracy. Civilizations accumulate gaming strategies until external reality delivers correction. The substrate varies. The physics does not.

The framework does not counsel despair. High-reliability organizations demonstrate that the ratchet can be interrupted. Leaders who understand the mechanism can change what gets selected for. Structures that create observation outside local selection pressure can make visible what the Cage obscures.

But the framework does counsel realism. Most organizations will not build adequate Mirrors. Most boards will continue choosing strategies that fit their fitness landscape rather than strategies that work. The theory predicts this. Time will tell whether it is accurate.

The Cage is the equilibrium state of compression and selection operating at scale. The forces that produce it do not relent because you understand them. They persist, and the question becomes what structures can offset their effects. Within those structures, clarity becomes possible. Without them, the drift continues regardless of intention or intelligence.

What can be done is not nothing. Gerstner did it at IBM. Mulally did it at Ford. Joly did it at Best Buy. The common thread is not genius but recognition—seeing the Cage for what it is and changing what it selects for.

Simple explanations outcompete complex ones. The story that leadership is the problem and new leadership is the solution is more legible than the story that selection environments generate outcomes regardless of who inhabits them. The fit story wins. Whether it is also the true story is a separate question.

This paper has argued that it is not. The evidence is suggestive. The propositions are testable. The falsification criteria are specified. What remains is systematic analysis.

The theory may be wrong. If it is, the tests will reveal it. But sometimes correct ideas spread despite complexity. Sometimes the Cage is seen. Seeing is the beginning. It is not the end. But you cannot move walls you cannot see.

References

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

Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. University of Chicago Press.

Crawford, V. P., & Sobel, J. (1982). Strategic information transmission. *Econometrica, 50*(6), 1431–1451.

DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review, 48*(2), 147–160.

Elkins, Z., Ginsburg, T., & Melton, J. (2009). *The endurance of national constitutions*. Cambridge University Press.

Gamson, W. A., & Scotch, N. A. (1964). Scapegoating in baseball. *American Journal of Sociology, 70*(1), 69–72.

Gerstner, L. V. (2002). *Who says elephants can't dance? Inside IBM's historic turnaround*. HarperBusiness.

Hambrick, D. C., & Mason, P. A. (1984). Upper echelons: The organization as a reflection of its top managers. *Academy of Management Review, 9*(2), 193–206.

Hannan, M. T., & Freeman, J. (1984). Structural inertia and organizational change. *American Sociological Review, 49*(2), 149–164.

Hughes, M. (2011). Do 70 per cent of all organizational change initiatives really fail? *Journal of Change Management, 11*(4), 451–464.

Joly, H. (2021). *The heart of business: Leadership principles for the next era of capitalism*. Harvard Business Review Press.

McEntire, J. (2025). Dysmemic pressure: Selection dynamics in organizational information environments. *arXiv preprint* arXiv:2512.14716.

McEntire, J. (2026). Compression, selection, and organizational self-deception: A unified theory. *arXiv preprint*.

Meindl, J. R., Ehrlich, S. B., & Dukerich, J. M. (1985). The romance of leadership. *Administrative Science Quarterly, 30*(1), 78–102.

Schepker, D. J., Kim, Y., Patel, P. C., Thatcher, S. M. B., & Campion, M. C. (2017). CEO succession, strategic change, and post-succession performance: A meta-analysis. *The Leadership Quarterly, 28*(6), 701–720.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal, 27*(3), 379–423.

Staw, B. M., Sandelands, L. E., & Dutton, J. E. (1981). Threat rigidity effects in organizational behavior: A multilevel analysis. *Administrative Science Quarterly, 26*(4), 501–524.

Vaughan, D. (1996). *The Challenger launch decision: Risky technology, culture, and deviance at NASA*. University of Chicago Press.

Weick, K. E., & Sutcliffe, K. M. (2007). *Managing the unexpected: Resilient performance in an age of uncertainty* (2nd ed.). Jossey-Bass.

Whaley, B. (1973). *Codeword Barbarossa*. MIT Press.

Wirtz, J. J. (2023). Are intelligence failures still inevitable? *International Journal of Intelligence and CounterIntelligence, 36*(4), 1112–1132.

Woods, K., Lacey, J., & Murray, W. (2006). Saddam's delusions: The view from the inside. *Foreign Affairs, 85*(3), 2–26.