

THE JUDGMENT LAYER

An Inductive Theory of Understanding Synthesis Failure in Large Enterprises

Shubham Agarwal

Founder and Chief Executive Officer, triNetra Research

www.trinetra.life

Submission Draft, July 2026

Conflict of Interest Statement: The author declares no financial or non-financial conflicts of interest. No external funding was received. All empirical material derives from publicly accessible regulatory filings, judicial records, legislative testimony, enforcement decisions, and peer-reviewed secondary analyses.

Keywords: judgment layer, understanding synthesis gap, organizational failure, institutional reasoning chain, fragmentation paradox, inductive theory building, organizational sensemaking, organizational memory, high-reliability organizations

ABSTRACT

Why do large enterprises fail recurrently in domains for which they possess sophisticated operational systems, even after system-level remediation? This empirical pattern, documented across fifteen global organizations spanning six sectors and six geographic regions with aggregate direct financial consequences exceeding USD 80 billion, constitutes an unresolved puzzle in organizational theory.

This paper proposes Judgment Layer Theory (JLT) to explain this pattern. The judgment layer is the institutional architecture, designed or default, governing how an enterprise synthesizes the collective outputs of its operational systems into documented reasoning at consequential cross-domain decisions. When this architecture is absent, synthesis is informal and leaves no reconstructable institutional artifact. This paper terms this condition the understanding synthesis gap.

The central theoretical mechanism is this: post-failure remediation addresses the system-level deficiencies visible in individual system records without addressing the synthesis gap, which operates above the system level and is therefore invisible to system-level assessment. The structural preconditions for same-category failure are consequently preserved, producing the recurrence pattern. This mechanism is logically distinct from those proposed by organizational learning theory (Argyris and Schon, 1978), high-reliability organization theory (Weick and Sutcliffe, 2001), information processing theory (Galbraith, 1977; Tushman and Nadler, 1978), agency theory (Jensen and Meckling, 1976), and the Viable System Model (Beer, 1972).

Drawing on inductive regulatory document analysis structured through the Gioia methodology (Gioia, Corley, and Hamilton, 2013) and Eisenhardt's (1989) multiple-case logic, the paper develops three falsifiable propositions and a four-category synthesis gap taxonomy. Convergent validation is provided by commercial decision intelligence platform development (Gartner Magic Quadrant, 2024), the EU AI Act's Articles 12 and 13 mandating decision-level traceability, and concurrent academic governance engineering research independently identifying cross-system coupling documentation as the unaddressed gap in decision trace schemas.

TABLE OF CONTENTS

1. Introduction: The Empirical Puzzle
 2. Theoretical Background: What Existing Theory Explains and Where It Stops
 3. Research Design: Methodological Rationale and Analytical Procedure
 4. Case Selection and Empirical Material
 5. Inductive Theory Development: From Codes to Constructs
 6. Judgment Layer Theory: Constructs, Causal Model, and Propositions
 7. Cross-Case Analysis: Literal and Theoretical Replication
 8. Taxonomy of Understanding Synthesis Failure
 9. Alternative Explanations and Boundary Conditions
 10. Discussion: Theoretical Contribution and Implications
 11. Limitations and Future Research Directions
 12. Conclusion
 13. References
 14. Appendix A: Case Profiles and Evidentiary Sources
-

1. INTRODUCTION: WHY SOPHISTICATED ORGANIZATIONS FAIL RECURRENTLY IN DOMAINS THEY KNOW

In October 2018 and March 2019, two Boeing 737 MAX aircraft crashed, killing 346 people. Boeing possessed sophisticated flight control engineering systems, airworthiness certification processes, and regulatory compliance infrastructure. Those systems generated data relevant to the flight control software at issue. Following a Deferred Prosecution Agreement, a USD 2.5 billion legal resolution, and two years of mandated compliance remediation, a door plug separated from a Boeing 737 MAX 9 mid-flight in January 2024. The category of failure was the same. The institutional synthesis of available technical and regulatory information into documented, auditable decision reasoning had not occurred at the relevant decision points on either occasion.

In 2016, Wells Fargo's unauthorized account creation scheme was publicly exposed after operating across its retail network for years. Wells Fargo possessed branch performance monitoring systems, internal audit frameworks, and compliance oversight infrastructure. Following an eight-year Federal Reserve consent order and fourteen regulatory enforcement actions, a mortgage loan modification calculation error resulted in 545 customers losing their homes. The systematic failure persisted through years of documented remediation.

In 2012 and 2013, Goldman Sachs underwrote three bond transactions for 1Malaysia Development Berhad. Goldman possessed transaction monitoring systems, anti-money laundering platforms, and multi-layered compliance review frameworks. Those systems generated data relevant to the risk profile of the transactions. The Federal Reserve Board subsequently determined that Goldman had failed to detect or prevent the scheme or to address obvious risk indicators. In 2026, Goldman settled a shareholder class action alleging it had misrepresented the integrity of its risk management systems. Total documented financial consequence: approximately USD 3.4 billion across fourteen years.

These are not isolated incidents. They exemplify a systematic pattern documented across all fifteen organizations in this paper: consequential failures occurring and recurring in domains for which the organization possessed relevant operational systems. The pattern persists across aerospace, banking, healthcare, automotive, energy, retail, and technology sectors. It persists across six geographic regions and multiple regulatory regimes. It persists under active regulatory supervision, following documented remediation, in organizations whose individual operational systems were assessed as functioning within their designed parameters.

The pattern constitutes an empirical puzzle for existing organizational theory. Organizational learning theory predicts that experienced organizations encode corrective routines following failure (Levitt and March, 1988; Crossan, Lane, and White, 1999). High-reliability organization theory predicts that mindful organizations detect and address developing failure conditions before they cascade (Weick and Sutcliffe, 2001; Roberts, 1990). Information processing theory predicts that organizations facing high task uncertainty design structures to route relevant information to decision points (Galbraith, 1977; Tushman and Nadler, 1978). Agency theory predicts that appropriate incentive structures align managerial behavior with principal objectives (Jensen and Meckling, 1976). Normal accident theory attributes recurrence to immutable system architecture (Perrow, 1984), which cannot explain recurrence following structural remediation. Collectively, these traditions predict that well-resourced organizations operating relevant systems under active regulatory oversight should not exhibit the systematic recurrence pattern documented in this paper.

The explanation the paper proposes is structural. The organizational governance architectures of the enterprises examined lacked a defined institutional component for synthesizing the collective outputs of their specialized operational systems into documented reasoning at the point of consequential cross-domain decisions. The paper terms this missing component the judgment layer. The paper terms its absence the understanding synthesis gap. This paper proposes that the synthesis gap is the mechanism through which post-failure remediation fails to prevent same-category

recurrence: remediation addresses the failure symptoms visible in individual system records, while the synthesis gap operates above individual systems at the cross-domain decision interface, invisible to system-level remediation, and therefore structurally preserved.

This is Judgment Layer Theory. It does not replace existing theories of organizational failure. It proposes a structural mechanism that those theories, together, do not account for and that explains the specific phenomenon of post-remediation same-category recurrence. The institutional reasoning chain, the documented artifact produced when judgment layer architecture is present, is proposed as the organizational governance analogue of a class of institutional artifacts that have historically enabled reliable institutional action in complex systems: the double-entry account in commercial governance, the judicial precedent in legal governance, and the standard operating procedure in manufacturing governance. Each converted informal, individually variable synthesis into a documented, auditable institutional artifact. The institutional reasoning chain proposes to do the same for cross-domain organizational decision synthesis.

This paper addresses one research question: What structural organizational property explains the documented pattern of consequential failure occurring and recurring in organizations that possess relevant operational systems and have applied system-level remediation?

The following section reviews the theoretical traditions most directly relevant to the judgment layer, identifying the cumulative gap that JLT addresses.

2. THEORETICAL BACKGROUND: WHAT EXISTING THEORY EXPLAINS AND WHERE IT STOPS

The literature review is organized by the specific explanatory problem each tradition addresses relative to the empirical puzzle in Section 1. For each tradition, we identify what it explains, the boundary of its explanatory reach, and what remains unexplained. The cumulative gap across traditions identifies the theoretical space that JLT occupies.

2.1 Organizational Learning Theory

Argyris and Schon (1978) distinguish single-loop learning, error correction within existing governing variables, from double-loop learning, revision of the governing variables themselves. Levitt and March (1988) extend this to the population level, documenting how organizations encode inferences from experience into routines that guide behavior. Huber (1991) identifies four subprocesses of organizational learning: knowledge acquisition, information distribution, information interpretation, and organizational memory. Crossan, Lane, and White (1999) propose the 4I framework connecting individual insight to institutional embedding of new knowledge.

What organizational learning theory explains: How organizations improve performance through experience when feedback mechanisms are functional. How learning failures occur when feedback is distorted, delayed, or absent.

Where it stops: Organizational learning theory presupposes that the reasoning chain connecting experience, interpretation, decision, and outcome is available for retrospective analysis and encoding. It does not theorize the structural condition in which the synthesis of multiple system outputs into an institutional decision occurs informally and leaves no artifact. When no reasoning chain exists, organizational learning theory's encoding mechanisms have nothing to operate on.

What remains unexplained: Why organizations that have encoded prior failures into revised systems fail again in the same category after remediation, in the absence of evidence that feedback mechanisms are themselves deficient.

2.2 High-Reliability Organization Theory

Weick and Sutcliffe (2001) propose that high-reliability organizations maintain collective mindfulness through five properties: preoccupation with failure, reluctance to simplify interpretations, sensitivity to operations, commitment to resilience, and deference to expertise. Roberts (1990) and LaPorte and Consolini (1991) document these properties in nuclear aircraft carriers and air traffic control systems. Reason (1990) proposes the Swiss cheese model, in which active failures combine with latent conditions to produce accidents when defensive barriers are simultaneously breached.

What HRO theory explains: How organizations in high-consequence environments maintain reliability through attentional and behavioral practices. How the collapse of collective mindfulness contributes to failure.

Where it stops: HRO theory identifies collective mindfulness as a process property. It does not theorize what happens when collective mindfulness produces no institutional artifact. Reason's (1990) latent conditions construct locates failure preconditions within individual system components. JLT locates failure preconditions in the space between systems, specifically in the absence of synthesis architecture at the inter-system interface.

What remains unexplained: Why organizations that have remediated system-level deficiencies and subsequently demonstrate attentional compliance with regulatory oversight fail again in the same category.

2.3 Normal Accident Theory

Perrow (1984) argues that accidents are emergent properties of systems characterized by interactive complexity and tight coupling. The theory is structurally deterministic: certain system configurations will produce accidents regardless of management quality.

What normal accident theory explains: Why certain technological system architectures are accident-prone by design. Why accidents emerge from the interaction of individually manageable failures.

Where it stops: Normal accident theory does not address the organizational governance layer above the technological system. It cannot account for post-remediation recurrence in organizations that have modified their technical architectures unless the theory posits that the remediated architecture

is as fundamentally accident-prone as the original.

What remains unexplained: Why the same category of failure recurs in organizations that have applied structural remediation to identified failure sources, and why that recurrence persists under active regulatory supervision.

2.4 Organizational Sensemaking

Weick (1995) proposes that organizational action is preceded by sensemaking: the retrospective construction of plausible interpretations of equivocal stimuli. Weick, Sutcliffe, and Obstfeld (2005) extend this to organizing as constituted by ongoing cycles of sensemaking, enactment, and retention. Maitlis and Christianson (2014) identify conditions under which organized sensemaking processes succeed and fail.

What sensemaking theory explains: How organizations construct interpretations of ambiguous situations. How the breakdown of collective sensemaking contributes to organizational failure.

Where it stops: Sensemaking theory is a process theory. It theorizes how interpretation occurs. It does not theorize whether the interpretation process produces an institutional artifact and what the organizational consequences of artifact absence are. The organizations examined in this paper produced sensemaking in the functional sense: they generated interpretations of system outputs and acted on them. What was absent was any institutional artifact of the sensemaking that would permit reconstruction, audit, or revision after the fact.

What remains unexplained: Why organizations whose sensemaking processes produce organized action cannot reconstruct the reasoning that produced that action when post-failure reconstruction is required.

2.5 Organizational Memory

Walsh and Ungson (1991) identify five organizational memory retention facilities: individuals, culture, transformations (standard operating procedures, rules, and routines), structures, and ecology. Argote (1999) documents how organizations retain and transfer knowledge. Nonaka and Takeuchi (1995) distinguish tacit from explicit knowledge and propose the SECI model.

What organizational memory theory explains: How organizations retain and retrieve knowledge from prior experience. How knowledge transfer succeeds and fails across organizational boundaries.

Where it stops: Walsh and Ungson's (1991) five retention facilities all operate within defined organizational units or domains. No facility in their framework captures the synthesis of knowledge across multiple system domains at the point of a consequential cross-system decision. The specific absence of a cross-system synthesis retention facility is not theorized in their framework.

What remains unexplained: Why organizations with functional memory in individual system domains fail to produce institutional memory of the synthesis across those domains at the point of cross-domain decisions.

2.6 Agency Theory and Corporate Governance

Jensen and Meckling (1976) attribute organizational governance failures to incentive misalignment between principals and agents. Fama and Jensen (1983) distinguish decision management from decision control. Blair (1995) extends the stakeholder dimension of agency analysis.

What agency theory explains: How incentive misalignment produces behavioral divergence from principal objectives. How governance mechanisms partially address this divergence.

Where it stops: Agency theory is an intentionality theory. It does not account for the structural condition in which agents acting with appropriate incentives and good intentions nevertheless fail to produce documented synthesis of system outputs into institutional reasoning chains.

What remains unexplained: Why governance failures of the synthesis type persist in organizations that have implemented standard agency-theoretic remedies, including board oversight, audit committees, compliance programs, and incentive realignment.

2.7 Information Processing Theory

Galbraith (1977) proposes that when task uncertainty exceeds an organization's information processing capacity, performance degrades. Organizations can respond by either reducing information processing demands through goal setting and slack resources, or increasing information processing capacity through investment in lateral relations, vertical information systems, and organizational design. Tushman and Nadler (1978) extend this to organizational design: they propose that organizations should design their structures to match information processing requirements with information processing capacity. March and Simon (1958) establish that organizations are information processors and that organizational design is fundamentally about routing information to the right decision points.

What information processing theory explains: How organizations design structures to route information to decision points in alignment with task uncertainty. Why misalignment between information processing requirements and organizational information processing capacity produces performance degradation. How organizational design choices affect the adequacy of information flow to consequential decisions.

Where it stops: Information processing theory addresses the routing of information to decision points. It does not address whether the processing of information at the decision point produces a documented institutional artifact. An organization can achieve optimal information routing design, satisfying Galbraith's prescriptions entirely, while its decision-makers process routed information informally and produce no institutional record of that processing. Goldman Sachs's transaction monitoring architecture routed AML and KYC outputs to relevant decision-makers. The information routing was adequate. The processing of that routed information into specific transaction authorization decisions left no reconstructable institutional chain. Information processing theory cannot explain this failure because the failure was not in routing capacity or design but in synthesis documentation architecture at the decision point.

What remains unexplained: Why organizations with formally adequate information routing structures nevertheless fail to produce documented institutional reasoning from the information that arrives at decision points, and what architectural design addresses this documentation gap that is not captured in information processing theory's routing and capacity framework.

2.8 Dynamic Capabilities and Knowledge Integration

Grant (1996) proposes that competitive advantage derives from the organization's ability to integrate specialized knowledge. Teece, Pisano, and Shuen (1997) develop dynamic capabilities theory. Eisenhardt and Martin (2000) further specify dynamic capabilities as organizational and strategic routines.

What dynamic capabilities theory explains: How organizations integrate, build, and reconfigure knowledge and competencies. How sensing, seizing, and reconfiguring capabilities enable organizational adaptation.

Where it stops: Dynamic capabilities theory addresses strategic-level knowledge integration. It does not address the operational-level failure to document synthesis of available system outputs into institutional reasoning chains at specific consequential decisions.

What remains unexplained: How organizations with documented dynamic capabilities in sensing and seizing market opportunities simultaneously fail to synthesize available operational system outputs into documented governance decisions.

2.9 The Cumulative Theoretical Gap

The traditions reviewed collectively address: individual cognition, collective interpretation, system architecture, incentive structures, knowledge retention, and strategic capability reconfiguration. No existing theory directly addresses the following structural condition:

An enterprise operates multiple specialized operational systems that function adequately within their respective domains and generate data material to a developing failure. The institutional synthesis of those systems' collective outputs into a documented reasoning chain at the point of the consequential decisions preceding the failure does not occur or occurs only informally. When the failure manifests and reconstruction is attempted, no synthesis artifact is available. Post-failure remediation addresses individual system deficiencies. The synthesis gap at the inter-system decision interface is not addressed. The same category of failure recurs.

This structural condition is the theoretical focus of JLT. It is logically distinct from, and not fully reducible to, any of the theoretical constructs reviewed above. The borrowing relationships between JLT constructs and existing constructs are specified explicitly in Section 6.1.

3. RESEARCH DESIGN: METHODOLOGICAL RATIONALE AND ANALYTICAL PROCEDURE

3.1 Methodological Approach and Justification

This paper employs inductive regulatory document analysis as its primary methodology, organized using a coding hierarchy adapted from the Gioia methodology (Gioia, Corley, and Hamilton, 2013) and structured through Eisenhardt's (1989) multiple-case theory-building logic.

The primary data source for this paper is the corpus of regulatory enforcement actions, judicial records, legislative testimony, and consent order documentation produced by regulatory and judicial bodies investigating the fifteen organizational failure events in the sample. This data type was selected for two reasons. First, regulatory and judicial bodies possess investigative powers and internal document access that no external researcher possesses, making their findings the most authoritative available characterizations of internal organizational governance properties. Second, these documents contain explicit assessments of organizational governance architecture adequacy that directly bear on the understanding synthesis gap construct. The methodology is therefore labeled as inductive regulatory document analysis: a systematic coding and analysis of regulatory intelligence to identify patterns in organizational governance architecture across the sample.

The coding hierarchy is adapted from the Gioia methodology (Gioia, Corley, and Hamilton, 2013), which provides a structured procedure for moving from first-order empirical concepts to second-order theoretical themes to aggregate theoretical dimensions. The Gioia procedure was designed for interview-based research in which informant language constitutes the first-order data. This paper adapts the procedure for regulatory document analysis, treating regulatory findings language as the first-order data. This is methodologically non-standard and is acknowledged as a limitation: regulatory characterizations of organizational behavior are produced for regulatory purposes by external actors, and may not fully represent the internal organizational reality they describe. This adaptation is justified by the absence of any other available data source with equivalent access to internal organizational governance records across fifteen global organizations spanning six sectors and six geographic regions.

The Eisenhardt (1989) multiple-case logic governs case selection, within-case analysis prior to cross-case comparison, and the search for literal and theoretical replication. These requirements provide methodological safeguards against post-hoc imposition of theory on selected cases and together produce a theory grounded in empirical material, systematically constructed, and expressed in falsifiable propositional form.

3.2 Case Selection Criteria

Cases were selected according to explicit criteria applied prior to analysis.

Inclusion Criteria

IC1 (Scale): The organization is among the largest 100 enterprises by revenue in its region, as documented in public financial reporting in the annual period immediately preceding the documented failure. This criterion ensures resource constraints do not explain the synthesis gap.

IC2 (Documentation Quality): The failure event is documented in at least two independent primary sources drawn from: regulatory enforcement actions, consent orders, or civil money penalty decisions; court judgments or conviction records; parliamentary, congressional, or senate testimony transcripts; or government audit reports. This criterion ensures failures are not contested and the evidentiary basis is robust.

IC3 (System Presence): The organization is documented to have operated operational systems in the domain in which failure occurred at the time of the failure event. This criterion is essential: the paper does not theorize resource-constrained failures.

IC4 (Consequence Magnitude): The failure event produced a directly quantifiable documented financial consequence of no less than USD 100 million. This criterion ensures analytical focus on consequential failures.

IC5 (Synthesis Evidence): Public documentation is sufficient to permit analysis of both the operational systems present and evidence regarding the presence or absence of documented institutional reasoning connecting system outputs to consequential decisions.

Exclusion Criteria

EC1: Organizations excluded where the primary documented cause of failure was attributed exclusively to external market forces, macroeconomic conditions, or geopolitical events beyond organizational control.

EC2: Organizations excluded where the primary regulatory or judicial finding identified individual criminal misconduct as the sole proximate cause, without systemic governance or institutional synthesis findings.

EC3: Organizations excluded where documentary evidence was insufficient to evaluate IC3 or IC5.

Fifteen organizations satisfied all inclusion and no exclusion criteria. The sample was closed at fifteen when theoretical saturation was observed. Specifically, no new first-order codes emerged from the twelfth, thirteenth, fourteenth, and fifteenth cases analyzed, which were Toyota Motor Corporation (automotive, Asia Pacific), Saudi Aramco (energy, Middle East), TotalEnergies (energy, Europe), and Ford Motor Company (automotive, North America). These four cases were analyzed last specifically because they represented sectors and regions already represented in the sample, providing the most conservative test of saturation: new codes would be least likely to emerge from cases in already-represented categories, and therefore the absence of new codes from these cases is a stronger saturation indicator than their absence from novel cases would be (Strauss and Corbin, 1998).

Theoretical Rationale for Sample Size

Fifteen cases exceeds the minimum recommended for multiple-case theory building (Eisenhardt, 1989, recommends four to ten cases). The larger sample was retained to enable theoretical replication across six governance traditions and six regulatory regimes. Statistical representativeness is neither claimed nor required: the analytical purpose is theoretical, not statistical, generalization (Yin, 2018).

3.3 Data Sources and Evidence Hierarchy

Tier 1 (Highest Weight): Regulatory enforcement actions, consent orders, and civil money penalty decisions; court judgments and conviction records; congressional, parliamentary, and senate hearing transcripts; government audit and investigation reports; securities regulatory filings.

Tier 2 (Corroborating): Corporate annual reports and proxy statements filed with securities regulators; independent auditor reports; third-party assurance reports filed with regulatory bodies.

Tier 3 (Contextual Only): Peer-reviewed academic analyses of specific cases; major news organization reporting corroborated by at least one Tier 1 or Tier 2 source.

A source was included only if corroborated by at least one Tier 1 or Tier 2 document. Encyclopedic reference sources and uncorroborated commercial analyses are excluded from all citations.

3.4 Analytical Procedure

Stage 1: Within-Case Analysis. For each organization, a structured case narrative was developed from primary sources prior to cross-case comparison, documenting: (a) operational systems present in the failure domain; (b) the failure event and its regulatory or judicial characterization; (c) evidence for or against a documented institutional reasoning chain; (d) financial and operational consequences; and (e) evidence of recurrence following remediation.

Stage 2: First-Order Coding. Open coding of case narratives produced first-order codes describing empirical observations in language close to the source data (Gioia, Corley, and Hamilton, 2013).

Stage 3: Second-Order Theming. First-order codes were aggregated into second-order theoretical themes through iterative analysis.

Stage 4: Aggregate Dimension Identification. Second-order themes were aggregated into theoretical dimensions representing the highest level of abstraction in the Gioia hierarchy.

Stage 5: Cross-Case Pattern Analysis. Following Eisenhardt (1989), cross-case patterns were identified by seeking literal replication and theoretical replication.

Stage 6: Proposition Formulation and Literature Confrontation. Theoretical propositions were formulated from cross-case patterns, then explicitly compared with existing theory to identify convergence, divergence, and JLT's incremental explanatory contribution.

4. CASE SELECTION AND EMPIRICAL MATERIAL

The fifteen cases exhibit the theoretical property of interest across diverse sectors and governance contexts. Table 1 presents the sample with primary source documentation.

Table 1: Case Sample

Organization	Region	Sector	IC3 System Evidence	IC4 Consequence	Primary Tier 1 Source
Boeing	North America	Aerospace	FAA ODA Program; MCAS design systems	USD 2.5B+	DOJ DPA Case 4:21-cr-00005-O
UnitedHealth Group	North America	Healthcare	Enterprise cybersecurity infrastructure	USD 1.6B profit impact	Senate Finance Hearing May 1 2024
Goldman Sachs	North America	Financial Services	Transaction monitoring; KYC; AML platforms	USD 2.9B resolution	Federal Reserve Board Order Oct 22 2020
Wells Fargo	North America	Financial Services	Branch monitoring; audit systems	14 consent orders; USD 1.95T asset cap	OCC Consent Order 2021
Walmart	North America	Retail	Pharmacy management; CSA compliance	USD 3.1B opioid resolution	SEC 8-K Oct 18 2024
Volkswagen	Europe	Automotive	Emissions testing; engineering certification	EUR 31B+ total cost	EPA Notice of Violation Sep 18 2015
Deutsche Bank	Europe	Financial Services	AML transaction monitoring	USD 186M Fed penalty; EUR 23M BaFin	Federal Reserve Order July 2024
HSBC	Europe	Financial Services	Automated transaction monitoring	GBP 63.9M FCA penalty; USD 7.5B cumulative	FCA Decision Notice Dec 17 2021
Shell	Europe and Africa	Energy	Pipeline integrity monitoring	Open liability; 50,000 affected	UK High Court Judgment June 20 2025
TotalEnergies	Europe and Africa	Energy	Sustainability reporting systems	French court conviction	Tribunal de Paris Oct 23 2025
Toyota	Asia Pacific	Automotive	Quality certification across subsidiaries	9%+ Japan sales decline	MLIT Investigation Records 2024
Alibaba	Asia Pacific	Technology	Merchant relationship management	CNY 18.2B fine	SAMR Decision April 10 2021
Glencore	Global	Extractives	Agent payment systems; AML compliance	USD 1.5B+ multi-jurisdictional	DOJ Guilty Plea May 24 2022
Saudi Aramco	Middle East	Energy	Environmental and safety monitoring	Abqaiq attack 5.7Mbb/d loss	Aramco Annual Reports 2024 and 2025

Organization	Region	Sector	IC3 System Evidence	IC4 Consequence	Primary Tier 1 Source
Ford	North America	Automotive	Quality management; supplier; warranty systems	USD 165M NHTSA fine; USD 5B+ recall costs	Ford 10-K FY2025; NHTSA Database

Recurrence was observed in eleven of fifteen cases. Table 2 documents the recurrence evidence.

Table 2: Recurrence Evidence

Organization	Initial Failure	Remediation Applied	Recurrence Event	Interval
Boeing	2018 to 2019 737 MAX crashes	2021 DPA compliance program	2024 door plug blowout	Approximately 2 years post-DPA
Wells Fargo	2016 unauthorized accounts	2018 Federal Reserve consent order	2021 mortgage modification errors: 545 home losses	Approximately 3 years
HSBC	2012 DOJ settlement	USD 1.9B settlement and reforms	2021 FCA penalty for same AML categories 2010 to 2018	8 years overlapping failures
Deutsche Bank	2015 and 2017 Fed consent orders	Mandated remediation programs	2024 USD 186M Fed penalty; 2025 EUR 23M BaFin penalty	7 to 9 years post-orders
Volkswagen	2015 EPA notice	USD 4.3B DOJ resolution 2017	2025 EUR 1.5B EU emissions compliance fine	10 years of ongoing failures
Toyota	2009 to 2010 mass recalls	Cultural and process remediation	2023 to 2024 Daihatsu and Toyota Industries falsification	Approximately 14 years
Goldman Sachs	2020 1MDB resolutions	DOJ DPA compliance	2026 USD 500M shareholder settlement	Approximately 6 years
Ford	2019 to 2022 quality commitments	Multiple management initiatives	2025 record 152 recalls; reissued recalls	Persistent across commitment cycles
Walmart	2009 DEA investigation	No documented comprehensive remediation	2022 USD 3.1B opioid settlement	13 years
AT&T	2023 FCC consent decree for vendor breach	Mandated data governance improvements	2024 two discrete breaches	Less than 1 year
Glencore	2022 multi-jurisdictional resolution	DOJ compliance monitor; DPA	2025 to 2027 individual criminal trials ongoing	3 to 5 years post-resolution

5. INDUCTIVE THEORY DEVELOPMENT: FROM CODES TO CONSTRUCTS

5.1 First-Order Codes

Open coding of the fifteen case narratives produced the following recurring first-order concepts, presented in the language of primary regulatory and judicial sources:

“Failed to assess and manage risk on a sufficiently holistic basis” (FCA, Goldman Sachs, 2020). “Failed to implement a compliance program capable of preventing violations” (DOJ, Boeing, 2024). “Transaction monitoring controls showed serious weaknesses across three key systems” (FCA, HSBC, 2021). “Still in the process of bringing Change Healthcare’s cybersecurity protections to company standards” (CEO Witty, Senate, 2024). “Did not possess a robust mechanism to verify that software updates were correctly installed” (Ford 10-K FY2025). “30 years of forged safety test results” (MLIT, Daihatsu, 2023). “Bribery scheme spanned more than a decade across seven countries” (DOJ, Glencore, 2022). “DEA investigation initiated 2009; settlement 13 years later” (SEC 8-K, Walmart, 2024). “Defeat device software detected when a car was undergoing official emissions testing” (EPA, VW, 2015). “Certification irregularities at Toyota Industries following those at Daihatsu” (Toyota statement, 2024).

5.2 Second-Order Theoretical Themes

Aggregation of first-order codes produced five second-order theoretical themes:

Theme 1: System Capability Decoupled from Institutional Synthesis.

Across the cases examined, individual operational systems functioned within their designed domains. Regulatory findings consistently distinguished the adequacy of individual systems from the absence of holistic, cross-system synthesis. The failure was not located within any individual system. It was located in the organizational space between systems and decisions.

Theme 2: Synthesis Operates Informally and Produces No Institutional Artifact.

In each case, human agents synthesized system outputs in the course of decision-making. No institutional record of that synthesis was produced in a form available for post-hoc reconstruction. When regulatory authorities or courts sought to examine the reasoning that preceded consequential decisions, it could not be reconstructed from available documentation.

Theme 3: Remediation Targets Systems, Not Synthesis Architecture.

Post-failure remediation in the cases examined focused on individual system upgrades, new system acquisitions, procedural amendments within individual systems, or enhanced staffing in specific compliance domains. The evidence does not indicate that any remediation program in the sample addressed the absence of institutional architecture for documenting cross-system synthesis.

Theme 4: The Recurrence Signature.

In eleven of fifteen cases, the same category of failure recurred following documented system-level remediation. This pattern is inconsistent with theories that locate failure cause within individual systems, as system-level remediation would address those causes.

Theme 5: Regulators and Courts Identify Holistic Synthesis Failure.

Multiple primary regulatory and judicial sources explicitly identify the failure of holistic, cross-system synthesis as a proximate finding. The FCA's finding that Goldman Sachs failed to assess risk on a sufficiently holistic basis, the DOJ's finding that Boeing's compliance program failed to prevent violations rather than merely documenting policies, and the FCA's finding that HSBC's transaction monitoring showed serious weaknesses across three separate systems in the same eight-year period collectively constitute independent regulatory identification of synthesis layer failure, providing triangulation for Theme 1.

5.3 Aggregate Theoretical Dimensions

The five second-order themes converge on three aggregate theoretical dimensions:

Dimension 1: The Fragmentation Condition.

Large enterprises structurally distribute operational knowledge across specialized systems. This distribution is a designed property of modern enterprise architecture, appropriate for individual system optimization. It produces an organizational condition in which no single system possesses the complete informational basis for consequential cross-domain decisions. This condition is a necessary antecedent of the synthesis gap but is not itself the synthesis gap.

Dimension 2: The Synthesis Absence Condition.

The institutional architecture for documenting cross-system synthesis at the point of consequential decisions is absent as the structural default in the organizations examined. Individual systems produce outputs. Those outputs flow to human decision-makers. The synthesis performed by decision-makers is informal, ephemeral, and produces no institutional artifact. This is the synthesis gap.

Dimension 3: The Recurrence Mechanism.

Post-failure remediation in the cases examined addresses visible system deficiencies without addressing the synthesis absence condition. Because the synthesis gap is not addressed, the structural preconditions for the same category of failure are preserved. New instances of the same category emerge from new combinations of conditions.

Figure 1 depicts the recommended coding hierarchy diagram for the published paper.

Figure 1 Description: A three-level hierarchy diagram showing first-order concepts (regulatory language from primary sources) at the base, aggregating upward to five second-order themes, then to three aggregate dimensions, then to the four theoretical constructs of JLT. Arrows indicate the analytical path from data to theory.

6. JUDGMENT LAYER THEORY: CONSTRUCTS, CAUSAL MODEL, AND PROPOSITIONS

6.1 Core Constructs

Construct 1: The Judgment Layer

Definition: The judgment layer is the organizational space between the collective knowledge produced by an enterprise's operational systems and the institutional decisions made from that knowledge. It is the locus where human agents synthesize fragmented system outputs into consequential decisions.

Key Property: Its analytically critical property is its documentation requirement: whether the institutional architecture requires synthesis to produce a documented artifact (judgment layer present) or permits synthesis to occur informally without documentation (synthesis gap). This is a binary structural property at the architectural design level, though it admits of degrees of adequacy at the implementation level.

Distinction from Existing Constructs: The judgment layer is related to Weick's (1995) sensemaking but is architecturally rather than process-focused: sensemaking asks how interpretation occurs; JLT asks whether interpretation produces a documented institutional trace. The judgment layer is related to Walsh and Ungson's (1991) retention facilities but represents a cross-system synthesis layer not captured by any of their five facilities.

Observable Indicators: Presence of documented cross-system synthesis at defined decision gates; ability of organizations to produce coherent accounts of cross-domain decision reasoning when required by regulatory or judicial authorities; existence of formal cross-system review protocols.

Construct 2: The Understanding Synthesis Gap

Definition: The understanding synthesis gap is the structural condition in which the judgment layer produces no documented reasoning chain connecting the collective outputs of operational systems to institutional decisions. It is distinct from cognitive understanding failure: decision-makers may possess adequate individual understanding. The gap refers specifically to the absence of institutional documentation in a form that permits reconstruction, audit, and organizational learning.

Key Property: Binary in its foundational form (present or absent) but admits of degrees in practice. The critical threshold for the purposes of this paper is the minimum documentation required to permit regulatory reconstruction of institutional reasoning.

Distinction from Existing Constructs: Related to Reason's (1990) latent conditions but distinct: Reason locates failure preconditions within individual systems; JLT locates them at the inter-system synthesis interface. Related to Huber's (1991) information interpretation failure but distinct: Huber addresses organizational interpretation of environmental signals; JLT addresses institutional documentation of cross-system synthesis within the organization.

Observable Indicators: Regulatory or judicial findings that the organization could not produce coherent decision reasoning accounts; internal audit findings regarding absent cross-system decision documentation; post-failure reconstruction failures documented in enforcement actions.

Construct 3: The Institutional Reasoning Chain

Definition: The institutional reasoning chain is the documented artifact connecting: (a) the specific system outputs available to decision-makers at the point of a consequential cross-domain decision; (b) the synthesis of those outputs into the collective understanding that informed the decision; (c) the decision itself, including its authorization chain; and (d) the expected and actual outcomes against the documented expectation.

Key Property: Its completeness is evaluable against a minimum standard: it must be sufficient to permit reconstruction of the decision logic by an independent party with access to the same system outputs available to the original decision-makers.

Distinction from Existing Constructs: Distinct from individual system audit trails, which document what a system did. The institutional reasoning chain documents what the organization understood from the collective outputs of multiple systems and why it decided as it did. Related to Argyris and Schon's (1978) action theory but distinct: espoused theory documents what organizations claim to do; the institutional reasoning chain documents the synthesis reasoning that actually produced specific decisions.

Observable Indicators: Formal decision memoranda documenting cross-system evidence synthesis; regulatory submissions demonstrating the chain of evidence considered; governance committee minutes documenting cross-domain synthesis at defined decision gates.

Construct 4: The Fragmentation Paradox

Definition: The fragmentation paradox is the organizational dynamic in which investment in individual operational systems increases the volume of data available to the judgment layer without proportionally increasing the institutional architecture for documenting the synthesis of that data. Under the fragmentation paradox, increased system investment, absent proportional synthesis architecture investment, increases rather than reduces institutional vulnerability in cross-system decision domains.

Key Property: A dynamic property that intensifies over time as system portfolios grow through organic development and acquisition. Most pronounced in organizations with high system counts, frequent acquisition activity, or geographically dispersed operations.

Distinction from Existing Constructs: Related to Grant's (1996) knowledge integration challenge but distinct: Grant's challenge concerns strategic knowledge integration for competitive advantage; the fragmentation paradox concerns the operational-level failure of synthesis documentation that persists even when strategic knowledge integration is functioning.

6.2 Causal Model

The complete causal model of JLT is presented below. Each arrow represents a theorized causal mechanism.

Figure 2 Description for published paper: A vertical flow diagram with eight levels. Level 1: Operational Systems (multiple, specialized, functionally fragmented). Arrow 1 label: System outputs flow to judgment layer. Level 2: Judgment Layer (synthesis occurs informally; no institutional architecture present). Arrow 2 label: Synthesis absence condition produces no documented artifact. Level 3: Understanding Synthesis Gap (structural absence of institutional reasoning chain). Arrow 3 label: Decisions made without documented synthesis. Level 4: Consequential Decision (cannot be reconstructed post-hoc). Arrow 4 label: Decision produces outcome. Level 5: Failure Event (organizational, financial, regulatory, or reputational consequence). Arrow 5 label: Post-failure remediation targets systems, not synthesis architecture. Level 6: System-Level Remediation (synthesis gap preserved). Arrow 6 label: Structural conditions for same category failure preserved. Level 7: Recurrence. Arrow 7 (conditional, labeled as moderator): If synthesis architecture is developed at Arrow 5, recurrence probability is reduced; feedback loop to organizational learning is created. Level 8 (conditional): Organizational Learning (conditional on institutional reasoning chain availability).

Arrow Explanations:

Arrow 1: Operational systems generate outputs that flow, through reporting and monitoring processes, to the human decision-makers who constitute the judgment layer. This flow is continuous and adequate within individual system domains.

Arrow 2: Because no institutional architecture documents the synthesis of system outputs at the judgment layer, synthesis occurs informally. No institutional artifact is produced. This is the structural condition constituting the understanding synthesis gap.

Arrow 3: Consequential decisions are made from informal synthesis, without documented reasoning chains. The decisions may be fully informed at the individual cognitive level. The institutional architecture for documenting that reasoning is absent.

Arrow 4: Decisions produce outcomes. When undocumented synthesis is deficient, the outcome may constitute a failure.

Arrow 5: Post-failure investigation identifies system-level deficiencies as the proximate cause visible in evidentiary records. Remediation addresses those deficiencies. The synthesis gap, which is not captured in individual system records, is not identified as requiring remediation.

Arrow 6: Because the synthesis gap is preserved, the structural preconditions for the same category of failure are preserved.

Arrow 7 (Conditional): If post-failure remediation includes the development of synthesis architecture, the institutional reasoning chain is produced at subsequent decision points. This creates an artifact from which organizational learning can proceed. Recurrence probability is correspondingly reduced.

6.3 Theoretical Propositions

The following three propositions are expressed in falsifiable form consistent with AMR and ASQ standards (Bacharach, 1989; Whetten, 1989).

Proposition 1 (Synthesis Gap and Failure Incidence):

Among organizations operating in cross-domain decision environments with documented operational systems in the relevant domains, the probability of consequential failure is positively associated with the degree of understanding synthesis gap in those domains, controlling for organizational size, industry, regulatory environment, and individual system investment levels.

Theoretical Logic: Flowing from Dimension 2 (Synthesis Absence Condition) and Theme 1 (System Capability Decoupled from Synthesis), consequential failures in cross-domain environments are produced not by the absence of relevant systems but by the absence of documented synthesis of their outputs.

Boundary Conditions: Applies in cross-domain decision environments where at least two operationally independent systems with relevant outputs are engaged by the decision. Less applicable to single-domain decisions.

Measurable Variables: The understanding synthesis gap is measurable through: completeness of cross-domain decision documentation; regulatory assessment of synthesis adequacy; organizational ability to produce reconstruction-adequate accounts of cross-domain decisions upon request.

Falsifiability Condition: Proposition 1 would be disconfirmed if comparative analysis of organizations with documented synthesis architecture and organizations without it shows no statistically significant difference in consequential failure rates in cross-system decision domains, controlling for the specified covariates.

Proposition 2 (Remediation Type and Recurrence):

Among organizations that have experienced a consequential failure and applied remediation, those that apply exclusively system-level remediation without addressing synthesis architecture will exhibit significantly higher rates of same-category failure recurrence within five years compared to organizations that also develop synthesis architecture as part of the remediation program.

Theoretical Logic: Flowing from Dimension 3 (Recurrence Mechanism) and Theme 3 (Remediation Targets Systems, Not Architecture), system-level remediation preserves the structural preconditions for recurrence. Organizations that address synthesis architecture eliminate the condition that enables new circumstances to produce the same failure category.

Boundary Conditions: Applies where the failure category involves cross-domain synthesis. Less applicable where the failure is traceable to a single system deficiency.

Measurable Variables: System-level remediation is operationally defined as remediation documented in regulatory consent orders or corporate disclosures that addresses individual system deficiencies without documented cross-system synthesis architecture development. Recurrence is a regulatory finding, consent order, or judicial decision in the same failure category within five years of documented remediation.

Falsifiability Condition: Proposition 2 would be disconfirmed if longitudinal analysis shows that organizations applying exclusively system-level remediation do not exhibit higher recurrence rates than those that also address synthesis architecture, over a five-year observation period.

Proposition 3 (Fragmentation Paradox):

Among organizations of comparable size and industry classification, the probability of an understanding synthesis gap in cross-domain decision environments is positively associated with the number of specialized operational systems maintained, *ceteris paribus*, unless organizational investment in documented synthesis architecture scales proportionally with system portfolio growth.

Theoretical Logic: Flowing from Dimension 1 (Fragmentation Condition) and the Fragmentation Paradox construct, each additional system increases data volume and complexity available to the judgment layer without automatically increasing synthesis documentation architecture.

Boundary Conditions: Applies after a minimum organizational complexity threshold at which cross-domain synthesis becomes routinely required for consequential decisions.

Measurable Variables: System count measurable through enterprise architecture inventories. Synthesis architecture investment measurable through governance documentation completeness, cross-domain decision protocol presence, and regulatory assessment of synthesis adequacy.

Falsifiability Condition: Proposition 3 would be disconfirmed if cross-sectional analysis shows no positive relationship between system count and understanding synthesis gap probability, controlling for organizational age, size, and industry complexity.

7. CROSS-CASE ANALYSIS: LITERAL AND THEORETICAL REPLICATION

7.1 Literal Replication

The understanding synthesis gap is predicted to be present in all fifteen cases given that all satisfy IC3 (documented systems in failure domain) and IC5 (insufficient documented institutional synthesis). The following structural pattern is observed across all fifteen cases:

Operational systems were present and functional in the failure domain. System-level data material to the failure was generated prior to the failure event. No documented institutional reasoning chain connecting system outputs to consequential decisions was available for post-failure reconstruction, as evidenced by the inability of organizations to produce such chains when required by regulatory, judicial, or legislative authorities. Post-failure remediation focused on individual system deficiencies. Same-category recurrence occurred in eleven of fifteen cases within the observation window.

This pattern constitutes literal replication across the full sample.

7.2 Theoretical Replication

Three theoretically predicted sub-patterns are identified:

Sub-Pattern A: Integration Boundary Synthesis Gap.

Theory predicts that organizations undergoing acquisition or managing subsidiary networks should exhibit synthesis gaps at the integration boundary: the acquiring or parent entity's synthesis architecture does not extend to the acquired or subsidiary entity within the timeframe required. This produces a temporally bounded, boundary-specific gap.

Observation: UnitedHealth Group's CEO testified that synthesis of Change Healthcare's security posture against UnitedHealth's own standards was incomplete 16 months post-acquisition. Toyota's parent-level governance failed to detect Daihatsu subsidiary falsification over a 30-year period. Both cases exhibit the predicted boundary-specific, temporally extended pattern, consistent with the theoretical prediction.

Sub-Pattern B: Distributed Network Policy-Execution Gap.

Theory predicts that geographically distributed organizations should exhibit synthesis gaps at the policy-execution interface: the synthesis architecture connecting corporate policy to distributed operational practice is absent.

Observation: Walmart's CSA compliance policy existed at the corporate level. Pharmacy-level operational data across 5,000 locations was generated for 13 years without triggering adequate corrective synthesis. Wells Fargo's sales practice policies existed at the corporate level. Branch-level behavioral data was generated without triggering adequate synthesis-level detection. Both cases exhibit the predicted distributed, policy-execution gap pattern.

Sub-Pattern C: Functional Synthesis Gap.

Theory predicts that in cases involving deliberate undisclosed practices, the synthesis gap functions in a pathological sense: informal synthesis permits decisions that would not survive documentation.

Observation: Volkswagen's defeat device software operated for years without triggering a documented institutional decision to disclose it to regulators. Glencore's agent payment arrangements were maintained for over a decade across seven countries without triggering

documented institutional decisions to discontinue or disclose them. Alibaba’s exclusive dealing arrangements were maintained without triggering documented institutional assessment of their antitrust implications. All three exhibit the predicted functional gap pattern.

Note: This sub-pattern does not invalidate JLT as a structural theory. The same architectural absence that enables accidental synthesis failure also enables deliberate synthesis gap maintenance. JLT encompasses both manifestations within the same theoretical construct.

7.3 Rival Explanations in the Cross-Case Evidence

Against Agency Theory: If synthesis gaps were primarily products of deliberate agent self-interest, they should be concentrated in cases involving documented fraudulent intent. Sub-Pattern C cases (Volkswagen, Glencore, Alibaba: three of fifteen) are the candidates most consistent with agency theory. The remaining twelve cases exhibit synthesis gaps in the absence of documented concealment motivation, which is inconsistent with agency theory as the primary explanation for the full sample.

Against HRO Theory: If synthesis gaps were primarily products of mindfulness failure, organizations with documented attentional compliance under regulatory monitoring should not exhibit recurrence. Of the eleven recurrence cases, the evidence suggests that nine occurred during periods of active regulatory supervision. Attentional compliance appears necessary but not sufficient to address synthesis gaps.

Against Normal Accident Theory: If synthesis gaps were primarily products of tight coupling and interactive complexity, they should be concentrated in technologically complex sectors. The synthesis gap pattern is observed equally in retail (Walmart), financial services (Goldman Sachs, Wells Fargo, HSBC, Deutsche Bank), technology (Alibaba), and industrial sectors, suggesting the pattern is not sector-specific.

8. TAXONOMY: MANIFESTATIONS OF JUDGMENT LAYER ABSENCE

The grounded analysis produces a four-category taxonomy of judgment layer absence manifestations, defined by the specific architectural interface at which the synthesis gap is most pronounced and by the regulatory or judicial signature it produces in the empirical record. These categories are not failure types in the conventional sense. They are architectural failure modes: patterns through which the absence of a defined institutional synthesis architecture expresses itself in observable organizational and regulatory outcomes. The categories are not mutually exclusive.

Table 3: Taxonomy of Understanding Synthesis Failure

Category	Definition	Structural Layer	Regulatory Signature	Primary Cases
1: Judgment Reconstruction Failure	Organization cannot reconstruct institutional reasoning that produced a consequential decision	Decision synthesis layer	Inability to produce coherent decision account when required by regulatory or judicial authority	Goldman Sachs; Boeing; Volkswagen; Glencore; Alibaba
2: Cross-System Integration Failure	Failure propagates across organizational boundaries because synthesis of cross-boundary system outputs was not documented	Cross-boundary synthesis layer	Post-acquisition failures; subsidiary governance gaps; vendor ecosystem breaches	UnitedHealth Group; Toyota; Walmart; Deutsche Bank
3: Policy-to-Execution Translation Failure	Policy stated at corporate level does not translate to consistent execution; synthesis connecting policy to execution is absent in both directions	Policy-execution interface	Consent order violations; documented divergence between policy and operational practice	Wells Fargo; Boeing; Walmart; Volkswagen; HSBC; TotalEnergies
4: Feedback Loop Failure	System outputs material to developing failure not synthesized into documented institutional understanding that produces corrective action	Feedback synthesis layer	Post-hoc analysis reveals failure was detectable from data already in organizational possession	Ford; HSBC; Deutsche Bank; Saudi Aramco; Toyota

Table 4: Construct Summary and Nomological Network

Construct	Definition	Key Property	Antecedents	Consequences	Relationship to Existing Constructs
Judgment Layer	Organizational space between collective system knowledge and institutional decisions	Presence or absence of synthesis documentation architecture	Organizational complexity; system specialization; acquisition growth	Understanding synthesis gap if documentation absent; institutional reasoning chain if documentation present	Extends Weick (1995) sensemaking; adds artifact requirement
Understanding Synthesis Gap	Structural absence of documented reasoning chain	Binary presence or absence; degree of completeness	Judgment layer without documentation architecture; fragmentation condition	Failure incidence; post-remediation recurrence; learning failure	Extends Walsh and Ungson (1991); identifies untheorized sixth retention facility
Institutional Reasoning Chain	Documented artifact connecting system outputs to decisions and outcomes	Reconstruction adequacy; regulatory verifiability	Synthesis documentation architecture in judgment layer	Organizational learning; regulatory compliance demonstration; failure prevention	Distinct from individual system audit trails; extends Argyris and Schon (1978)
Fragmentation Paradox	Dynamic in which system investment increases synthesis gap absent proportional architecture investment	Intensifies with system count growth and acquisition activity	System portfolio growth; absent synthesis architecture investment	Increased understanding synthesis gap probability	Distinct from Grant (1996) knowledge integration; applies at operational decision level

9. ALTERNATIVE EXPLANATIONS AND BOUNDARY CONDITIONS

9.1 Systematic Engagement with Alternative Explanations

Table 5: Alternative Explanations and JLT's Incremental Explanatory Contribution

Theory	Core Claim	Empirical Fit with Sample	JLT's Incremental Contribution
Agency Theory (Jensen and Meckling, 1976)	Incentive misalignment produces behavioral divergence	High fit for 3 of 15 Sub-Pattern C cases	Explains 12 of 15 cases where concealment motivation is absent; identifies structural synthesis gap independent of incentive alignment
Organizational Learning (Argyris and Schon, 1978; Levitt and March, 1988)	Organizations encode experience into revised routines	Moderate fit; explains some initial failures	Addresses why learning fails even when feedback mechanisms are present: no synthesis artifact exists from which to learn
HRO Theory (Weick and Sutcliffe, 2001)	Mindfulness collapse produces failure	Moderate fit; applies where attentional failure is documented	Addresses why organizations with compliance attestations under regulatory monitoring still exhibit recurrence
Normal Accident Theory (Perrow, 1984)	Complex coupled systems produce emergent failure	Low fit for post-remediation recurrence	Explains recurrence without requiring immutable system architecture
Sensemaking Theory (Weick, 1995)	Equivocality breakdown produces failure	Moderate fit	Addresses what happens when sensemaking produces action but leaves no institutional artifact
Organizational Memory (Walsh and Ungson, 1991)	Retention failure produces knowledge loss	Moderate fit	Identifies the specific cross-system synthesis retention facility absent from Walsh and Ungson's five-facility framework
Dynamic Capabilities (Teece et al., 1997)	Sensing and seizing failures produce maladaptation	Low fit for operational failures	Addresses the operational decision-level gap that dynamic capabilities theory does not reach
Institutional Theory (DiMaggio and Powell, 1983)	Isomorphic pressures produce symbolic compliance	Moderate fit	Provides mechanism explaining why symbolic compliance persists: synthesis gap makes actual compliance unverifiable internally
Behavioral Theory (Cyert and March, 1963)	Bounded search and satisficing produce suboptimal outcomes	Moderate fit	Synthesis architecture extends the effective search boundary of bounded decision-makers
Complex Adaptive Systems (Holland, 1992)	Emergent properties produce unpredictable outcomes	Low fit as primary explanation	Identifies synthesis architecture as a designed counter to emergent synthesis failure in complex organizational systems

9.2 Boundary Conditions

Table 6: Boundary Conditions for JLT

Condition Type	Specification
Scope: Organizational Size	Most applicable to large, complex organizations operating multiple specialized operational systems. Less applicable to small organizations where single-domain decisions dominate.
Scope: Decision Type	Applies to decisions requiring synthesis across the outputs of two or more operationally independent systems. Less applicable to single-domain decisions.
Scope: Consequence Stakes	Most relevant to consequential decisions with significant demonstrable consequences. Less directly applicable to low-stakes decisions where synthesis gap costs are negligible.
Necessary Condition 1	Multiple operational systems generating relevant outputs in the failure domain.
Necessary Condition 2	Cross-domain decisions requiring synthesis of those outputs.
Necessary Condition 3	Absence of documented synthesis architecture at the relevant decision interfaces.
Sufficient Conditions for Failure	Necessary conditions above, plus: developing failure conditions detectable from available system outputs, and absence of independent regulatory or audit mechanisms performing the synthesis function.
Where JLT May Not Apply	Single-domain operational architectures; failures traceable exclusively to external shocks; small organizations where the judgment layer is functionally equivalent to the individual decision-maker.
Moderating Conditions	Organizational size; system portfolio size and growth rate; acquisition frequency; geographic distribution of operations; regulatory environment specificity.

10. DISCUSSION: THEORETICAL CONTRIBUTION AND IMPLICATIONS

10.1 What JLT Contributes to Organizational Theory

JLT makes four theoretically distinct contributions.

Contribution o: The Judgment Layer as a Missing Enterprise Architectural Component.

Before specifying JLT's contributions to organizational failure theory, its primary contribution should be stated precisely: JLT identifies the judgment layer as a missing architectural component in contemporary enterprise governance design. Just as double-entry bookkeeping identified a missing institutional artifact in medieval commercial governance, and just as standard operating procedures identified a missing replication artifact in industrial manufacturing governance, JLT proposes that modern complex enterprises lack a defined architectural component for cross-domain synthesis documentation at consequential decision interfaces.

This architectural contribution is prior to the failure theory contributions that follow. The failure theory contributions explain what happens when the architectural component is absent. The architectural contribution identifies what the component is and where it belongs in enterprise

design.

Contribution 1: A Novel Construct Logically Prior to Existing Failure Mechanisms.

The judgment layer construct introduces an organizational-level architectural variable that is logically prior to the process-level mechanisms theorized in sensemaking, organizational learning, and HRO theory. Each of those theories assumes that the organizational decision space produces some form of institutional artifact, whether encoded routines, revised mental models, or collective attentional outputs. JLT identifies the structural condition in which no such artifact is produced and proposes this condition is the structural default, not the exception, in large complex organizations.

This reframes the level of analysis. Prior theories ask how processes within the decision space succeed or fail. JLT asks whether the decision space is designed to produce institutional artifacts at all. The evidence examined in this paper suggests it typically is not.

Contribution 2: A Mechanism Explaining Post-Remediation Recurrence.

The recurrence pattern observed across eleven of fifteen cases constitutes the strongest evidence for JLT's incremental explanatory contribution. Existing theories explain initial failure adequately. They do not explain why organizations that have identified failure causes, applied remediation, and demonstrated compliance to regulatory authorities fail again in the same category.

JLT explains this through a specific mechanism: system-level remediation addresses failure symptoms visible in individual system records. The synthesis gap, operating at the inter-system layer above individual systems, is not visible in individual system records and is therefore not addressed. The structural preconditions for same-category failure are consequently preserved.

Contribution 3: The Fragmentation Paradox as a Named Organizational Dynamic.

The fragmentation paradox names and theorizes a previously unnamed dynamic: the tendency of system investment, absent proportional synthesis architecture investment, to increase rather than decrease institutional vulnerability in cross-domain decision environments. This dynamic has practical significance for organizational governance and investment that has not been captured in prior frameworks.

10.2 The “So What” Question

If JLT is correct, three significant implications follow for organizational theory.

First, the appropriate unit of analysis for organizational failure research must expand beyond individual systems, processes, and cognitive biases to include the institutional architecture of the judgment layer itself. Papers in organizational learning, HRO theory, and sensemaking that treat this architecture as given may be studying the quality of the content within a vessel while the vessel itself is absent.

Second, post-failure remediation research should investigate whether remediation programs address the synthesis architecture rather than only individual system deficiencies. The eleven recurrence cases in this paper suggest that this distinction is not currently made in either organizational practice or academic analysis of remediation effectiveness.

Third, governance theory should incorporate synthesis architecture adequacy as a governance property distinct from, and complementary to, the properties addressed by current governance frameworks, including COSO, ISO 31000, and standard audit committee charters.

10.3 Implications for Enterprise Architecture Design

JLT's most direct practical contribution is to enterprise governance architecture design. Contemporary enterprise architecture frameworks (Zachman, 1987; TOGAF; COBIT; COSO) define architectural layers from technology infrastructure through data management through application systems through business process through business strategy. None of these frameworks includes a defined architectural layer for cross-domain institutional synthesis documentation.

JLT proposes that enterprise governance architects should define the judgment layer as an explicit component, specifying: (1) which decision domains require cross-domain synthesis documentation; (2) what the minimum viable content of an institutional reasoning chain is for each decision domain; (3) what governance roles are accountable for synthesis documentation; (4) what audit mechanisms verify synthesis documentation adequacy; and (5) how synthesis documentation feeds back into organizational learning processes.

This design specification is analogous to how enterprise architects specify information architecture requirements (what data must be captured, in what format, with what access controls) for operational systems. The judgment layer specification addresses the analogous requirements for cross-domain synthesis documentation at consequential decision interfaces.

10.4 Implications for Regulatory Practice

The regulatory direction observed across multiple jurisdictions in the cases examined, toward requiring demonstration of holistic, cross-system governance rather than individual system compliance, is theoretically well-founded under JLT. JLT provides the theoretical foundation for making this regulatory direction explicit, systematic, and measurable.

An important implication for corporate governance law requires explicit statement. The Delaware Caremark standard (In re Caremark International, Inc. Derivative Litigation, 1996) requires boards of directors to maintain information and reporting systems sufficient to enable the board to be informed of organizational risk and compliance. The Caremark standard is satisfied by the existence of adequate information systems and adequate information flow to the board. JLT proposes a governance requirement that is logically distinct from and more specific than Caremark: not that information reach the board, but that the synthesis of information from multiple operational systems into specific institutional decisions at defined decision interfaces produce a reconstructable reasoning chain. The Boeing governance failure is illustrative: the Delaware Chancery Court found

that Boeing's board did not maintain safety oversight at the highest level of the corporate governance structure, despite the existence of information systems and board-level reporting. The failure was not that information did not reach the board. The failure was that no committee charter assigned responsibility for overseeing the synthesis function connecting operational safety systems to institutional safety decisions. The Caremark standard does not require synthesis architecture. JLT implies it should. This distinction is a specific contribution to corporate governance law scholarship and constitutes a recommendation for Caremark doctrine development.

10.5 Implications for Investment Due Diligence

The aggregate documented direct financial consequence of understanding synthesis failure across the fifteen cases examined exceeds USD 80 billion. Standard investment due diligence and credit rating frameworks do not assess synthesis architecture as a governance property. JLT implies that synthesis architecture assessment should be incorporated as a material governance risk factor in investment due diligence and credit risk frameworks.

10.6 Practitioner and Regulatory Convergent Validation

A theory's claim to identify a genuine structural gap is strengthened when independent actors, operating without knowledge of the theoretical framework, converge on the same gap from different directions. Three independent streams of convergent evidence have emerged concurrent with the development of this paper.

Stream 1: Commercial Platform Development

Cloverpop, a San Francisco-based enterprise decision intelligence platform recognized in Gartner's 2024 Market Guide for Decision Intelligence Platforms and the inaugural Gartner Magic Quadrant for Decision Intelligence Platforms, has independently arrived at a product architecture that closely corresponds to the judgment layer construct. Cloverpop characterizes its core offering as the Decision Layer, described explicitly as "the layer that connects AI intelligence to organizational action" and positioned as a missing component in existing enterprise architecture. Its Decision System of Record is defined as "a structured log of decisions made, the context and reasoning behind them, and the outcomes that followed," characterized as "the institutional memory that enables organizations to systematically improve decision quality over time, something no ERP, CRM, or BI platform can do" (Cloverpop, 2026a). Cloverpop states directly that "the decision infrastructure that captures how your business thinks and decides does not exist yet in most enterprises" (Cloverpop, 2026b). This practitioner diagnosis is structurally identical to JLT's understanding synthesis gap construct, expressed in commercial product language rather than theoretical terms.

Aera Technology, a decision intelligence platform serving large operational enterprises and also recognized in Gartner's decision intelligence platform evaluations, addresses an adjacent architectural layer through its Decision Data Model, which "records every decision made in Aera Decision Cloud along with the context in which the decision was made, a record of actions taken, and the resulting outcome" and provides "a complete audit trail" enabling "real-time monitoring of

decisions across your enterprise” (Aera Technology, 2026). Aera’s architecture addresses automated operational decisions within defined workflows rather than the cross-domain governance synthesis decisions that constitute JLT’s primary scope, but its recognition of decision context and reasoning documentation as an absent enterprise architecture component is convergent with JLT’s diagnosis.

The commercial emergence of this product category provides practitioner validation that the understanding synthesis gap is not merely an academic construct. The decision intelligence platform market was valued at USD 15.22 billion in 2024, with projected growth to USD 36.34 billion by 2030 (Grand View Research, cited in AI CERTS, 2026), suggesting that market actors have independently identified and are responding to the architectural gap that JLT theorizes.

Stream 2: Regulatory Mandate

The European Union Artificial Intelligence Act (EU AI Act), which entered into force on August 1, 2024, establishes through Articles 12 and 13 what one analysis characterizes as “decision-level traceability” requirements, explicitly distinguishing this from system-level activity logging: “traditional logging approaches focus on system activity, inputs, outputs, timestamps. What regulators increasingly expect, however, is a structured record of decisions that connects” inputs, reasoning pathways, decision logic, and outcomes. The analysis concludes that “this shifts governance from monitoring systems to documenting decisions” (AI Governance Desk, 2026). The EU AI Act thus constitutes regulatory codification of the institutional reasoning chain requirement for AI-assisted decisions in high-risk domains. JLT proposes that the institutional reasoning chain requirement applies to all consequential cross-domain institutional decisions, not only AI-assisted ones. The regulatory direction is consistent with JLT’s theoretical claim and provides a natural institutional experiment for testing Proposition 1: organizations that develop EU AI Act Article 12-compliant decision traceability architecture are, by definition, constructing a domain-specific version of the judgment layer. Comparative analysis of those organizations against organizations that have not developed such architecture would provide one of the first direct empirical tests of Proposition 1.

Stream 3: Independent Academic Convergence

A 2026 preprint in the governance evidence literature proposes a Decision Trace Schema for governance evidence in real-time risk systems, identifying six required properties for decision trace records. The authors note that “five of the six properties are independently proposed across separate source clusters in the decision tracing literature, confirming convergent demand for these governance evidence requirements.” They further identify the sixth property, cross-system coupling defined as decision boundary, as “not identified as a distinct requirement in any source examined,” representing the unique gap their design approach addresses (Solozobov, 2026). This independently confirms JLT’s most important structural claim: cross-system synthesis documentation is the specific architectural property that existing frameworks do not yet address. Academic research is converging on the same gap from a computer science and governance engineering direction simultaneously with JLT’s organizational theory development.

These three independent convergences, from commercial platform development, from regulatory mandate, and from adjacent academic research, do not constitute theoretical validation of JLT's causal claims. They do constitute evidence that the understanding synthesis gap is a real and consequential organizational property recognized by actors operating entirely independently of this paper's theoretical development. This convergent recognition strengthens the face validity of the construct and the practical significance of the theory's implications.

11. LIMITATIONS AND FUTURE RESEARCH DIRECTIONS

11.1 Limitations

Causal Inference: This paper establishes a consistent structural pattern through inductive theory building. It does not establish statistical causation. The theoretical propositions are inductive conjectures requiring empirical testing. Claims regarding causal relationships are theoretical proposals.

Sample Composition: The fifteen cases are drawn from large, publicly documented organizations. The theory may not generalize to smaller organizations, to organizations in jurisdictions with limited public regulatory documentation, or to organizations whose failures have not produced regulatory or judicial records of sufficient specificity.

Survivorship: The analysis examines organizations that experienced documented failures. It does not include organizations with documented synthesis architecture that did not fail, which would constitute the comparison group required to test Proposition 1 directly.

Public Documentation Constraint: Observable evidence for the understanding synthesis gap is constrained by what appears in public regulatory and judicial records.

Reverse Causation: The current research design cannot rule out the reverse causation interpretation: rather than synthesis gap causing failure, failure may cause synthesis gap through post-failure document destruction, institutional opacity activated to manage legal liability, or regulatory-imposed confidentiality. In several Sub-Pattern C cases (Volkswagen, Glencore, Alibaba), the evidence is consistent with both the structural default interpretation (Type I synthesis gap: architecture was never designed) and the reverse causation interpretation (synthesis records existed but were not disclosed or were destroyed). Future research using prospective governance assessment designs (assessing synthesis architecture before failure events) rather than retrospective inference from failure records would substantially address this limitation.

Saudi Aramco Case: Analyzed exclusively from corporate annual reports and assurance filings, in the absence of independent multi-jurisdictional regulatory enforcement documentation of the type available for other cases. This reduces the evidentiary weight of this case relative to others.

Methodological Reconstruction: The analytical procedure described is a reconstruction of the theory-building process. The iterative nature of inductive theory building means early analytical decisions influenced later ones in ways not fully captured in a strictly sequential account.

11.2 Future Research Directions

Direction 1 (Testing Proposition 1): Construct a comparative sample of organizations with and without documented synthesis architecture operating in cross-domain decision environments, and measure consequential failure rates. This provides the first direct empirical test of Proposition 1. The EU AI Act's Article 12 decision traceability mandate, effective August 2026 for high-risk AI systems, provides a natural institutional experiment: organizations that develop Article 12-compliant decision logging architecture are constructing domain-specific judgment layer instances. Longitudinal comparison of Article 12-compliant versus non-compliant organizations on governance failure rates in AI-assisted decision domains would constitute a tractable empirical test with existing regulatory enforcement data as the outcome measure.

Direction 2 (Testing Proposition 2): Track organizations following consequential failures over five-year post-remediation windows, comparing same-category recurrence rates between organizations that address synthesis architecture and those that do not. Regulatory enforcement databases, consent order records, and judicial databases provide suitable longitudinal data sources.

Direction 3 (Testing Proposition 3): Develop a survey instrument measuring synthesis architecture completeness and operational system count across a large organizational sample. Test the predicted positive relationship controlling for size, industry, and regulatory environment.

Direction 4 (Instrument Development): Develop and validate a measurement instrument for synthesis architecture completeness, operationalizing the judgment layer construct for organizational surveys, regulatory assessments, and investment due diligence. Instrument development should follow established scale development procedures (Churchill, 1979; DeVellis, 2016).

Direction 5 (Experimental Studies): Conduct organizational simulation experiments in which teams with and without synthesis documentation protocols make cross-domain decisions under time pressure, measuring decision quality, reconstruction adequacy, and learning from failure. Experimental methods would permit causal inference that field studies cannot provide.

Direction 6 (Construct Operationalization): Develop and validate measures for all four JLT constructs, including the understanding synthesis gap, the institutional reasoning chain completeness, and the fragmentation paradox intensity. Validated scales would enable large-sample empirical tests of the three propositions.

Direction 7 (Decision Intelligence Platform Research): The commercial emergence of decision intelligence platforms (Cloverpop, Aera Technology, ACTICO, and others in Gartner's inaugural Magic Quadrant for Decision Intelligence Platforms) provides a population of organizations that have begun constructing judgment layer architecture. Longitudinal study of these organizations

relative to non-adopter control groups in the same industries would provide empirical evidence for Proposition 2 and would enable the first direct test of whether judgment layer architecture reduces post-failure recurrence rates in practice. This research direction bridges the current theory paper with the confirmatory empirical research program required to establish JLT's causal claims.

12. CONCLUSION

This paper introduces Judgment Layer Theory as a structural explanation for a specific empirical puzzle: why consequential failures occur and recur in organizations that demonstrably possess relevant operational systems and have applied system-level remediation following prior failures.

The proposed explanation is structural rather than behavioral, intentional, or systemic. It locates the failure mechanism in the organizational architecture connecting operational systems to institutional decisions, specifically in the structural absence of documented synthesis at the inter-system decision interface. This structural condition, the understanding synthesis gap, is proposed as the mechanism through which failure propagates, resists system-level remediation, and recurs.

The theory is developed inductively from systematic analysis of fifteen global enterprise cases spanning six sectors and six geographic regions, with aggregate documented direct financial consequences exceeding USD 80 billion. It is expressed in three falsifiable propositions, a four-category failure taxonomy, and explicit boundary conditions. It engages systematically with ten alternative theoretical traditions.

The central theoretical proposition is that the judgment layer is a first-class organizational architectural construct whose design properties are a primary structural determinant of institutional vulnerability to cross-domain failure and post-remediation recurrence. This proposition does not claim that system quality is irrelevant. It claims that system quality is a necessary but not sufficient condition for organizational reliability in cross-domain decision environments. The sufficient condition requires, in addition, institutional architecture that documents the synthesis of system outputs into traceable reasoning chains at the point of consequential decisions. That architecture is what the judgment layer construct identifies. Its absence is what JLT proposes as a structurally undertheorized mechanism of consequential organizational failure.

13. REFERENCES

- Argyris, C., and Schon, D. A. (1978). *Organizational learning: A theory of action perspective*. Addison-Wesley.
- Argote, L. (1999). *Organizational learning: Creating, retaining, and transferring knowledge*. Kluwer Academic Publishers.

- Bacharach, S. B. (1989). Organizational theories: Some criteria for evaluation. *Academy of Management Review*, 14(4), 496-515.
- Bazerman, M. H., and Moore, D. A. (2009). *Judgment in managerial decision making* (7th ed.). John Wiley and Sons.
- Beer, S. (1972). *Brain of the firm: The managerial cybernetics of organization*. Allen Lane.
- Blair, M. M. (1995). *Ownership and control: Rethinking corporate governance for the twenty-first century*. Brookings Institution Press.
- Churchill, G. A. (1979). A paradigm for developing better measures of marketing constructs. *Journal of Marketing Research*, 16(1), 64-73.
- Crossan, M. M., Lane, H. W., and White, R. E. (1999). An organizational learning framework: From intuition to institution. *Academy of Management Review*, 24(3), 522-537.
- Cyert, R. M., and March, J. G. (1963). *A behavioral theory of the firm*. Prentice-Hall.
- DeVellis, R. F. (2016). *Scale development: Theory and applications* (4th ed.). SAGE Publications.
- DiMaggio, P. J., and Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, 48(2), 147-160.
- Eisenhardt, K. M. (1989). Building theories from case study research. *Academy of Management Review*, 14(4), 532-550.
- Eisenhardt, K. M., and Graebner, M. E. (2007). Theory building from cases: Opportunities and challenges. *Academy of Management Journal*, 50(1), 25-32.
- Eisenhardt, K. M., and Martin, J. A. (2000). Dynamic capabilities: What are they? *Strategic Management Journal*, 21(10-11), 1105-1121.
- Fama, E. F., and Jensen, M. C. (1983). Separation of ownership and control. *Journal of Law and Economics*, 26(2), 301-325.
- Gioia, D. A., Corley, K. G., and Hamilton, A. L. (2013). Seeking qualitative rigor in inductive research: Notes on the Gioia methodology. *Organizational Research Methods*, 16(1), 15-31.
- Glaser, B. G., and Strauss, A. L. (1967). *The discovery of grounded theory: Strategies for qualitative research*. Aldine.
- Grant, R. M. (1996). Toward a knowledge-based theory of the firm. *Strategic Management Journal*, 17(S2), 109-122.
- Holland, J. H. (1992). *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control, and artificial intelligence*. MIT Press.

- Huber, G. P. (1991). Organizational learning: The contributing processes and the literatures. *Organization Science*, 2(1), 88-115.
- Jensen, M. C., and Meckling, W. H. (1976). Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4), 305-360.
- Kahneman, D., and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263-291.
- Klein, G. A. (1998). *Sources of power: How people make decisions*. MIT Press.
- LaPorte, T. R., and Consolini, P. M. (1991). Working in practice but not in theory: Theoretical challenges of high-reliability organizations. *Journal of Public Administration Research and Theory*, 1(1), 19-48.
- Levitt, B., and March, J. G. (1988). Organizational learning. *Annual Review of Sociology*, 14, 319-340.
- Locke, K., Golden-Biddle, K., and Feldman, M. S. (2008). Perspective: Making doubt generative: Rethinking the role of doubt in the research process. *Organization Science*, 19(6), 907-918.
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1), 71-87.
- March, J. G., and Simon, H. A. (1958). *Organizations*. John Wiley and Sons.
- Maitlis, S., and Christianson, M. (2014). Sensemaking in organizations: Taking stock and moving forward. *Academy of Management Annals*, 8(1), 57-125.
- Monks, R. A. G., and Minow, N. (2011). *Corporate governance* (5th ed.). John Wiley and Sons.
- Nahapiet, J., and Ghoshal, S. (1998). Social capital, intellectual capital, and the organizational advantage. *Academy of Management Review*, 23(2), 242-266.
- Nonaka, I., and Takeuchi, H. (1995). *The knowledge-creating company: How Japanese companies create the dynamics of innovation*. Oxford University Press.
- Oliver, C. (1991). Strategic responses to institutional processes. *Academy of Management Review*, 16(1), 145-179.
- Perrow, C. (1984). *Normal accidents: Living with high-risk technologies*. Basic Books.
- Reason, J. (1990). *Human error*. Cambridge University Press.
- Roberts, K. H. (1990). Some characteristics of one type of high reliability organization. *Organization Science*, 1(2), 160-176.
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69(1), 99-118.

Strauss, A., and Corbin, J. (1998). Basics of qualitative research: Techniques and procedures for developing grounded theory (2nd ed.). SAGE Publications.

Teece, D. J., Pisano, G., and Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7), 509-533.

In re Caremark International, Inc. Derivative Litigation. (1996). 698 A.2d 959. Delaware Court of Chancery.

Walsh, J. P., and Ungson, G. R. (1991). Organizational memory. *Academy of Management Review*, 16(1), 57-91.

Weick, K. E. (1995). Sensemaking in organizations. SAGE Publications.

Weick, K. E., and Sutcliffe, K. M. (2001). Managing the unexpected: Assuring high performance in an age of complexity. Jossey-Bass.

Weick, K. E., Sutcliffe, K. M., and Obstfeld, D. (2005). Organizing and the process of sensemaking. *Organization Science*, 16(4), 409-421.

Whetten, D. A. (1989). What constitutes a theoretical contribution? *Academy of Management Review*, 14(4), 490-495.

Yin, R. K. (2018). Case study research and applications: Design and methods (6th ed.). SAGE Publications.

Aera Technology. (2026). Decision Data Model. Retrieved July 2026 from <https://www.aeratechnology.com/decision-data-model>

AI Governance Desk. (2026). EU AI Act Articles 12 and 13 explained: Decision traceability and audit compliance. Retrieved July 2026 from <https://aigovernancedesk.com/eu-ai-act-articles-12-13-decision-traceability/>

Cloverpop. (2026a). The decision layer: The missing piece in your enterprise AI strategy. Retrieved July 2026 from <https://www.cloverpop.com/resources/the-decision-layer-the-missing-piece-in-your-enterprise-ai-strategy>

Cloverpop. (2026b). What you get in the Cloverpop Decision Intelligence Platform. Retrieved July 2026 from <https://www.cloverpop.com/blog/what-you-get-in-the-cloverpop-decision-intelligence-platform>

European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union.

Gartner. (2024). Market Guide for Decision Intelligence Platforms. Gartner Research.

Solozobov, O. (2026). Decision trace schema for governance evidence in real-time risk systems. arXiv preprint arXiv:2604.09296. <https://doi.org/10.48550/arXiv.2604.09296>

14. APPENDIX A: CASE PROFILES AND EVIDENTIARY SOURCES

All Tier 1 and Tier 2 sources are cited in full below, organized by case.

A.1 Boeing Company

United States Department of Justice. (January 2021). Deferred prosecution agreement. Case 4:21-cr-00005-O, Northern District of Texas.

United States Department of Justice. (May 2024). Notice of breach of deferred prosecution agreement. Case 4:21-cr-00005-O, Northern District of Texas.

United States Department of Justice. (May 2025). Non-prosecution agreement press release.

United States Senate Commerce Committee. (October 29, 2019). Hearing record: Aviation safety and the FAA's oversight of Boeing. Congressional Record.

Boeing Company. (February 2026). Annual report on Form 10-K for fiscal year ended December 31, 2025. SEC EDGAR. CIK 0000037996.

A.2 UnitedHealth Group

United States Senate Finance Committee. (May 1, 2024). Chairman Wyden opening statement: Hearing on the Change Healthcare cyberattack and UnitedHealth Group's response. Senate Finance Committee Records.

Witty, A. (May 1, 2024). Testimony before the House Energy and Commerce Subcommittee.

UnitedHealth Group. Annual report on Form 10-K for fiscal year ended December 31, 2024. SEC EDGAR.

A.3 Goldman Sachs

Board of Governors of the Federal Reserve System. (October 22, 2020). Order to cease and desist and order of assessment of a civil money penalty issued upon consent. www.federalreserve.gov/newsevents/pressreleases/enforcement20201022a.htm

Financial Conduct Authority and Prudential Regulation Authority. (October 23, 2020). FCA and PRA fine Goldman Sachs International for risk management failures in relation to 1MDB. www.fca.org.uk/news/press-releases/fca-pra-fine-goldman-sachs-international-risk-management-failures-1mdb

A.4 Wells Fargo

Wells Fargo and Company. (March 17, 2025). Current report on Form 8-K. SEC EDGAR. CIK 0000072971.

Office of the Comptroller of the Currency. (2021). Consent order and civil money penalty: In the

matter of Wells Fargo Bank, N.A. OCC Enforcement Actions Database.

Board of Governors of the Federal Reserve System. (February 2018). Consent cease and desist order.

A.5 Walmart

Walmart Inc. (October 18, 2024). Current report on Form 8-K. SEC EDGAR. <https://stock.walmart.com/sec-filings/all-sec-filings/content/0000104169-24-000166/ex991-10182024.htm>

A.6 Volkswagen Group

United States Environmental Protection Agency. (September 18, 2015). Notice of violation: Clean Air Act Section 203(a)(3)(B). EPA Enforcement Actions.

United States Department of Justice. (January 11, 2017). Volkswagen AG agrees to plead guilty and pay 4.3 billion in criminal and civil penalties. DOJ Press Release.

Braunschweig Regional Court, Germany. (May 2025). Criminal conviction: Four former Volkswagen managers.

A.7 Deutsche Bank

Board of Governors of the Federal Reserve System. (July 2024). Consent order and assessment of civil money penalty: Deutsche Bank AG.

Deutsche Bank AG. (2025). Annual report on Form 20-F for fiscal year ended December 31, 2024. SEC EDGAR. CIK 0001159508.

A.8 HSBC Holdings

Financial Conduct Authority. (December 17, 2021). Decision notice: HSBC Bank PLC. www.fca.org.uk/news/press-releases/fca-fines-hsbc-bank-plc-deficient-transaction-monitoring-controls

HSBC Bank PLC. (2024). Form SBSE-A/A. SEC EDGAR. CIK 0001140465.

A.9 Shell

UK High Court of Justice. (June 20, 2025). Judgment on preliminary issues of Nigerian law: Okpabi and others v. Shell PLC and Renaissance Africa Energy Company Limited. Mrs Justice May.

A.10 TotalEnergies

Tribunal Judiciaire de Paris. (October 23, 2025). Judgment: Greenpeace France and others v. TotalEnergies SE and TotalEnergies Electricite et Gaz France.

TotalEnergies SE. (March 26, 2025). Form 6-K: Mozambique LNG investigation disclosure. SEC EDGAR. CIK 0000879764.

A.11 Toyota Motor Corporation

Ministry of Land, Infrastructure, Transport and Tourism, Japan. (2023 to 2024). Investigation records: Daihatsu Motor Co. safety certification falsification. MLIT Press Releases.
Toyota Motor Corporation. (January 24, 2024). Recall notification and corporate statement regarding Toyota Industries certification irregularities. Toyota Global Newsroom.

A.12 Alibaba Group

State Administration for Market Regulation, People's Republic of China. (April 10, 2021). Administrative penalty decision: Alibaba Group Holding Limited. SAMR Enforcement Records.
State Administration for Market Regulation. (August 2024). Confirmation of Alibaba rectification completion. SAMR Website.

A.13 Glencore

United States Department of Justice. (May 24, 2022). Glencore International AG guilty plea: FCPA conspiracy. Southern District of New York Press Release.
Commodity Futures Trading Commission. (May 24, 2022). CFTC order: Glencore International AG, Glencore Ltd, and Chemoil Corporation.
UK Serious Fraud Office. (June 21, 2022). Glencore Energy (UK) Ltd: Conviction on seven counts of bribery.

A.14 Saudi Aramco

Saudi Aramco. (2025). Annual report 2024: Sustainability section. www.aramco.com
Saudi Aramco. (2025). Annual report 2024: Governance section. www.aramco.com
Saudi Aramco. (2026). Annual report 2025. www.aramco.com

A.15 Ford Motor Company

Ford Motor Company. (February 2025). Annual report on Form 10-K for fiscal year ended December 31, 2024. SEC EDGAR. CIK 0000037996.
Ford Motor Company. (February 2026). Annual report on Form 10-K for fiscal year ended December 31, 2025. SEC EDGAR. CIK 0000037996.
National Highway Traffic Safety Administration. (2025). NHTSA recall database. www.nhtsa.gov/vehicle-safety/recalls

The author declares this paper presents original research not previously published. All empirical material derives from publicly accessible sources cited above.

Shubham Agarwal, Founder and Chief Executive Officer, triNetra Research, www.trinetra.life, July 2026