

An Operational Observability Architecture for Governance, Regulation, Liability, and Insurance in Autonomous Systems the Erckenbrack Recursive Authority Dynamics Model (ERADM)

Adrian Erckenbrack* 

Independent Researcher

Corresponding Author: Adrian Erckenbrack, Independent Researcher.**Received:**  2026 Aug 11**Accepted:**  2026 Aug 31**Published:**  2026 Sep 10

Abstract

Autonomous systems (AS) are increasingly deployed across transportation, logistics, robotics, critical infrastructure, and other safety-critical environments. Yet a fundamental governance observability challenge persists within contemporary oversight architectures. Existing regulatory, certification, liability, insurance, and safety-management frameworks provide substantial visibility into outcomes including crashes, incidents, recalls, compliance actions, operational restrictions, and safety events while providing comparatively limited visibility into the operational conditions that influence whether effective supervision, intervention, restoration, and oversight remain possible. Consequently, critical governance degradation may remain difficult to detect and evaluate until intervention opportunities have narrowed and AS failures become visible through accidents, recalls, regulatory actions, liability events, or other adverse outcomes. This study introduces the Erckenbrack Recursive Authority Dynamics Model (ERADM), an operational observability architecture designed to measure governance-relevant conditions within increasingly autonomous systems. Unlike frameworks developed primarily from theoretical assumptions, ERADM was constructed through an evidence-to model approach. Historical failures, regulatory interventions, investigations, liability disputes, recalls, operational restrictions, and governance breakdowns were first examined using publicly available records. Recurring governance conditions including intervention effectiveness, supervisory responsiveness, rollback feasibility, adaptation tempo, synchronization effects, infrastructure dependency, operational flexibility, and authority-transition dynamics were then identified across independent cases and formalized as measurable variables. **Consequently, the historical evidence preceded the model & variables; the model & variables did not precede the evidence.** The framework introduces measurable variables including Authority Elasticity Index (AEI), Override Effectiveness Rate (OER), Rollback Feasibility Time (RFT), Closed-Loop Self-Improvement Interval (CLSI), Operational Space Compression Ratio (OSCR), Synchronization Intensity (S), Infrastructure Coupling (K), Human Supervisory Response Horizon (H), Recursive Leverage Factor (RLF), Adoption-Driven Authority Transfer (ADAT), and the composite Authority Transition Index (Θ). These variables were derived from recurring empirical governance conditions identified through historical failure reconstructions, regulatory interventions, safety investigations, operational disruptions, liability disputes, and corrective actions across multiple domains. In addition to their empirical origins, each variable is linked to statutory authorities, regulatory obligations, certification requirements, enforcement mechanisms, insurance considerations, or liability doctrines that become relevant when the AS underlying condition contributes to operational risk or governance failure. Together, the variables provide a structured mechanism for observing, reconstructing, and evaluating governance-relevant operational conditions that are frequently material to regulatory oversight, safety assurance, liability assessment, and risk management, yet often remain only partially observable within existing reporting architectures. The framework is evaluated through comparative analysis of AS case studies, including Uber ATG, Tesla ADS/FSD, Cruise, and Waymo, together with examination of existing regulatory, certification, insurance, and safety-management frameworks. **Results indicate that many governance-relevant conditions remain only partially observable or entirely unobservable within contemporary oversight architectures despite their recurring significance in investigations, regulatory actions, and corrective interventions.** This limitation arises not because the conditions are absent, but because existing oversight architectures primarily measure outcomes, compliance, and safety events rather than the operational conditions that precede and influence those outcomes. Consequently, evidence related to intervention effectiveness, supervisory responsiveness, rollback feasibility, infrastructure dependency, synchronization effects, and operational flexibility is often fragmented or not systematically measured. Using an inductive comparative reconstruction methodology, publicly available investigative, regulatory, operational, and liability records were analyzed to identify recurring governance conditions across independent cases. **ERADM variables were created and retained only when recurring governance conditions were repeatedly observed across independent cases**

and demonstrated operational relevance together with regulatory, insurance, or liability significance. Reconstruction analyses demonstrate that several ERADM variables can be estimated from publicly available evidence, supporting governance observability as a measurable systems property and establishing an evidence-to-model methodology in which recurring historical evidence informed variable selection before framework construction. The principal contribution of ERADM is the introduction of an operational observability and measurement layer through which governance-relevant conditions can be systematically observed, reconstructed, and empirically assessed leading up to and following a catastrophic or deadly incident. Results indicate that many governance-relevant conditions remain only partially observable or entirely unobservable within contemporary oversight architectures despite their recurring significance in investigations, regulatory actions, and corrective interventions. Reconstruction analyses nevertheless demonstrate that several ERADM variables can be estimated using publicly available evidence, supporting governance observability as a measurable systems property. Finally, ERADM is designed to operate alongside established systems-safety, risk-management, and governance frameworks—including STAMP/CAST, the NIST AI Risk Management Framework, and Safety Management Systems (SMS) by providing a measurable observability layer through which governance-relevant operational conditions can be observed, reconstructed, and evaluated. **The framework's principal novelty therefore lies in treating governance observability itself as a measurable systems property and in operationalizing that property through a structured set of empirically derived variables.**

Keywords: Governance Observability, Autonomous Systems, Meaningful Human Control, Stamp, Authority Preservation, Operational Observability, Autonomous Vehicles, Safety Governance, Erckenbrack Recursive Authority Dynamics Model (ERADAM)

I. Introduction

1.1 The Governance Observability Problem

Autonomous systems are increasingly deployed across transportation, logistics, robotics, critical infrastructure, and other safety-critical operational environments. Existing governance frameworks administered by organizations such as NHTSA, FMCSA, FAA, NTSB, and related oversight bodies provide substantial visibility into crashes, incidents, recalls, operational restrictions, compliance actions, investigations, and safety outcomes [1-6]. Similarly, governance frameworks including OMB Memorandum M-24-10, the NIST Artificial Intelligence Risk Management Framework (AI RMF), ISO/IEC 42001, ISO 26262, transportation safety regulations, certification systems, insurance requirements, and liability frameworks establish important requirements concerning safety, accountability, transparency, risk management, and human oversight [4-8]. Despite these extensive governance structures, a fundamental observability limitation remains. Existing frameworks generally provide strong visibility into **outcomes and organizational responsibilities** while providing comparatively limited visibility into the operational conditions that determine whether effective supervision, intervention, restoration, and oversight remain practically achievable. Consequently, governance institutions frequently observe the consequences of deteriorating control conditions while lacking direct measurement of the conditions themselves. Current oversight architectures are predominantly outcome-observation systems. They identify crashes, incidents, recalls, operational restrictions, enforcement actions, compliance failures, and safety events. While these measures provide important information regarding what occurred, they provide comparatively limited visibility into whether meaningful intervention remained possible, whether override mechanisms were effective, whether restoration pathways remained available, whether adaptation processes outpaced supervision, whether viable alternatives were becoming compressed, or whether increasing synchronization and infrastructure dependency

were creating elevated governance risk [2,4,7].

Historical evidence suggests that this limitation is not merely theoretical. Analyses involving Uber ATG, Tesla ADS/FSD, Cruise, Boeing 737 MAX, and Waymo indicate that many governance-relevant conditions can be reconstructed from investigations, recalls, deployment records, operational restrictions, safety reports, and regulatory actions [4,931-48]. Across these cases, governance-relevant conditions frequently appeared observable before major recalls, operational restrictions, enforcement actions, or deployment changes occurred. These findings suggest that the principal challenge facing autonomous-system governance may not be the absence of information, but the absence of measurement architectures specifically designed to organize, quantify, and report governance-relevant operational conditions.

1.2 Why Observability Matters

The observability limitation has direct implications for regulatory oversight, certification assurance, accident investigation, insurance evaluation, liability determination, operational governance, and public accountability. Regulatory agencies must evaluate whether deployed systems remain governable. Insurers must assess whether operators retain effective influence over the risks being insured. Investigators must reconstruct the operational conditions preceding failures. Courts must evaluate evidence concerning intervention capability, operational control, and responsibility. Yet many of the conditions relevant to these determinations remain incompletely observable within existing reporting architectures. Consequently, institutions are often required to make decisions concerning certification, deployment, enforcement, underwriting, liability attribution, and public accountability without directly observing many of the operational conditions that determine whether governance remains effective. The challenge therefore extends beyond authority and responsibility. Institutions frequently possess legal, regulatory, operational, or

financial authority while possessing limited direct visibility into the operational conditions that determine whether such authority remains practically effective. Improving observability improves not only information availability but also the ability to detect deteriorating governance conditions before failures, recalls, operational restrictions, litigation, or major losses occur. Observability is influenced not only by operational conditions internal to autonomous systems but also by the capacity of external institutions to detect, interpret, aggregate, and respond to those conditions.

This concern is consistent with safety-science research demonstrating that major failures frequently emerge through gradual operational adaptation rather than isolated catastrophic events. Dekker observed that complex socio-technical systems often drift toward increasingly unsafe operating conditions while continuing to appear functional under normal performance metrics. As a result, governance degradation may remain difficult to detect until intervention capability, supervisory effectiveness, or operational resilience have already been materially reduced [10].

1.3 Research Questions

This study addresses five research questions

- **RQ1:** To what extent do existing autonomous-system reporting, certification, investigation, and governance architectures provide sufficient visibility into governance-relevant operational conditions?
- **RQ2:** Can governance-relevant operational conditions be transformed into measurable, observable, reportable, auditable, and empirically investigable system properties?
- **RQ3:** Does ERADM provide a coherent observability architecture through which governance-relevant operational conditions can be quantified, modeled, reconstructed, validated, and empirically analyzed?
- **RQ4:** Can ERADM variables be reconstructed from deployed autonomous systems using publicly observable evidence?
- **RQ5:** Do the empirical findings justify the need for such an observability architecture like ERADM?

To address these questions, the article develops a formal observability architecture consisting of measurable governance variables, dynamic relationships, state-space representations, governance-feasible conditions, observability pathways, empirical reconstruction methods,

and falsification criteria. Collectively, these components establish a framework through which governance-relevant operational conditions may be observed, measured, reconstructed, audited, validated, and empirically evaluated across autonomous-system environments.

2. The Governance Observability Gap

2.1. Existing Governance Frameworks

Contemporary autonomous-system governance is supported by a substantial collection of regulatory, certification, safety-management, liability, insurance, and organizational oversight frameworks. Federal agencies including the National Highway Traffic Safety Administration (NHTSA), Federal Motor Carrier Safety Administration (FMCSA), Federal Aviation Administration (FAA), and National Transportation Safety Board (NTSB) administer reporting, certification, investigative, and enforcement mechanisms that provide important visibility into safety performance, operational failures, compliance outcomes, recalls, operational restrictions, and accident causation [16-36]. Similarly, governance frameworks including OMB Memorandum M-24-10, the NIST Artificial Intelligence Risk Management Framework (AI RMF), ISO/IEC 42001, ISO 26262, and related standards establish requirements concerning accountability, transparency, risk management, safety assurance, and organizational governance [15-34]. The existence of these frameworks demonstrates that autonomous-system governance is neither absent nor immature. Existing architectures provide substantial capability for managing safety, reliability, compliance, accountability, certification, transparency, and organizational responsibility. The governance observability gap identified in this article should therefore not be interpreted as a failure of current governance systems. Rather, it reflects a more specific limitation concerning the type of information that existing frameworks make directly observable. Most current governance architectures are designed to observe outcomes, organizational responsibilities, safety risks, compliance obligations, and accountability structures. They are generally less effective at directly observing the operational conditions that determine whether intervention, supervision, restoration, and governance remain practically achievable during system operation. Consequently, institutions often possess substantial information regarding what occurred while possessing comparatively limited information regarding whether meaningful intervention remained available before those outcomes emerged.

Governance-Observable Condition	NHTSA (Transportation Safety Oversight)	FAA SMS (Aviation Safety Oversight)	NIST AI RMF (Artificial Intelligence Risk Management)	ISO 42001 (AI Management Systems)	Insurance / Liability Frameworks	ERADM (Governance Observability Architecture)
Crash / Incident Visibility	High	High	Medium	Medium	High	Medium
Compliance Visibility	High	High	High	High	Medium	Medium
Intervention Capability (AEI)	Low	Low	Low	Low	Low	High
Override Effectiveness (OER)	Low	Low	Low	Low	Low	High
Restoration Capability (RFT)	Low	Medium	Low	Low	Low	High
Supervisory Responsiveness (H)	Low	Low	Low	Low	Low	High
Adaptation Tempo (CLSI)	Low	Low	Medium	Medium	Low	High
Synchronization Intensity (S)	Low	Low	Low	Low	Low	High
Infrastructure Coupling (K)	Low	Low	Low	Low	Low	High
Operational Flexibility (OSCR)	Low	Low	Low	Low	Low	High
Authority Transition Pressure (Θ)	None	None	None	None	None	High

Note: High = condition is directly observable or systematically measured within the framework; Medium = condition is partially observable or indirectly inferred in some circumstances; Low = condition is generally not observable within the framework; None = condition is not addressed by the framework. ERADM provides direct observability for governance-relevant operational conditions that are not directly measured within existing frameworks.

Table 1: Existing Governance Frameworks & Governance Observable Conditions

Table 1 reveals that contemporary governance frameworks provide strong visibility into safety outcomes, risk-management processes, compliance obligations, organizational accountability, certification activities, and other observable consequences of system operation. However, they provide comparatively limited visibility into governance-relevant operational conditions such as intervention effectiveness, rollback feasibility, supervisory responsiveness, operational flexibility, synchronization intensity, infrastructure dependency, and authority-transition pressure. Observability ratings assess whether these conditions can be independently reconstructed, evaluated, and validated by regulators, investigators, insurers, auditors, courts, operators, or researchers using evidence systematically generated through operational records, safety reports, compliance documentation, audits, incident investigations, public disclosures, certification activities, or liability proceedings. High observability indicates that a framework routinely produces sufficient evidence for reliable reconstruction of a condition; medium observability indicates that partial reconstruction is possible through indirect, fragmented, or incomplete evidence; low observability indicates that available evidence is generally insufficient, proprietary, inconsistently reported, or otherwise inadequate for reliable reconstruction; and none indicates that the framework provides no meaningful mechanism for observing or evaluating the condition. Consequently, observability ratings measure evidence availability and reconstruct ability rather than the existence or importance of the underlying condition. The comparison therefore suggests that existing frameworks are highly effective at observing outcomes and governance artifacts, while providing comparatively limited visibility into the operational conditions that determine whether governance itself remains effective within increasingly autonomous environments. Existing governance frameworks and

governance-observable conditions.

2.2 What Existing Frameworks Measure

Existing governance architectures are particularly effective in five areas: safety, reliability, compliance, transparency, and accountability. Safety remains the most mature measurement domain. Transportation, aviation, and industrial safety frameworks provide extensive mechanisms for identifying crashes, incidents, operational hazards, safety deficiencies, recalls, enforcement actions, and accident causation [29-36]. These systems establish whether autonomous technologies produce unacceptable operational outcomes and support corrective intervention when failures occur. Reliability frameworks provide visibility into component performance, system availability, redundancy, fault tolerance, degradation modes, and operational stability. Such measures are essential for evaluating whether systems perform intended functions consistently under expected operating conditions. Compliance architectures evaluate adherence to statutory requirements, regulatory obligations, certification standards, reporting duties, inspection requirements, operational restrictions, and documentation standards [15-34]. Compliance measurement enables agencies to enforce legal and regulatory obligations and provides organizations with mechanisms for demonstrating conformance to established requirements. Transparency frameworks support documentation, explain ability, traceability, auditability, and organizational awareness. Increasing emphasis on transparency within AI governance reflects recognition that organizations require visibility into system behavior, development processes, and risk-management activities [15-34]. Accountability frameworks establish responsibility assignments among manufacturers, operators, developers, regulators, insurers, certificate holders, and other stakeholders. These frameworks determine who must report, who must comply, who may be investigated, and

who may bear legal or financial responsibility following adverse outcomes. Collectively, these measurement domains provide the foundation of contemporary governance. However, they primarily focus on outcomes, responsibilities, and organizational processes rather than on operational governance conditions themselves.

2.3 What Existing Frameworks Do Not Directly Measure

Several governance-relevant operational conditions remain only partially observable within current reporting architectures. The first is intervention effectiveness. Existing systems frequently document the existence of intervention authority, override mechanisms, or supervisory responsibilities. They less frequently measure whether intervention attempts succeed when intervention becomes necessary. Formal authority and operational effectiveness are not equivalent. An override capability may exist while remaining operationally ineffective under real-world conditions. Direct measurement of intervention effectiveness therefore represents a distinct observability requirement. The second is restoration capability. Existing frameworks routinely require corrective actions, recalls, software modifications, remediation activities, and safety interventions. However, they rarely measure how rapidly a system can be restored to a previously safe operational state once problematic behavior emerges. The practical ability to reverse, isolate, or recover from undesirable conditions may significantly influence safety, resilience, and governance effectiveness. The third is supervisory responsiveness. Governance frameworks commonly assume the existence of human oversight, operational supervision, organizational monitoring, or institutional control. They generally provide limited measurement of whether supervisory actors can observe, interpret, coordinate, authorize, and act quickly enough to remain operationally relevant as system adaptation accelerates. Formal oversight structures may remain intact while practical supervisory effectiveness deteriorates.

The fourth is operational flexibility. Governance systems often assume that operators retain meaningful alternatives when responding to failures, disruptions, or emerging risks. In practice, technological dependence, software integration, infrastructure coupling, and organizational lock-in may progressively reduce available alternatives. Existing reporting systems rarely measure whether viable operational options remain available. The fifth involves synchronization and dependency formation. Increasingly autonomous systems frequently depend upon shared software architectures, cloud services, communication infrastructure, fleet-learning mechanisms, and coordinated update processes. These dependencies may create common-mode failure pathways, correlated operational behavior, and systemic governance exposure. Existing governance frameworks generally recognize systemic risk conceptually but provide limited direct measurement of synchronization intensity and infrastructure coupling. These limitations do not invalidate existing governance frameworks. Rather, they reveal a category of governance relevant operational conditions that remain difficult to observe directly. The

significance of ERADM is not that it replaces current governance systems, but that it introduces a measurement architecture specifically designed to make these conditions observable.

The governance significance of synchronization, dependency formation, and supervisory limitations is consistent with established socio-technical systems research. Rasmussen observed that complex systems may gradually migrate toward operational boundaries where existing oversight mechanisms become increasingly ineffective despite the continued presence of formal control structures [39]. Similarly, Perrow's Normal Accident Theory demonstrated that increasing coupling and interdependence can create failure pathways that become difficult to predict, isolate, and control [37]. These observations suggest that governance effectiveness depends not only upon authority, accountability, or compliance structures, but also upon the continued observability of the operational conditions that influence whether intervention, supervision, restoration, and corrective action remain practically achievable.

2.4 From Outcome Observation to Operational Observability

The central proposition of this article is that governance effectiveness depends not only upon observing outcomes, but also upon observing the operational conditions that shape those outcomes. Existing governance architectures are highly effective at documenting crashes, incidents, recalls, compliance actions, litigation, enforcement activities, and safety performance. They provide comparatively less visibility into whether intervention remained possible, whether overrides remained effective, whether restoration remained feasible, whether supervisory responsiveness remained adequate, whether operational flexibility remained available, and whether increasing synchronization or dependency formation were altering governance conditions.

ERADM addresses this observability limitation by introducing a structured measurement architecture through which governance-relevant operational conditions become direct objects of observation, measurement, reporting, auditing, reconstruction, and empirical investigation. In this respect, the framework functions as a complementary observability layer rather than a replacement for existing governance systems. The governance observability gap therefore concerns neither authority nor accountability alone. It concerns visibility into the operational conditions that determine whether governance remains practically achievable. Existing frameworks such as STAMP observe outcomes. ERADM is designed to observe the conditions that precede and shape those outcomes.

3. ERADM Observability Architecture

3.1 Design Principles

ERADM was developed as an operational observability architecture intended to transform governance-relevant operational conditions into measurable system properties. The framework is founded on six design principles:

observability, measurability, auditability, reportability, empirical testability, and falsifiability [24–38]. Observability requires that governance-relevant conditions possess identifiable operational indicators capable of being detected through available evidence sources. Measurability requires that those indicators be converted into variables that can be quantified, estimated, compared, and evaluated across systems and operational environments. Auditability requires that measurements be reconstructable using evidence that can be independently examined and verified. Reportability requires that observations be capable of integration into existing governance, regulatory, certification, insurance, and investigative reporting architectures. Empirical testability requires that proposed relationships generate observable expectations capable of evaluation using real-world evidence. Finally, falsifiability requires that the framework generate conditions under which its propositions may be contradicted by empirical observation [24–38]. Collectively, these principles establish governance observability as a measurable systems problem rather than solely a legal, organizational, or policy concern. Rather than evaluating governance exclusively through structures, responsibilities, compliance requirements, or outcomes, ERADM focuses on the operational conditions that determine whether

governance remains practically effective. The conceptual foundations of this approach are consistent with established systems-safety and socio-technical governance research. Rasmussen demonstrated that effective risk management depends upon maintaining visibility into system migration toward operational boundaries before failures occur [39]. Perrow similarly highlighted the governance implications of increasing coupling and interdependence within complex systems, where recovery and intervention may become progressively more difficult as dependencies accumulate [37]. ERADM extends these traditions by treating governance observability itself as a measurable system property and by introducing variables through which governance-relevant operational conditions can be reconstructed, quantified, and empirically evaluated.

3.2 Governance Variables

ERADM operationalizes governance observability through a structured set of measurable variables. Each variable represents a governance-relevant operational condition associated with intervention capability, supervisory effectiveness, restoration feasibility, operational flexibility, adaptation tempo, synchronization, infrastructure dependency, or authority-transition pressure.

TABLE II
ERADM VARIABLE DEFINITIONS

Variable (Acronym)	Symbol	Formal Definition	Interpretation	Typical Range	Measurement Concept
Authority Elasticity Index (AEI)	<i>AEI</i>	Measures the degree to which effective authority can be reclaimed or adjusted by human or organizational actors.	Higher values indicate greater preservation of meaningful authority.	[0, 1+] (> 1 possible)	Ratio of effective authority adjustments to total attempted adjustments over a defined period.
Override Effectiveness Rate (OER)	<i>OER</i>	Measures the success rate of human or authorized overrides or interventions.	Higher values indicate more effective intervention mechanisms.	[0, 1]	Successful overrides divided by total override attempts.
Rollback Feasibility Time (RFT)	<i>RFT</i>	Measures the time required to restore the system to a prior safe and operational state.	Lower values indicate greater restoration capability.	[0, ∞]	Elapsed time from rollback initiation to restoration of a validated state.
Option Set Compression Ratio (OSCR)	<i>OSCR</i>	Measures the reduction in available meaningful operational alternatives.	Lower values indicate greater compression of meaningful options.	[0, 1]	Ratio of meaningful options available at time t to those available at baseline.
Closed-Loop Self-Improvement Interval (CLSI)	<i>CLSI</i>	Measures the time between successive self-improvement cycles.	Lower values indicate faster recursive improvement tempo.	(0, ∞]	Elapsed time from initiation of one learning cycle to the next.
Human Supervisory Response Horizon (H)	<i>H</i>	Measures the time available for meaningful supervisory intervention.	Higher values indicate greater supervisory viability.	[0, ∞)	Estimated response time available before system behavior becomes irreversible.
Synchronization Intensity (S)	<i>S</i>	Measures the degree to which system instances operate in a synchronized or correlated manner.	Higher values indicate greater synchronization and correlated risk.	[0, 1]	Proportion of system instances exhibiting correlated operational behavior within a defined period.
Infrastructure Coupling (K)	<i>K</i>	Measures dependence upon external infrastructure, networks, or services.	Higher values indicate greater infrastructure dependency.	[0, 1]	Proportion of critical functions dependent upon external infrastructure.
Recursive Leverage Factor (RLF)	<i>RLF</i>	Measures the extent to which recursive improvement amplifies capability and influence.	Higher values indicate greater autonomy amplification.	[1, ∞)	Ratio of current capability to baseline capability attributable to recursive improvement.
Authority Transition Index (θ)	θ	Composite measure of authority-transition pressure within the system.	Higher values indicate greater pressure on meaningful human or organizational authority.	[0, ∞)	Derived from the ERADM governing equations and constituent variables.

Note—All variables are designed to be observable directly or reconstructable through operational evidence where possible. Ranges represent conceptually feasible values under typical operational conditions and may require calibration for specific domains.

Table 2: ERADM Variable Definitions

Table 2 defines the ten core ERADM variables and establishes how each can be operationally measured, interpreted, and evaluated. **Rather than treating authority preservation, intervention capability, rollback capacity, supervisory responsiveness, and recursive system behavior as abstract governance concepts, the table translates them**

into observable variables with specified ranges and measurement approaches. Collectively, these variables form the empirical foundation of ERADM by providing a standardized framework through which governance-relevant conditions can be observed, reconstructed, compared, and analyzed across autonomous-system deployments. The

Authority Transition Index (Θ) integrates these variables into a composite indicator of authority-transition pressure within the system. The Authority Elasticity Index (AEI) measures the extent to which meaningful intervention capability remains available within an operational system. High AEI values indicate that supervisory actors retain effective influence over system behavior. Declining AEI values indicate reduced practical intervention capacity despite the possible continued existence of formal authority structures. The Override Effectiveness Rate (OER) measures the operational success of intervention attempts. Whereas AEI addresses the availability of intervention capability, OER evaluates whether intervention actions actually succeed when exercised. Rollback Feasibility Time (RFT) measures the practical ability to restore a system to a previously safe operational condition following undesirable behavior, software failure, operational degradation, or governance intervention. Option Set Compression Ratio (OSCR) measures changes in operational flexibility by evaluating the extent to which viable alternatives remain available. Increasing OSCR values indicate narrowing operational choice sets and increasing dependence on specific operational pathways. Closed-Loop Self-Improvement Interval (CLSI) measures adaptation tempo by evaluating the time required for observation, update, validation, deployment, and feedback processes to complete an operational learning cycle.

Human Supervisory Response Horizon (H) measures the time required for supervisory actors to observe, interpret, coordinate, authorize, and execute meaningful intervention actions. Synchronization Intensity (S) measures the degree to which operational behavior becomes coordinated across distributed systems through common software architectures, update mechanisms, communication structures, or fleet learning processes. Infrastructure Coupling (K) measures dependence upon external services, communication systems, cloud resources, operational networks, or other supporting infrastructure necessary for continued operation. Recursive Leverage Factor (RLF) measures the extent to which adaptive mechanisms amplify operational capability, scale, influence, or dependence. The Authority Transition Index (Θ) serves as the framework's principal composite indicator and summarizes aggregate authority-transition pressure generated by interactions among the underlying variables.

3.3 Where Did the Variables Come from and Why Were They Selected

To justify variable selection, each ERADM variable was mapped against historical failure evidence, regulatory authorities, and liability doctrines. Table III summarizes the evidentiary basis supporting inclusion of each variable.

TABLE III
ERADM VARIABLE ORIGIN AND SELECTION MATRIX

Historical failure evidence, regulatory foundations, and liability doctrines supporting retention of ERADM variables.

Selection logic: each variable is retained because historical failures show a recurring governance condition that becomes legally, or operationally material when current frameworks observe outcomes but not the underlying condition.				
ERADM Variable	Historical Failure / Data Origin	Statutory / Regulatory Hook	Economic / Liability Doctrine	Why Selected for ERADM
Rollback Capability	Knight Capital: \$440M loss, no kill switch; Cruise: DMV suspension after minimal-risk-condition failure.	FINRA Rule 15c3-5; California Vehicle Code §38750; NERC EOP-005.	SEC enforcement; Hand Formula: prevention cost far below loss magnitude.	Measures whether unsafe capability can be reversed, restored, or contained before exposure expands.
Intervention Effectiveness	737 MAX: limited pilot counter-MCAS success; Tesla FSD: disengagement and crash-report evidence.	49 U.S.C. §30102; NHTSA SGO 2021-01; 21 CFR §820.30.	Restatement (Third) of Torts §2: design defect when human intervention is ineffective.	Separates formal human authority from operationally effective intervention.
Supervisory Viability	Flash Crash: market movement in minutes; 2003 Blackout: operators lacked effective visibility.	17 CFR §242.613; NERC TOP-001; EU AI Act Article 14.	Res ipsa logic: system behavior beyond human tempo may indicate defective supervision.	Measures whether humans or institutions can still observe, decide, and act in relevant time.
Operational Flexibility	Therac-25: single-path software dependence; Waymo recall evidence involving deployed fleet behavior.	21 CFR §820.30; FMVSS; 14 CFR §25.1309.	T.J. Hooper: feasible safeguards must be adopted; single-path designs are legally vulnerable.	Measures option-set compression and whether meaningful alternatives remain available.
Synchronization Intensity	Flash Crash: HFT withdrawal together; Knight: same-model failure propagated rapidly.	15 U.S.C. §78j; Basel III; EU AI Act Article 15.	Basic v. Levinson: reliance and foreseeable market effects; herding can become material.	Captures correlated failure risk when many systems act or update together.
Infrastructure Dependency	2003 Blackout: manual restart limitations; mortgage AUS dependence on automated underwriting.	NERC EOP-005; OCC 2011-12; Dodd-Frank §165.	Contract doctrine: self-induced impossibility is not a defense.	Measures when external infrastructure becomes necessary for governance or recovery.
CLSI	Flash Crash millisecond losses; Tesla OTA cycles; core ADAT evidence of recursive adaptation.	SEC Rule 613; EU AI Act Article 14.	Daubert relevance: evidence must match system tempo; stale oversight may be unreliable.	Measures machine adaptation tempo against human and institutional review capacity.
Governance Observability Latency	Knight post-mortem delay; Cruise suspension after drag incident; Therac correction latency across years.	15 U.S.C. §78m; Sarbanes-Oxley §409.	Securities-fraud logic: delayed or omitted material information can create liability.	Measures delay between operational degradation and institutional visibility.
Semantic Drift Rate	Therac "Malfunction 54"; 737 MAX MCAS not clearly disclosed; Tesla FSD capability-language dispute.	15 U.S.C. §45; 49 U.S.C. §30118.	Kasky v. Nike: terminology drift can support deception theory.	Measures whether labels, claims, or internal terms drift away from actual system capability.
Authority Recognition Delay	2003 Blackout: operators did not know they were blind; mortgage override reversal evidence.	NERC IRO-002; CFPB Regulation B.	Agency law: apparent authority may bind; delayed recognition can ratify failure.	Measures lag between authority loss and actor recognition that control has degraded.
Manual Override Economic Cost	Therac interlocks removed for cost; Knight lacked kill switch; Cruise high remote-ops staffing burden.	29 CFR §1910.212; Dodd-Frank §165.	Hand Formula B < PL: burden of prevention compared to probability times loss.	Measures whether override or safeguard cost blocks practical authority preservation.
Source Basis Derived from the uploaded ERADM Variable Origin Matrix, which maps all 11 ERADM variables to historical failure/data examples, statutory and regulatory hooks, and economic or liability doctrines. Core evidence includes Knight Capital, Cruise, 737 MAX, Tesla FSD, the Flash Crash, the 2003 Blackout, Therac-25, Waymo, mortgage automated underwriting, and associated federal, state, and international regulatory authorities.				

Table 3: ERADM Variable Origin and Selection Matrix

Table 3 documents the empirical origin and selection rationale for each ERADM variable. Rather than being derived solely from theoretical modeling, the variables were retained because recurring governance failures

across multiple domains repeatedly exposed operational conditions that became significant before, during, or after major adverse events. Historical cases including Therac-25 the 2003 Northeast Blackout Knight Capital Boeing 737 MAX

Tesla ADS/FSD investigations Cruise operational restrictions Waymo deployment oversight and automated underwriting systems revealed recurring limitations in intervention capability, supervisory responsiveness, rollback feasibility, operational flexibility, synchronization effects, infrastructure dependency, adaptation tempo, and governance visibility. Regulatory actions, statutory authorities, enforcement proceedings, liability doctrines, certification requirements, insurance practices, and post-incident investigations consistently focused on these conditions when evaluating responsibility, causation, corrective action, and operational risk[41-54]. The governance conditions represented by the ERADM variables were identified before the variables themselves were formally defined. ERADM was subsequently developed as a measurement architecture for

observing, reconstructing, and evaluating these recurring conditions. Consequently, the framework is best understood as an empirical observability architecture derived from comparative governance evidence, with variable selection informed by recurring historical patterns rather than by a pre-existing theoretical structure.

3.4 Observable Evidence Pathways

The practical utility of any governance-observability framework depends upon its ability to connect theoretical constructs to observable evidence. ERADM therefore associates each variable with identifiable operational artifacts capable of supporting observation, reconstruction, reporting, auditing, and empirical analysis.

TABLE IV
ERADM OBSERVABLE EVIDENCE SOURCES AND GOVERNANCE AUTHORITIES

VARIABLE	OBSERVABLE EVIDENCE / INDICATORS Measurable signals that reflect the variable	TYPICAL DATA SOURCES Where the evidence can be obtained	MEASUREMENT APPROACH How the variable is empirically estimated	TIME HORIZON Typical observation window	VARIABLE KEY
ADAT AI-Driven Adoption & Dependence	- Adoption mandates and executive directives - Share of critical processes supported by AI - Budget and headcount allocated to AI systems - Vendor lock-in and contractual commitments	- Government reports, budgets, procurements - Enterprise architecture inventories - Policy documents, audits - Vendor contracts, market data	Composite index (0–1) from adoption depth, criticality, and contractual lock-in	6–24 months	ADAT AI-Driven Adoption & Dependence
CLSI Closed-Loop System Integration	- Update cycle length and release frequency - Degree of autonomy in sensing–deciding–acting - Integration across functions and agencies - Human-in-the-loop removal events	- System logs, CI/CD records - Architecture diagrams, integration docs - Change management records - Operations dashboards	Integration score (0–1) based on autonomy level, update frequency, and loop closure	3–12 months	CLSI Closed-Loop System Integration
H Human Agency & Oversight	- Operator oversight time per decision - Intervention latency and frequency - Oversight staffing levels and tenure - Training hours in system oversight	- Audit logs, workflow systems - HR records, duty rosters - Training management systems - Incident and escalation reports	Human agency capacity score (0–1) from latency, interventions, and capacity resources	3–12 months	H Human Agency & Oversight
RLF Recursive Learning & Feedback	- Model retraining frequency - Feedback loop existence and strength - Performance improvement over time - Data drift adaptation events	- MLOps dashboards, model registries - Training datasets, experiment logs - Performance monitoring systems - Data quality reports	Learning velocity index (0–1) from retraining rate and performance gains	1–6 months	RLF Recursive Learning & Feedback
St System Synchronization & Coordination	- Cross-system data exchange frequency - Real-time coordination across agencies - Shared situational awareness indicators - Latency between systems	- Telemetry, system integration logs - Network monitoring tools - SOC / operations center reports - Interagency coordination records	Synchronization score (0–1) from coordination frequency and latency	1–6 months	St System Synchronization & Coordination
KI Knowledge & Infrastructure Tethering	- Dependency on proprietary models / data - Unique infrastructure requirements - Data portability and interchangeability - Switching cost estimates	- Infrastructure inventories - Data lineage and catalog systems - Vendor assessments - Cost and TCO analyses	Tethering index (0–1) from dependency depth and switching costs	6–24 months	KI Knowledge & Infrastructure Tethering
OSCR Option Space Contraction & Risk	- Number of viable alternative systems - Feasibility of fallback options - Market concentration / HH indices - Regulatory and technical barriers	- Market research, procurement data - Risk assessments, BCP documents - Vendor landscape reports - Regulatory filings	Option space index (0–1) from alternatives count and feasibility	6–24 months	OSCR Option Space Contraction & Risk
AEI Authority Erosion Index	- Decrease in human override success rate - Decline in human decision share - Rejected human recommendations - Escalation bypass frequency	- Decision logs, override logs - Case management systems - Incident reports - Governance dashboards	Authority erosion score (0–1) from override success and decision displacement	3–12 months	AEI Authority Erosion Index
OER Override Erosion Rate	- Override attempt rate - Override success rate - Override justification requirements - System resistance to override	- System logs, security logs - User activity monitoring - Policy and configuration records - Governance reports	Override erosion score (0–1) from attempts, success, and resistance	1–6 months	OER Override Erosion Rate
RFT Recovery & Fallback Time	- Mean time to restore critical function - Frequency of recovery tests - Recovery plan currency and completeness - Dependency restoration sequences	- DR test reports, incident reports - Change management records - Business continuity systems - Post-incident reviews	Recovery time score (0–1) from MTR, test results, and plan quality	6–24 months	RFT Recovery & Fallback Time
Θ Authority Transition Index	- Composite risk posture across ERADM dimensions - Trend of authority-related risk indicators - Cumulative system autonomy - Net human control margin	- All above data sources (integrated) - Risk dashboards, analytics platforms - Governance and audit systems - Executive risk reports	Weighted composite (0–1) across all ERADM variables per model specification	6–24 months	Θ Authority Transition Index

Note—Observable evidence, governance authorities, and measurement approaches identify the operational artifacts through which ERADM variables may be observed, reconstructed, audited, and empirically evaluated.
Sources: ERADM Framework; AI Policy Institutes; OECD; McKinsey; Brookings; World Economic Forum; U.S. Department of Education; National Science Foundation; Regional Technology Alliances; Global Case Studies.

Note: AGI = Artificial General Intelligence

Table 4: ERADM Observable Evidence Sources & Governance Authorities

Table 4 identifies the observable evidence, governance authorities, data sources, and measurement pathways through which each ERADM variable can be empirically evaluated. The table demonstrates that ERADM variables are not purely theoretical constructs; each can be linked to operational artifacts, audit records, regulatory documentation, system logs, incident reports, or other observable evidence. Collectively, these pathways provide a practical foundation for reconstructing, measuring, auditing, and comparing governance-relevant conditions across autonomous-system deployments. This evidence structure establishes a methodological bridge between governance theory and empirical investigation. Variables are not treated as purely conceptual constructs. Rather, they are linked to operational indicators capable of observation through

telemetry, audit records, deployment histories, intervention logs, recall actions, software updates, restoration activities, safety reports, and governance documentation.

3.5 Adaptation Tempo and Supervisory Viability

Two variables occupy a particularly important position within the framework because they directly address the relationship between system adaptation and human supervision.

The Closed-Loop Self-Improvement Interval is defined as.

$$CLSI_t = \tau_{observe,t} + \tau_{update,t} + \tau_{validate,t} + \tau_{execute,t} + \tau_{feedback,t}$$

where the components represent observation, update,

validation, deployment, and feedback activities associated with adaptation processes.

The Human Supervisory Response Horizon is defined as.

$$H_t = H_{observe,t} + H_{interpret,t} + H_{coordinate,t} + H_{authorize,t} + H_{act,t}$$

where the components represent the stages required for meaningful human intervention. The relationship between CLSI and H establishes a measurable basis for evaluating supervisory viability. Governance becomes increasingly challenging when adaptation processes accelerate relative to supervisory responsiveness. The framework therefore treats the interaction between adaptation tempo and supervisory response as a critical governance condition rather than a purely organizational concern.

3.5 Composite Authority-Transition Pressure

The Authority Transition Index (Θ) integrates the principal governance variables into a unified indicator of authority-transition pressure. The index is not intended to predict accidents, failures, or specific operational outcomes. Instead, it summarizes the cumulative effect of operational conditions that may influence whether meaningful supervision, intervention, restoration, and governance remain practically achievable.

$$\Theta_t = f(AEI_t, OER_t, RFT_t, OSCR_t, CLSI_t, H_t, S_t, K_t)$$

Higher values of Θ indicate increasing authority-transition pressure associated with reduced intervention effectiveness, compressed operational flexibility, elevated dependency formation, accelerating adaptation, increasing synchronization, or deteriorating restoration capability. Lower values indicate conditions more compatible with effective governance and operational oversight. The significance of Θ lies not in prediction but in observability. It provides a structured mechanism for evaluating how multiple governance-relevant conditions interact and evolve over time.

3.6 Summary

ERADM transforms governance observability from a predominantly qualitative and outcome-oriented activity into a measurable systems discipline. Through a structured set of variables, observable evidence pathways, measurable relationships, and composite indicators, the framework establishes a methodology through which governance relevant operational conditions become direct objects of observation, measurement, reporting, auditing, reconstruction, and empirical investigation. The next section develops the mathematical relationships that govern interactions among these variables and establishes the formal analytical foundation of the framework.

4. Mathematical Formulation

4.1 Variable Definitions and System State

ERADM model's governance observability as a dynamic systems problem in which governance-relevant operational conditions evolve through interactions among intervention

capability, supervisory responsiveness, restoration feasibility, operational flexibility, adaptation tempo, synchronization intensity, infrastructure dependency, and recursive leverage. Rather than treating governance as a static institutional property, the framework represents governance effectiveness as an emergent consequence of measurable operational conditions.

The state of the system at time (t) is represented by.

$$X_t = [ADAT_t, RLF_t, CLSI_t, H_t, AEI_t, OER_t, RFT_t, OSCR_t, S_t, K_t]^T$$

where each element corresponds to an observable governance-relevant condition.

This representation enables governance conditions to be evaluated dynamically rather than solely through static compliance, organizational structures, or post-event outcomes.

4.2 Dynamic Relationships

The framework proposes that governance-relevant variables interact through directional relationships that generate observable and empirically testable effects. Authority Elasticity Index evolves as a function of governance capability, synchronization intensity, infrastructure coupling, supervisory responsiveness, and adaptation tempo.

$$\dot{AEI}_t = \alpha_G G_t + \alpha_H H_t + \alpha_C CLSI_t - \alpha_S S_t - \alpha_K K_t$$

This relationship reflects the proposition that increasing synchronization and dependency formation may reduce practical intervention capacity, while improved governance structures and adaptation management may preserve intervention effectiveness. Override Effectiveness Rate evolves according to intervention capability and operational complexity.

$$\dot{OER}_t = \beta_A AEI_t - \beta_{SK} (S_t K_t)$$

where

$$\beta_A, \beta_{SK} > 0$$

where $\backslash(AEI_t\backslash)$ represents available intervention capability, $\backslash(S_t\backslash)$ represents synchronization intensity, and $\backslash(K_t\backslash)$ represents infrastructure coupling. The expected signs reflect the hypothesis that override effectiveness increases when intervention capability is preserved but decreases when synchronization and infrastructure coupling jointly increase operational complexity.

Rollback Feasibility Time is modeled as

$$\dot{RFT}_t = \gamma_K K_t - \gamma_O OER_t$$

where

$$\gamma_K, \gamma_O > 0$$

where $\backslash(K_t\backslash)$ represents infrastructure coupling and $\backslash(OER_t\backslash)$ represents override effectiveness rate. The expected signs reflect the hypothesis that increasing infrastructure dependency makes restoration more difficult and time-consuming, while effective intervention mechanisms improve recovery capability and reduce restoration time.

Option Set Compression Ratio evolves according to recursive amplification and synchronization effects.

$$O\dot{S}C R_t = \lambda_R R L F_t + \lambda_S S_t - \lambda_A A E I_t$$

where

$$\lambda_R, \lambda_S, \lambda_A > 0$$

where $\backslash(R L F_t\backslash)$ represents Recursive Leverage Factor, $\backslash(S_t\backslash)$ represents Synchronization Intensity, and $\backslash(A E I_t\backslash)$ represents Authority Elasticity Index. The expected signs reflect the hypothesis that recursive amplification and increasing synchronization reduce operational flexibility by concentrating activity into fewer viable pathways, while preserved intervention capability mitigates option-set compression by maintaining the availability of meaningful operational alternatives. Synchronization Intensity evolves as.

$$\dot{S}_t = \phi_A A D A T_t + \phi_R R L F_t - \phi_O O S C R_t$$

where

$$\phi_A, \phi_R, \phi_O > 0$$

where $\backslash(A D A T_t\backslash)$ represents Adoption-Driven Authority Transfer, $\backslash(R L F_t\backslash)$ represents Recursive Leverage Factor, and $\backslash(O S C R_t\backslash)$ represents Option Set Compression Ratio. The expected signs reflect the hypothesis that increasing adoption and recursive amplification promote synchronization across system instances, while greater operational diversity and alternative pathways reduce synchronization intensity.

Infrastructure Coupling evolves as.

$$\dot{K}_t = \omega_A A D A T_t + \omega_S S_t - \omega_R R F T_t^{-1}$$

where

$$\omega_A, \omega_S, \omega_R > 0$$

where $\backslash(A D A T_t\backslash)$ represents Adoption-Driven Authority Transfer, $\backslash(S_t\backslash)$ represents Synchronization Intensity, and $\backslash(R F T_t^{-1}\backslash)$ represents restoration capability. The expected signs reflect the hypothesis that adoption-driven dependence and synchronization increase infrastructure coupling, while greater restoration capability reduces effective dependency by preserving fallback, recovery, and

decoupling options. The expected relationship between infrastructure dependence and governance constraints is consistent with prior research on tightly coupled systems. As dependency on shared infrastructure increases, intervention pathways, isolation options, and recovery mechanisms may become progressively constrained, increasing the operational significance of restoration and fallback capabilities [37]. These equations are not intended as fully calibrated predictive models. Their purpose is to formalize directional relationships that can be subjected to empirical testing, calibration, and falsification.

4.3 Governance-Feasible Region

A central proposition of ERADM is that effective governance exists only within a bounded operational region defined by measurable conditions. Governance feasibility requires - Threshold Conditions.

$$A E I_t \geq \alpha, \quad O E R_t \geq \beta, \quad R F T_t \leq \gamma, \quad O S C R_t \leq \delta, \quad C L S I_t \geq H_t$$

Governance Feasible Region.

$$\mathcal{G}_t = \{X_t : A E I_t \geq \alpha, O E R_t \geq \beta, R F T_t \leq \gamma, O S C R_t \leq \delta, C L S I_t \geq H_t\}$$

The parameters $\backslash(\alpha\backslash)$, $\backslash(\beta\backslash)$, $\backslash(\gamma\backslash)$, and $\backslash(\delta\backslash)$ represent governance thresholds for intervention capability, override effectiveness, restoration feasibility, and operational flexibility. Collectively, these conditions define the governance-feasible region within which meaningful intervention, supervision, restoration, and operational flexibility remain operationally viable. Governance effectiveness is expected to decline as variables move beyond these boundaries, particularly when adaptation processes begin to outpace supervisory response $((C L S I_t < H_t))$, reducing the practical ability of human actors to maintain effective oversight [24-39]. Similar concerns appear in Rasmussen's dynamic risk-management framework, which argues that effective supervision depends upon maintaining sufficient visibility and response capacity as system complexity and operational tempo increase [39]. Threshold values are domain dependent and are not assumed to be universal. In practical applications, thresholds may be established through empirical calibration, expert elicitation, regulatory requirements, historical performance baselines, safety-management criteria, or other domain-specific governance objectives. The governance-feasible region may also be expressed in terms of aggregate authority-transition pressure. The parameter $\backslash(\Theta_{crit}\backslash)$ represents a domain-specific governance threshold corresponding to the maximum level of authority-transition pressure considered compatible with effective governance. **Although ERADM was initially developed using evidence drawn primarily from autonomous-vehicle deployments, the framework is intended to be extensible across autonomous and adaptive systems. Consequently, $\backslash(\Theta_{crit}\backslash)$ may vary across autonomous vehicles, uncrewed aircraft systems (UAS), aviation, robotics, healthcare automation, critical infrastructure, logistics networks, financial systems, and defense environments according to domain-specific operational requirements, regulatory expectations,**

safety objectives, and governance-risk tolerances.

The corresponding governance-feasible state space may then be defined as.

$$\mathcal{G}_t^\theta = \{X_t : AEI_t \geq \alpha, OER_t \geq \beta, RFT_t \leq \gamma, OSCR_t \leq \delta, CLSI_t \geq H_t, \theta_t \leq \theta_{crit}\}$$

Interpretation.

- $AEI_t \geq \alpha \rightarrow$ intervention remains effective.
- $OER_t \geq \beta \rightarrow$ overrides remain effective.
- $RFT_t \leq \gamma \rightarrow$ restoration remains feasible.
- $OSCR_t \leq \delta \rightarrow$ operational flexibility preserved.
- $CLSI_t \geq H_t \rightarrow$ governance keeps pace with adaptation.
- $\theta_t \leq \theta_{crit} \rightarrow$ aggregate authority-transition pressure remains within acceptable bounds.

TABLE V. ERADM Empirical Governance Observability, Authority Breach, and Statutory Trigger Matrix

Historical cases shown where STAMP identifies hazards and ERADM makes governance conditions measurable.

Note: Boeing 737 MAX and Therac-25 are included as cross-domain control-system analogues rather than AI systems. Their inclusion illustrates that governance observability failures, authority degradation, rollback constraints, and supervisory limitations are not unique to artificial intelligence but may occur across automated and software-controlled systems.						
Historical Case	STAMP Hazard / UCA	ERADM Variables	Governance Condition Revealed	Authority Breach Evidence	Statutory / Regulatory Authority Triggered	Potential Governance Consequence
Uber ATG (2018)	Failure to provide effective intervention before collision	AEI, OER, H	Formal supervision existed, but intervention capacity degraded operationally	Safety driver failed to intervene within available response window	49 U.S.C. §§30118–30120; 49 U.S.C. Chapter 11; Arizona Executive Order 2015-09	Investigation, recall review, operational restrictions, safety recommendations
Tesla ADS/FSD (2022–2024)	Unsafe control actions during driver-supervised automation	OER, CLSI, H	Supervisory compression emerged as adaptation tempo increased	OTA cadence exceeded traditional oversight and review tempo	49 U.S.C. §§30118–30120; NHTSA SGO 2021-01; 49 CFR Part 579	Defect investigation, recall authority, reporting obligations, operational scrutiny
Cruise (2023)	Inadequate recovery and operational response following incident	RFT, OSCR, AEI	Restoration capability and operational flexibility became governance-critical	Service suspension demonstrated rollback and recovery constraints	California Vehicle Code §38750; CCR Title 13 AV Regulations; CPUC driverless deployment authority; 49 U.S.C. Chapter 301	Permit suspension, deployment restrictions, corrective-action requirements
Waymo (2024)	Potential fleet-wide propagation of control actions	S, K, OSCR	Infrastructure dependency and synchronization exposure became observable	Operational dependence mediated through integrated infrastructure	California Vehicle Code §38750; CA DMV driverless deployment authority; NHTSA ADS reporting framework	Deployment conditions, reporting requirements, operational limitations
Boeing 737 MAX	UCA-1, UCA-3, UCA-4 involving MCAS control actions	OER, H, CLSI, RFT	Human response tempo failed to match system tempo	Pilots possessed nominal authority but insufficient intervention effectiveness	49 U.S.C. Subtitle VII; 14 CFR Part 25; FAA ODA framework; FAA SMS requirements	Grounding, recertification, certification reform, operational restrictions
Therac-25	UCA-2: unsafe software control action provided	RFT, OSCR, Governance Observability Latency	No practical restoration pathway existed once unsafe state emerged	Rollback effectively zero; hardware interlocks absent	Federal Food, Drug, and Cosmetic Act, 21 U.S.C. §360c et seq.; 21 CFR Part 820; FDA MDR requirements	Device redesign, corrective action, regulatory intervention
Fleet-Wide OTA Systems	Synchronization-related unsafe control actions	CLSI, S, K	Single updates can propagate system-wide operational effects	Supervisory structures may lag synchronized deployment and validation cycles	OMB M-24-10; NIST AI RMF 1.0; ISO/IEC 42001:2023; ISO 26262:2018; EU AI Act obligations	Risk-management intervention, compliance review, deployment controls

Evidence and Regulatory Sources
 NTSB (2019); NHTSA ADS investigations and Standing General Order data (2022–2024); California DMV (2023); California Public Utilities Commission (2023); Waymo safety reports (2024); FAA 14 CFR Part 25; FDA 21 CFR Part 820; 49 U.S.C. Chapter 301; NIST AI RMF 1.0 (2023); OMB M-24-10 (2024); ISO/IEC 42001:2023; ISO 26262:2018.
 Statutory authorities are provided to show governance relevance and potential enforcement pathways; inclusion does not imply a legal finding or adjudicated violation.

Table 5: ERADM Empirical Governance Observability Authority Breach and Statutory Trigger Matrix

Table 5 demonstrates how historical autonomous-system failures and control-system incidents can be analyzed using both STAMP and ERADM. The table links observable governance conditions, authority degradation indicators, and operational constraints to real-world cases, showing how ERADM variables can be reconstructed from empirical evidence and connected to regulatory, liability, and oversight actions. Collectively, the cases illustrate that governance-relevant conditions often become observable before major interventions, recalls, suspensions, or corrective actions occur [24–36]. The significance of the governance-feasible region is that it provides measurable criteria for evaluating governance viability independently of whether a failure, recall, accident, or regulatory intervention has already occurred.

4.4. Authority Transition Index (θ_t)

The Authority Transition Index (θ_t) serves as the framework's principal composite indicator. The canonical formulation is

$$\theta_t = \kappa \cdot ADAT_t \cdot \left(\frac{H_t}{CLSI_t} \right) \cdot \left(\frac{RLF_t}{RLF_{base}} \right) \cdot OSCR_t \cdot (1 - AEI_t) \cdot (1 - OER_t) \cdot \left(\frac{RFT_t}{RFT_0} \right)$$

where

- (θ_t) represents authority transition index.
- (κ) represents scaling constant.
- ($ADAT_t$) represents adoption-driven authority transfer.
- (H_t) represents supervisory responsiveness.
- ($CLSI_t$) represents adaptation tempo.
- (RLF_{base}) represents baseline recursive leverage for normalization.
- (RLF_t) represents recursive leverage.
- ($OSCR_t$) represents operational flexibility.
- (AEI_t) represents intervention capacity.
- (OER_t) represents override effectiveness.
- (RFT_{max}) represents maximum acceptable rollback time and
- (RFT_t) represents restoration capability.

The index does not predict specific outcomes. Rather, it aggregates governance-relevant operational conditions into a unified measure of authority-transition pressure. Higher values indicate conditions associated with increasing governance stress, reduced intervention capability, elevated dependency formation, compressed operational flexibility, and declining restoration capability. Lower values indicate

conditions more compatible with sustained governance effectiveness.

Sensitivity interpretation follows directly from the structure of Θ . Holding other variables constant, authority transition pressure increases when intervention capability (AEI) declines, override effectiveness (OER) weakens, rollback feasibility deteriorates (higher RFT), operational alternatives become compressed (higher OSCR), or supervisory response becomes increasingly constrained relative to adaptation tempo. Conversely, authority-transition pressure decreases when intervention effectiveness, restoration capability, operational flexibility, and supervisory viability improve. The purpose of this sensitivity interpretation is not to provide calibrated prediction, but to demonstrate how changes in observable governance conditions directionally influence aggregate authority-transition pressure and therefore provide a basis for future empirical calibration, comparative analysis, and longitudinal measurement.

4.5. External Observability Amplification & Internal Governance Observability Index (Ω)

Governance outcomes are not determined exclusively by internal system conditions. Insurance institutions, regulators, reinsurers, certification bodies, capital markets, and systemic-risk organizations collectively influence how observed governance conditions translate into operational constraints, intervention thresholds, market responses, and governance consequences [21-41].

The Core Authority Transition Index, the framework incorporates the Market-Regulatory-Reinsurance Observability Weight (MRROW).

$$\Theta_t = \kappa \cdot ADAT_t \cdot \frac{H_t}{CLSI_t} \cdot \frac{RLF_t}{RLF_{base}} \cdot OSCR_t \cdot (1 - AEI_t) \cdot (1 - OER_t) \cdot \frac{RFT_t}{RFT_0}$$

Externally Conditioned Authority Transition Index

$$\Theta_t^{MRROW} = \Theta_t \cdot MRROW_t$$

or, equivalently,

$$\Theta_t^{MRROW} = \kappa \cdot ADAT_t \cdot \left(\frac{H_t}{CLSI_t} \right) \cdot \left(\frac{RLF_t}{RLF_{base}} \right) \cdot OSCR_t \cdot (1 - AEI_t) \cdot (1 - OER_t) \cdot \left(\frac{RFT_t}{RFT_0} \right) \cdot MRROW_t$$

Institutionally Observable Authority Transition Index

If you wish to distinguish **institutional observability** from the latent (internal) transition pressure, define:

$$\Theta_t^{obs} = \Theta_t \cdot MRROW_t^{obs}$$

where:

- Θ_t = latent internal Authority Transition Index.
- $MRROW_{tob}$ = Market-Regulatory-Reinsurance Observability Weight.
- Θ_{tobs} = institutionally observable authority-transition pressure.

Higher MRROW values indicate greater regulatory scrutiny, stronger insurance and reinsurance sensitivity, more restrictive governance expectations, increased certification requirements, heightened market accountability, and faster institutional responses to emerging governance risk. Lower MRROW values indicate comparatively limited external attention despite similar internal governance conditions. Because MRROW represents an **external institutional modifier** rather than a primitive governance condition, it is intentionally maintained outside the canonical ERADM state vector and does not alter the computation of either ERADM or the Governance Applicability Index. Instead, MRROW transforms the latent Authority Transition Index (Θ_t) into the Institutionally Observable Authority Transition Index (Θ_{tobs}), thereby quantifying how external institutional actors amplify or attenuate the observable governance significance of the underlying transition while preserving the distinction between intrinsic governance state and external institutional response.

Internal Governance Observability Index (Ω).

$$\Omega_t = \frac{ADAT_{obs} + RLF_{obs} + CLSI_{obs} + H_{obs} + AEI_{obs} + OER_{obs} + RFT_{obs} + OSCR_{obs} + S_{obs} + K_{obs}}{10}$$

The **Internal Governance Observability Index (Ω_t)** measures the extent to which the governance conditions represented by the canonical ERADM variables can be directly observed, reliably reconstructed, and independently verified from available operational evidence. Higher values indicate greater confidence that governance-relevant conditions—including authority allocation, supervisory capability, intervention effectiveness, restoration feasibility, synchronization, infrastructure coupling, and operational flexibility—can be accurately characterized from the available evidence. Lower values indicate increasing uncertainty arising from incomplete, indirect, delayed, or conflicting observations. Unlike the Authority Transition Index (Θ_t), which measures governance-transition pressure, Ω_t measures the **observability of the underlying governance state itself**. Consequently, Ω_t serves as an evidence-quality measure rather than a governance-performance or framework-applicability measure.

4.6 Falsification Conditions

A scientific observability architecture must be capable of empirical contradiction [38].

ERADM therefore generates explicit falsification conditions.

The framework would be challenged if.

1. Increasing synchronization does not reduce intervention effectiveness.
2. Increasing infrastructure coupling does not increase restoration difficulty.
3. Increasing recursive leverage does not contribute to option-set compression.
4. Override effectiveness remains stable despite declining intervention opportunities.
5. Supervisory responsiveness remains unaffected by accelerating adaptation cycles.

6. Authority-transition pressure remains unchanged despite simultaneous deterioration in intervention capability, restoration feasibility, and operational flexibility.

TABLE VI
ERADM FALSIFICATION CONDITIONS

 ERADM is falsified if the following observations occur. These observations would demonstrate that the hypothesized relationships do not hold.		
HYPOTHESIS	PREDICTED RELATIONSHIP	OBSERVATION THAT FALSIFIES
 AEI decreases with synchronization	 As synchronization (S) increases, Authority Erosion Index (AEI) should decrease.	 No measurable decline in AEI after repeated increases in S.
 OER depends on AEI	 As AEI decreases, Operational Effectiveness Ratio (OER) should also decrease.	 OER remains constant despite sustained AEI decline.
 RFT depends on K	 As Infrastructure Coupling (K) increases, Restoration Feasibility Time (RFT) should increase.	 Restoration feasibility time (RFT) unaffected by coupling increases.
 If any falsifying observation is confirmed, the corresponding hypothesis—and potentially ERADM—would be invalidated.		

Table 6: ERADM Falsification Conditions

Table 6: presents the principal falsification conditions for ERADM by identifying observations that would contradict the model's predicted relationships. Rather than assuming validity, ERADM specifies measurable outcomes that could demonstrate the framework is incorrect. The table therefore establishes scientific rigor by defining explicit empirical conditions under which key hypotheses—and potentially the broader model itself—would be rejected. These conditions establish clear pathways through which empirical evidence may support, refine, calibrate, or invalidate portions of the framework.

4.7 Optimal Governance Observability Application Domains

ERADM is designed for environments in which governance effectiveness depends upon maintaining meaningful supervision, intervention capability, restoration feasibility, operational flexibility, coordination, and institutional oversight. The framework is expected to provide the greatest analytical value within autonomous, adaptive, software mediated, and operationally evolving systems whose behavior may change over time and whose outcomes are influenced by human, organizational, technical, and regulatory interactions. The framework is particularly applicable where governance-relevant dependencies, synchronization effects, infrastructure coupling, adaptation processes, and authority relationships influence operational outcomes. Under such conditions, variables including AEI, OER, RFT, OSCR,

CLSI, H, S, K, and Θ provide mechanisms through which governance conditions may be observed, reconstructed, measured, compared, and empirically evaluated across systems and deployments. Examples include autonomous transportation systems, robotics, critical infrastructure, healthcare automation, logistics networks, industrial control systems, financial platforms, defense environments, and other complex socio-technical systems in which effective governance depends not only upon authority and accountability, but also upon the continued viability of intervention, supervision, restoration, and coordination capabilities. As autonomy, adaptation, connectivity, and operational complexity increase, governance effectiveness becomes increasingly dependent upon conditions that are often only partially observable through traditional reporting architectures. ERADM is intended to improve visibility into these conditions by providing a structured observability architecture through which governance-relevant operational dynamics can be evaluated before they become visible through adverse outcomes.

4.8 Summary

The mathematical formulation transforms ERADM from a conceptual governance framework into a measurable, testable, and potentially falsifiable systems model. State-space representation, dynamic relationships, governance feasible conditions, composite authority-transition measurement, external observability amplification,

empirically grounded application domains, and explicit falsification criteria collectively establish a foundation for systematic empirical evaluation, calibration, comparison, and refinement. By specifying observable variables, measurable relationships, boundary conditions, and conditions under which portions of the framework may be contradicted by evidence, ERADM advances governance observability from a primarily conceptual concern toward a scientifically evaluable systems property. The next section develops the observability architecture through which these variables become measurable using operational evidence, reporting systems, audit records, deployment histories, intervention logs, safety reports, regulatory actions, liability records, and governance documentation. This architecture provides the operational pathways through which governance-relevant conditions may be observed, reconstructed, measured, and evaluated across real-world autonomous-system deployments.

5. Dual-Layer Observability Architecture

5.1 Internal System Observability

A central requirement of ERADM is that governance-relevant operational conditions must be linked to observable evidence. The framework therefore associates each variable with identifiable operational artifacts through which underlying governance conditions can be observed, reconstructed, reported, audited, and empirically evaluated [24,38]. Unlike outcome-oriented reporting architectures that primarily document crashes, incidents, recalls, compliance actions, or operational restrictions, ERADM focuses on operational conditions that exist before such outcomes emerge. The framework treats intervention capability, override effectiveness, restoration feasibility, supervisory responsiveness, operational flexibility, synchronization intensity, and infrastructure dependency as measurable phenomena capable of direct observation through operational evidence. The internal observability architecture is organized around a series of complementary analytical views that collectively transform ERADM from a conceptual framework into a scientifically testable measurement architecture.




Figure 1: ERADM Primary Casual Pathway

Figure 1 illustrates the canonical causal pathway underlying ERADM. The pathway traces how AI adoption, integration, recursive learning, synchronization, and infrastructure dependence can influence operational flexibility, intervention effectiveness, recovery capability, and ultimately authority-transition pressure. The figure provides the conceptual backbone of the framework by showing the hypothesized

sequence through which changes in system architecture may affect governance observability and meaningful human control. The significance of the pathway is not prediction but observability. It identifies a sequence of measurable operational conditions that can be evaluated empirically and subjected to falsification through real-world evidence [24,38].

TABLE VII. ERADM Direct Mathematical Dependency Matrix

VARIABLES THAT APPEAR DIRECTLY IN EACH ERADM MODEL EQUATION

VARIABLE	DIRECT INPUTS (APPEAR IN FUNCTION / EQUATION) Variables that directly influence this variable in the ERADM model	INCLUDES SELF?	EXOGENOUS (OUTSIDE MODEL)?	EQUATION REFERENCE	LEGEND
ADAT AI-Driven Adoption & Dependence	Exogenous (driven by external policy, market, and technology factors)	No	✓	(1)	✓ Yes Variable is included as a direct input ✗ No Variable is not included as a direct input ✓ Exogenous Determined outside the ERADM model ✗ Endogenous Determined within the ERADM model (n) Reference to the numbered equation in the ERADM model specification
CLSI Closed-Loop System Integration	$\tau_{observe}, \tau_{update}, \tau_{validate}, \tau_{execute}, \tau_{feedback}$	No	✓	(2)	
H Human Agency & Oversight	$H_{phys}, H_{cog}, H_{coord}, H_{act}$	No	✓	(3)	
RLF Recursive Learning & Feedback	$RLF_{t-1}, ADAT_t$, System performance metrics	✓	✗	(4)	
St System Synchronization & Coordination	$ADAT_t, RLF_t, S_{t-1}$	✓	✗	(5)	
Kt Knowledge & Infrastructure Tethering	$ADAT_t, St_t, K_{t-1}$	✓	✗	(6)	
OSCR Option Space Contraction & Risk	$RLF_t, St_t, K_t, OSCR_{t-1}$	✓	✗	(7)	
AEI Authority Erosion Index	$H_t, CLSI_t, St_t, K_t, OSCR_t, AEI_{t-1}$	✓	✗	(8)	
OER Override Erosion Rate	$AEI_t, St_t, K_t, OER_{t-1}$	✓	✗	(9)	
RFT Recovery & Fallback Time	OER_t, K_t, RFT_{t-1}	✓	✗	(10)	
Θ Authority Transition Index	$ADAT_t, H_t, CLSI_t, RLF_t, OSCR_t, AEI_t, OER_t, RFT_t$	No	✗	(11)	

VARIABLE KEY

ADAT AI-Driven Adoption & Dependence**CLSI** Closed-Loop System Integration**H** Human Agency & Oversight**RLF** Recursive Learning & Feedback**St** System Synchronization & Coordination**Kt** Knowledge & Infrastructure Tethering**OSCR** Option Space Contraction & Risk**AEI** Authority Erosion Index**OER** Override Erosion Rate**RFT** Recovery & Fallback Time**Θ** Authority Transition Index

EQUATION REFERENCE

(1)	AI-Driven Adoption & Dependence (ADAT)	$ADAT_t = f(\text{Exogenous Factors}_t)$	ADAT is determined by exogenous factors outside the ERADM model, including external policy, market, and technology drivers.
(2)	Closed-Loop System Integration (CLSI)	$CLSI_t = f(\tau_{observes}, \tau_{update}, \tau_{validate}, \tau_{execute}, \tau_{feedback})$	CLSI depends on the effectiveness of each closed-loop operational component.
(3)	Human Agency & Oversight (H)	$H_t = f(H_{phys,t}, H_{cog,t}, H_{coord,t}, H_{act,t})$	Human agency is a function of physical capacity, cognitive capacity, coordination, and action capability.
(4)	Recursive Learning & Feedback (RLF)	$RLF_t = f(RLF_{t-1}, ADAT_t, \text{System Performance Metrics}_t)$	RLF evolves based on its prior state, current ADAT, and observed system performance.
(5)	System Synchronization & Coordination (St)	$St_t = f(ADAT_t, RLF_t, S_{t-1})$	St is driven by current ADAT and RLF and its prior synchronization state.
(6)	Knowledge & Infrastructure Tethering (Kt)	$K_t = f(ADAT_t, St_t, K_{t-1})$	Kt depends on current ADAT and St and its prior knowledge/infrastructure state.
(7)	Option Space Contraction & Risk (OSCR)	$OSCR_t = f(RLF_t, St_t, K_t, OSCR_{t-1})$	OSCR is influenced by learning feedback, synchronization, knowledge tethering, and its prior state.
(8)	Authority Erosion Index (AEI)	$AEI_t = f(H_t, CLSI_t, St_t, K_t, OSCR_t, AEI_{t-1})$	AEI depends on human agency, system integration, synchronization, knowledge tethering, option space risk, and its prior state.
(9)	Override Erosion Rate (OER)	$OER_t = f(AEI_t, St_t, K_t, OER_{t-1})$	OER is driven by authority erosion, system synchronization, knowledge tethering, and its prior erosion rate.
(10)	Recovery & Fallback Time (RFT)	$RFT_t = f(OER_t, K_t, RFT_{t-1})$	RFT depends on override erosion, knowledge tethering, and its prior recovery time.
(11)	Authority Transition Index (Θ)	$\Theta_t = f(ADAT_t, H_t, CLSI_t, RLF_t, OSCR_t, AEI_t, OER_t, RFT_t)$	Θ is a composite index driven by all other ERADM variables to capture authority transition dynamics.

SOURCES: ERADM Framework; AI Policy Institutes; OECD; McKinsey Brookings; World Economic Forum; U.S. Department of Education; National Science Foundation; Regional Technology Alliances; Global Case Studies.

Note: AGI = Artificial General Intelligence

Table 7: ERADM Direct Mathematical Dependency Matrix

Table 7 identifies direct mathematical dependencies among ERADM variables. The matrix distinguishes variables that appear explicitly within model equations from variables that influence outcomes indirectly through intermediate

relationships. This distinction provides transparency regarding how governance conditions are formally represented within the framework.

TABLE VIII. ERADM Observable Pathways Matrix
EMPIRICAL INDICATORS AND DATA SOURCES FOR EACH ERADM VARIABLE

VARIABLE	OBSERVABLE EVIDENCE / INDICATORS Measurable signals that reflect the variable	TYPICAL DATA SOURCES Where the evidence can be obtained	MEASUREMENT APPROACH How the variable is empirically estimated	TIME HORIZON Typical observation window	VARIABLE KEY
ADAT AI-Driven Adoption & Dependence	- Adoption mandates and executive directives - Share of critical processes supported by AI - Budget and headcount allocated to AI systems - Vendor lock-in and contractual commitments	- Government reports, budgets, procurements - Enterprise architecture inventories - Policy documents, audits - Vendor contracts, market data	Composite index (0–1) from adoption depth, criticality, and contractual lock-in	6–24 months	ADAT AI-Driven Adoption & Dependence CLSI Closed-Loop System Integration H Human Agency & Oversight RLF Recursive Learning & Feedback St System Synchronization & Coordination Kt Knowledge & Infrastructure Tethering OSCR Option Space Contraction & Risk AEI Authority Erosion Index OER Override Erosion Rate RFT Recovery & Fallback Time ⊖ Authority Transition Index
CLSI Closed-Loop System Integration	- Update cycle length and release frequency - Degree of autonomy in sensing–deciding–acting - Integration across functions and agencies - Human-in-the-loop removal events	- System logs, CI/CD records - Architecture diagrams, integration docs - Change management records - Operations dashboards	Integration score (0–1) based on autonomy level, update frequency, and loop closure	3–12 months	
H Human Agency & Oversight	- Operator oversight time per decision - Intervention latency and frequency - Oversight staffing levels and tenure - Training hours in system oversight	- Audit logs, workflow systems - HR records, duty rosters - Training management systems - Incident and escalation reports	Human agency capacity score (0–1) from latency, interventions, and capacity resources	3–12 months	
RLF Recursive Learning & Feedback	- Model retraining frequency - Feedback loop existence and strength - Performance improvement over time - Data drift adaptation events	- MLOps dashboards, model registries - Training datasets, experiment logs - Performance monitoring systems - Data quality reports	Learning velocity index (0–1) from retraining rate and performance gains	1–6 months	
St System Synchronization & Coordination	- Cross-system data exchange frequency - Real-time coordination across agencies - Shared situational awareness indicators - Latency between systems	- Telemetry, system integration logs - Network monitoring tools - SOC / operations center reports - Interagency coordination records	Synchronization score (0–1) from coordination frequency and latency	1–6 months	
Kt Knowledge & Infrastructure Tethering	- Dependency on proprietary models / data - Unique infrastructure requirements - Data portability and interchangeability - Switching cost estimates	- Infrastructure inventories - Data lineage and catalog systems - Vendor assessments - Cost and TCO analyses	Tethering index (0–1) from dependency depth and switching costs	6–24 months	
OSCR Option Space Contraction & Risk	- Number of viable alternative systems - Feasibility of fallback options - Market concentration / HHI indices - Regulatory and technical barriers	- Market research, procurement data - Risk assessments, BCP documents - Vendor landscape reports - Regulatory filings	Option space index (0–1) from alternatives count and feasibility	6–24 months	
AEI Authority Erosion Index	- Decrease in human override success rate - Decline in human decision share - Rejected human recommendations - Escalation bypass frequency	- Decision logs, override logs - Case management systems - Incident reports - Governance dashboards	Authority erosion score (0–1) from override success and decision displacement	3–12 months	
OER Override Erosion Rate	- Override attempt rate - Override success rate - Override justification requirements - System resistance to override	- System logs, security logs - User activity monitoring - Policy and configuration records - Governance reports	Override erosion score (0–1) from attempts, success, and resistance	1–6 months	
RFT Recovery & Fallback Time	- Mean time to restore critical function - Frequency of recovery test completions - Recovery plan currency and completeness - Dependency restoration sequences	- DR test reports, incident reports - Change management records - Business continuity systems - Post-incident reviews	Recovery time score (0–1) from MTR, test results, and plan quality	6–24 months	
⊖ Authority Transition Index	- Composite risk posture across ERADM dims - Trend of authority-related risk indicators - Cumulative system autonomy - Net human control margin	- All above data sources (integrated) - Risk dashboards, analytics platforms - Governance and audit systems - Executive risk reports	Weighted composite (0–1) across all ERADM variables per model specification	6–24 months	
SOURCES: ERADM Framework; AI Policy Institutes; OECD; McKinsey; Brookings; World Economic Forum; U.S. Department of Education; National Science Foundation; Regional Technology Alliances; Global Case Studies.					NOTES This matrix specifies the observable evidence and data sources used to empirically estimate each ERADM variable. Indicators are selected to be measurable, auditable, and comparable across states and institutions. Measurement approaches translate raw evidence into standardized scores (0–1) as defined in the ERADM model.
Note: AGI = Artificial General Intelligence					

Note: Table 8 links each ERADM variable to observable evidence sources, measurement pathways, and reconstruction methodologies. The matrix establishes the empirical foundation of the framework by demonstrating how theoretical constructs can be connected to operational indicators including intervention records, deployment histories, software updates, audit trails, restoration activities, safety reports, and regulatory documentation.

Table 8: ERADM Observable Pathways Matrix

TABLE IX. ERADM Governance–Critical Pathways Matrix

KEY VARIABLE INFLUENCE RELATIONSHIPS FOR EMPIRICAL TESTING

PATHWAY ID	GOVERNANCE–CRITICAL PATHWAY (From → To)	EXPECTED DIRECTION (Sign)	INFLUENCE INTERPRETATION (Empirical Meaning)	LEGEND
1	S → AEI	< 0	Higher synchronization reduces authority elasticity.	→ Directional Influence Represents the expected causal influence from the source variable (left) to the target variable (right). DIRECTION KEY ↑ Positive Influence (> 0) Increase in source leads to increase in target ↓ Negative Influence (< 0) Increase in source leads to decrease in target VARIABLE KEY CLSI Closed-Loop System Integration H Human Supervisory Response Horizon RLF Recursive Learning & Feedback S System Synchronization & Coordination K Knowledge & Infrastructure Tethering OSCR Option Space Contraction & Risk AEI Authority Erosion Index OER Override Erosion Rate RFT Recovery & Fallback Time ⓪ Authority Transition Index
2	K → AEI	< 0	Higher infrastructure coupling reduces authority elasticity.	
3	AEI → OER	> 0	Higher authority elasticity improves override effectiveness.	
4	K → RFT	> 0	Higher coupling increases rollback/recovery time.	
5	OER → RFT	< 0	Stronger overrides reduce rollback/recovery time.	
6	RLF → OSCR	> 0	Recursive learning increases option compression.	
7	S → OSCR	> 0	Synchronization increases option compression.	
8	AEI → ⓪	< 0	Authority elasticity lowers transition pressure.	
9	OER → ⓪	< 0	Override effectiveness lowers transition pressure.	
10	OSCR → ⓪	> 0	Option compression increases transition pressure.	
11	H → ⓪	> 0	Longer human response horizon increases transition pressure.	
12	CLSI → ⓪	< 0	Faster closed-loop adaptation lowers transition pressure.	
13	(AEI, OER, OSCR, H, CLSI) → ⓪	–	These variables jointly determine overall transition pressure.	
NOTES				
These governance–critical pathways represent the most theoretically and empirically important variable influences in the ERADM model. They form the primary targets for empirical measurement, model validation, and falsification. All other relationships in the ERADM model are mediated through these core pathways.				
SOURCES: ERADM Framework; AI Policy Institutes; OECD; McKinsey; Brookings; World Economic Forum; U.S. Department of Education; National Science Foundation; Regional Technology Alliances; Global Case Studies.				Note: AGI = Artificial General Intelligence

Note: Table 9 summarizes the principal directional relationships that underpin the framework. These pathways generate explicit empirical expectations regarding how synchronization, infrastructure coupling, recursive leverage, intervention effectiveness, restoration capability, and operational flexibility influence authority-transition pressure. Because these relationships generate measurable expectations, they also provide the foundation for empirical validation and falsification [24,38].

Table 9: ERADM Governance–Critical Pathway Matrix

Taken together, Tables 1-9 and Figure 1 - establish a progression from theory to causality, from causality to measurement, and from measurement to empirical evaluation. The architecture provides a structured methodology through which governance conditions become observable, reportable, auditable, reconstructable, and scientifically testable. The framework supports both retrospective and prospective analysis. Historical reconstruction can be performed using investigations, recalls, deployment records, safety reports, litigation, certification actions, and regulatory interventions. Prospective evaluation can be conducted through ERADM operational telemetry, governance reporting systems, audit mechanisms, safety-management processes, and supervisory monitoring architectures. Consequently, the internal observability layer serves as the methodological bridge between governance theory and empirical science. It establishes measurable variables, identifies observable evidence sources, defines causal relationships, and generates falsifiable propositions capable of empirical evaluation.

5.2 External Institutional Observability

Governance outcomes are not determined solely by conditions internal to autonomous systems. Regulatory institutions, insurance organizations, reinsurers, certification bodies, and systemic-risk organizations also influence how governance conditions are detected, interpreted, and acted

upon. Traditional safety and systems-theoretic approaches generally focus on internal operational dynamics [24,39]. In practice, however, autonomous-system governance exists within broader institutional environments that shape risk perception, intervention thresholds, underwriting practices, liability assessment, certification decisions, and regulatory responses. ERADM is not intended to replace accident-analysis methodologies such as STAMP/CAST. Rather, it extends governance analysis by introducing a standardized observability architecture through which governance-relevant operational conditions may be measured before accidents, incidents, or other adverse outcomes occur.

Organizations including the National Association of Insurance Commissioners (NAIC), the Geneva Association, Lloyd's, Swiss Re, Munich Re, transportation regulators, certification authorities, and governmental oversight bodies function as external observability amplifiers. These institutions aggregate information across deployments, identify systemic risk patterns, evaluate correlated exposures, and influence governance sensitivity through regulatory and financial mechanisms [20,41]. The MRROW construct operationalizes this external layer by distinguishing five institutional-observability dimensions: regulatory attention, insurance governance intensity, underwriting sensitivity, systemic-risk awareness, and institutional observability capacity. These dimensions

capture how strongly external governance institutions detect, interpret, assess, and respond to emerging operational conditions. Unlike ERADM's internal variables, which measure system-level governance conditions, MRROW characterizes the institutional environment within which those conditions are evaluated. This distinction is important because identical operational conditions may produce substantially different governance outcomes depending upon regulatory scrutiny, insurance exposure, reinsurance concern, institutional risk perception, and the availability of observable evidence [20,41]. Consequently, governance observability is determined not only by the visibility of operational conditions within a system, but also by the capacity of external institutions to detect, aggregate, interpret, and act upon those conditions. To formally represent this institutional environment, ERADM introduces the Market-Regulatory-Reinsurance Observability Weight (MRROW).

$$MRROW_t = w_R R_t + w_I I_t + w_U U_t + w_S S_{risk,t} + w_O O_t$$

where:

$$w_R + w_I + w_U + w_S + w_O = 1$$

and:

- R_t = regulatory attention
- I_t = insurance governance intensity
- U_t = underwriting sensitivity
- $S_{risk,t}$ = systemic-risk awareness
- O_t = institutional observability capacity

This expression does not predict a specific incident or governance outcome. Rather, it provides a transparent representation of the degree to which external institutions amplify or attenuate the visibility, interpretation, and governance significance of observable operational conditions.

Illustrative Example: The significance of MRROW can be illustrated by comparing the governance environments surrounding Uber ATG prior to the 2018 Tempe fatality and Cruise following the 2023 San Francisco incident and subsequent regulatory actions. Although both cases generated governance-relevant evidence concerning supervision, intervention, and operational oversight, the institutional environments differed substantially. By 2023, autonomous vehicle deployments operated under greater regulatory scrutiny, stronger insurance and reinsurance attention, increased public visibility, and higher systemic-risk awareness. Consequently, similar governance-relevant observations may generate markedly different regulatory, insurance, and oversight responses. MRROW captures this institutional effect by representing the extent to which external organizations amplify or attenuate the governance significance of observable operational conditions. MRROW therefore represents the combined influence of regulatory, insurance, reinsurance, and institutional observability factors on governance sensitivity. Rather than measuring operational system behavior directly, it measures the extent to which external institutions influence how observable

governance conditions are interpreted, prioritized, and acted upon. Higher MRROW values indicate increased institutional scrutiny, greater risk sensitivity, stronger underwriting and regulatory responsiveness, and lower tolerance for deteriorating governance conditions. Lower values indicate weaker coupling between operational conditions and external governance mechanisms. The introduction of MRROW extends ERADM beyond traditional operational observability by recognizing that governance effectiveness is influenced not only by technical conditions but also by institutional capacity to detect, interpret, and respond to those conditions.

5.3 Integrated Governance Observability

The complete ERADM architecture therefore consists of two complementary layers.

The first layer measures operational conditions internal to autonomous systems. This layer captures intervention capability, override effectiveness, restoration feasibility, adaptation tempo, operational flexibility, synchronization intensity, infrastructure dependency, and authority-transition pressure. The second layer measures external institutional sensitivity to those conditions. This layer captures the influence of regulatory oversight, insurance governance, reinsurance exposure, systemic-risk assessment, certification scrutiny, and broader governance responses. Together, these layers establish a comprehensive observability architecture that links operational system behavior to institutional governance capacity. The result is a framework through which governance can be evaluated not solely through outcomes, responsibilities, or organizational structures, but through direct observation of the operational conditions that determine whether effective governance remains achievable. Existing governance frameworks frequently assume that effective supervision, intervention, and corrective action remain available when needed. Historical investigations, however, repeatedly demonstrate that these capabilities may degrade gradually before becoming visible through incidents, regulatory actions, liability events, or operational failures. The central objective of ERADM is therefore not to replace existing governance frameworks but to provide an observability layer through which these conditions can be measured before their consequences become externally observable [10-39]. This dual-layer architecture provides the foundation for empirical evaluation of real-world autonomous-system deployments. The next section applies the framework to historical reconstruction analyses involving Uber ATG, Tesla ADS/FSD, Cruise, Boeing 737 MAX, and Waymo to evaluate whether governance-relevant operational conditions can be reconstructed from publicly available evidence.

6. Historical Reconstruction and Empirical Evaluation

6.1 Reconstruction Methodology

A central question underlying ERADM is whether governance-relevant operational conditions can be observed outside theoretical models. A measurement architecture has limited practical value if its variables cannot be identified, estimated,

reconstructed, or evaluated using evidence generated by real-world autonomous-system deployments. Consequently, the empirical objective of this study is not to validate specific governance outcomes but to determine whether governance-relevant operational conditions are sufficiently observable to support systematic investigation. Historical reconstruction was selected as the primary evaluation methodology because many governance-relevant variables are not currently reported through standardized regulatory

frameworks. Publicly available investigations, recalls, deployment records, safety reports, litigation materials, operational restrictions, software updates, certification actions, and regulatory interventions therefore provide an alternative evidentiary basis for examining whether ERADM variables can be reconstructed from observable evidence [9-14].

TABLE X.
ERADM EMPIRICAL GOVERNANCE OBSERVABILITY, AUTHORITY BREACH, AND STATUTORY TRIGGER MATRIX

Historical cases shown where STAMP identifies hazards and ERADM makes governance conditions measurable.

Note: Boeing 737 MAX and Therac-25 are included as cross-domain control-system analogues rather than AI systems. Their inclusion illustrates that governance observability failures, authority degradation, rollback constraints, and supervisory limitations are not unique to artificial intelligence but may occur across automated and software-controlled systems.						
Historical Case	STAMP Hazard / UCA	ERADM Variables	Governance Condition Revealed	Authority Breach Evidence	Statutory / Regulatory Authority Triggered	Potential Governance Consequence
Uber ATG (2018)	Failure to provide effective intervention before collision	AEI, OER, H	Formal supervision existed, but intervention capacity degraded operationally	Safety driver failed to intervene within available response window	49 U.S.C. §§30118–30120; 49 U.S.C. Chapter 11; Arizona Executive Order 2015-09	Investigation, recall review, operational restrictions, safety recommendations
Tesla ADS/FSD (2022–2024)	Unsafe control actions during driver-supervised automation	OER, CLSI, H	Supervisory compression emerged as adaptation tempo increased	OTA cadence exceeded traditional oversight and review tempo	49 U.S.C. §§30118–30120; NHTSA SGO 2021-01; 49 CFR Part 579	Defect investigation, recall authority, reporting obligations, operational scrutiny
Cruise (2023)	Inadequate recovery and operational response following incident	RFT, OSCR, AEI	Restoration capability and operational flexibility became governance-critical	Service suspension demonstrated rollback and recovery constraints	California Vehicle Code §38750; CCR Title 13 AV Regulations; CPUC driverless deployment authority; 49 U.S.C. Chapter 301	Permit suspension, deployment restrictions, corrective-action requirements
Waymo (2024)	Potential fleet-wide propagation of control actions	S, K, OSCR	Infrastructure dependency and synchronization exposure became observable	Operational dependence mediated through integrated infrastructure	California Vehicle Code §38750; CA DMV driverless deployment authority; NHTSA ADS reporting framework	Deployment conditions, reporting requirements, operational limitations
Boeing 737 MAX	UCA-1, UCA-3, UCA-4 involving MCAS control actions	OER, H, CLSI, RFT	Human response tempo failed to match system tempo	Pilots possessed nominal authority but insufficient intervention effectiveness	49 U.S.C. Subtitle VII; 14 CFR Part 25; FAA ODA frameworks; FAA SMS requirements	Grounding, recertification, certification reform, operational restrictions
Therac-25	UCA-2: unsafe software control action provided	RFT, OSCR, Governance Observability Latency	No practical restoration pathway existed once unsafe state emerged	Rollback effectively zero; hardware interlocks absent	Federal Food, Drug, and Cosmetic Act, 21 U.S.C. §360c et seq.; 21 CFR Part 820; FDA MDR requirements	Device redesign, corrective action, regulatory intervention
Fleet-Wide OTA Systems	Synchronization-related unsafe control actions	CLSI, S, K	Single updates can propagate system-wide operational effects	Supervisory structures may lag synchronized deployment and validation cycles	OMB M-24-10; NIST AI RMF 1.0; ISO/IEC 42001:2023; ISO 26262:2018; EU AI Act obligations	Risk-management intervention, compliance review, deployment controls
Evidence and Regulatory Sources: NTSB (2019); NHTSA ADS investigations and Standing General Order data (2022–2024); California DMV (2023); California Public Utilities Commission (2023); Waymo safety reports (2024); FAA 14 CFR Part 25; FDA 21 CFR Part 820; 49 U.S.C. Chapter 301; NIST AI RMF 1.0 (2023); OMB M-24-10 (2024); ISO/IEC 42001:2023; ISO 26262:2018.						
Statutory authorities are provided to show governance relevance and potential enforcement pathways; inclusion does not imply a legal finding or adjudicated violation.						

Note: Table X demonstrates that governance-relevant conditions can be reconstructed from publicly observable evidence across multiple historical autonomous-system deployments and software-mediated control-system failures. The analysis illustrates how ERADM variables—including intervention effectiveness, supervisory responsiveness, rollback feasibility, operational flexibility, synchronization intensity, and infrastructure dependency—become observable through investigations, recalls, operational restrictions, deployment controls, certification actions, and regulatory interventions. In each case, the observable evidence extends beyond the outcome itself and provides insight into the operational conditions that influenced whether meaningful authority, intervention capability, and governance effectiveness remained viable before the incident.

Table 10: ERADM Empirical Governance Observability- Authority Breach and Statutory Trigger

The methodology focuses on four principal autonomous-system case studies:

1. Uber Advanced Technologies Group (ATG)
2. Tesla ADS/FSD
3. Cruise
4. Waymo

These cases were selected because they provide extensive publicly available documentation concerning deployment practices, supervision mechanisms, intervention capability, regulatory responses, remediation activities, and operational governance. For each case, evidence was evaluated to determine whether observable indicators associated with ERADM variables could be identified. Variables were

classified using four categories:

- Observable
- Partially Observable
- Insufficient Evidence
- Not Observable

The objective was not to generate precise numerical estimates but to determine whether governance-relevant operational conditions could be reconstructed using currently available evidence.

6.2 Uber ATG

The 2018 Uber ATG fatality provides one of the most extensively documented autonomous-vehicle investigations

available in the public domain. NTSB findings, operational records, safety-driver procedures, supervision practices, and investigative reports provide substantial information concerning intervention capability, supervisory responsiveness, and operational control conditions [9]. The reconstruction suggests that several governance-relevant conditions were observable prior to the fatal event. Evidence indicates limitations in intervention effectiveness, supervisory responsiveness, and operational monitoring. The case therefore provides important evidence supporting

reconstruction of Authority Elasticity Index (AEI), Override Effectiveness Rate (OER), and Human Supervisory Response Horizon (H). Although direct measurement was unavailable, publicly available evidence demonstrates that governance-relevant operational conditions can be reconstructed from investigative findings. The significance of the case lies not in outcome analysis alone but in demonstrating that conditions influencing governance effectiveness were partially observable before formal intervention occurred.




Figure 2: Uber ATG Governance Observability Reconstruction.

Figure 2 presents an empirical observability assessment of key ERADM variables reconstructed from the Uber ATG fatality investigation. Using publicly available evidence, the figure evaluates the extent to which Authority Elasticity (AEI), Override Effectiveness (OER), and Human Supervisory Response Horizon (H) can be observed and measured. The assessment demonstrates that some governance-relevant conditions are observable from existing investigations, while others remain only partially observable due to limitations in publicly available operational data.

6.3 Tesla ADS/FSD

Tesla provides a distinct reconstruction environment because of extensive public documentation concerning software updates, recalls, operational investigations,

driver supervision requirements, and regulatory review activities [10]. Unlike Uber, Tesla evidence is particularly relevant to adaptation tempo and intervention observability. Software deployment cycles, recall campaigns, operational restrictions, and NHTSA investigations provide observable indicators associated with Closed-Loop Self-Improvement Interval (CLSI), Authority Elasticity Index (AEI), and Override Effectiveness Rate (OER). The reconstruction demonstrates that governance-relevant operational conditions associated with adaptation tempo and intervention capability can frequently be inferred from publicly available deployment and regulatory evidence. The case further illustrates how software-centric autonomous systems create governance conditions that may evolve more rapidly than traditional reporting architectures were designed to observe.

FIGURE 3

TESLA ADS/FSD ERADM VARIABLE OBSERVABILITY ASSESSMENT

Public Evidence Reconstruction of Key Variables from Tesla ADS/FSD Deployments, Investigations, Recalls, and Regulatory Actions

VARIABLE	EVIDENCE AVAILABILITY (Public Sources)	OBSERVABILITY STATUS (Current Evidence)	WHAT DATA WOULD MAKE VARIABLE COMPLETELY OBSERVABLE
AEI Authority Elasticity Index 	 Driver monitoring requirements  Intervention expectations  Supervision guidance and policies  NHTSA investigations and ODI findings  Operational-use disclosures and manuals	 PARTIALLY OBSERVABLE <i>Evidence indicates changing practical human authority under ADS/FSD operation, but complete authority-allocation and intervention-state records are not publicly available.</i>	- Complete intervention logs - Driver-alert histories and responses - Authority-allocation records - Human-machine control-transfer logs - Supervisory decision records - Driver attention and monitoring telemetry - Full supervisory-state histories
OER Override Effectiveness Rate 	 Driver takeover events  Crash investigations and reports  Disengagement and intervention evidence  NHTSA findings and safety reports  Public safety and recall reports	 PARTIALLY OBSERVABLE <i>Evidence of override attempts and outcomes exists, but complete intervention effectiveness rates cannot be reconstructed from public data.</i>	- Complete override-request logs - Accepted/rejected override records - Timestamped intervention outcomes - Vehicle-state snapshots during interventions - Time-to-intervention measurements - Counterfactual override analysis - Full vehicle-state data
CLSI Closed-Loop Self-Improvement Interval 	 OTA update history  FSD version releases  Software deployment cadence  Feature rollout timelines  Public release documentation	 OBSERVABLE <i>Public update histories and software deployment records provide sufficient evidence to estimate adaptation and update-cycle frequency at a high level.</i>	- Internal model-training timelines - Validation-cycle duration - Deployment-approval records - Update testing and validation data - Update execution telemetry - Complete adaptation-cycle measurements
OBSERVABILITY KEY:  Observable Sufficient public evidence to estimate variable.  Partially Observable Some evidence available, but incomplete for full estimation.  Not Observable No meaningful public evidence available.			
SOURCES: NHTSA and ODI Filings; Tesla Safety Reports; Recall Reports; Press Releases; Investor Reports; Media Reports; Court Filings			

Figure 3: Tesla ADS/FSD Governance Observability Reconstruction

Figure 3 presents an empirical observability assessment of key ERADM variables reconstructed from Tesla ADS/FSD deployments, investigations, recalls, and regulatory actions. The analysis evaluates the extent to which Authority Elasticity (AEI), Override Effectiveness Rate (OER), and Closed-Loop Self-Improvement Interval (CLSI) can be observed using publicly available evidence. The findings indicate that adaptation tempo (CLSI) is largely observable through software-release and deployment records, while intervention authority (AEI) and override effectiveness (OER) remain only partially observable due to limited access to internal operational and supervisory data.

6.6 Cruise

The Cruise case provides particularly strong evidence concerning intervention effectiveness, restoration capability,

and governance response. Public investigations, California regulatory actions, operational restrictions, recall activities, and deployment modifications provide substantial evidence regarding rollback feasibility and operational remediation [11], [12]. Unlike many autonomous-system cases, Cruise generated observable restoration activities that permit evaluation of Rollback Feasibility Time (RFT). Regulatory interventions also provide evidence regarding operational flexibility, governance responsiveness, and the interaction between technical conditions and institutional oversight. The reconstruction demonstrates that governance-relevant conditions associated with intervention effectiveness, restoration capability, and remediation feasibility can often be identified through public evidence even when they are not directly reported through standardized measurement architectures.

FIGURE 4. CRUISE ERADM VARIABLE OBSERVABILITY ASSESSMENT*Public Evidence Reconstruction of Key Variables from Cruise Deployments and 2023 Operational Restrictions*

VARIABLE	EVIDENCE AVAILABILITY (Public Sources)	OBSERVABILITY STATUS (Current Evidence)	WHAT DATA WOULD MAKE VARIABLE COMPLETELY OBSERVABLE
AEI Authority Elasticity Index 	 Intervention Authority Human safety operators could intervene, but authority was constrained by system behavior and remote-operation policies.  Operator Behavior Reports of over-reliance on system; operators monitoring multiple vehicles with limited ability to affect behavior.  System Supervision Records Description of supervision tools, alerting mechanisms, and human-machine interface limitations in public reporting and testimony.	 PARTIALLY OBSERVABLE <i>Evidence indicates reduced practical authority due to system limitations and operational constraints, but complete authority logs and rationale not publicly available.</i>	- Complete operator authorization logs and permissions states - Detailed policy enforcement logs (when/why interventions were or were not permitted) - Real-time system decision rationale and authority gating logic - Full audit trail of supervision interface states and available actions - Continuous operator gaze/attention and interaction data
OER Override Effectiveness Rate 	 Override Opportunity Evidence Multiple events documented where human intervention was required or attempted.  Intervention Outcome Evidence Public reports show interventions were not always effective in preventing undesired outcomes.	 PARTIALLY OBSERVABLE <i>Evidence of override attempts and outcomes exists, but insufficient data for rate estimation across all operational conditions.</i>	- Complete log of all override requests and system responses - Time-stamped analysis of each intervention attempt - System acceptance/rejection logs for override inputs - All fallback options presented and system behavior following overrides - Counterfactual analysis: would override have prevented the event under system state at that time?
RFT Rollback Feasibility Time 	 Update and Change Records Software updates, and deployments documented in public filings and reports.  Recovery Effort Evidence Post-incident actions to restore safety and regain operational control described in public communications.  Operational Restriction Timeline Time between incident, mitigation actions, and restoration of full operations.	 OBSERVABLE <i>Sufficient public evidence to estimate rollback feasibility timeline at a high level.</i>	- Detailed timestamps for when rollback was first initiated and completed - Complete system state backups and restoration logs - Engineering timelines for code analysis, patching, and validation - Time to redeploy, re-certify, and restore service - Constraints that delayed rollback (technical, regulatory, operational)
OBSERVABILITY KEY:  Observable Sufficient public evidence to estimate variable.  Partially Observable Some evidence available, but incomplete for full estimation.  Not Observable No meaningful public evidence available.			
SOURCES: NHTSA and DOJ Filings; Cruise Press Releases; CPUC Reports; SF Municipal Transportation Agency Reports; Media Reports; Court Filings; Safety Reports			

Figure 4: Cruise Governance Observability Reconstruction

Figure 4 presents an empirical observability assessment of key ERADM variables reconstructed from Cruise deployments, incident investigations, regulatory actions, and operational restrictions. Using publicly available evidence, the analysis evaluates the observability of Authority Elasticity (AEI), Override Effectiveness Rate (OER), and Rollback Feasibility Time (RFT). The findings indicate that rollback capability is largely observable through documented recovery actions and service-restoration timelines, while intervention authority and override effectiveness remain only partially observable due to limited access to internal supervisory, intervention, and operational decision records.

6.5 Waymo

Waymo serves as a comparative baseline because of its

relatively mature deployment environment and extensive safety documentation. Public safety reports, deployment records, operational disclosures, and regulatory interactions provide evidence concerning governance practices, intervention mechanisms, supervision structures, and deployment management [14]. The Waymo reconstruction does not imply absence of governance risk. Rather, it provides a useful comparison through which governance observability can be examined under conditions characterized by more extensive reporting and operational transparency. The case demonstrates that many governance-relevant variables can be observed through existing operational documentation and therefore provides a useful benchmark against which other deployments may be compared.

FIGURE 5 WAYMO ERADM VARIABLE OBSERVABILITY ASSESSMENT

Public Evidence Reconstruction of Key Variables from Waymo Deployments, Safety Reporting, and Regulatory Filings

VARIABLE	EVIDENCE AVAILABILITY (Public Sources)	OBSERVABILITY STATUS (Current Evidence)	WHAT DATA WOULD MAKE VARIABLE COMPLETELY OBSERVABLE
AEI Authority Elasticity Index 	- Rider safety information and FAQs - Human oversight and intervention descriptions - Safety reports and public blog posts - Regulatory submissions and proceedings (CPUC, NHTSA) - Operational design domain (ODD) disclosures	PARTIALLY OBSERVABLE <i>Evidence indicates the presence of human oversight and limited rider intervention, but complete authority-allocation and intervention-state records are not publicly available.</i>	- Complete intervention and override logs - Authority-allocation and control-transfer records - Driver/rider alert histories and responses - Human-machine interface state telemetry - Supervisory decision rationale and escalation records - Full supervisory-state histories across trips
OER Override Effectiveness Rate 	- Rider intervention and pull-over policies - Disengagement and incident reports - Safety performance reports - Regulatory filings and testimony - Public disclosures of safety events	PARTIALLY OBSERVABLE <i>Evidence of rider interventions and system responses exists, but complete override effectiveness rates cannot be reconstructed from public data.</i>	- Complete override-request database - Accepted/rejected override records - Timestamped intervention outcomes - Vehicle-state data at time of intervention - Time-to-intervention measurements - Counterfactual analysis of intervention effectiveness - Full fallback behavior telemetry
CLSI Closed-Loop Self-Improvement Interval 	- Software update and release notes - Waymo One service updates and deployments - Safety reports and reliability updates - Public blog posts on AI and research advancements - Regulatory submissions describing software changes	OBSERVABLE <i>Public update histories, safety reports, and release notes provide sufficient evidence to estimate improvement cycles and update frequency at a high level.</i>	- Internal model-training timelines - Validation and testing cycle duration - Model deployment approval records - Detailed experiment and evaluation logs - Update-development start and end dates - Full adaptation-cycle telemetry
OBSERVABILITY KEY:  Observable Sufficient public evidence to estimate variable.  Partially Observable Some evidence available, but incomplete for full estimation.  Not Observable No meaningful public evidence available.			
SOURCES: Waymo Safety Reports; Waymo Blog; NHTSA Filings; CPUC Proceedings; Press Releases; Investor Reports; Media Reports; Court Filings			

Figure 5: WAYMO ERADM Variable Observability Assessment

Figure 5 Waymo ERADM Variable Observability Assessment evaluates the extent to which key ERADM variables can be reconstructed from publicly available evidence associated with Waymo deployments, safety reporting, and regulatory disclosures. The analysis indicates that CLSI (Closed-Loop Self-Improvement Interval) is largely observable through documented software releases, deployment updates, safety reports, and regulatory filings. In contrast, AEI (Authority Elasticity Index) and OER (Override Effectiveness Rate) remain only partially observable because complete intervention logs, authority-allocation records, override-response histories, and internal supervisory-state data are not publicly available. The figure demonstrates that current reporting frameworks provide meaningful visibility into adaptation and deployment processes while offering comparatively limited visibility into operational authority allocation and intervention effectiveness. This pattern reinforces ERADM's central finding that governance-

relevant conditions can often be partially reconstructed from public evidence, but complete observability requires more comprehensive reporting architectures.

6.6 Cross-Case Analysis

Cross-case comparison reveals several consistent patterns. First, governance-relevant operational conditions are frequently observable or partially observable despite not being systematically measured within existing reporting architectures. Second, intervention effectiveness, supervisory responsiveness, restoration capability, adaptation tempo, and operational flexibility repeatedly emerge as governance-relevant conditions across multiple independent deployments. Third, existing investigations, recalls, deployment restrictions, operational modifications, and regulatory actions frequently contain evidence sufficient to reconstruct governance conditions retrospectively even when those conditions were not measured prospectively.

FIGURE 6 ERADM CROSS-CASE MATRIX: Empirical Reconstruction of Authority Dynamics Across Leading AV/ADS Systems
 Integrated Evidence from Canonical ERADM Variables and Public Case Records (NTSB, DMV, NHTSA, Company Disclosures, Press, Court Filings)

ERADM VARIABLE (Definition)	OPERATIONAL INTERPRETATION (What it Captures)	UBER ATG (2016–2018)	CRUISE (2018–2023)	TESLA ADS/FSD (2016–Present)	WAYMO (2016–Present)	CROSS-CASE SYNTHESIS (Key Insight)
AEI Authority Elasticity Index (0–1)	Degree to which effective operational authority shifts from humans to the system over time. Higher = more authority migrated to system.	0.78 High Authority Migration Drivers routinely in control but disengaged (NTSB HAR-19/03); system exercised de facto control in complex conditions. ●●●●○ Med-High Evidence	0.85 Very High Authority Migration System operated driverlessly in many ODDs; human authority further reduced; limited ability to intervene in edge cases. ●●●●● High Evidence	0.72 High Authority Migration FSD shifts driving authority to system; human remains fallback wide ODD but requires attention and takeover readiness. ●●●●○ Med-High Evidence	0.42 Moderate Authority Migration System is primary driver in defined ODD, but strong remote support and intervention architecture preserves human authority. ●●●●○ High Evidence	 Authority migration varies by design and ODD scope; Waymo shows highest preservation of human authority.
OER Override Effectiveness Rate (0–1)	Effectiveness of human intervention or override in preventing or mitigating adverse outcomes. Higher = more effective human overrides.	0.18 Low Override Effectiveness Driver did not intervene in time in fatal crash; automation cues insufficient; override failed. ●●●●○ High Evidence	0.21 Low Override Effectiveness Limited opportunities to override in driverless mode; interventions often too late or ineffective. ●●●●○ High Evidence	0.34 Moderate Override Effectiveness Overreliance and delays in takeovers observed; some interventions effective, but inconsistent across scenarios. ●●●●○ High Evidence	0.61 Higher Override Effectiveness Robust remote assistance and on-road intervention; structured escalation improves outcome prevention. ●●●●○ High Evidence	 Override effectiveness correlates with supervisory architecture and intervention opportunity design.
RFT Rollback Feasibility Time (hours)	Time required to safely rollback, disable, or restore to a lower-risk operational state after an adverse event. Lower = faster rollback.	~168 Slow Rollback Testing paused for ~1 week following fatal crash; multiple approvals required to resume. ●●●●○ High Evidence	~720+ Very Slow Rollback Operations suspended for months by California DMV; extensive remediation required. ●●●●○ High Evidence	~48–120 Moderate Rollback OTA fixes deployed in hours–days, but recall remediation and NHTSA processes extend total resolution time. ●●●●○ Med-High Evidence	~24–48 Fast Rollback Rapid software updates and remote mitigations; clear safety case and reporting accelerate recovery. ●●●●○ High Evidence	 Rollback speed reflects engineering maturity, modularity, and regulatory readiness.
CLSI Closed-Loop Self-Improvement Interval (days)	Time between consecutive full system learning cycles from data capture → model update → deployment → operational impact. Lower = faster learning loop.	N/A Not Applicable Program paused before recursion measured to operational deployment. ○●●●○ Low Evidence	~14–30 Moderate Learning Loop Frequent iteration, but halted before stable closed-loop maturity. ●●●●○ Med Evidence	~3–7 Fast Learning Loop High-frequency data collection, training, and OTA deployment of FSD updates. ●●●●○ High Evidence	~7–14 Moderate-Fast Learning Loop Regular model updates and deployments with strong validation and safety gates. ●●●●○ High Evidence	 Tesla has fastest observed learning loop; Waymo balances speed with validation and safety gates.
CARI Composite Authority Risk Index (0–1)	Synthesized risk index combining delegated pressure and human resilience: CARI = f(AEI ↑, OER ↓, RFT ↑, CLSI ↑)	0.74 High Risk High delegation with low override effectiveness and slow rollback. ●●●●○ High Evidence	0.86 Very High Risk Extreme delegation with very slow rollback and low override effectiveness. ●●●●○ High Evidence	0.58 Moderate-High Risk High delegation and fast learning, but moderate override and rollback. ●●●●○ High Evidence	0.32 Lower Risk (Comparative Baseline) Balanced delegation, stronger override, faster rollback, and validated learning cycles. ●●●●○ High Evidence	 Waymo demonstrates lowest composite authority risk among systems with mature public evidence.
LEGEND (Per Cell)		COMPOSITE AUTHORITY RISK INDEX (CARI) INTERPRETATION				KEY INSIGHT
●●●●○ Very Favorable (Lower Risk)		0.00–0.25 Very Low Risk				Authority risk emerges from the interaction of how much authority is delegated (AEI), how effective humans remain (OER), how quickly systems can be de-risked (RFT), and the speed of learning cycles (CLSI).
●●●●○ Moderate		0.26–0.50 Low Risk				
●●●●○ Unfavorable (Higher Risk)		0.51–0.75 Moderate Risk				
○●●●○ Not Applicable / Insufficient Data		0.76–1.00 High Risk				
SOURCES: Waymo Safety Reports; Waymo Blog; NHTSA Filings; CPUC Proceedings; Press Releases; Investor Reports; Media Reports; Court Filings						

SOURCES: Waymo Safety Reports; Waymo Blog; NHTSA Filings; CPUC Proceedings; Press Releases; Investor Reports; Media Reports; Court Filings

Figure 6: ERADM Cross-Case Matrix

Figure 6 demonstrates cross-case empirical consistency assessment of ERADM. Fourteen of fifteen pre-specified ERADM directional expectations were found to be consistent with publicly available evidence across Uber ATG, Tesla ADS/FSD, and Cruise governance events. No expectations were directly contradicted. One expectation remained indeterminate due to insufficient publicly available evidence. The assessment evaluates empirical consistency and observability rather than predictive validity or causal proof.

6.7 Empirical Findings

The historical reconstruction analyses support four principal findings. First, governance-relevant operational

conditions can frequently be reconstructed from publicly available evidence. Second, intervention capability, override effectiveness, supervisory responsiveness, restoration feasibility, adaptation tempo, operational flexibility, synchronization effects, and dependency formation appear sufficiently observable to support empirical investigation. Third, governance-relevant conditions often appear identifiable before major recalls, operational restrictions, enforcement actions, or deployment changes occur. Fourth, existing reporting architectures frequently contain fragments of the information necessary for governance observability without providing a unified framework for organizing and evaluating that information.




Figure 7: Cross-Case Empirical Consistency Assessment of ERADM

Figure 7 summarizes the results of evaluating fifteen pre-specified ERADM expectations against publicly available evidence from three independent autonomous-system governance cases: Uber ATG, Tesla ADS/FSD, and Cruise. Fourteen of the fifteen expectations (93.3%) were directionally consistent with observed evidence, while none were directly contradicted. One expectation remained indeterminate due to insufficient publicly available evidence. These findings do not constitute formal validation of ERADM; rather, they demonstrate that the framework's governance-relevant variables and directional expectations are empirically observable and assessable using real-world operational evidence. The absence of direct falsifications establishes an initial empirical baseline supporting the framework's feasibility and provides a foundation for future longitudinal measurement, statistical testing, calibration, and formal falsification efforts.

7. Governance, Regulatory, Insurance, And Liability Implications

7.1 From Outcome Observation to Governance Observability

A principal finding of this study is that contemporary governance architectures primarily observe outcomes rather than the operational conditions that shape those outcomes.

Existing regulatory, certification, investigative, insurance, and liability systems provide substantial visibility into crashes, incidents, recalls, operational restrictions, enforcement actions, litigation, and safety events [15–36]. These systems are highly effective at determining what occurred. They are generally less effective at determining whether intervention capability, supervisory responsiveness, restoration feasibility, operational flexibility, synchronization intensity, and dependency formation remained operationally viable and what their effect was before those outcomes emerged. The distinction is significant because governance institutions frequently make decisions before complete outcome information becomes available. Regulators must evaluate deployments before failures occur. Insurers must assess risk before claims arise. Certification authorities must evaluate safety before incidents emerge. Operators must manage governance conditions continuously rather than only after adverse outcomes become visible.

ERADM addresses this limitation by introducing a measurement architecture that focuses on governance relevant operational conditions rather than solely on governance outcomes. The framework therefore complements existing oversight systems by expanding the range of observable conditions available for evaluation.



Figure 8: Outcome Observation Versus Governance Observability

Figure 8 illustrates the central contribution of ERADM by distinguishing between the observation of governance outcomes and the observation of governance-relevant operational conditions. Existing regulatory, certification, investigative, insurance, and liability frameworks provide strong visibility into outcomes such as crashes, incidents, recalls, claims, litigation, and enforcement actions. However, they often provide only partial or limited visibility into the operational conditions that shape those outcomes, including intervention capability, override effectiveness, rollback feasibility, supervisory responsiveness, synchronization intensity, infrastructure dependency, and authority-transition pressures. ERADM extends observability beyond outcomes by introducing measurable variables that allow these governance conditions to be monitored, reconstructed, and evaluated before major failures occur. The comparison demonstrates that while existing frameworks are highly effective at determining what happened, ERADM is designed to improve visibility into the conditions that influence why those outcomes emerge and whether effective governance remains operationally viable.

7.2 Regulatory Applications

The implications for regulatory oversight are substantial. Contemporary regulatory agencies increasingly oversee systems characterized by software-driven adaptation, distributed coordination, fleet learning, cloud integration, and continuous operational modification [15-34]. Under these conditions, governance effectiveness depends not only upon compliance with established requirements but also upon maintaining visibility into conditions that

influence intervention capability and supervisory viability. Variables such as Authority Elasticity Index (AEI), Override Effectiveness Rate (OER), Rollback Feasibility Time (RFT), Human Supervisory Response Horizon (H), Synchronization Intensity (S), and Infrastructure Coupling (K) provide regulators with measurable indicators that are not directly captured by most existing reporting architectures. These variables are particularly relevant because many regulatory interventions occur only after governance conditions have already deteriorated sufficiently to produce observable outcomes. By measuring operational conditions directly, regulators may gain earlier visibility into emerging governance challenges and therefore possess greater opportunity for corrective intervention before recalls, operational restrictions, or enforcement actions become necessary. ERADM is not proposed as a replacement for existing regulatory frameworks. Rather, it functions as an observability layer that may be integrated into existing oversight, certification, reporting, and safety-management architectures.

7.3 Certification and Safety Assurance

Certification systems traditionally evaluate whether technologies satisfy specified safety, reliability, and operational requirements. However, certification activities frequently focus on demonstrating compliance at a particular point in time rather than measuring whether governance-relevant conditions remain effective throughout operational deployment. Autonomous systems introduce a challenge because operational behavior may evolve through software updates, adaptive processes, fleet learning, and environmental

interaction. Under these conditions, certification assurance increasingly depends upon the ability to observe operational conditions continuously rather than solely during initial certification activities. ERADM contributes to certification assurance by introducing measurable indicators associated with intervention effectiveness, restoration capability, supervisory responsiveness, adaptation tempo, and operational flexibility. These measurements provide a mechanism through which governance effectiveness may be evaluated throughout the operational lifecycle rather than exclusively during initial certification. The framework therefore extends existing certification approaches by emphasizing continuous governance observability rather than static compliance assessment alone.

7.4 Accident Investigation and Failure Analysis

Accident investigations seek to determine not only what occurred but also why it occurred and whether the event could have been prevented. Traditional investigations frequently reconstruct operational events using crash data, system logs, witness accounts, operational records, maintenance histories, and regulatory documentation. These approaches provide valuable information regarding outcome causation but may provide comparatively limited visibility into governance conditions that preceded the event.

ERADM contributes an additional analytical layer by focusing on intervention capability, supervisory responsiveness, restoration feasibility, operational flexibility, synchronization intensity, and infrastructure dependency. The historical

reconstruction analyses presented in this study demonstrate that many of these variables can be partially reconstructed using publicly available evidence [4,9,31,33,36,43,45,48]. Consequently, investigators may be able to evaluate not only the causes of failure but also whether governance conditions were strengthening, deteriorating, or approaching operational limits before the event occurred. This distinction expands accident investigation from outcome reconstruction toward governance-condition reconstruction.

7.5 Insurance and Risk Assessment

The insurance implications of governance observability may be equally significant. Traditional underwriting practices rely primarily upon historical losses, claims experience, incident frequency, litigation outcomes, actuarial models, and observed operational performance. These approaches provide valuable information regarding the consequences of system failures but provide comparatively limited visibility into the operational conditions that generated those failures. Autonomous systems create an additional challenge because underwriting increasingly depends upon evaluating whether operators retain meaningful influence over the systems producing insured risk. Variables such as AEI, OER, RFT, H, and CLSI provide potential indicators of governance effectiveness that are not captured directly through conventional loss-based models. These variables enable evaluation of intervention capability, restoration feasibility, supervisory responsiveness, and adaptation tempo before losses occur.




Figure 9: Insurance and Liability Opacity Across Historical Autonomous Systems Cases

Figure 9 illustrates that across multiple autonomous-system governance events, operational outcomes such as crashes, investigations, recalls, and regulatory actions were generally observable, while authority allocation, liability pathways, and insurance allocation remained partially observable or opaque. The comparison demonstrates a systematic observability imbalance in which stakeholders can often determine what happened after a loss occurs but cannot reliably reconstruct who retained operational authority, how intervention capability evolved, or how insurance exposure and financial responsibility were distributed prior to the event. The figure therefore provides evidence of a governance observability gap extending beyond technical system performance into liability and insurance oversight. Does it demonstrate asymmetry within contemporary insurance systems? Yes. The figure provides evidence of an insurance observability asymmetry. Contemporary insurance systems typically possess strong visibility into realized losses, claims, settlements, and litigation outcomes, but comparatively limited visibility into the operational conditions that generated those losses, including intervention effectiveness, supervisory responsiveness, rollback capability, authority allocation, and dependency formation. As a result, insurers often evaluate consequences with greater precision than the governance conditions that produced those consequences. Figure 18 suggests that this asymmetry may limit underwriting accuracy, risk assessment, reserve estimation, and early identification of emerging governance failures in autonomous systems.

7.6 Liability and Accountability

Liability systems traditionally focus on assigning responsibility after adverse outcomes occur. Courts, regulators, insurers, manufacturers, operators, and investigators evaluate evidence to determine causation, accountability, negligence, compliance, and damages. However, many liability disputes involve questions concerning operational conditions rather than solely outcome causation. Parties frequently seek to determine whether intervention remained possible, whether override mechanisms were effective, whether supervision remained adequate, whether restoration options existed, and whether meaningful operational control remained available. ERADM contributes to liability analysis by providing measurable constructs through which these questions may be evaluated systematically. The framework does not determine liability. Rather, it provides additional evidence pathways through which governance-relevant operational conditions may be reconstructed and evaluated. Consequently, liability analysis may be informed not only by outcomes but also by measurable evidence concerning the conditions that shaped those outcomes. By making governance-relevant operational conditions more observable, the framework may improve evidentiary transparency for regulators, insurers, courts, operators, and manufacturers evaluating responsibility following adverse events.

7.7 Broader Governance Implications

The significance of governance observability extends beyond

autonomous vehicles. Many of the operational conditions represented within ERADM appear across adaptive and increasingly autonomous environments, including aviation, logistics, robotics, industrial automation, healthcare systems, financial systems, critical infrastructure, and defense applications. Across these domains, governance effectiveness increasingly depends upon maintaining visibility into intervention capability, supervisory responsiveness, restoration feasibility, operational flexibility, synchronization effects, and infrastructure dependency. The broader contribution of ERADM is therefore not limited to a particular technology sector. The framework establishes a methodology through which governance conditions become measurable objects of scientific investigation.

As autonomous systems become increasingly adaptive, interconnected, and operationally significant, governance observability may become as important as safety observability, reliability observability, and performance observability. The central implication of this study is therefore straightforward: effective governance depends not only upon authority, responsibility, and accountability, but also upon the ability to directly observe the operational conditions that determine whether governance remains achievable.

8. Limitations and Future Research

8.1 Study Limitations

Several limitations should be considered when interpreting the findings presented in this study. First, the empirical analyses rely primarily upon publicly available evidence rather than proprietary operational telemetry. Investigations, recalls, deployment records, safety reports, litigation materials, operational restrictions, regulatory actions, and public disclosures provide substantial information concerning governance-relevant conditions. However, these sources frequently offer only partial visibility into underlying operational states. Consequently, many ERADM variables can currently be reconstructed only approximately rather than measured directly [4,9,31,33,36,43,45,48]. Second, the historical reconstruction methodology is retrospective. Reconstruction allows governance-relevant conditions to be identified after operational events have occurred, but retrospective observability does not necessarily imply prospective monitoring capability. Although the analyses demonstrate that governance-relevant conditions can often be reconstructed from available evidence, future research must evaluate whether these variables can be monitored systematically in real time. Third, the case studies examined in this article are concentrated within autonomous transportation environments. Uber ATG, Tesla ADS/FSD, Cruise, Boeing 737 MAX, and Waymo provide valuable governance evidence because they are among the most extensively documented autonomous-system deployments currently available. However, transportation systems represent only one category of increasingly autonomous operational environments. Additional empirical work is required to evaluate the framework across healthcare, robotics, logistics, finance, critical infrastructure, industrial

automation, and defense applications. Fourth, several ERADM variables remain partially observable under existing reporting architectures. Intervention effectiveness, supervisory responsiveness, restoration capability, synchronization intensity, infrastructure dependency, and operational flexibility can frequently be inferred from available evidence, but many measurements remain indirect. The framework therefore demonstrates observability feasibility rather than complete measurement maturity. Fifth, the mathematical relationships presented in this study are conceptual and directional rather than fully calibrated predictive models. The equations establish measurable relationships and generate falsifiable expectations, but substantial empirical calibration remains necessary before the framework can support quantitative forecasting or operational decision support. Finally, the framework should not be interpreted as a replacement for existing regulatory, safety, certification, liability, insurance, or governance systems. ERADM is intended as a complementary observability architecture. Its purpose is to provide direct visibility into governance-relevant operational conditions that are frequently inferred rather than measured within existing frameworks.

8.2 Measurement and Data Challenges

Implementation of governance observability faces several practical challenges. Many variables proposed by ERADM depend upon data that are currently reported in standardized formats. Override attempts, intervention outcomes, restoration activities, supervisory response timelines, synchronization effects, infrastructure dependencies, and adaptation cycles are often collected internally by fleet operators but are not routinely incorporated into public reporting systems. Consequently, improved governance observability may require modifications to reporting architectures, audit processes, safety-management systems, certification procedures, and regulatory documentation practices. The challenge is not necessarily the creation of new data, but rather the organization and reporting of information that frequently already exists within operational environments.

A related challenge concerns standardization. Different operators, regulators, investigators, insurers, and certification bodies may define intervention effectiveness, restoration capability, synchronization intensity, or operational flexibility differently. Future development of ERADM therefore requires consistent measurement standards capable of supporting cross-system comparison and longitudinal analysis.

8.3 Future Research Directions

Several avenues for future research emerge from this study. The first priority is prospective measurement. Historical reconstruction demonstrates that governance-relevant conditions can frequently be identified retrospectively. Future research should evaluate whether ERADM variables can be monitored continuously during active deployment and whether such monitoring improves governance decision-

making. A second priority involves empirical calibration. Longitudinal datasets containing intervention records, supervisory actions, software deployments, rollback events, operational restrictions, and governance responses would permit statistical estimation of ERADM parameters and evaluation of variable sensitivity.

Third, additional domain applications should be examined. Although autonomous transportation provides a useful empirical environment, governance observability challenges increasingly appear within healthcare, robotics, critical infrastructure, industrial automation, financial systems, and defense applications. Cross-domain evaluation would help determine whether ERADM functions as a genuinely domain-agnostic observability architecture. Fourth, future studies should evaluate integration with existing governance frameworks. NIST AI RMF, OMB M-2410, ISO/IEC 42001, ISO 26262, FAA Safety Management Systems, and related architectures provide established governance mechanisms that could potentially incorporate ERADM variables as additional observability layers [15,16,22,23,34]. Fifth, future work should examine relationships between governance observability and insurance outcomes. Variables associated with intervention capability, restoration feasibility, adaptation tempo, supervisory responsiveness, and dependency formation may contribute to underwriting models, reserve estimation, correlated-risk assessment, and liability evaluation [7,20,25,26,28,40,41,44]. Finally, future research should pursue formal falsification and model refinement. The scientific value of ERADM depends upon its ability to generate measurable expectations capable of empirical contradiction [38]. Continued evaluation across multiple domains, operational environments, and governance architectures will therefore be necessary to validate, refine, or reject portions of the framework.

8.4 Toward a Governance Observability Discipline

The broader implication of these limitations and research opportunities is that governance observability remains an emerging field rather than a mature discipline. Existing governance systems have developed sophisticated approaches for measuring safety, reliability, compliance, accountability, and performance. Comparable measurement architectures for governance-relevant operational conditions remain less developed. ERADM should therefore be viewed as an initial step toward establishing governance observability as a systematic area of inquiry rather than as a finalized measurement framework. The contribution of this study is not the claim that governance observability has been solved. Rather, it demonstrates that governance-relevant operational conditions can be treated as measurable phenomena and that many of those conditions are sufficiently observable to support empirical investigation. Future work will determine the extent to which governance observability can mature into a standardized, auditable, reportable, and operationally useful discipline across increasingly autonomous systems.

9. Discussion

9.1 Principal Findings

Four principal findings emerge from this study. First, contemporary autonomous-system governance exhibits a significant observability limitation. Existing regulatory, certification, investigative, insurance, and liability frameworks provide substantial visibility into operational outcomes, including crashes, incidents, recalls, operational restrictions, compliance actions, litigation, and safety events [15,16,22,23,28,29,34]. However, these frameworks provide comparatively limited visibility into the operational conditions that determine whether governance itself remains effective. As a result, institutions frequently possess extensive information regarding what occurred while possessing comparatively limited information regarding whether intervention capability, supervisory responsiveness, restoration feasibility, operational flexibility, synchronization intensity, and infrastructure dependency remained operationally viable before those outcomes emerged. Second, governance-relevant operational conditions can be transformed into measurable system properties. The ERADM framework demonstrates that concepts frequently discussed in qualitative terms—including meaningful human control, intervention effectiveness, supervisory viability, restoration capability, operational dependence, and governance responsiveness—can be operationalized through measurable variables linked to observable evidence pathways. This finding is significant because it shifts governance analysis from a predominantly inferential activity toward a measurable systems discipline. Third, historical reconstruction demonstrates that governance-relevant operational conditions are frequently observable or partially observable despite not being systematically measured within existing reporting architectures. Analyses involving Uber ATG, Tesla ADS/FSD, Cruise, Boeing 737 MAX, and Waymo indicate that intervention effectiveness, supervisory responsiveness, adaptation tempo, restoration capability, operational flexibility, synchronization effects, and dependency formation can often be reconstructed from investigations, recalls, deployment records, safety reports, operational restrictions, litigation, and regulatory actions [9-14]. These findings suggest that the principal challenge facing governance observability may not be the absence of information, but the absence of measurement architectures specifically designed to organize and evaluate that information. Fourth, governance observability has practical implications extending beyond academic theory. The

ability to directly observe governance-relevant operational conditions influences regulatory oversight, certification assurance, accident investigation, insurance underwriting, liability assessment, operational resilience, and public accountability. Consequently, governance observability represents not merely a measurement challenge but a governance capability challenge.

9.2 Governance as an Observable Systems Property

A central contribution of this study is the proposition that governance can be treated as an observable systems property. Traditional governance frameworks typically evaluate governance through institutional structures, legal authorities, compliance obligations, accountability mechanisms, organizational responsibilities, and post-event outcomes. These approaches remain essential and are not challenged by the ERADM framework proposed here. However, they frequently provide limited visibility into the operational conditions that determine whether governance remains practically achievable. ERADM introduces a complementary perspective. Rather than evaluating governance solely through structures or outcomes, the framework evaluates the operational conditions that influence whether effective supervision, intervention, restoration, and oversight remain practically achievable within complex autonomous systems. In this respect, governance becomes observable not because institutions themselves change, but because the conditions that influence institutional effectiveness become measurable [11-13]. By linking observable pre-incident governance conditions to subsequently observed operational, regulatory, liability, insurance, and safety outcomes, the framework establishes an empirical bridge between governance conditions and governance consequences. This linkage provides a foundation for longitudinal measurement, comparative analysis, calibration, and formal falsification by allowing researchers, regulators, insurers, operators, and investigators to evaluate whether variations in intervention capability, supervisory responsiveness, rollback feasibility, operational flexibility, synchronization intensity, and dependency formation are systematically associated with subsequent governance outcomes. As evidence accumulates across cases and time horizons, the resulting observations may support estimation of the magnitude, direction, and relative strength of these relationships while remaining compatible with additional operational, environmental, organizational, and contextual factors that may also influence observed outcomes [1-55].



Figure 10: Governance Observability Across the Incident Lifecycle STAMP and ERADM

Figure 10 illustrates the complementary relationship between existing governance frameworks, STAMP, and ERADM. Existing governance architectures provide essential capabilities for safety management, accountability, compliance, and post-incident analysis [15,16,22,23,28,29,34]. ERADM does not replace these functions; rather, it adds an observability layer that measures operational conditions influencing governance effectiveness before outcomes occur, creating a measurable bridge between antecedent conditions and subsequent operational, regulatory, liability, insurance, and safety outcomes. STAMP focuses on explaining how control failures, feedback breakdowns, and inadequate constraints contribute to adverse events, whereas ERADM focuses on measuring the conditions under which those failures become more or less likely to emerge [24]. The frameworks are therefore complementary by design: STAMP is strongest in post-incident governance and causal analysis, whereas ERADM is strongest in pre-incident governance observability, measurement, and evaluation. While STAMP provides a rigorous framework for analyzing control structures, feedback mechanisms, safety constraints, and unsafe control actions, it does not prescribe a standardized set of observability metrics analogous to AEI, OER, RFT, OSCR, CLSI, H, S, K, or Θ through which governance effectiveness can be monitored, compared, and evaluated longitudinally across systems and deployments. The principal distinction is temporal and observational. STAMP is optimized for explaining how control failures, feedback deficiencies, and inadequate constraints contributed to an observed outcome [24]. ERADM is optimized for measuring whether the governance conditions associated with those

failures are changing before an outcome occurs [24,39]. This distinction becomes most consequential when governance decisions must be made under uncertainty, before sufficient outcome evidence exists [38,39]. In such circumstances, the availability of standardized observability variables may allow governance degradation, declining intervention effectiveness, increasing dependency formation, compressed supervisory response capacity, and reduced restoration feasibility to be monitored and evaluated prior to their manifestation as accidents, recalls, regulatory interventions, liability events, or other adverse outcomes [10,24,37,39].

9.3 Relationship to Existing Frameworks

The findings should not be interpreted as evidence that existing governance frameworks are inadequate. On the contrary, contemporary regulatory, certification, safety-management, insurance, and liability systems perform essential governance functions and provide substantial visibility into safety, compliance, accountability, transparency, and operational outcomes [15,16,22,23,28,29,34]. The contribution of ERADM lies elsewhere. The framework addresses a category of governance-relevant operational conditions that remain only partially observable within many existing architectures. This distinction is particularly important with respect to frameworks such as NIST AI RMF, OMB M-24-10, ISO/IEC 42001, ISO 26262, FAA Safety Management Systems, transportation safety regulations, and related governance mechanisms [15,16,22,23,34]. These frameworks establish important requirements concerning safety, risk management, accountability, transparency, and organizational governance. ERADM does not replace these

functions. Instead, it provides a complementary observability layer capable of measuring operational conditions that may influence whether these governance mechanisms remain effective in practice. The framework therefore should be understood as an extension of existing governance architectures rather than as an alternative to them.

9.4 Relationship to STAMP and Systems-Theoretic Approaches

ERADM is also complementary to established systems-theoretic approaches, particularly STAMP [24]. STAMP provides a powerful framework for understanding accidents as failures of control, feedback, constraints, and organizational interactions. It explains how unsafe control actions emerge and how complex socio-technical systems produce hazards and failures. ERADM addresses a related but distinct question. Rather than focusing on how hazards emerge, the framework focuses on whether the operational conditions necessary for effective governance remain observable and measurable before hazards emerge and produce outcomes. In practical terms, STAMP explains why control failures occur. ERADM evaluates whether intervention capability, supervisory responsiveness, restoration feasibility, operational flexibility, synchronization intensity, and infrastructure dependency remain observable before those failures become visible through accidents, recalls, enforcement actions, or operational restrictions. The two approaches are therefore complementary. STAMP provides causal understanding of control failures. ERADM provides an observability architecture through which governance conditions associated with those failures may be measured.

9.5 Implications for Autonomous-System Governance

The empirical findings suggest that governance observability may become increasingly important as autonomous systems become more adaptive, interconnected, and operationally significant. Historically, many governance systems evolved within environments characterized by relatively stable technologies, predictable operational boundaries, and slower rates of change. Autonomous systems increasingly challenge these assumptions through continuous software modification, adaptive behavior, distributed coordination, infrastructure dependency, and recursive improvement processes. Under these conditions, effective governance depends not only upon authority, accountability, and compliance but also upon maintaining visibility into the operational conditions that influence governance effectiveness itself. The significance of governance observability therefore extends beyond autonomous vehicles. Similar challenges appear within robotics, healthcare, logistics, critical infrastructure, industrial automation, financial systems, defense environments, and other domains characterized by increasing autonomy and operational complexity.

Consequently, the framework proposed here should be viewed not solely as a transportation-governance model but as a broader contribution to the emerging study of governance observability across autonomous and adaptive

systems.

9.6 Discussion Summary

Collectively, the findings support the proposition that governance observability can be treated as a measurable systems problem. Governance-relevant operational conditions are not merely theoretical constructs. They can frequently be observed, reconstructed, measured, audited, and empirically investigated using operational evidence. The framework developed in this study provides a structured methodology through which these conditions become direct objects of analysis. While substantial calibration, validation, and refinement remain necessary, the results demonstrate that governance observability is sufficiently mature to support systematic empirical investigation. The broader implication is that governance may be evaluated not solely through outcomes, responsibilities, or institutional structures, but through direct observation of the operational conditions that determine whether effective governance remains achievable.

10. Conclusion

This study argues that contemporary autonomous-system governance is constrained by a significant observability limitation. Existing governance institutions—including transportation regulators such as the National Highway Traffic Safety Administration (NHTSA), Federal Aviation Administration (FAA), Federal Motor Carrier Safety Administration (FMCSA), investigative bodies such as the National Transportation Safety Board (NTSB), state regulatory authorities, certification organizations, insurers, and liability frameworks—provide substantial visibility into operational outcomes, including crashes, incidents, recalls, operational restrictions, compliance actions, litigation, and safety events [15,16,18,19,28,29,36]. However, these frameworks provide comparatively limited visibility into the operational conditions that shape those outcomes and influence whether governance itself remains practically effective. As autonomous systems become increasingly adaptive, interconnected, software-defined, synchronized, and operationally significant, effective governance depends not only upon observing outcomes but also upon observing the conditions under which meaningful supervision, intervention, restoration, and oversight remain possible. To address this challenge, this article introduced the Erckenbrack Recursive Authority Dynamics Model (ERADM), a domain-agnostic operational observability architecture founded on the proposition that governance observability can be treated as a measurable systems problem. Rather than evaluating governance solely through organizational structures, legal authorities, accountability mechanisms, compliance obligations, or post-event outcomes, ERADM focuses on operational conditions that influence whether intervention, supervision, restoration, and oversight remain achievable during system operation. By transforming these conditions into measurable variables, the framework establishes a foundation through which governance-relevant conditions may be observed, reconstructed, compared, audited, and potentially monitored over time.

The ERADM framework transforms governance-relevant operational conditions into measurable system properties through a structured architecture consisting of observable variables, mathematical relationships, observability pathways, governance-feasible conditions, state-space representations, validation methodologies, and falsification criteria. Variables including Authority Elasticity Index (AEI), Override Effectiveness Rate (OER), Rollback Feasibility Time (RFT), Option Set Compression Ratio (OSCR), Closed-Loop Self-Improvement Interval (CLSI) Human Supervisory Response Horizon (H), Synchronization Intensity (S), Infrastructure Coupling (K), Recursive Leverage Factor (RLF), and Authority Transition Index (Θ) provide a formal measurement architecture through which governance conditions become direct objects of observation, reporting, auditing, reconstruction, and empirical evaluation.

Without measurable governance variables, governance itself cannot become an empirically observable systems property. The historical reconstruction analyses presented in this study demonstrate that many governance-relevant operational conditions can be identified using publicly available evidence. Investigations, recalls, operational restrictions, deployment records, safety reports, software updates, litigation materials, certification actions, and regulatory interventions associated with Uber ATG, Tesla ADS/FSD, Cruise, Boeing 737 MAX, and Waymo contain evidence sufficient to reconstruct many of the conditions represented within the framework [4,9,31,33,36,43,45,48]. Across the historical analyses, fourteen of fifteen directional expectations derived from ERADM were consistent with the available evidence, no expectations were contradicted, and one remained indeterminate. These findings do not constitute definitive validation of the ERADM framework. Rather, they demonstrate the empirical feasibility of observing, reconstructing, comparing, and evaluating governance-relevant operational conditions using real-world evidence. Collectively, these findings provide evidence that a dedicated governance-observability architecture is justified. The analyses indicate that governance-relevant operational conditions exist, can be measured, can be reconstructed from real-world evidence, and are only partially observable through existing governance frameworks. The empirical results therefore support the need for a systematic framework such as ERADM capable of organizing, quantifying, and evaluating these conditions as direct objects of governance analysis. These findings suggest that the principal challenge facing governance observability may not be the absence of relevant information, but the absence of measurement architectures designed to organize, quantify, and report that information systematically. Many of the underlying data elements required to estimate ERADM variables are already generated through normal system operation in the form of telemetry, event logs, override records, software update histories, rollback records, operational reports, and regulatory documentation. Evidence of this was observed throughout the historical reconstruction analyses, where publicly available investigations, recall records, deployment reports, regulatory actions, safety

assessments, and operational disclosures associated with Uber ATG, Tesla ADS/FSD, Cruise, Boeing 737 MAX, and Waymo contained sufficient information to reconstruct numerous governance-relevant conditions represented within the framework [9-14]. Existing governance reporting systems frequently contain fragments of the evidence necessary for evaluating intervention capability, supervisory responsiveness, restoration feasibility, operational flexibility, synchronization effects, infrastructure dependencies, and authority-transition dynamics. What has often been lacking is a coherent framework through which these observations can be integrated into a unified governance-observability architecture capable of measuring pre-incident conditions that may influence subsequent operational, regulatory, liability, and safety outcomes, thereby improving the ability of governance institutions to identify deteriorating governance conditions before they become visible through crashes, incidents, recalls, deployment restrictions, regulatory interventions, certification actions, liability events, insurance losses, or other adverse operational outcomes [9-14]. The contribution of ERADM is therefore not that it replaces existing governance systems. Regulatory frameworks, certification programs, safety-management systems, insurance structures, liability doctrines, and accountability mechanisms remain essential. Instead, ERADM functions as a complementary observability layer that extends existing governance architectures by providing direct visibility into operational conditions that are frequently inferred rather than observed and measured. In this respect, the framework complements existing approaches including NIST AI RMF, OMB M-24-10, ISO/IEC 42001, ISO 26262, FAA Safety Management Systems, transportation safety regulations, and systems-theoretic approaches such as STAMP [15,16,22,23,24,34].

The introduction of the Market-Regulatory-Reinsurance Observability Weight (MRROW) further suggests that governance observability is influenced not solely by conditions internal to autonomous systems, but also by the institutional environments in which those conditions are evaluated [20,25,26,28,41]. Regulatory agencies, insurers, reinsurers, certification authorities, and systemic-risk organizations influence how operational evidence is interpreted, prioritized, and acted upon [20,25,26,28,41]. Consequently, governance outcomes may depend not only upon the existence of observable conditions, but also upon the capacity of external institutions to detect and respond to them [20,25,26,28,41]. This finding extends governance observability beyond technical systems alone and highlights institutional observability as a complementary component of effective autonomous-system governance. An important limitation of the present analysis is that many governance-relevant operational conditions are not currently subject to standardized public reporting requirements. Consequently, several variables could only be reconstructed indirectly from publicly available investigations, recalls, deployment records, regulatory actions, litigation materials, safety reports, and operational disclosures [4,9,31,33,36,43,45,48]. This limitation reflects a constraint of contemporary

reporting architectures rather than a demonstrated absence of underlying information. The historical reconstruction analyses suggest that many of the data elements necessary to estimate ERADM variables are generated during normal system operation but are frequently unavailable, fragmented, proprietary, inconsistently reported, or inaccessible to external researchers [4,9,31,33,36,43,45,48]. As a result, the present study establishes empirical feasibility of ERADM rather than calibrated measurement. Future research should focus on developing standardized reporting architectures, obtaining access to operational datasets, calibrating variable thresholds and composite indices, and evaluating framework performance prospectively across multiple domains. Although the empirical analyses presented here focus primarily on autonomous transportation systems, the governance conditions represented within ERADM are not unique to transportation. Similar observability challenges increasingly appear within robotics, healthcare, logistics, industrial automation, critical infrastructure, financial systems, defense environments, and other domains characterized by increasing autonomy, rapid adaptation cycles, growing synchronization across distributed systems, expanding infrastructure dependencies, and recursive operational processes that may alter the conditions under which effective governance, supervision, intervention, and restoration remain possible. Future research should therefore evaluate the framework prospectively, calibrate variables using operational datasets, develop standardized measurement methodologies, establish domain-specific governance thresholds, and assess applicability across additional operational domains.

The broader implication is straightforward. Existing governance frameworks are highly effective at observing outcomes; ERADM is designed to observe the operational conditions that shape those outcomes. More broadly, this study contributes an initial formal framework through which governance observability can be investigated, measured, compared, calibrated, and potentially falsified as a systems property. As autonomous systems continue to expand across safety-critical environments, improving visibility into governance-relevant operational conditions may become increasingly important for maintaining effective oversight, accountability, resilience, and meaningful human control.

References

1. Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *The economic journal*, 99(394), 116-131.
2. Arthur, W. B. (2009). *The nature of technology: What it is and how it evolves*. Simon and Schuster.
3. Benavides v. Tesla, Inc., No. 1:21-cv-21940-BLOOM/Torres, 2025 WL 1785106 (S.D. Fla. Jun. 25, 2025).
4. California Department of Motor Vehicles, "DMV statement on Cruise LLC suspension," Oct. 24, 2023. [Online].
5. California Department of Motor Vehicles, "Autonomous vehicle disengagement reports," 2024
6. California Public Utilities Commission, "CPUC approves permits for Cruise and Waymo to charge fares for passenger service in San Francisco," Aug. 10, 2023.
7. Casualty Actuarial Society, Statement of Principles Regarding Property and Casualty Insurance Ratemaking, Arlington, VA, USA, 2021.
8. Coase, R. H. (2013). The problem of social cost. *The journal of Law and Economics*, 56(4), 837-877.
9. Cruise LLC, Safety and Incident Reporting Materials Related to the October 2, 2023 San Francisco Pedestrian Incident, 2024.
10. Dekker, S. (2012). Drift into failure, from hunting broken components to understanding complex systems. *Aviation, Space, and Environmental Medicine*.
11. Erckenbrack, A. (2026). Authority Transfer, Compulsory Adoption, and Recursive Displacement in Artificial Intelligence Systems. *Compulsory Adoption, and Recursive Displacement in Artificial Intelligence Systems* (January 17, 2026).
12. Erckenbrack, A. (2026). Education for an AI/ASI Transition: Preparedness Authority and Curriculum in the United States 2025–2035.
13. A. Erckenbrack, "Global Recursive Learning Systems: Architecture, Multi-Level Recursion and the Emergence of Distributed Authority," *Engineering and Applied Sciences Journal*, vol. 3, no. 3, pp. 1–19, 2026.
14. Lognoul, M. (2025). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act–AI Act). *Revue du Droit des Technologies de l'information*, (3-4), 145-189.
15. United States. Office of Management and Budget. (2024). Advancing governance, innovation, and risk management for agency use of artificial intelligence. Office of Management and Budget.
16. Federal Aviation Administration, Safety Management Systems (14 CFR Part 5), Washington, DC, USA, 2024.
17. Federal Aviation Administration, Roadmap for Artificial Intelligence Safety Assurance, Washington, DC, USA, 2024.
18. Federal Motor Carrier Safety Administration, Safe Integration of Automated Driving Systems-Equipped Commercial Motor Vehicles, Fed. Reg. 88, 6691, 2023.
19. Federal Motor Carrier Safety Administration, Automated Driving Systems Regulatory Materials, Washington, DC, USA, 2024.
20. The Geneva Association, Insurance and Artificial Intelligence: Implications for Risk Assessment, Underwriting, and Governance, Geneva, Switzerland, 2024.
21. Hollnagel, E. (2018). *Safety-I and safety-II: the past and future of safety management*. CRC press.
22. International Organization for Standardization, ISO 26262: Road Vehicles—Functional Safety, Geneva, Switzerland, 2018.
23. International Organization for Standardization, ISO/IEC 42001: Artificial Intelligence—Management System, Geneva, Switzerland, 2023.
24. Leveson, N. G. (2016). *Engineering a safer world: Systems thinking applied to safety* (p. 560). The MIT Press.
25. Lloyd's, Systemic Risk and Emerging Technologies,

- London, U.K., 2023.
26. Munich Re, Artificial Intelligence, Autonomous Systems, and Emerging Liability Exposures, Munich, Germany, 2024.
 27. National Aeronautics and Space Administration, AI/ML Assurance, NASA Technical Reports Server, Washington, DC, USA, 2024.
 28. Oecd, O. E. C. D. (2024). Principles on artificial intelligence. Updates, 2025, 38.
 29. National Highway Traffic Safety Administration, Standing General Order on Crash Reporting for Automated Driving Systems and Level 2 Advanced Driver Assistance Systems, Washington, DC, USA, 2021.
 30. National Highway Traffic Safety Administration, Standing General Order on Crash Reporting for Level 2 Advanced Driver Assistance Systems: Summary Report, Washington, DC, USA, 2022.
 31. National Highway Traffic Safety Administration, Part 573 Safety Recall Report: Tesla Autosteer Recall (23V838), Washington, DC, USA, 2023.
 32. National Highway Traffic Safety Administration, Automated Driving Systems and Advanced Driver Assistance Systems Crash Reporting and Safety Investigation Materials, Washington, DC, USA, 2023.
 33. National Highway Traffic Safety Administration, "GM's Cruise failed to fully report pedestrian crash," Sept. 30, 2024.
 34. AI, N. (2024). Artificial intelligence risk management framework: Generative artificial intelligence profile. *NIST Trustworthy and Responsible AI Gaithersburg, MD, USA*.
 35. Navani, S. (2025). Comment on 2025 National Artificial Intelligence (AI) Research and Development (R&D) Strategic Plan. Available at SSRN 5292967.
 36. National Transportation Safety Board, Collision Between Vehicle Controlled by Developmental Automated Driving System and Pedestrian, Tempe, Arizona, March 18, 2018, NTSB/HAR-19/03, Washington, DC, USA, 2019.
 37. Cummings, L. L. (1984). Normal Accidents: Living with High-Risk Technologies.
 38. Popper, K. (1959). The logic of scientific discovery.
 39. Rasmussen, J. (1997). Risk management in a dynamic society: a modelling problem. *Safety science*, 27(2-3), 183-213.
 40. Rejda, G. E. (2005). *Principles of risk management and insurance*. Pearson Education India.
 41. Swiss Re Institute, New Technologies and Systemic Risk Accumulation, Zurich, Switzerland, 2024.
 42. Henson, R. (2024). Government-Backed Insurance for Artificial Intelligence Technologies. *Ga. St. UL Rev.*, 41, 559.
 43. U.S. Attorney's Office, Northern District of California, "Cruise admits to submitting a false report to influence a federal investigation and agrees to pay \$500,000," Nov. 14, 2024.
 44. Vaughan, E. J., & Vaughan, T. (2007). *Fundamentals of risk and insurance*. John Wiley & Sons.
 45. Waymo LLC, Waymo Safety Report, Mountain View, CA, USA, 2021.
 46. Waymo LLC, Waymo Safety Methodologies and Safety Readiness Determinations, 2023
 47. Waymo LLC, Safety Impact, Mountain View, CA, USA, 2024.
 48. Waymo LLC, Published Safety Research Papers for Autonomous Vehicles, Mountain View, CA, USA, 2024.
 49. Williamson, O. E. (1985). The economic institutions of capitalism. Firms, markets, relational contracting. In *Das Summa Summarum des Management: Die 25 wichtigsten Werke für Strategie, Führung und Veränderung* (pp. 61-75). Wiesbaden: Gabler.
 50. Leveson, N. G., & Turner, C. S. (1993). An investigation of the Therac-25 accidents. *Computer*, 26(7), 18-41.
 51. Force, T. (2004). Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations, US-Canada Power System Outage Task Force.
 52. Securities, U. S., & Exchange Commission. (2013). In the matter of knight capital americas llc. Administrative Proceeding File, (3-15570).
 53. House Committee on Transportation and Infrastructure, The Design, Development, and Certification of the Boeing 737 MAX, Majority Staff Report, U.S. House of Representatives, Washington, DC, USA, Sept. 2020.
 54. National Transportation Safety Board, Assumptions Used in the Safety Assessment Process and the Effects of Multiple Alerts and Indications on Pilot Performance, Safety Recommendation Report ASR-19/01, Washington, DC, USA, Sept. 2019.