



# AI Governance Maturity and Digital Resilience: A Multi-Level Model of Trustworthy AI Implementation

Navya Sri Maddukuri, Hari Nagakoteswar Tripurari

<sup>1</sup>Graduate Researcher Faulkner University

<sup>2</sup>Graduate Research Assistant Dakota State University

**Abstract-** As artificial intelligence systems become embedded in core business processes, organizations increasingly require governance mechanisms that translate high-level ethical principles into operational capabilities capable of preventing, detecting, and recovering from AI-related failures. This study develops and validates an AI Governance Maturity Index (AGMI) spanning six dimensions — policy formalization, risk monitoring, model documentation, human oversight, incident response, and continuous auditing — and examines its relationship to organizational digital resilience. Drawing on survey data, archival AI incident records, and six in-depth case studies from 341 firms deploying high-impact AI systems across six industry sectors, the study tests whether governance maturity reduces AI failure frequency and severity and improves post-incident recovery. Regression results indicate that AGMI is significantly associated with lower incident frequency ( $\beta = -0.58, p < .001$ ), lower incident severity ( $\beta = -0.49, p < .001$ ), shorter mean time to recovery ( $\beta = -8.92$  hours,  $p < .001$ ), and higher Digital Resilience Index scores ( $\beta = 0.11, p < .001$ ), with continuous auditing and human oversight intensity emerging as the dimensions most strongly associated with resilience outcomes and exhibiting significant complementary interaction effects. A five-tier maturity comparison reveals a tenfold difference in mean time to recovery between Tier 1 (Ad Hoc) and Tier 5 (Adaptive) firms (78.4 versus 9.8 hours). A twelve-month governance investment pilot across 48 firms demonstrates that combined investment in continuous auditing, cross-functional incident response, standardized documentation, tiered oversight, and automated risk telemetry produces a 1.54-point AGMI gain and a 3.67-incident annual reduction, substantially exceeding the effect of any single investment. Six case studies of firms experiencing significant AI incidents illustrate the governance-resilience feedback loop through which incident response activates targeted maturity advancement. The paper contributes a multi-level (macro-meso-micro-outcome) model linking responsible AI practices to digital resilience, a validated maturity instrument, and a practical roadmap for firms seeking to scale AI deployment while maintaining trust, regulatory compliance, and organizational accountability.

**Keywords-** AI governance, digital resilience, trustworthy AI, maturity model, incident response, model documentation, human oversight, continuous auditing, information systems governance, responsible AI.

## I. Introduction

Artificial intelligence systems have transitioned from peripheral analytical tools to load-bearing components of core business processes — underwriting credit decisions, triaging healthcare cases, managing supply chains, screening job applicants, and increasingly, operating with substantial autonomy across enterprise workflows. This



transition has been accompanied by a parallel proliferation of high-profile AI failures: biased credit-scoring models triggering regulatory action, clinical decision-support tools producing erroneous recommendations, generative AI systems producing offensive or legally consequential outputs in customer-facing channels, and autonomous systems causing operational disruptions when their behavior diverges from expected parameters in ways that went undetected until material harm occurred.

In response, a substantial body of work has emerged articulating ethical principles for trustworthy AI — frameworks emphasizing fairness, accountability, transparency, and human oversight (Floridi et al., 2018; OECD, 2024). However, a persistent gap exists between the articulation of such principles and their operationalization as organizational capabilities. Floridi (2019) explicitly identified this 'ethics shirking' risk: organizations may adopt the language of AI ethics — publishing principles, forming ethics boards, issuing public commitments — without developing the underlying operational infrastructure (documented processes, monitoring systems, incident response capabilities, audit mechanisms) required to translate these principles into consistent practice. This gap between principle and practice has direct consequences for organizational resilience: when AI systems fail — and the evidence suggests they will, given the inherent uncertainty and complexity of machine learning systems operating in dynamic environments (Power, 2007) — organizations without operationalized governance capabilities are poorly positioned to detect failures promptly, respond effectively, and recover without lasting damage to operations, regulatory standing, and stakeholder trust.

The information systems governance literature provides substantial theoretical resources for understanding organizational governance capability development (Tallon et al., 2013; Tiwana & Konsynski, 2010), and the digital resilience literature (Madnick et al., 2017; Linkov et al., 2018) provides frameworks for understanding organizational recovery from technology-related disruptions. However, these literatures have not yet been integrated into a model that specifically addresses AI governance as a multi-dimensional organizational capability and tests its relationship to the resilience outcomes — incident frequency, severity, and recovery speed — that ultimately determine whether AI deployment at scale is sustainable from a risk management perspective.

This study addresses this gap through four research questions: (RQ1) What dimensions constitute AI governance maturity as an operational organizational capability, and how can this maturity be validly measured? (RQ2) Does AI governance maturity predict AI incident frequency, severity, and recovery speed (digital resilience outcomes), controlling for firm characteristics and AI system risk exposure? (RQ3) Which governance dimensions — and which combinations of dimensions — are most strongly associated with resilience outcomes? (RQ4) How does the governance-resilience relationship operate dynamically — that is, how do AI incidents themselves catalyze governance maturity advancement, and can deliberate governance investment accelerate this advancement?

Drawing on survey data, archival AI incident records, and six in-depth case studies from 341 firms deploying high-impact AI systems across six industry sectors (financial



services, healthcare, technology/SaaS, retail/e-commerce, manufacturing, and insurance), this study makes four contributions. First, it develops and validates the AI Governance Maturity Index (AGMI), a six-dimensional, five-level instrument grounded in established IT governance and IS maturity model traditions. Second, it provides large-sample empirical evidence on the relationship between governance maturity and digital resilience outcomes, including a five-tier maturity comparison revealing the magnitude of resilience differences across maturity levels. Third, it develops a multi-level (macro-meso-micro-outcome) conceptual model situating AGMI within the broader regulatory and market context that shapes governance investment incentives. Fourth, it provides intervention-based evidence — via a twelve-month governance investment pilot — demonstrating that targeted governance investment produces measurable maturity advancement and resilience improvement within a practically relevant timeframe.

## II. Theoretical Background

### From AI Ethics Principles to Operational Governance

The proliferation of AI ethics frameworks over the past decade — including the OECD AI Principles (OECD, 2024), the EU's trustworthy AI guidelines, and numerous corporate and academic frameworks (Floridi et al., 2018) — converges on a recurring set of high-level principles: fairness, transparency, accountability, privacy, safety, and human oversight. Floridi (2019) provided an influential critique of this principle-centric approach, identifying five risks of 'unethical' AI practice that principle adoption alone does not address, including ethics shirking (adopting ethical rhetoric without operational substance), ethics bluewashing (presenting an organization as more ethically AI-oriented than it is), and ethics lobbying (using ethical commitments to forestall binding regulation while avoiding operational change).

Mökander and Floridi's (2021) subsequent work on ethics-based auditing represents a key conceptual bridge between principle-level AI ethics and operational governance: ethics-based auditing reframes ethical AI principles as auditable criteria against which specific AI systems and organizational processes can be assessed, creating a pathway from abstract principle to concrete, verifiable organizational practice. Raji et al.'s (2020) internal algorithmic auditing framework similarly operationalizes accountability as a structured organizational process — comprising design-stage, pre-deployment, and post-deployment audit activities — rather than as an abstract value commitment. This study's AGMI framework builds directly on this operationalization tradition, translating the high-level principles of trustworthy AI into six measurable organizational capability dimensions (policy formalization, risk monitoring, model documentation, human oversight, incident response, and continuous auditing) that collectively span the pre-deployment, deployment, and post-incident phases of the AI system lifecycle.

### Dynamic Capabilities and Governance as Organizational Capability

Teece et al.'s (1997) dynamic capabilities framework — comprising sensing, seizing, and reconfiguring capabilities — provides a theoretical foundation for conceptualizing AI governance not as a static compliance function but as an organizational capability that itself must evolve in response to a changing technological and regulatory environment. Applied to AI governance, sensing capabilities correspond to risk



monitoring and continuous auditing (detecting emerging risks and performance issues); seizing capabilities correspond to incident response (mobilizing organizational resources to address identified issues); and reconfiguring capabilities correspond to the feedback mechanisms through which policy, documentation standards, and oversight structures are revised based on operational experience (Zollo & Winter, 2002).

This dynamic capabilities framing has a direct empirical implication that distinguishes this study's approach from purely cross-sectional governance assessments: if AI governance maturity is genuinely a dynamic capability, it should not only predict resilience outcomes (RQ2) but should itself be responsive to organizational learning from incidents — a proposition this study tests through the case study analyses (Section 5) documenting AGMI changes following significant AI incidents, and through the governance investment pilot (Section 4.6) demonstrating that deliberate investment accelerates maturity advancement beyond the rate of incident-driven learning alone.

### **Digital Resilience and the Resilience-Risk Distinction**

The digital resilience literature draws an important conceptual distinction between risk management — activities aimed at preventing adverse events — and resilience — organizational capacity to maintain or rapidly restore essential functions when adverse events occur despite prevention efforts (Linkov et al., 2018; Weick & Sutcliffe, 2007). Madnick et al. (2017) operationalize digital resilience along dimensions including the capacity to anticipate, monitor, respond to, and learn from disruptions — a structure that maps closely onto this study's AGMI dimensions (risk monitoring as anticipation and monitoring capacity; incident response as response capacity; continuous auditing and documentation as learning capacity).

This resilience framing is theoretically important because it implies that governance maturity's value does not depend on its capacity to prevent all AI incidents — an unrealistic standard given the inherent complexity and environmental dependence of AI systems (Power, 2007) — but on its capacity to reduce the severity and duration of incidents that do occur, even when occurrence itself is not fully preventable. This study's Digital Resilience Index (DRI) and Mean Time to Recovery (MTTR) measures directly operationalize this resilience-centric framing, complementing (rather than substituting for) the incident frequency and severity measures that more directly reflect risk-prevention outcomes (Sharma et al., 2025).

### **A Multi-Level Model of AI Governance and Digital Resilience**

Synthesizing the preceding theoretical perspectives, this study proposes a multi-level conceptual model (Figure 1) that situates AI governance maturity within a broader macro-meso-micro-outcome structure. At the macro level, regulatory frameworks (e.g., the EU AI Act), industry self-regulatory bodies, and market-level pressures (investor ESG expectations, competitive dynamics) establish the external accountability context that shapes organizational incentives for governance investment. At the meso level, the AGMI's six dimensions represent organizational-level governance architecture — the policies, processes, and capabilities that an organization develops in response to (and sometimes ahead of) macro-level pressures. At the micro level, these meso-level capabilities are instantiated for specific AI systems through system-level controls — risk classifications, oversight assignments, documentation artifacts, and monitoring



configurations specific to individual deployed systems. Finally, at the outcome level, the cumulative effect of micro-level controls, aggregated across an organization's AI system portfolio, produces the resilience and trust outcomes — incident frequency, severity, recovery speed, and stakeholder trust — that this study's empirical analyses examine.

Figure 1. A Multi-Level Model of AI Governance Maturity and Digital Resilience

| Macro Level:<br>Regulatory &<br>Market Context   | Meso Level:<br>Organizational<br>Governance<br>Architecture  | Micro Level:<br>System-Specific<br>Controls  | Outcome Level:<br>Resilience & Trust  |
|--|--|--|---|
| Drivers: <ul style="list-style-type: none"> <li>• AI-specific regulation (EU AI Act, sector rules)</li> <li>• Industry self-regulatory bodies</li> <li>• Investor ESG/AI-risk disclosure expectations</li> <li>• Competitive AI adoption pressure</li> </ul> Effect:<br>Sets the baseline expectations and external accountability pressures that shape meso-level governance investment | AGMI Dimensions: <ul style="list-style-type: none"> <li>• Policy formalization</li> <li>• Risk monitoring</li> <li>• Model documentation</li> <li>• Human oversight</li> <li>• Incident response</li> <li>• Continuous auditing</li> </ul> Organizational capabilities that translate macro-level expectations into operational practice | System-Level Controls: <ul style="list-style-type: none"> <li>• Per-system risk classification</li> <li>• System-specific oversight assignment</li> <li>• Model card completeness</li> <li>• Drift and performance monitoring</li> <li>• Escalation pathway clarity</li> </ul> Where meso-level AGMI capabilities are instantiated for individual AI systems | Resilience Outcomes: <ul style="list-style-type: none"> <li>• Lower incident frequency</li> <li>• Lower incident severity</li> <li>• Faster recovery (MTTR)</li> <li>• Higher Digital Resilience Index</li> </ul> Trust Outcomes: <ul style="list-style-type: none"> <li>• Stakeholder trust score</li> <li>• Reduced regulatory inquiry likelihood</li> <li>• Feedback loop to macro-level legitimacy</li> </ul> |

Note. AGMI = AI Governance Maturity Index. The model proposes that macro-level context shapes meso-level governance investment incentives, meso-level AGMI capabilities are instantiated through micro-level system controls, and the aggregate quality of micro-level controls across a firm's AI portfolio determines outcome-level resilience and trust, which in turn feed back to macro-level legitimacy (e.g., via reduced regulatory inquiry rates, Table 8).

### Hypothesized Relationships

Based on the theoretical model presented in Figure 1 and the preceding literature review, this study formulates the following hypotheses, tested in Section 4. H1: AI Governance Maturity (AGMI) is negatively associated with AI incident frequency and severity. H2: AGMI is negatively associated with Mean Time to Recovery (MTTR) and positively associated with the Digital Resilience Index (DRI). H3: Continuous Auditing and Human Oversight Intensity exhibit significant complementary (positive) interaction effects on resilience outcomes, beyond their independent associations —



reflecting the dynamic capabilities logic (Section 2.2) in which sensing (auditing) and seizing (oversight-enabled response) capabilities are mutually reinforcing rather than independently sufficient. H4: AI incidents catalyze subsequent AGMI advancement (a within-firm, pre/post-incident comparison), and this incident-driven advancement is accelerated by deliberate governance investment beyond the rate observed in firms relying on incident-driven learning alone (Jagatha et al., 2025).

### **Governance Signaling, Legitimacy, and the Documentation-Practice Gap**

A further theoretical consideration relevant to this study's design concerns the relationship between governance signals — the documented policies, model cards, and audit artifacts that constitute much of what is externally observable about an organization's AI governance — and governance practice, the actual organizational behaviors these artifacts are intended to represent. The institutional theory literature on organizational legitimacy (Suchman, 1995, as applied in adjacent IS governance contexts) suggests that organizations facing external accountability pressure (Section 2.4's macro level) may respond by producing governance signals — policy documents, audit reports, model documentation — that satisfy external legitimacy requirements with varying degrees of fidelity to underlying organizational practice, a phenomenon related to but distinct from Floridi's (2019) ethics bluewashing concept in that signal-practice gaps may arise from genuine organizational complexity (the difficulty of ensuring that documented policies are uniformly implemented across a large, decentralized AI system portfolio) rather than deliberate misrepresentation.

This signal-practice gap consideration has direct methodological implications for this study's design, discussed further in Section 3.2: the AGMI's reliance on a combination of self-reported practice descriptions and archival documentation review (model cards, policy documents, incident response playbooks) means that the instrument may be differentially sensitive to governance signals versus governance practice depending on which AGMI dimensions are assessed. Policy Formalization and Model Documentation — dimensions whose maturity levels (Table 1) are substantially defined in terms of the existence and standardization of documentary artifacts — may be more readily assessable through archival review and therefore more subject to potential signal-practice gaps, whereas Human Oversight Intensity and Incident Response — dimensions whose maturity levels are substantially defined in terms of behavioral and procedural practices that may or may not be fully reflected in documentation — may be less directly observable through documentation alone and more dependent on the self-reported and interview-based components of this study's measurement approach.

This consideration motivates this study's archival-validation subsample design (Section 3.2,  $n = 118$ ): by independently verifying self-reported AGMI scores against archival documentation for a subset of firms, this study can assess the degree to which self-reported and archivally-observable governance maturity converge, providing an empirical check on the signal-practice gap concern. The strong observed convergence ( $r = 0.87$ , Section 4.1) suggests that, at least for the firms in this study's validation subsample, self-reported AGMI scores are not substantially inflated relative to archivally-verifiable practice — though this study cannot rule out that firms agreeing to archival validation may be systematically different (e.g., more confident in the accuracy of their self-assessment, and therefore more likely to have smaller signal-



practice gaps) from firms that did not participate in validation, a limitation discussed further in Section 6.4.

### III. Research Methodology

#### Research Design

This study employs a sequential mixed-methods design integrating three data sources: (1) a cross-sectional survey of AI governance practices administered to senior technology, risk, and compliance executives at 341 firms deploying high-impact AI systems (defined as AI systems whose failure could plausibly cause material financial, operational, legal, or reputational harm) across six industry sectors; (2) archival AI incident records spanning a 36-month period (2023–2026), sourced from firm disclosures, regulatory filings, and a proprietary AI incident database maintained by a risk-analytics partner organization, covering incident frequency, severity classifications, and recovery timelines for 326 of the 341 surveyed firms; and (3) six in-depth case studies of firms that experienced significant, publicly documented AI incidents during the study period, each comprising pre- and post-incident AGMI assessments (12-month interval) and structured interviews with risk and compliance leadership.

The sample of 341 firms spans six industry sectors: Financial Services ( $n = 68$ ), Healthcare ( $n = 54$ ), Technology/SaaS ( $n = 61$ ), Retail/E-Commerce ( $n = 58$ ), Manufacturing ( $n = 52$ ), and Insurance ( $n = 48$ ). Firms were included if they had deployed at least one AI system meeting the high-impact definition above for a minimum of 12 months prior to survey administration, ensuring sufficient operational history for incident records and resilience measures to be meaningful.

#### The AI Governance Maturity Index (AGMI)

The AGMI was developed through a multi-stage process combining literature-derived dimension identification (Section 2.1–2.2), expert panel review (14 AI governance practitioners and academics), and confirmatory factor analysis. The instrument comprises six dimensions — Policy Formalization, Risk Monitoring, Model Documentation, Human Oversight, Incident Response, and Continuous Auditing — each assessed across five maturity levels (Ad Hoc, Initiated, Established, Integrated, Adaptive), presented in full in Table 1. Each dimension was operationalized through 4–6 survey items assessing specific organizational practices corresponding to each maturity level, with respondents' practice descriptions mapped to maturity levels through a combination of self-assessment (against level descriptors provided in the survey instrument) and, for a validation subsample ( $n = 118$ ), independent archival verification by the research team based on firm-provided documentation (policy documents, model cards, incident response playbooks).

Confirmatory factor analysis confirmed a six-factor structure with excellent fit ( $CFI = 0.96$ ,  $RMSEA = 0.051$ ,  $SRMR = 0.058$ ), and the composite AGMI score (a CFA-weighted average across the six dimensions) demonstrated strong composite reliability ( $\omega = 0.92$ ). Comparison of self-assessed AGMI scores with archivally-verified scores in the validation subsample showed strong agreement ( $r = 0.87$ ), supporting the validity



of the self-assessment approach for the full sample while providing a verified-subsample robustness check reported in Table 6.

### **Outcome Measures**

AI Incident Frequency was measured as the annualized count of AI-related incidents (defined as any AI system behavior causing material operational disruption, erroneous decisions affecting customers or stakeholders, regulatory inquiry, or public complaint) recorded in archival incident data over the 36-month observation period. Incident Severity Index (0–10 scale) was constructed from a weighted composite of incident characteristics including financial impact, regulatory consequence, customer impact scope, and media attention, coded by two independent raters (inter-rater reliability ICC = 0.85) for each recorded incident and averaged at the firm level.

Mean Time to Recovery (MTTR, hours) was measured as the time elapsed between incident detection and verified resolution (operational normalization), averaged across a firm's recorded incidents 0/0/00 0:00:00 AM. The Digital Resilience Index (DRI, 0–1) was constructed as a composite survey-based measure assessing organizational capacity to anticipate, monitor, respond to, and learn from AI-related disruptions, following the Madnick et al. (2017) resilience dimension structure (composite reliability  $\omega = 0.90$ ). The Stakeholder Trust Score (STS, 1–7) was a survey-based measure assessing internal and external stakeholder confidence in the organization's AI systems and governance ( $\omega = 0.88$ ). The Regulatory Inquiry Indicator was a binary measure of whether the firm had received a formal regulatory inquiry related to an AI system during the observation period.

### **Analytical Strategy**

The primary analytical strategy employed hierarchical multiple regression to test H1–H3, with AGMI, Human Oversight Intensity, Continuous Auditing Score, and their interaction as predictors of Incident Frequency, Incident Severity, MTTR, and DRI, controlling for AI System Risk Tier (count of high-risk AI systems deployed), firm size, and industry fixed effects. For H4, a within-firm pre/post design compared AGMI scores 12 months before and after a significant AI incident for the six case study firms (Section 5), supplemented by a quasi-experimental analysis of the governance investment pilot (Section 4.6), in which 48 firms implementing a combination of governance investments were compared to a matched comparison group on AGMI and resilience outcome changes over a 12-month period. Robustness checks (Section 4.5) addressed alternative specifications, sample restrictions, and a placebo test examining whether post-period AGMI predicts prior-period incident frequency (which a purely confound-driven account would predict, but a causal-adjacent account would not).

### **AI Incident Classification Taxonomy**

Given this study's reliance on archival AI incident records as a primary data source, a structured incident classification taxonomy was developed to ensure consistent severity coding across the diverse incident types represented in the sample. Incidents were classified along four dimensions: origin (whether the incident arose from model behavior, data quality, integration/deployment infrastructure, or third-party dependency), impact scope (whether effects were contained to internal operations, affected external customers/users, or extended to broader stakeholders such as



regulators or the public), detection mechanism (whether the incident was detected via automated monitoring, internal human observation, external complaint, or regulatory/audit discovery), and consequence severity (financial impact, regulatory consequence, and reputational impact, each rated on a 0–10 scale and averaged to form the composite Incident Severity Index reported in Table 2).

This four-dimensional classification serves two purposes beyond the construction of the Incident Severity Index. First, it enables the detection-mechanism dimension to be cross-referenced against AGMI's Continuous Auditing and Risk Monitoring dimensions, providing a direct behavioral validation of these AGMI dimensions' theoretical function: if Continuous Auditing maturity genuinely reflects an organization's capacity for automated incident detection, firms with higher Continuous Auditing scores should show a higher proportion of incidents detected via automated monitoring (versus external complaint or regulatory discovery) — a pattern confirmed in supplementary analysis (not presented in full tabular form) showing that Tier 4–5 firms detect 68% of incidents via automated monitoring, compared to 19% for Tier 1–2 firms, with the remainder detected via human observation, external complaint, or regulatory/audit discovery (Suryawanshi et al., n.d.). Second, the origin dimension's identification of third-party dependency as a distinct incident origin category provided the initial empirical signal — corroborated by the Firm Epsilon case study (Section 5) — motivating this study's discussion of a potential seventh AGMI dimension addressing third-party AI governance specifically (Section 6.4).

All incident classifications were conducted by two independent raters trained on the classification taxonomy using a calibration set of 40 incidents not included in the final analytic sample; inter-rater reliability for the composite Incident Severity Index was  $ICC = 0.85$  (Section 3.3), with dimension-specific reliabilities ranging from  $ICC = 0.79$  (reputational impact, the most subjective dimension) to  $ICC = 0.91$  (financial impact, the most objectively verifiable dimension). Disagreements were resolved through discussion and, where necessary, reference to additional source documentation (e.g., regulatory filings providing more detailed incident descriptions than initial firm disclosures).

## IV. Results

### The AGMI Framework and Sample Distribution

Table 1 presents the full AGMI framework: six dimensions, each specified across five maturity levels from Ad Hoc to Adaptive. The framework's level descriptors were derived from the expert panel review process described in Section 3.2 and represent a synthesis of academic governance frameworks (Mökander et al., 2021; Raji et al., 2020) and practitioner maturity models (Gartner, 2025; IBM Institute for Business Value, 2024).

Table 1. The AI Governance Maturity Index (AGMI): Six Dimensions Across Five Maturity Levels



| Dimension                   | Level 1 Ad Hoc                        | Level 2 Initiated                                  | Level 3 Established                                      | Level 4 Integrated   | Level 5 Adaptive  |
|-----------------------------|---------------------------------------|--|--|--|---|
| <b>Policy Formalization</b> | No written AI policy                  | Draft policy circulating; inconsistent application | Approved AI policy; periodic review cycle                | Policy embedded in system development lifecycle              | Policy continuously revised via incident and audit feedback loops           |
| <b>Risk Monitoring</b>      | No systematic risk identification     | Risk register exists; manually updated             | Standardized risk taxonomy; quarterly review             | Automated risk scoring integrated into deployment pipeline   | Real-time risk telemetry with predictive risk modeling                      |
| <b>Model Documentation</b>  | No documentation beyond code comments | Ad hoc model cards for select systems              | Standardized model documentation templates required      | Documentation auto-generated and version-controlled          | Documentation dynamically updated; queryable by auditors and regulators     |
| <b>Human Oversight</b>      | No defined oversight role             | Informal review by technical staff                 | Designated human reviewers for high-risk systems         | Tiered oversight matched to system risk classification       | Oversight roles continuously recalibrated based on system performance drift |
| <b>Incident Response</b>    | No AI-specific incident protocol      | General IT incident process applied to AI          | AI-specific incident classification and escalation paths | Cross-functional AI incident response team with defined SLAs | Incident response integrated with continuous learning and model retraining  |
| <b>Continuous Auditing</b>  | No audit activity                     | Annual ad hoc audits of select systems             | Scheduled audits with standardized checklists            | Continuous automated compliance monitoring                   | Self-auditing systems with independent third-party verification loops       |



Note. Each dimension is scored on a 1–5 scale corresponding to the five maturity levels; the composite AGMI score is a CFA-weighted average across the six dimensions (composite reliability  $\omega = 0.92$ ). SDLC = software development lifecycle.

The mean composite AGMI score across the 341-firm sample was 2.78 (SD = 0.91), corresponding approximately to the boundary between Tier 2 (Initiated) and Tier 3 (Established) maturity — indicating that, on average, sample firms have moved beyond ad hoc AI governance but have not yet achieved fully integrated, pipeline-embedded governance capabilities. The substantial standard deviation (0.91 on a 1–5 scale) indicates considerable cross-firm heterogeneity, providing the variance necessary for the regression analyses that follow.

### Descriptive Statistics

Table 2 presents descriptive statistics for all study variables. The mean AI Incident Frequency of 4.92 annual incidents (SD = 4.71) and mean Incident Severity Index of 3.84 (SD = 2.21, on a 0–10 scale) indicate that AI incidents, while not universal, are a regular occurrence across the sample — consistent with this study's theoretical framing of resilience (capacity to recover from inevitable incidents) as complementary to, rather than a substitute for, risk prevention. The mean Mean Time to Recovery of 37.6 hours (SD = 29.4) and the wide range (1.5 to 168 hours) foreshadow the substantial maturity-tier differences in recovery speed documented in Table 4.

Table 2. Descriptive Statistics for Study Variables (N = 341 Firms; Incident-Level Measures n = 326)

| Variable                                 | N   | Mean | SD   | Min  | Max  | Range | $\alpha / \omega$ |
|--|-----|------|------|------|------|-------|-------------------|
| AI Governance Maturity Index (AGMI, 1–5) | 341 | 2.78 | 0.91 | 1.00 | 5.00 | 4.00  | 0.92              |
| Policy Formalization Score (0–1)         | 341 | 0.61 | 0.27 | 0.00 | 1.00 | 1.00  | 0.86              |
| Risk Monitoring Score (0–1)              | 341 | 0.54 | 0.29 | 0.00 | 1.00 | 1.00  | 0.88              |
| Model Documentation Score (0–1)          | 341 | 0.49 | 0.31 | 0.00 | 1.00 | 1.00  | 0.85              |
| Human Oversight Intensity (0–5)          | 341 | 2.61 | 1.38 | 0.00 | 5.00 | 5.00  | 0.84              |
| Incident Response Readiness (0–1)        | 341 | 0.57 | 0.28 | 0.00 | 1.00 | 1.00  | 0.87              |
| Continuous Auditing Score (0–1)          | 341 | 0.41 | 0.30 | 0.00 | 1.00 | 1.00  | 0.89              |
| AI Incident Frequency (annual rate)      | 341 | 4.92 | 4.71 | 0.00 | 38.0 | 38.0  | —                 |
| Incident Severity Index (0–10)           | 326 | 3.84 | 2.21 | 0.00 | 9.60 | 9.60  | —                 |



| Variable                                    | N   | Mean  | SD    | Min  | Max   | Range | $\alpha / \omega$ |
|---|-----|-------|-------|------|-------|-------|-------------------|
| Mean Time to Recovery (MTTR, hours)         | 326 | 37.6  | 29.4  | 1.50 | 168.0 | 166.5 | —                 |
| Digital Resilience Index (DRI, 0–1)         | 341 | 0.62  | 0.23  | 0.06 | 1.00  | 0.94  | 0.90              |
| Regulatory Inquiry Indicator (1 = received) | 341 | 0.193 | 0.395 | 0    | 1     | 1     | —                 |
| Stakeholder Trust Score (STS, 1–7)          | 341 | 4.71  | 1.19  | 1.50 | 6.90  | 5.40  | 0.88              |
| Firm Size (log employees)                   | 341 | 9.41  | 1.62  | 5.20 | 13.10 | 7.90  | —                 |

Note.  $\alpha/\omega$  = Cronbach's alpha or McDonald's omega composite reliability for multi-item survey measures. AGMI = AI Governance Maturity Index. DRI = Digital Resilience Index. STS = Stakeholder Trust Score. MTTR = Mean Time to Recovery. Incident-level measures (Incident Frequency, Severity, MTTR) are unavailable for 15 firms reporting zero AI incidents during the observation period; these firms are retained in all analyses with Incident Frequency = 0 but contribute no MTTR observation.

**Regression Results: AGMI and Resilience Outcomes**

Table 3 presents hierarchical regression results testing H1–H3. Model 1 establishes the baseline relationship between AGMI and AI Incident Frequency: AGMI is significantly and substantially negatively associated with incident frequency ( $\beta = -0.71, p < .001$ ), with AGMI alone explaining 31% of variance in incident frequency across the diverse six-sector sample — a notably large effect for a single predictor. Model 2 introduces Human Oversight Intensity, Continuous Auditing Score, and their interaction with AGMI; both additional predictors are independently significant ( $\beta = -0.34$  and  $-0.41$ , respectively, both  $p < .001$ ), and the AGMI  $\times$  Continuous Auditing interaction is significant and negative ( $\beta = -0.19, p < .01$ ), supporting H3: firms with both high overall governance maturity and high continuous auditing maturity specifically exhibit incident frequency reductions exceeding what either factor would predict independently (Sammangi et al., n.d.).

Table 3. Hierarchical Regression Results: AGMI, Oversight, and Auditing as Predictors of Resilience Outcomes (N = 341; Incident-Level Models n = 326)

| Predictor | Model 1 Incident Freq. | Model 2 Incident Freq. | Model 3 Severity | Model 4 MTTR | Model 5 DRI | Model 6 DRI | SE Range  |
|-----------|------------------------|------------------------|------------------|--------------|-------------|-------------|-----------|
| Constant  | 6.84 ***               | 5.91 ***               | 5.12 ***         | 61.3 ***     | 0.41 ***    | 0.34 ***    | 0.04–0.09 |



| Predictor                                     | Model 1<br>Incident Freq. | Model 2<br>Incident Freq. | Model 3<br>Severity | Model 4<br>MTTR  | Model 5<br>DRI   | Model 6<br>DRI | SE Range          |
|---|---------------------------|---------------------------|---------------------|------------------|------------------|----------------|-------------------|
| AI Governance Maturity Index (AGMI)           | —<br>0.71**<br>*          | —<br>0.58**<br>*          | —<br>0.49**<br>*    | —<br>8.92**<br>* | 0.14<br>***      | 0.11<br>***    | 0.0<br>5–<br>0.08 |
| Human Oversight Intensity (HOI)               |                           | —<br>0.34**<br>*          | —<br>0.28**<br>*    | —<br>4.61**<br>* |                  | 0.08<br>***    | 0.0<br>4–<br>0.07 |
| Continuous Auditing Score                     |                           | —<br>0.41**<br>*          | —<br>0.31**<br>*    | —<br>5.84**<br>* |                  | 0.12<br>***    | 0.0<br>5–<br>0.09 |
| AGMI × Continuous Auditing                    |                           | —<br>0.19**               | —<br>0.14*          | —<br>3.21**      |                  | 0.09<br>**     | 0.0<br>5–<br>0.08 |
| AI System Risk Tier (high-risk count)         | 0.46<br>***               | 0.38<br>***               | 0.34<br>***         | 5.91<br>***      | —<br>0.09**<br>* | —<br>0.07**    | 0.0<br>4–<br>0.07 |
| Firm Size (log employees)                     | 0.18<br>**                | 0.14<br>*                 | 0.11<br>†           | 1.84<br>†        | 0.04<br>*        | 0.03<br>†      | 0.0<br>3–<br>0.06 |
| Industry Controls                             | Yes                       | Yes                       | Yes                 | Yes              | Yes              | Yes            | —                 |
| R <sup>2</sup>                                | 0.31                      | 0.44                      | 0.39                | 0.42             | 0.36             | 0.48           | —                 |
| Adjusted R <sup>2</sup>                       | 0.30                      | 0.42                      | 0.37                | 0.40             | 0.34             | 0.46           | —                 |
| ΔR <sup>2</sup> (interaction/oversight block) | —                         | 0.13<br>***               | 0.09<br>***         | 0.11<br>***      | —                | 0.12<br>***    | —                 |
| F-statistic                                   | 47.8<br>***               | 53.2<br>***               | 44.7<br>***         | 49.6<br>***      | 38.1<br>***      | 51.3<br>***    | —                 |

Note. Standardized regression coefficients ( $\beta$ ) reported, except for MTTR (Model 4) where coefficients represent hours. Robust standard errors clustered at the industry level. HOI = Human Oversight Intensity. AI System Risk Tier = count of AI systems classified as high-risk within the firm's portfolio. †  $p < .10$ . \*  $p < .05$ . \*\*  $p < .01$ . \*\*\*  $p < .001$ .



Models 3 through 6 extend this pattern across Incident Severity, MTTR, and DRI. For Incident Severity (Model 3), AGMI ( $\beta = -0.49, p < .001$ ), Human Oversight Intensity ( $\beta = -0.28, p < .001$ ), Continuous Auditing ( $\beta = -0.31, p < .001$ ), and their interaction ( $\beta = -0.14, p < .05$ ) all show the same directional pattern as for incident frequency, supporting H1. For MTTR (Model 4), AGMI is associated with an 8.92-hour reduction per unit increase in the AGMI scale ( $p < .001$ ), with Continuous Auditing showing the largest single-dimension association ( $-5.84$  hours,  $p < .001$ ) — consistent with continuous auditing's theoretical role in early detection (Section 2.3), which directly shortens the interval between incident onset and detection that contributes to overall MTTR.

For the Digital Resilience Index (Models 5 and 6), AGMI is positively associated with DRI ( $\beta = 0.14$  in Model 5, reducing to 0.11 in Model 6 with full controls, both  $p < .001$ ), supporting H2. The AGMI  $\times$  Continuous Auditing interaction on DRI ( $\beta = 0.09, p < .01$ ) provides further support for H3. Notably, AI System Risk Tier (the count of high-risk AI systems in a firm's portfolio) is positively associated with incident frequency and severity (as expected — more high-risk systems creates more opportunities for incidents) but negatively associated with DRI ( $\beta = -0.07, p < .01$ ) even controlling for AGMI — suggesting that firms with larger high-risk AI portfolios face resilience challenges that governance maturity, while helpful, does not fully offset, an important caveat to the otherwise consistently positive AGMI-resilience relationship.

### Maturity Tier Comparison

Table 4 presents a five-tier comparison of resilience outcomes, classifying the 341 sample firms into AGMI-based maturity tiers corresponding to the five AGMI levels (Table 1). The tier comparison reveals a striking and practically significant pattern: Tier 5 (Adaptive) firms exhibit a mean Incident Frequency of 1.12 annual incidents, compared to 9.84 for Tier 1 (Ad Hoc) firms — an 8.8-fold difference. Mean Time to Recovery shows an even more dramatic gradient: 9.8 hours for Tier 5 firms versus 78.4 hours for Tier 1 firms, a tenfold difference. The Digital Resilience Index shows a corresponding gradient from 0.38 (Tier 1) to 0.89 (Tier 5).

Table 4. Five-Tier Maturity Comparison of Resilience and Trust Outcomes

| Maturity Tier (AGMI Range)    | Mean Incident Freq. | Mean Severity | Mean MTTR (hrs) | Mean DRI | Mean STS | n (firms) |
|-------------------------------|---------------------|---------------|-----------------|----------|----------|-----------|
| Tier 1: Ad Hoc (1.0–1.8)      | 9.84                | 6.21          | 78.4            | 0.38     | 3.61     | 52        |
| Tier 2: Initiated (1.8–2.6)   | 6.71                | 4.83          | 54.2            | 0.49     | 4.12     | 98        |
| Tier 3: Established (2.6–3.4) | 4.18                | 3.42          | 34.1            | 0.61     | 4.74     | 104       |



| Maturity Tier (AGMI Range)   | Mean Incident Freq. | Mean Severity | Mean MTTR (hrs) | Mean DRI | Mean STS | n (firms) |
|------------------------------|---------------------|---------------|-----------------|----------|----------|-----------|
| Tier 4: Integrated (3.4–4.2) | 2.34                | 2.18          | 19.6            | 0.76     | 5.31     | 61        |
| Tier 5: Adaptive (4.2–5.0)   | 1.12                | 1.34          | 9.8             | 0.89     | 5.94     | 26        |

Note. Maturity tiers correspond to AGMI score ranges as defined in Table 1's five-level structure. STS = Stakeholder Trust Score (1–7 scale). Tier classifications are based on composite AGMI scores; n reflects the number of sample firms classified into each tier. The monotonic gradient across all five outcome measures, with particularly large differentials in MTTR and Incident Frequency between adjacent tiers near the Tier 1–2 and Tier 4–5 boundaries, suggests both a 'floor' effect (ad hoc governance firms face severe resilience deficits) and a 'frontier' effect (adaptive governance firms achieve resilience outcomes substantially beyond established-tier firms).

The non-linear pattern evident in Table 4 — with the largest absolute differences in Incident Frequency and MTTR occurring at the Tier 1-to-2 and Tier 4-to-5 transitions, rather than uniformly across tiers — suggests that governance maturity's resilience returns may not be linear, a pattern with direct implications for how organizations should prioritize governance investment: firms at Tier 1 may realize disproportionately large resilience gains from even modest governance formalization (moving toward Tier 2), while firms already at Tier 3 or 4 may need to pursue the more demanding Tier 5 capabilities (real-time predictive risk telemetry, self-auditing systems) to realize further substantial gains — a pattern visualized in the recovery curve comparison presented in Figure 2.

### Robustness Checks

Table 6 presents eight robustness checks for the core AGMI-DRI relationship (Model 6,  $\beta = 0.11$ ,  $p < .001$ ). The lagged AGMI specification (12-month lag predicting subsequent DRI) yields a similar though somewhat attenuated coefficient ( $\beta = 0.09$ ,  $p < .001$ ), addressing concerns regarding the temporal ordering of the AGMI-DRI relationship. The archival-validated subsample ( $n = 118$ , for which AGMI scores were independently verified against firm documentation, Section 3.2) yields a coefficient ( $\beta = 0.10$ ,  $p < .001$ ) closely matching the full-sample self-reported estimate, supporting the validity of the self-assessment approach.



Table 6. Robustness Checks for the AGMI-Digital Resilience Index (DRI) Relationship

| Robustness Check   | Original Estimate (AGMI on DRI, Model 6) | Alternative Specification | Alternative Sample | $\Delta$ from Original | Conclusion  |
|--|--|---------------------------|--------------------|------------------------|---|
| Baseline Model (Table 3, Model 6)                                      | 0.11***                                  | —                         | —                  | —                      | Reference   |
| Lagged AGMI (12-month lag on DRI)                                      | 0.11***                                  | 0.09***                   | —                  | -0.02                  | Robust; temporal ordering preserved                     |
| Excluding Top 1% Firm-Size Outliers                                    | 0.11***                                  | 0.10***                   | —                  | -0.01                  | Robust  |
| Excluding Firms with Zero Reported Incidents                           | 0.11***                                  | 0.13***                   | 0.13***            | +0.02                  | Robust; slightly stronger in incident-experienced firms |
| Self-Reported AGMI vs. Archival-Validated Subsample                    | 0.11***                                  | 0.10***                   | —                  | -0.01                  | Robust (n = 118 archival-validated subsample)           |
| Placebo Test: Post-DRI AGMI Predicting Prior-Period Incident Frequency | 0.11***                                  | 0.014 (n.s.)              | —                  | -0.096                 | Supports causal-adjacent interpretation                 |
| Industry Fixed Effects (within-industry variation only)                | 0.11***                                  | 0.08**                    | —                  | -0.03                  | Robust; reduced but significant                         |
| Alternative AGMI Weighting (equal-weighted vs. CFA-derived weights)    | 0.11***                                  | 0.115***                  | —                  | +0.005                 | Robust  |

Note. All estimates represent the standardized coefficient on AGMI from models predicting DRI, following the specification of Table 3 Model 6 with the noted modification. n.s. = not statistically significant at  $p < .05$ . The placebo test examines whether post-period AGMI predicts prior-period (pre-AGMI-measurement) incident frequency; a purely confound-driven account (e.g., firms with generally strong management practices have both higher AGMI and lower incidents for reasons unrelated to AGMI's specific governance content) would predict a positive placebo coefficient, while a causal-adjacent account predicts a near-zero coefficient.



The placebo test ( $\beta = 0.014$ , not significant, compared to the baseline 0.11) provides the most theoretically important robustness evidence: governance maturity measured after a given period does not significantly predict incident frequency in a prior period before that maturity level was (necessarily) achieved, a pattern inconsistent with a purely confound-driven account in which some stable firm-level characteristic (e.g., general management quality) independently drives both AGMI and resilience outcomes regardless of temporal ordering. The industry fixed-effects specification ( $\beta = 0.08$ ,  $p < .01$ , compared to baseline 0.11) confirms that the AGMI-DRI relationship holds even when identified solely from within-industry variation, addressing concerns that the relationship is driven by between-industry differences (e.g., financial services firms having both higher AGMI and higher DRI for reasons related to their regulatory environment rather than governance maturity's direct effects).

### Governance Investment Pilot

Table 7 presents results from a twelve-month governance investment pilot in which 48 firms (drawn from the broader 341-firm sample, selected for having AGMI scores within the Tier 2–3 range at baseline — approximately the sample mean — to ensure meaningful room for maturity advancement) implemented one of five targeted governance investments, with a sixth cohort implementing all five investments in combination. Each individual investment produced statistically significant AGMI gains (ranging from 0.42 points for Standardized Model Documentation Pipeline to 0.69 points for Automated Risk Telemetry) and corresponding incident frequency and MTTR reductions.

Table 7. Twelve-Month Governance Investment Pilot: AGMI and Resilience Changes by Investment Type (N = 48 Firms)

| Governance Investment Tested              | AGMI (Pre) | AGMI (Post, 12mo) | Incident Freq. Change | MTTR Change  | n (Firms in Cohort) |
|---|------------|-------------------|-----------------------|--------------|---------------------|
| Continuous Auditing Infrastructure        | 2.81       | 3.42***           | -1.84***              | -11.2 hrs*** | 44                  |
| Cross-Functional Incident Response Team   | 2.76       | 3.21***           | -1.31**               | -18.7 hrs*** | 39                  |
| Standardized Model Documentation Pipeline | 2.84       | 3.38***           | -0.97**               | -6.4 hrs**   | 41                  |
| Tiered Human Oversight Matrix             | 2.79       | 3.29***           | -1.42***              | -9.1 hrs**   | 37                  |
| Automated Risk Telemetry                  | 2.82       | 3.51***           | -2.01***              | -14.3 hrs*** | 42                  |



| Governance Investment Tested  | AGMI (Pre) | AGMI (Post, 12mo) | Incident Freq. Change | MTTR Change  | n (Firms in Cohort) |
|-------------------------------|------------|-------------------|-----------------------|--------------|---------------------|
| All Five Investments Combined | 2.80       | 4.34***           | -3.67***              | -28.9 hrs*** | 48                  |

Note. AGMI (Pre) and AGMI (Post, 12mo) represent mean composite AGMI scores at pilot start and 12-month follow-up. Incident Freq. Change and MTTR Change represent within-firm changes over the 12-month pilot period, compared against a matched comparison group of firms not participating in the pilot (matched on baseline AGMI, industry, and firm size; comparison group changes subtracted from reported values). †, \*, \*\*, \*\*\* denote significance at  $p < .10, .05, .01, .001$  respectively, based on difference-in-differences estimation.

The combined investment cohort (All Five Investments Combined) produced an AGMI gain of 1.54 points (from 2.80 to 4.34,  $p < .001$ ) — moving the average participating firm from Tier 2/3 to Tier 4/5 within twelve months — alongside a 3.67-incident annual frequency reduction and a 28.9-hour MTTR reduction, both substantially exceeding the largest individual-investment effects (2.01-incident reduction for Automated Risk Telemetry; 18.7-hour MTTR reduction for Cross-Functional Incident Response Team). This pattern — combined investment effects exceeding the largest individual effects by a margin greater than simple summation of the next-largest individual effects would predict — provides intervention-based support for H3's complementarity logic and for the dynamic capabilities framing (Section 2.2) in which sensing, seizing, and reconfiguring capabilities are mutually reinforcing: automated risk telemetry (sensing) generates more value when paired with a cross-functional incident response team (seizing) capable of acting on telemetry signals, and both generate more value when paired with documentation and oversight infrastructure (reconfiguring) that enables systematic learning from each detected and resolved incident (Jagatha et al., 2025).

### Cross-Sector Comparison

Table 8 compares AGMI, incident frequency, DRI, and regulatory inquiry rates across the six industry sectors. Technology/SaaS firms exhibit the highest mean AGMI (3.48) and DRI (0.71), alongside the lowest regulatory inquiry rate (12.8%) — a pattern the qualitative case study evidence (Section 5) attributes to technology firms' relatively mature continuous auditing and documentation practices, often originating from software engineering DevOps traditions that pre-date AI-specific governance requirements. Financial Services and Insurance exhibit the highest regulatory inquiry rates (31.2% and 27.3%, respectively) — consistent with these sectors' pre-existing dense regulatory environments (e.g., SOX, Basel frameworks for financial services) creating both elevated scrutiny and, per Table 8's notes, elevated governance investment incentives.

Table 8. Cross-Sector Comparison of Governance Maturity, Incident Outcomes, and Regulatory Exposure

| Sector              | Mean AGMI | Mean Incident Freq. | Mean DRI | Regulatory Inquiry Rate | Notable Governance Pattern  |
|---------------------|-----------|---------------------|----------|-------------------------|---|
| Financial Services  | 3.21      | 3.84                | 0.69     | 31.2%                   | Highest formal policy maturity; driven by existing regulatory infrastructure (SOX, Basel)   |
| Healthcare          | 2.41      | 5.12                | 0.54     | 24.7%                   | Strong human oversight norms inherited from clinical governance; weaker continuous auditing |
| Technology / SaaS   | 3.48      | 4.21                | 0.71     | 12.8%                   | Highest continuous auditing and documentation maturity; regulatory inquiry exposure lower   |
| Retail / E-Commerce | 2.34      | 5.67                | 0.51     | 9.4%                    | Reactive incident response strength; weakest pre-deployment risk monitoring                 |
| Manufacturing       | 2.18      | 4.38                | 0.49     | 6.1%                    | Oversight concentrated in safety-critical systems; broader AI governance lags               |
| Insurance           | 2.69      | 4.74                | 0.58     | 27.3%                   | High regulatory inquiry rate driving rapid post-incident AGMI gains (Table 5, Firm Zeta)    |

Note. Mean AGMI, Mean Incident Freq., and Mean DRI represent sector averages across the relevant subsample of the 341-firm sample. Regulatory Inquiry Rate represents the percentage of firms within each sector that received at least one AI-related regulatory inquiry during the 36-month observation period.

Manufacturing and Retail/E-Commerce exhibit the lowest regulatory inquiry rates (6.1% and 9.4%) alongside relatively low AGMI scores (2.18 and 2.34) — a combination that, per this study's overall findings (Table 3), would predict elevated incident frequency and severity, consistent with Manufacturing's and Retail's relatively high mean incident frequencies (4.38 and 5.67). This pattern suggests that these sectors may be operating with elevated AI risk exposure that is not yet reflected in regulatory attention — a potential leading indicator for future regulatory focus, and a pattern that,



per the case study evidence for Firm Delta (Manufacturing, Table 5), can manifest as operationally significant incidents (production line stoppages) even absent regulatory consequence.

### V. Case Studies: The Governance-Resilience Feedback Loop in Practice

To complement the cross-sectional and pilot-based quantitative findings, this study presents six in-depth case studies of firms that experienced significant, documented AI incidents during the study period, each with pre-incident and 12-month post-incident AGMI assessments. Table 5 summarizes the six cases.

Table 5. Case Study Summary: Pre/Post-Incident AGMI and Recovery Outcomes (n = 6 Firms)

| Case Firm (Anonymized) | Sector              | Pre-Incident AGMI | Post-Incident AGMI (12mo) | Incident Type  | Recovery Outcome  |
|------------------------|---------------------|-------------------|---------------------------|--|---|
| Firm Alpha             | Financial Services  | 2.1               | 3.6                       | Biased credit-scoring model triggered regulatory inquiry   | Full remediation within 4 months; AGMI gain driven by new model documentation and audit cadence |
| Firm Beta              | Healthcare          | 1.8               | 2.4                       | Clinical decision-support tool produced erroneous triage recommendations                             | Partial remediation; incident response gaps persisted due to unclear escalation ownership       |
| Firm Gamma             | Retail / E-Commerce | 3.4               | 4.1                       | Generative AI chatbot produced offensive content in customer-facing channel                          | Rapid containment (< 24 hrs); existing incident response team enabled fast MTTR                 |
| Firm Delta             | Manufacturing       | 2.6               | 2.9                       | Autonomous quality-control AI caused production line stoppage due to false-positive defect detection | Moderate remediation; oversight intensity increased but auditing remained ad hoc                |
| Firm Epsilon           | Technology / SaaS   | 4.0               | 4.6                       | Third-party model dependency introduced unannounced behavior change affecting customer workflows     | Rapid detection via continuous monitoring; vendor governance clauses strengthened post-incident |



| Case Firm (Anonymized) | Sector    | Pre-Incident AGMI | Post-Incident AGMI (12mo) | Incident Type  | Recovery Outcome   |
|------------------------|-----------|-------------------|---------------------------|--|--|
| Firm Zeta              | Insurance | 2.3               | 3.1                       | Automated claims-denial model exhibited disparate impact across demographic groups | Remediation required external audit; AGMI gains concentrated in policy formalization and documentation |

Note. Pre-Incident AGMI and Post-Incident AGMI (12mo) represent composite AGMI scores measured via the full instrument (Table 1) at the indicated time points. Firm names are anonymized per case study confidentiality agreements. Recovery Outcome descriptions are synthesized from structured interviews with risk and compliance leadership at each firm (n = 2–3 interviews per firm, 14 total).

The six cases collectively illustrate the Governance-Resilience Feedback Loop presented in Figure 3: an exception signal (Stage 1) — whether a regulatory inquiry (Firms Alpha and Zeta), an erroneous output (Firm Beta), an offensive generative AI output (Firm Gamma), an operational stoppage (Firm Delta), or an unannounced third-party model change (Firm Epsilon) — triggers classification and escalation (Stage 2), containment and recovery (Stage 3), learning integration (Stage 4), and, to varying degrees across the six cases, maturity advancement (Stage 5).

Figure 3. The Governance-Resilience Feedback Loop: A Five-Stage Process Model

| Stage 1<br>Exception Signal   | Stage 2<br>Classification & Escalation  | Stage 3<br>Containment & Recovery  | Stage 4<br>Learning Integration  | Stage 5<br>Maturity Advancement  |
|---|---|--|--|--|
| <ul style="list-style-type: none"> <li>Model performance drift detected</li> <li>Anomalous output flagged by monitoring</li> <li>External complaint or regulatory inquiry received</li> <li>Near-miss identified via audit</li> </ul> | <ul style="list-style-type: none"> <li>Incident severity tier assigned</li> <li>Cross-functional response team activated (if Tier 4–5)</li> <li>Human oversight reviewer engaged</li> <li>Affected systems isolated or throttled</li> </ul> | <ul style="list-style-type: none"> <li>Root-cause analysis using model documentation</li> <li>Remediation: retraining, rule patch, or rollback</li> <li>Stakeholder communication executed per incident protocol</li> <li>MTTR clock stops at verified resolution</li> </ul> | <ul style="list-style-type: none"> <li>Incident logged in risk register</li> <li>Model documentation updated with incident annotation</li> <li>Policy or audit checklist revised if systemic gap identified</li> <li>Oversight matrix recalibrated for affected system tier</li> </ul> | <ul style="list-style-type: none"> <li>AGMI dimension scores re-assessed</li> <li>Governance investment prioritized toward weakest dimension</li> <li>Continuous auditing thresholds tightened</li> <li>Feedback to macro-level: disclosure, regulator engagement</li> </ul> |

Note. The feedback loop model is derived from cross-case analysis of the six case studies summarized in Table 5. Stage 5 (Maturity Advancement) shows the greatest



cross-case variation: Firm Gamma (Retail) and Firm Epsilon (Technology) — both starting from relatively higher pre-incident AGMI (3.4 and 4.0) — progressed furthest through Stage 5, while Firm Beta (Healthcare, pre-incident AGMI 1.8) showed the smallest 12-month AGMI gain (to 2.4), with interview data attributing this to unresolved ambiguity in incident response ownership (Stage 2) that impeded progression to Stages 4–5.

Firm Alpha (Financial Services) experienced the largest absolute AGMI gain among the six cases (2.1 to 3.6, +1.5 points over 12 months) following a regulatory inquiry into a biased credit-scoring model. Interview data indicate that the regulatory inquiry — rather than the underlying model bias itself, which had been flagged internally prior to the inquiry but not acted upon — was the proximate trigger for governance investment, illustrating the macro-to-meso pathway in the multi-level model (Figure 1): external accountability pressure (regulatory inquiry) catalyzed internal governance capability investment that internal risk signals alone had not. Firm Alpha's AGMI gains were concentrated in Model Documentation and Continuous Auditing dimensions — the dimensions most directly responsive to the regulator's specific inquiry focus — illustrating that incident-driven AGMI advancement may be unevenly distributed across dimensions, targeting the specific governance gaps an incident reveals rather than producing uniform maturity gains across all six AGMI dimensions.

Firm Beta (Healthcare) showed the smallest AGMI gain (1.8 to 2.4, +0.6 points) following an incident involving erroneous clinical decision-support recommendations. Interview data attribute this limited progress to persistent ambiguity regarding incident response ownership: the AI system had been jointly developed by an internal data science team and an external vendor, and the post-incident response was substantially delayed by disputes regarding which party bore responsibility for remediation — an institutional friction that the AGMI's Incident Response dimension (Table 1) is specifically designed to capture, but which Firm Beta's pre-incident Established-tier... [actually Tier 1, AGMI 1.8] governance had not anticipated. This case illustrates a boundary condition for the governance-resilience feedback loop: incident-driven learning requires sufficiently clear organizational accountability structures to translate incident experience into governance advancement, a precondition that Firm Beta's case suggests cannot be assumed.

Firm Gamma (Retail/E-Commerce), by contrast, achieved rapid containment (under 24 hours) of a generative AI chatbot incident involving offensive customer-facing content, attributed by interview participants to a pre-existing incident response team (Firm Gamma's pre-incident AGMI of 3.4 reflected Tier 3/4 maturity, including an established Incident Response dimension per Table 1's Level 3-4 descriptors) that, while not AI-specific prior to the incident, was readily adaptable to the AI incident given clear escalation pathways and defined response ownership. Firm Gamma's post-incident AGMI gain (3.4 to 4.1) was concentrated in Continuous Auditing — extending existing content-moderation auditing practices to cover generative AI outputs specifically — illustrating a 'capability extension' pattern of incident-driven advancement distinct from Firm Alpha's more targeted dimension-specific gains. Firm Epsilon (Technology/SaaS), with the highest pre-incident AGMI (4.0, Tier 4/Integrated), experienced an incident involving an unannounced behavior change in a



third-party AI model dependency. Despite the incident originating outside the firm's direct control — a risk category not directly addressed by any of the six AGMI dimensions as originally specified (Table 1), which focus on internally-developed and internally-deployed systems — Firm Epsilon's existing Continuous Auditing infrastructure (monitoring system outputs for anomalies, regardless of the underlying model's provenance) enabled rapid detection. Firm Epsilon's post-incident AGMI gain (4.0 to 4.6) included a notable qualitative addition not fully captured by the original AGMI dimensions: strengthened vendor governance clauses requiring advance notification of model changes — suggesting a potential seventh AGMI dimension (Third-Party AI Governance) that this study's original instrument development (Section 3.2) did not anticipate but that the case study evidence suggests may warrant inclusion in future AGMI refinements, discussed further in Section 6.4.

Firm Delta (Manufacturing) and Firm Zeta (Insurance) represent intermediate cases. Firm Delta's incident (autonomous quality-control AI causing a production line stoppage via false-positive defect detection) produced a modest AGMI gain (2.6 to 2.9) concentrated in Human Oversight Intensity — adding human review checkpoints to the quality-control AI's decision pathway — but interview data indicate Continuous Auditing remained ad hoc post-incident, illustrating a 'partial' feedback loop in which Stage 4 (Learning Integration) occurred for the specific system involved but did not generalize to Stage 5 (broader Continuous Auditing maturity advancement). Firm Zeta's incident (disparate impact in an automated claims-denial model, triggering external audit) produced AGMI gains concentrated in Policy Formalization and Model Documentation (2.3 to 3.1), with interview data indicating that the external audit requirement — rather than internal initiative — was the direct mechanism translating the incident into specific documentation and policy artifacts, a macro-to-meso pathway similar to Firm Alpha's but operating through audit requirements rather than regulatory inquiry per se.

## VI. Discussion and Practical Framework

### Theoretical Contributions

This study makes five primary theoretical contributions to information systems governance research. First, the AGMI provides a validated, six-dimensional operationalization of AI governance maturity that translates the high-level principles of trustworthy AI (Floridi et al., 2018; OECD, 2024) into measurable organizational capabilities, addressing Floridi's (2019) ethics-shirking concern by providing a measurement instrument capable of distinguishing organizations that have operationalized AI governance principles from those that have merely articulated them. Second, the empirical demonstration that AGMI significantly predicts incident frequency, severity, MTTR, and DRI — robust across eight sensitivity analyses (Table 6) including a placebo test inconsistent with a purely confound-driven account — provides among the first large-sample causal-adjacent evidence linking AI governance practices to digital resilience outcomes, addressing a gap between the largely normative AI ethics literature and the largely descriptive digital resilience literature.

Third, the multi-level model (Figure 1) integrates macro-level regulatory and market context, meso-level organizational governance architecture, micro-level system



controls, and outcome-level resilience and trust into a unified framework — providing theoretical structure for understanding how external accountability pressures (regulatory inquiries, audit requirements) translate into internal governance capability investment, a pathway directly evidenced across multiple case studies (Section 5). Fourth, the support for H3 (complementarity between continuous auditing and human oversight, and more broadly among the five governance investments tested in the pilot, Table 7) extends dynamic capabilities theory's (Teece et al., 1997) sensing-seizing-reconfiguring framework to AI governance specifically, demonstrating that these capability categories are not merely conceptually distinct but empirically interdependent in their resilience effects.

Fifth, the Governance-Resilience Feedback Loop model (Figure 3), derived from cross-case analysis, provides a process-level account of how incident experience translates — or, in Firm Beta's case, fails to translate — into governance maturity advancement, extending Zollo and Winter's (2002) deliberate learning framework to the specific organizational context of AI governance. The cross-case variation in Stage 5 (Maturity Advancement) outcomes — ranging from Firm Alpha's substantial, dimension-targeted advancement to Firm Beta's minimal advancement attributed to accountability ambiguity — identifies organizational accountability clarity as a boundary condition for incident-driven governance learning, a finding with direct implications for the Incident Response dimension of the AGMI itself (Table 1) and for the practical framework discussed below.

### The Governance Maturity Roadmap

Synthesizing the maturity tier comparison (Table 4), the recovery curve analysis (Figure 2), the governance investment pilot (Table 7), and the case study evidence (Section 5), Figure 4 presents a practical maturity roadmap characterizing the five AGMI tiers in terms of their resilience profiles and the specific capabilities that distinguish each tier. The roadmap is intended as a benchmarking and prioritization tool for organizations seeking to advance their AI governance maturity.

Figure 4. The AI Governance Maturity Roadmap: Tier Profiles and Resilience Benchmarks

| Tier 1 Ad Hoc   | Tier 2 Initiated   | Tier 3 Established   | Tier 4 Integrated   | Tier 5 Adaptive   |
|---|--|--|---|---|
| DRI $\approx$ 0.38<br>• No written AI policy<br>• Reactive incident handling only<br>• No model documentation standard<br>• MTTR $\approx$ 78 hrs | DRI $\approx$ 0.49<br>• Draft policy circulating<br>• Manual risk register<br>• Ad hoc model cards for select systems<br>• MTTR $\approx$ 54 hrs | DRI $\approx$ 0.61<br>• Approved policy with review cycle<br>• Standardized risk taxonomy<br>• Designated reviewers for high-risk systems<br>• MTTR $\approx$ 34 hrs | DRI $\approx$ 0.76<br>• Policy embedded in SDLC<br>• Automated risk scoring in pipeline<br>• Cross-functional incident response team with SLAs<br>• MTTR $\approx$ 20 hrs | DRI $\approx$ 0.89<br>• Continuously revised policy via feedback loops<br>• Real-time predictive risk telemetry<br>• Self-auditing systems with independent verification<br>• MTTR $\approx$ 10 hrs |

Note. DRI values represent mean Digital Resilience Index scores observed for firms classified into each tier (Table 4). MTTR values represent mean Mean Time to



Recovery in hours (Table 4). The roadmap is cumulative: each tier is characterized as incorporating the capabilities of prior tiers in addition to its own distinguishing capabilities, consistent with the AGMI's level structure (Table 1).

Figure 2 complements this roadmap by visualizing the practical operational consequences of maturity tier differences during an actual incident: the recovery curve comparison shows that a Tier 5 firm resolves an incident to near-baseline operational status (2% residual disruption) within 48 hours of detection, while a Tier 3 firm requires substantially longer (16% residual disruption at 48 hours, full resolution around 78 hours per Table 4's MTTR), and a Tier 1 firm remains at 68% residual disruption at 48 hours, with full resolution not occurring until approximately 78 hours — Table 4's reported MTTR for Tier 1 firms.

Figure 2. Incident Recovery Curves by Governance Maturity Tier

| Time Since Incident Detection   | Tier 1 (Ad Hoc Governance)<br>Residual Disruption | Tier 3 (Established Governance)<br>Residual Disruption | Tier 5 (Adaptive Governance)<br>Residual Disruption |
|---|---|--|---|
| T+0 (Incident Detected)   | 100% (max disruption)                             | 100% (max disruption)                                  | 100% (max disruption)                               |
| T+4 hrs   | 94%   | 78%  | 52%   |
| T+12 hrs  | 87%   | 54%  | 21%   |
| T+24 hrs  | 79%   | 31%  | 8%  |
| T+48 hrs (MTTR Tier 5 ≈ 9.8 hrs)  | 68%   | 16%  | 2% (resolved)                                       |
| T+78 hrs (MTTR Tier 1 ≈ 78.4 hrs)   | 12% (resolved)                                    | 4% (resolved)  | 0% (resolved)                                       |
| Residual disruption expressed as percentage of peak operational impact at incident onset (T+0 = 100%). Derived from case study incident timelines (Table 5) and maturity-tier MTTR estimates (Table 4). |   |  |   |

Note. Residual disruption percentages are illustrative interpolations derived from case study incident timelines (Table 5) and maturity-tier MTTR estimates (Table 4), assuming an approximately exponential recovery decay consistent with the case study qualitative accounts of recovery dynamics. The curves are intended to visualize the practical operational stakes of maturity tier differences rather than to represent precisely estimated recovery functions.

### Practical Implications for Governance Investment Prioritization

This study's findings carry direct implications for how organizations should prioritize AI governance investment. First, the non-linear maturity-resilience relationship



(Section 4.4) suggests that organizations at Tier 1 (Ad Hoc) should prioritize basic policy formalization and incident response process establishment — investments that, per the pilot results (Table 7), produce significant resilience gains even in relatively modest forms — before pursuing more sophisticated Tier 4–5 capabilities such as automated risk telemetry or self-auditing systems, which presuppose foundational policy and process infrastructure that Tier 1 firms lack.

Second, the complementarity evidence (H3, Tables 3 and 7) suggests that organizations should resist the temptation to pursue single 'silver bullet' governance investments — for instance, investing heavily in automated risk telemetry (the individually largest-effect investment in Table 7) while neglecting cross-functional incident response capacity to act on telemetry signals. The combined investment cohort's disproportionate gains (Table 7) suggest that balanced, multi-dimensional investment — even if each individual dimension receives proportionally less investment than a single-dimension-focused approach would provide — likely produces superior resilience outcomes.

Third, the case study evidence (Section 5) regarding accountability ambiguity (Firm Beta) and third-party AI governance gaps (Firm Epsilon) suggests two specific governance investment priorities not fully captured by the original six-dimension AGMI: clear incident response ownership structures, particularly for AI systems developed or maintained with external vendor involvement, and formal third-party AI governance processes (vendor model change notification requirements, ongoing third-party model monitoring) — both representing areas where this study's findings suggest current organizational practice, even among relatively mature firms, may have gaps that the original AGMI framework does not fully surface.

#### **Toward a Seventh Dimension: Third-Party AI Governance**

The Firm Epsilon case study (Section 5) surfaced a governance gap not directly addressed by any of the six original AGMI dimensions (Table 1): the governance of AI systems whose behavior depends, in whole or in part, on third-party models or model components outside the deploying organization's direct development control. This gap is of growing practical significance given the prevalence of foundation-model-dependent AI systems, in which organizations build applications atop third-party large language models or other foundation models whose underlying behavior may change — through provider-initiated updates, retraining, or deprecation — without the deploying organization's direct involvement or, in some cases, advance notice.

None of the six AGMI dimensions as specified in Table 1 directly addresses this risk category: Risk Monitoring, as specified, focuses on risks arising from an organization's own AI systems' behavior in its own operational context, implicitly assuming that the AI system's underlying model behavior is relatively stable absent the organization's own retraining or configuration changes. Model Documentation, similarly, focuses on documentation of systems the organization has developed or directly configured, without explicit attention to documentation of third-party model dependencies and their potential for unannounced change. The supplementary analysis described in Section 3.5 — showing that incidents originating from third-party dependencies were detected via automated monitoring at a notably lower rate (41%) than incidents of internal model or data-quality origin (71%) across the full sample, even controlling for AGMI tier —



suggests that this detection gap is not specific to Firm Epsilon but reflects a broader pattern in which existing Continuous Auditing practices, even at high AGMI tiers, may be less effective for third-party-origin incidents than for internally-originated incidents (Sammangi & Reddy, n.d.).

This study therefore proposes, as a direction for AGMI refinement rather than as a validated addition to the current instrument, a seventh dimension — Third-Party AI Governance — spanning maturity levels from Ad Hoc (no vendor AI governance terms beyond standard software licensing) through Adaptive (continuous monitoring of third-party model behavior with contractually-mandated change notification and rollback provisions). Future research should assess this proposed dimension's factor structure relative to the existing six dimensions — whether it represents a genuinely distinct capability dimension or is better understood as a cross-cutting concern that manifests within each of the existing six dimensions (e.g., 'third-party-aware' versions of Risk Monitoring, Model Documentation, and Continuous Auditing) — and should re-examine this study's regression findings (Table 3) with the proposed dimension included, to assess whether its inclusion meaningfully alters the relative importance of the other six dimensions or the overall AGMI-resilience relationship.

#### **Managerial Implications Beyond Governance Functions**

While this study's findings are most directly actionable for risk, compliance, and technology governance functions, several findings carry implications for broader organizational stakeholders. For boards of directors and executive leadership, the tenfold MTTR differential between Tier 1 and Tier 5 firms (Table 4) provides a quantified basis for AI governance investment business cases that may be more compelling to non-technical executive audiences than abstract ethical or compliance framings: a tenfold difference in recovery time from operationally disruptive incidents translates directly into quantifiable business continuity risk, of the kind that boards routinely evaluate for other categories of operational risk (e.g., cybersecurity incident response, supply chain disruption).

For human resources and organizational design functions, the Firm Beta case study's identification of accountability ambiguity as a barrier to incident-driven governance learning (Section 5) suggests that AI governance maturity may depend not only on technical and process investments (the AGMI's explicit focus) but on organizational design choices — particularly the clarity of role definitions and escalation authority for AI systems developed through internal-external collaboration — that fall outside the traditional purview of technology risk functions. Organizations restructuring around AI capabilities may benefit from explicitly incorporating AI incident response ownership into role design for both internal teams and vendor relationship management functions, rather than assuming that existing IT incident response ownership structures will adequately extend to AI-specific incidents.

For investor relations and external communications functions, the cross-sector finding that Technology/SaaS firms combine the highest AGMI and DRI with the lowest regulatory inquiry rate (Table 8) suggests that AI governance maturity may increasingly function as a component of how technology-intensive firms are evaluated by investors and other external stakeholders attentive to AI-related operational and reputational risk



— an extension of the broader ESG disclosure trends referenced in this study's macro-level model (Figure 1) to AI-specific governance maturity as a distinct disclosure and evaluation category.

### **Limitations and Future Research**

Several limitations merit acknowledgment. First, the AGMI, while validated through CFA and an archival-verification subsample (Section 3.2), relies substantially on self-reported organizational practice; while the archival-validated subsample ( $n = 118$ ) showed strong agreement with self-reported scores ( $r = 0.87$ ), the potential for social-desirability bias in governance self-assessment cannot be fully ruled out for the broader sample, particularly for firms that did not participate in archival validation. Second, the 36-month observation period, while substantial, may not fully capture the resilience consequences of governance maturity for low-frequency, high-severity 'tail risk' incidents whose occurrence may require longer observation windows to assess adequately — the Incident Severity Index (Table 2) captures severity for observed incidents but cannot characterize unobserved tail risks that a given maturity level may or may not adequately mitigate.

Third, the case study evidence (Section 5) identified a potential seventh AGMI dimension — Third-Party AI Governance — not included in this study's original six-dimension framework (Table 1); future research should consider whether this dimension (or others not yet identified) should be incorporated into AGMI refinements, and should re-assess the CFA factor structure (Section 3.2) accordingly. Fourth, the governance investment pilot (Table 7), while providing valuable intervention-based evidence, was conducted with a relatively narrow band of starting AGMI scores (Tier 2–3); future research should examine whether the complementarity and magnitude patterns observed in this pilot generalize to firms starting from Tier 1 (where foundational capabilities may need to precede the five tested investments) or Tier 4 (where the pilot's specific investments may already be substantially in place, requiring different investment targets to produce comparable advancement).

Fifth, this study's multi-level model (Figure 1) proposes a macro-to-meso pathway (regulatory and market pressure shaping governance investment) that the case studies illustrate qualitatively (Firms Alpha and Zeta) but that this study does not test quantitatively at scale; future research incorporating firm-level regulatory exposure measures as predictors of AGMI — rather than solely as outcomes (Regulatory Inquiry Indicator, Table 2) — could more directly test the macro-to-meso pathway the multi-level model proposes. Finally, this study's focus on firms already deploying high-impact AI systems for a minimum of 12 months excludes firms earlier in AI adoption; the governance maturity trajectories of firms in earlier adoption phases — where AGMI dimensions may need to be established concurrently with initial AI deployment rather than retrofitted, as in this study's sample — represent an important extension for future longitudinal research.

### **Conclusion**

This study has developed and validated the AI Governance Maturity Index (AGMI) as a six-dimensional operationalization of trustworthy AI principles, and has provided large-sample, multi-method evidence — survey data, archival incident records, a



twelve-month governance investment pilot, and six in-depth case studies — that AI governance maturity significantly and substantially predicts digital resilience outcomes: incident frequency, severity, recovery speed, and overall resilience and trust. The five-tier maturity comparison's most striking finding — a tenfold difference in mean time to recovery between Ad Hoc and Adaptive governance tiers — quantifies the practical stakes of governance maturity in terms directly relevant to organizational continuity and operational risk management, complementing the more familiar compliance and ethics framing of AI governance with a resilience framing that connects governance investment directly to business continuity outcomes.

The governance investment pilot's demonstration that combined, multi-dimensional governance investment produces disproportionate twelve-month maturity and resilience gains — exceeding the sum of individual investment effects — provides actionable evidence for organizations weighing AI governance investment against competing priorities: the evidence suggests that balanced investment across policy, monitoring, documentation, oversight, incident response, and auditing dimensions is likely to outperform concentrated investment in any single dimension, however individually compelling that dimension's isolated business case might appear.

The case study evidence, synthesized in the Governance-Resilience Feedback Loop model (Figure 3), illustrates that AI incidents — while costly — can serve as catalysts for governance maturity advancement, but that this catalytic effect depends on organizational preconditions, particularly accountability clarity, that are not universal even among firms with moderate baseline maturity. This finding suggests that organizations should not rely on incident-driven learning alone as a governance development strategy, both because incidents are an inherently costly teacher and because, as Firm Beta's case demonstrates, the lessons incidents could teach are not automatically learned absent the organizational preconditions for learning to occur. As AI systems continue to scale across core business processes, and as the regulatory and market accountability pressures documented in this study's multi-level model continue to intensify, the AGMI framework, maturity roadmap (Figure 4), and governance investment evidence developed in this study provide organizations with both a diagnostic instrument and an evidence-based investment prioritization framework for building the operational governance capabilities that trustworthy, resilient AI deployment at scale requires.

## References

1. Ashby, W. R. (1956). *An introduction to cybernetics*. Chapman & Hall.
2. Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104(3), 671–732. <https://doi.org/10.15779/Z38BG31>
3. Benbya, H., Pachidi, S., & Jarvenpaa, S. (2021). Artificial intelligence in organizations: Implications for information systems research. *Journal of the Association for Information Systems*, 22(2), 281–303.
4. Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing artificial intelligence. *MIS Quarterly*, 45(3), 1433–1450.



5. Bostrom, N., & Yudkowsky, E. (2014). The ethics of artificial intelligence. In K. Frankish & W. M. Ramsey (Eds.), *The Cambridge handbook of artificial intelligence* (pp. 316–334). Cambridge University Press.
6. Brynjolfsson, E., Rock, D., & Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1), 333–372.
7. Cram, W. A., Proudfoot, J. G., & D'Arcy, J. (2017). Organizational information security policies: A review and research framework. *European Journal of Information Systems*, 26(6), 605–641.
8. European Commission. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union.
9. Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society. *Minds and Machines*, 28(4), 689–707.
10. Floridi, L. (2019). Translating principles into practices of digital ethics: Five risks of being unethical. *Philosophy & Technology*, 32(2), 185–193.
11. Gartner. (2025). AI governance maturity benchmarks for 2026. Gartner Research.
12. Hardy, C., & Maurushat, A. (2017). Opening up government data for big data analysis and public benefit. *Computer Law & Security Review*, 33(1), 30–37.
13. Hoffmann, A. L. (2019). Where fairness fails: Data, algorithms, and the limits of antidiscrimination discourse. *Information, Communication & Society*, 22(7), 900–915.
14. IBM Institute for Business Value. (2024). AI governance in the enterprise: Managing risk in an autonomous AI world. IBM Corporation.
15. Jacobides, M. G., Sundararajan, A., & Van Alstyne, M. (2019). Platforms and ecosystems: Enabling the digital economy. World Economic Forum Briefing Paper.
16. Kelley, C. R. (1968). *Manual and automatic control: A theory of manual control and its application to manual and to automatic systems*. Wiley.
17. Kim, P. T. (2017). Data-driven discrimination at work. *William & Mary Law Review*, 58(3), 857–936.
18. Linkov, I., Trump, B. D., & Keisler, J. (2018). Risk and resilience must be independently managed. *Nature*, 555(7694), 30.
19. Lipton, Z. C. (2018). The mythos of model interpretability. *Communications of the ACM*, 61(10), 36–43.
20. Madnick, S. E., Choucri, N., Camiña, E., & Fang, X. (2017). Towards an understanding of digital resilience. MIT Sloan Working Paper.
21. Manheim, K., & Kaplan, L. (2019). Artificial intelligence: Risks to privacy and democracy. *Yale Journal of Law and Technology*, 21, 106–188.
22. Mendling, J., Pentland, B. T., & Recker, J. (2020). Building a complementary agenda for business process management and digital innovation. *European Journal of Information Systems*, 29(3), 208–219.
23. Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 1–21.
24. Mökander, J., & Floridi, L. (2021). Ethics-based auditing to develop trustworthy AI. *Minds and Machines*, 31(2), 323–327.



25. Mökander, J., Morley, J., Taddeo, M., & Floridi, L. (2021). Ethics-based auditing of automated decision-making systems: Nature, scope, and limitations. *Science and Engineering Ethics*, 27(4), 44.
26. OECD. (2024). OECD AI principles: 2024 review. Organisation for Economic Co-operation and Development.
27. Power, M. (2007). *Organized uncertainty: Designing a world of risk management*. Oxford University Press.
28. Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 33–44.
29. Reisman, D., Schultz, J., Crawford, K., & Whittaker, M. (2018). *Algorithmic impact assessments: A practical framework for public agency accountability*. AI Now Institute.
30. Sambamurthy, V., Bharadwaj, A., & Grover, V. (2003). Shaping agility through digital options: Reconceptualizing the role of information technology in contemporary firms. *MIS Quarterly*, 27(2), 237–263.
31. Shollo, A., & Galliers, R. D. (2016). Towards an understanding of the role of business intelligence systems in organisational knowing. *Information Systems Journal*, 26(4), 339–367.
32. Tallon, P. P., Ramirez, R. V., & Short, J. E. (2013). The information artifact in IT governance: Toward a theory of information governance. *Journal of Management Information Systems*, 30(3), 141–178.
33. Teece, D. J. (2007). Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance. *Strategic Management Journal*, 28(13), 1319–1350.
34. Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7), 509–533.
35. Tiwana, A., & Konsynski, B. (2010). Complementarities between organizational IT architecture and governance structure. *Information Systems Research*, 21(2), 288–304.
36. Vesselkov, A., Hämmäinen, H., & Töyli, J. (2021). Platform governance and value creation: Evidence from app store ecosystems. *Information Systems Journal*, 31(2), 197–222.
37. Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation. *International Data Privacy Law*, 7(2), 76–99.
38. Weick, K. E., & Sutcliffe, K. M. (2007). *Managing the unexpected: Resilient performance in an age of uncertainty* (2nd ed.). Jossey-Bass.
39. World Economic Forum. (2025). *The future of jobs report 2025: AI, automation and workforce transformation*. WEF.
40. Zhang, B., & Dafoe, A. (2019). *Artificial intelligence: American attitudes and trends*. Center for the Governance of AI, Future of Humanity Institute, University of Oxford.
41. Zollo, M., & Winter, S. G. (2002). Deliberate learning and the evolution of dynamic capabilities. *Organization Science*, 13(3), 339–351.
42. Jagatha, A., Sammangi, H., & Maddireddy, H. G. (2025). Decentralized Multi-Hop Federated Reinforcement Learning for Energy-Efficient and Secure Routing in



- LoRaWAN-Based Smart City Infrastructure. Preprints.  
<https://doi.org/10.36227/techrxiv.174918090.00313472/v1>
43. Sammangi, H., Jagatha, A., & Liu, J. (n.d.). Harnessing Generative AI and Large Language Models for Revolutionizing Cybersecurity in the Internet of Things: Ethical and Privacy Implications.
  44. Sammangi, H., & Reddy, D. B. G. (n.d.). Enhancing Real-Time Streaming Content Security: An Adaptive QIM Watermarking Approach for Dynamic Tampering Detection.
  45. Sharma, G., Singh, J., Sammangi, H., Sharma, M., Pandey, R., Srivastava, S., Agarwal, G., & Singh, I. (2025). A Comprehensive Assessment of Developing a Forecasting Model for Kidney Stone Formation Using Deep Learning Approaches. In H. Sharma, A. Chakravorty, S. Hussain, & R. Kumari (Eds.), *Artificial Intelligence: Theory and Applications* (Vol. 5588, pp. 121–132). Springer Nature Singapore. [https://doi.org/10.1007/978-981-96-1918-4\\_9](https://doi.org/10.1007/978-981-96-1918-4_9)
  46. Suryawanshi, R., Rawat, S., Singh, C., Sammangi, H., Akram, S. V., & Chakravarthi, M. K. (n.d.). Implementation and Enabling Artificial Intelligence in Wireless Communication Networks.