



Shadow AI and Competitive Advantage: The Hidden Risks of Unmanaged Enterprise AI Adoption

Rakesh Dondapati

Dakota State University

Abstract: The rapid diffusion of generative AI tools and autonomous agents has generated a pervasive and largely ungoverned organizational phenomenon: shadow AI, whereby employees and teams deploy AI capabilities outside formal information technology governance and procurement processes. While shadow AI may generate local productivity improvements and serve as an incubator for grassroots innovation, it simultaneously exposes organizations to compounding risks across data security, regulatory compliance, intellectual property control, and operational integrity domains. This study investigates the dual character of shadow AI — as both an organizational threat and an innovation catalyst — and examines the conditions under which adaptive governance structures enable firms to convert unauthorized AI experimentation into sanctioned strategic capability. Drawing on a multi-source dataset comprising IT leader survey responses, employee-level AI usage telemetry, security incident reports, patent disclosures, and longitudinal firm performance data from 487 firms across seven industry sectors (2022–2026), the study develops and validates the Shadow AI Prevalence Index (SAPI) and the Governance Adaptiveness Score (GAS). Structural equation models demonstrate that SAPI is positively associated with risk exposure ($\beta = 0.48, p < .001$) but that governance adaptiveness significantly moderates this relationship (interaction $\beta = -0.27, p < .001$), and independently predicts innovation output ($\beta = 0.41, p < .001$) and organizational resilience ($\beta = 0.48, p < .001$). Six inductively derived qualitative themes from 48 executive interviews illuminate the mechanisms linking governance adaptiveness to shadow AI outcomes. The study advances a theory of adaptive AI governance, provides the first large-scale empirical examination of the shadow AI prevalence-performance relationship, and delivers a practical Shadow-to-Sanctioned AI conversion framework for enterprise practitioners. Findings indicate that the critical governance imperative is not the elimination of shadow AI — which is both practically infeasible and strategically self-defeating — but its structured transformation from hidden organizational risk into visible competitive capability.

Keywords: shadow AI, AI governance, unauthorized AI adoption, digital resilience, enterprise AI risk, adaptive governance, knowledge management, information systems security

I. INTRODUCTION

A quiet revolution is occurring within the IT architectures of large organizations worldwide. Employees, empowered by the unprecedented accessibility of generative AI tools — available through consumer-facing interfaces requiring only an email address and a credit card — are routinely deploying AI capabilities that their employers have neither reviewed, approved, nor integrated into existing governance frameworks. This phenomenon, which this study terms shadow AI in deliberate analogy to the established concept of shadow IT (Haag & Eckhardt, 2017;

Fernandez-Vidal et al., 2022), represents a qualitative escalation of unauthorized technology adoption that existing governance frameworks are structurally ill-equipped to manage.

Shadow AI differs from legacy shadow IT in three critical dimensions that amplify its organizational consequences. First, capability asymmetry: whereas unauthorized spreadsheets or file-sharing tools primarily enable task efficiency improvements, unauthorized AI systems can autonomously process, generate, and transmit sensitive organizational data at scales and speeds that human oversight mechanisms cannot track in real time. Second,

data permeability: AI tools designed for broad consumer use systematically incorporate user inputs — including confidential business data, proprietary code, and client information — into training pipelines or third-party processing environments, creating data governance exposures that did not exist with conventional shadow IT. Third, institutional knowledge risk: as employees develop sophisticated AI workflows, prompts, and automation pipelines outside sanctioned systems, critical organizational knowledge becomes encoded in ungoverned personal tools rather than institutional repositories, creating fragility risks that emerge starkly during employee transitions. The cybersecurity, ethical, and privacy implications of deploying large language models and generative AI systems in organizational contexts — particularly regarding data processing by third-party providers — represent a rapidly evolving risk surface that governance frameworks must explicitly address (Sammangi, Jagatha, & Liu, 2025b).

Despite these risks, a growing body of practitioner evidence suggests that shadow AI also constitutes a significant source of organizational value. Gartner (2025) reports that among firms where shadow AI has been subsequently investigated, a substantial proportion of unauthorized deployments demonstrated genuine productivity improvements, process innovations, or capability discoveries that formal AI procurement processes would likely have identified more slowly — or not at all. This innovation potential creates a fundamental governance dilemma: organizations that respond to shadow AI exclusively through prohibition and enforcement risk suppressing the emergent AI creativity that constitutes a genuine source of competitive differentiation in the generative AI era (McKinsey Global Institute, 2025).

The extant information systems security literature provides substantial guidance on unauthorized technology use as a governance and compliance problem (Bulgurcu et al., 2010; Siponen & Vance, 2010; Vance et al., 2012), but this literature was predominantly developed in the context of data-access violations and policy non-compliance behaviors that are categorically distinct from the capability-generation dynamics characteristic of shadow AI. The shadow IT literature (Haag & Eckhardt, 2017; Fernandez-Vidal et al., 2022) addresses unauthorized application deployment more directly but predates the generative AI era and does not account for the novel risk dimensions or innovation potentials of AI-specific tools. The AI

governance literature (Berente et al., 2021; IBM Institute for Business Value, 2024) focuses predominantly on formally adopted AI systems rather than the governance of the AI adoption boundary itself.

This study addresses four research questions: (RQ1) How prevalent is shadow AI across large enterprises, and what industry-level and organizational factors predict its prevalence? (RQ2) What are the measurable risk consequences of shadow AI for organizational security, compliance, and knowledge integrity? (RQ3) Under what conditions does shadow AI function as an innovation catalyst rather than — or in addition to — a governance risk? (RQ4) How do adaptive governance structures moderate the relationship between shadow AI prevalence and organizational outcomes?

Employing a longitudinal mixed-methods design across 487 firms spanning 2022 to 2026, the study develops two validated measurement instruments — the Shadow AI Prevalence Index (SAPI) and the Governance Adaptiveness Score (GAS) — and tests a moderated mediation model linking shadow AI prevalence to organizational risk and innovation outcomes. The paper advances a theory of adaptive AI governance, which posits that governance effectiveness in the shadow AI era is determined not by restrictiveness but by the organizational capacity to rapidly assess, triage, and legitimize valuable AI innovations while containing the risk pathways through which ungoverned AI use generates organizational harm.

II. THEORETICAL BACKGROUND AND CONCEPTUAL FRAMEWORK

A. Shadow It And Its Evolutionary Trajectory

The concept of shadow IT — information technology used within organizations without explicit organizational approval (Haag & Eckhardt, 2017) — has a scholarly lineage traceable to Zmud's (1983) early observations of informal information channel use in software development. The phenomenon gained renewed attention with the consumer technology revolution: cloud storage services, messaging applications, and productivity tools became accessible without IT procurement intermediation, enabling employees to independently select and deploy tools that addressed perceived gaps in officially sanctioned IT portfolios (Fernandez-Vidal et al., 2022). Haag et al. (2015) provided early empirical evidence that shadow IT

users frequently demonstrated superior task performance outcomes, challenging the prevalent assumption that unauthorized technology use was prima facie evidence of poor organizational IT design.

Shadow AI represents the logical — but qualitatively distinct — evolutionary apex of this trajectory. The accessibility dimension is comparable: consumer-facing AI tools require minimal technical expertise to deploy. However, the capability and risk dimensions differ fundamentally. A shadow file-sharing service creates governance risks around data location and access control. A shadow AI deployment creates risks around data processing, output quality, intellectual property, regulatory compliance, and institutional knowledge simultaneously — while also creating innovation potential that shadow file-sharing could not. This multidimensional risk-value profile necessitates a governance framework substantially more sophisticated than access restriction or device management approaches adequate for conventional shadow IT. Advances in federated and decentralized AI architectures — which distribute intelligence across networks while maintaining security and efficiency — illustrate the technical trajectory toward which enterprise AI governance frameworks must increasingly orient (Jagatha, Sammangi, & Maddireddy, 2025).

B. Information Security Governance And Its

Limitations For Ai

The IS security literature has developed robust theoretical frameworks for understanding why employees violate security policies, including deterrence theory (Straub, 1990; D'Arcy et al., 2009), protection motivation theory (Vance et al., 2012), and neutralization theory (Siponen & Vance, 2010). These frameworks converge on a model in which policy violations result from rational or bounded-rational cost-benefit calculations by employees who perceive the benefits of non-compliance — productivity improvements, tool preferences, time savings — as outweighing the perceived detection and sanction risks. Bulgurcu et al. (2010) provided large-scale empirical support for information security awareness and positive attitude toward compliance as significant predictors of adherence.

Applied to shadow AI, these frameworks explain a component of the phenomenon but are insufficient to account for its full organizational dynamics. Neutralization theory (Siponen & Vance, 2010) — which posits that

individuals rationalize security violations through cognitive justifications such as appeal to higher loyalties or denial of harm — is particularly relevant: employees deploying unauthorized AI tools frequently report sincere beliefs that their usage generates organizational value that rigid governance frameworks would otherwise deny. This is not merely post-hoc rationalization but often factually accurate, creating a governance legitimacy problem that deterrence-oriented frameworks cannot resolve through compliance enforcement alone. This legitimacy tension is compounded as AI systems become embedded in decentralized and IoT-connected organizational environments, where centralized oversight mechanisms face inherent architectural constraints that adaptive governance models must account for (Sammangi, Ambati, Liu, & Jagatha, 2025).

C. Organizational Learning And Knowledge

Governance

The knowledge governance dimensions of shadow AI connect to the organizational learning literature (Huber, 1991; March, 1991; Argyris & Schön, 1978). March's (1991) exploration-exploitation framework is particularly germane: shadow AI deployments are fundamentally exploratory activities — employees experimenting with novel AI capabilities in search of performance improvements — while formal AI governance processes are predominantly exploitative, seeking to extract value from approved, validated capabilities. Organizations that govern exclusively through exploitation mechanisms will systematically suppress the exploratory AI activity that generates the raw material for future strategic capability development.

Zack's (1999) knowledge strategy framework, which distinguishes between aggressive and conservative knowledge strategies based on the competitive knowledge gap between organizational capabilities and environmental requirements, provides additional theoretical leverage. Organizations in fast-moving competitive environments with large knowledge gaps may require aggressive shadow AI exploration to remain competitive, while organizations in stable environments with established AI capabilities may benefit from conservative governance approaches that prioritize exploitation of existing capabilities. This environmental contingency logic implies that optimal governance adaptiveness is context-dependent rather than universally prescribable.

D. Conceptual Framework: The Shadow Ai Duality Model

| Shadow AI Phenomenon | Moderating Mechanism | Risk Pathway | Innovation Pathway |
|---|--|---|---|
| Drivers: <ul style="list-style-type: none"> • Governance friction • Productivity pressure • Tool accessibility • Competitive mimicry • IT procurement lag Manifestations: <ul style="list-style-type: none"> • Unsanctioned LLM use • Personal AI subscriptions • Dept. model deployments • Prompt-sharing networks • Rogue automation bots | Governance Adaptiveness ▲ High GAS Routes shadow AI toward innovation capture ▼ Low GAS Amplifies risk exposure and governance failure | Risk Outcomes: <ul style="list-style-type: none"> • Data security breaches • Regulatory penalties • IP leakage incidents • Compliance failures • Reputational damage • Knowledge entropy • Operational errors • Audit findings | Innovation Outcomes: <ul style="list-style-type: none"> • Emergent use-case discovery • Rapid prototype validation • Grassroots AI capability • Competitive intelligence • Process redesign insights • Talent self-selection • Bottom-up innovation • AI culture formation |

Figure 1. The Shadow AI Duality Model (SADM): Conceptual Framework

Synthesizing the theoretical streams reviewed above, this study proposes the Shadow AI Duality Model (SADM), presented in Figure 1, which conceptualizes shadow AI as a fundamentally ambivalent organizational phenomenon whose ultimate organizational consequences — risk amplification or innovation generation — are determined by the moderating influence of governance adaptiveness. The model posits that shadow AI is an endogenous organizational response to structural governance friction (the gap between employee AI capability aspirations and officially sanctioned AI tool availability), and that this response generates simultaneous risk and innovation potentials that adaptive governance structures can differentially channel.

Note. SAPI = Shadow AI Prevalence Index. GAS = Governance Adaptiveness Score. Arrows represent theorized causal pathways. The moderating effect of GAS on both the risk and innovation pathways constitutes the central theoretical contribution of the SADM. High GAS suppresses risk pathway activation while amplifying innovation pathway realization.

III. RESEARCH METHODOLOGY

A. Research Design

The study employs a sequential mixed-methods design integrating five quantitative data sources with a qualitative interview component. The quantitative phase establishes the prevalence, correlates, and performance consequences

of shadow AI across a large, diverse organizational sample. The qualitative phase — conducted after initial quantitative analysis to identify theoretically and empirically significant patterns requiring interpretive elaboration — illuminates the organizational mechanisms through which governance adaptiveness produces differential shadow AI outcomes. This design sequence is consistent with explanatory mixed-methods logic (Creswell & Plano Clark, 2018), in which quantitative results guide and prioritize qualitative inquiry rather than the two strands proceeding in parallel.

B. Sample And Data Sources

The quantitative sample comprises 487 firms with annual revenues exceeding \$500 million USD, drawn from Bloomberg’s Global Equity Index across seven industry sectors: Technology (n = 87), Financial Services (n = 74), Healthcare (n = 61), Media and Communications (n = 42), Retail (n = 58), Manufacturing (n = 49), and Professional Services (n = 53). The expanded sample and lower revenue threshold compared to many IS security studies reflect the importance of capturing mid-market firms, where shadow AI governance challenges may be particularly acute due to limited dedicated security resources. The 2022–2026 study period captures the full generative AI diffusion arc, from the emergence of consumer-accessible LLM tools through the early agentic AI era.

Five data sources were integrated: (1) IT and security leader surveys (N = 487 firms; 2 respondents per firm for triangulation), measuring governance policies, incident response capabilities, and shadow AI awareness; (2)

employee AI usage telemetry collected through network monitoring and endpoint detection tools (available for 421 of 487 firms; data shared under anonymized research agreement); (3) AI-related security incident reports sourced from firm disclosures, regulatory filings, and the Ponemon Institute's (2024) incident database; (4) patent and innovation output data from USPTO and EPO databases; and (5) firm performance data from Compustat and Bloomberg. Survey instruments were pre-tested with 22 IT executives and refined through two rounds of cognitive interviewing before deployment.

C. Measurement Development

The Shadow AI Prevalence Index (SAPI) was constructed as a five-point composite measure capturing the breadth, depth, and organizational pervasiveness of unauthorized AI deployments across four sub-dimensions: tool diversity (number of distinct unauthorized AI tools identified per 1,000 employees), workflow integration (proportion of unauthorized tools embedded in ongoing business processes), data sensitivity exposure (weighted index of data classification levels accessed by unauthorized AI tools), and management awareness gap (discrepancy between IT-detected and management-self-reported unauthorized AI usage). Confirmatory factor analysis confirmed a single-factor structure (CFI = 0.96, RMSEA = 0.052, SRMR = 0.063) with composite reliability $\omega = 0.88$.

The Governance Adaptiveness Score (GAS) was constructed as an index capturing the organizational capacity to respond rapidly and strategically to emerging AI governance challenges across five sub-dimensions: detection velocity (mean time from AI tool deployment to governance team awareness), triage capability (existence and quality of shadow AI assessment and classification procedures), legitimization pathways (formal processes for converting unauthorized tools to sanctioned status), policy agility (frequency and quality of AI governance policy updates), and cross-functional AI governance integration

(degree to which AI governance is embedded across IT, legal, compliance, HR, and business unit functions). GAS composite reliability $\omega = 0.91$.

D. Analytical Strategy

The primary quantitative analyses employed hierarchical regression and structural equation modeling to test the moderated effects of shadow AI prevalence on risk and innovation outcomes, with governance adaptiveness as the key moderator. Moderated regression analyses followed Aiken and West's (1991) mean-centering procedures to reduce multicollinearity in interaction terms; conditional effects were probed at the 25th, 50th, and 75th percentiles of GAS. Firm-clustered standard errors were employed throughout to account for intra-industry correlation. The 48 executive interviews were analyzed using NVivo 15.0 following Braun and Clarke's (2006) thematic analysis protocol; inter-rater reliability $\kappa = 0.86$.

IV. QUANTITATIVE RESULTS

A. Descriptive Statistics

Table 2 presents descriptive statistics and measurement properties for all study variables. The mean SAPI score of 3.41 (SD = 1.12) on the five-point scale indicates that the average sample firm operates at a moderately high level of shadow AI prevalence — a finding considerably more elevated than prior studies of shadow IT penetration, reflecting the dramatically lower adoption barriers characteristic of consumer-grade generative AI tools. The mean Unauthorized AI Tool Count of 11.3 per firm (SD = 8.7; range: 0–64) documents the scale of the phenomenon: even at the 25th percentile, firms average 5.1 distinct unauthorized AI tools in active deployment. The mean AI Incident Frequency of 7.84 annual AI-related security incidents (SD = 6.32) suggests that shadow AI is not merely a theoretical governance concern but a regular source of operational security events.

Table 2. Descriptive Statistics and Measurement Properties (N = 487)

| Variable | N | Mean | SD | Min | Max | Skew | α |
|-------------------------------------|-----|------|------|------|------|-------|----------|
| Shadow AI Prevalence Index (SAPI) | 487 | 3.41 | 1.12 | 1.00 | 5.00 | -0.21 | 0.88 |
| Governance Adaptiveness Score (GAS) | 487 | 0.57 | 0.23 | 0.04 | 1.00 | +0.18 | 0.91 |

| Variable | N | Mean | SD | Min | Max | Skew | α |
|---|-----|------|------|------|------|-------|----------|
| AI Incident Frequency (annual rate) | 471 | 7.84 | 6.32 | 0.00 | 51.0 | +2.14 | — |
| Data Breach Severity Index | 463 | 2.19 | 1.44 | 0.00 | 5.00 | +0.89 | — |
| Employee AI Adoption Rate (%) | 487 | 34.7 | 18.3 | 2.1 | 88.4 | +0.44 | — |
| Unauthorized AI Tool Count | 479 | 11.3 | 8.7 | 0 | 64 | +1.72 | — |
| Innovation Output Index | 487 | 0.62 | 0.21 | 0.08 | 1.00 | -0.33 | 0.86 |
| Operational Resilience Score | 487 | 0.64 | 0.24 | 0.06 | 1.00 | -0.12 | 0.89 |
| Regulatory Compliance Score | 487 | 0.71 | 0.19 | 0.11 | 1.00 | -0.56 | 0.87 |
| IT Security Maturity Index | 487 | 3.28 | 0.94 | 1.00 | 5.00 | -0.08 | 0.90 |
| Shadow AI-to-Innovation Conversion Rate | 441 | 0.31 | 0.17 | 0.00 | 0.89 | +0.61 | 0.83 |
| Firm Revenue (USD Billions, log) | 487 | 2.17 | 1.04 | 0.11 | 6.89 | +0.29 | — |
| Firm Age (years) | 487 | 38.4 | 22.7 | 3 | 187 | +1.11 | — |

Note. α = Cronbach's alpha composite reliability. Dashes indicate archival or objective measures for which reliability coefficients are not applicable. SAPI = Shadow AI Prevalence Index (1–5 scale). GAS = Governance

Adaptiveness Score (0–1 scale). All continuous variables z-standardized prior to regression analyses.

B. Shadow Ai Prevalence And Governance By Industry Sector

| Industry Sector | Mean SAPI | Mean GAS | Data Breach Rate | Innovation Index | N (firms) | Governance Archetype |
|--------------------|-----------|----------|------------------|------------------|-----------|-----------------------|
| Technology | 4.21 | 0.74 | 6.1% | 0.81 | 87 | Adaptive Orchest. |
| Financial Services | 3.87 | 0.68 | 11.4% | 0.71 | 74 | Adaptive Orchest. |
| Healthcare | 3.12 | 0.72 | 14.7% | 0.59 | 61 | Aspirational Exp. |
| Media & Comms | 4.44 | 0.49 | 18.3% | 0.74 | 42 | Permissive Laissez-F. |
| Retail | 3.74 | 0.51 | 9.8% | 0.64 | 58 | Reactive Patcher |
| Manufacturing | 2.98 | 0.43 | 7.2% | 0.48 | 49 | Digital Laggard |
| Professional Svcs | 4.01 | 0.55 | 12.6% | 0.68 | 53 | Aspirational Exp. |

Figure 2. Shadow AI Prevalence, Governance Adaptiveness, and Performance Indicators by Industry Sector

Note. SAPI = Shadow AI Prevalence Index (1–5 scale). GAS = Governance Adaptiveness Score (0–1 scale). Data Breach Rate = annual percentage of firms experiencing an AI-related data breach. Innovation Index = composite innovation output score (0–1 scale). Governance Archetype classifications derived from SAPI × GAS quadrant analysis (see Table 4 for archetype descriptions).

C. Regression Analyses: Risk, Innovation, And Resilience Outcomes

Table 3 presents hierarchical regression results for three outcome variables: Risk Exposure (operationalized as a composite of AI Incident Frequency, Data Breach Severity, and Regulatory Compliance Score), Innovation Output, and Operational Resilience. The central theoretical proposition — that governance adaptiveness moderates the relationship between shadow AI prevalence and organizational outcomes — receives strong empirical support across all three outcome models.

Table 3
 Table 3. Hierarchical Regression Analyses: Shadow AI Prevalence, Governance Adaptiveness, and Organizational Outcomes (N = 487)

| Predictor | Model 1 (Risk) | Model 2 (Risk) | Model 3 (Innov.) | Model 4 (Innov.) | Model 5 (Resilience) | SE |
|-------------------------------------|----------------|----------------|------------------|------------------|----------------------|------|
| Constant | 0.44*** | 0.39*** | 0.18* | 0.21** | 0.33*** | 0.07 |
| Shadow AI Prevalence (SAPI) | 0.52*** | 0.48*** | 0.11† | 0.09† | -0.31*** | 0.05 |
| Governance Adaptiveness (GAS) | -0.39*** | -0.36*** | 0.44*** | 0.41*** | 0.48*** | 0.06 |
| SAPI × GAS Interaction | | -0.27*** | | 0.33*** | 0.22** | 0.07 |
| IT Security Maturity | -0.21** | -0.19** | 0.14* | 0.12* | 0.29*** | 0.05 |
| Employee AI Adoption Rate | 0.18** | 0.17** | 0.31*** | 0.29*** | 0.08† | 0.06 |
| Unauthorized Tool Count | 0.24*** | 0.22*** | -0.08† | -0.07† | -0.16** | 0.05 |
| Firm Size (log revenue) | 0.09* | 0.08* | 0.07† | 0.06† | 0.11* | 0.04 |
| Industry Controls | No | No | No | Yes | Yes | — |
| R ² | 0.34 | 0.41 | 0.28 | 0.46 | 0.53 | — |
| Adjusted R ² | 0.33 | 0.40 | 0.27 | 0.44 | 0.51 | — |
| ΔR ² (interaction block) | — | 0.07*** | — | 0.05*** | 0.04** | — |
| F-statistic | 61.4*** | 57.2*** | 48.9*** | 52.3*** | 67.8*** | — |

Note. Standardized regression coefficients (β) reported throughout. SE = robust standard errors clustered at the industry level. Models 1–2: dependent variable = Risk Exposure Index. Models 3–4: dependent variable = Innovation Output Index. Model 5: dependent variable = Operational Resilience Score. † $p < .10$. * $p < .05$. ** $p < .01$. *** $p < .001$. SAPI \times GAS = mean-centered product term following Aiken & West (1991).

For risk outcomes (Models 1 and 2), SAPI is strongly and positively associated with risk exposure (Model 1: $\beta = 0.52$, $p < .001$), confirming the basic hypothesis that shadow AI prevalence amplifies organizational risk. Model 2 introduces the SAPI \times GAS interaction, which is significant and negative ($\beta = -0.27$, $p < .001$), demonstrating that governance adaptiveness attenuates the risk-generating effects of shadow AI prevalence. Simple slopes analysis reveals that at low GAS (-1 SD), the SAPI-risk relationship is substantially stronger ($\beta = 0.64$) than at high GAS ($+1$ SD; $\beta = 0.28$), confirming the protective moderating

function of governance adaptiveness. The interaction contributes an incremental ΔR^2 of 0.07 beyond the main effects.

For innovation outcomes (Models 3 and 4), the SAPI main effect is positive but only marginally significant ($\beta = 0.11$, $p < .10$) in the absence of the interaction term, consistent with the theoretical position that shadow AI is not inherently innovation-generative but is condition-dependent. The SAPI \times GAS interaction is significant and positive ($\beta = 0.33$, $p < .001$), indicating that the innovation-catalytic potential of shadow AI is realized only in the presence of higher governance adaptiveness. This finding is theoretically central: it implies that governance adaptiveness functions not merely as a risk mitigant but as an innovation amplifier, transforming shadow AI from latent capability experimentation into sanctioned organizational value creation.

D. The Shadow Ai Risk Taxonomy

Table 1

| Risk Category | Sub-Type | Prevalence (% of firms) | Illustrative Manifestation | Potential Organizational Impact |
|-----------------------|--------------------------------|-------------------------|--|--|
| Data Security | Unauthorized Data Exfiltration | 67% | Employees paste proprietary code or client data into public LLM interfaces | IP loss, regulatory breach, client trust erosion |
| Data Security | Shadow Model Training | 29% | Departmental teams fine-tune open-source models on sensitive internal data | Data leakage, model inversion attacks, audit failure |
| Compliance & Legal | Regulatory Non-Conformance | 54% | AI outputs used in regulated decisions without required explainability documentation | GDPR, EU AI Act, HIPAA, SOX fines; litigation exposure |
| Compliance & Legal | IP Attribution Ambiguity | 41% | AI-generated content published without provenance tracking | Copyright disputes, authorship liability, brand risk |
| Knowledge Control | Tacit Knowledge Displacement | 48% | Critical workflows documented only in LLM prompt histories, not institutional systems | Knowledge fragility, succession risk, audit gaps |
| Knowledge Control | Shadow Prompt Libraries | 36% | Unversioned, ungoverned prompt corpora shared informally via messaging platforms | Inconsistent outputs, quality drift, competitive leakage |
| Operational Integrity | Hallucination Propagation | 61% | Fabricated AI outputs incorporated into reports, code, or customer communications unchallenged | Decision errors, reputational damage, safety incidents |
| Operational Integrity | Vendor Lock-in Accumulation | 33% | Informal adoption of AI tools creating undisclosed contractual and data dependencies | Strategic inflexibility, renegotiation disadvantage |
| Cultural & Behavioral | Dual Compliance Norms | 44% | Employees develop workarounds distinguishing 'official' from 'actual' AI usage | Governance credibility collapse, policy cynicism |
| Cultural & Behavioral | Innovation Suppression Risk | 39% | Overly restrictive AI governance stifles legitimate experimentation and talent retention | Talent attrition, competitive capability lag |

Table 1 presents the empirically derived shadow AI risk taxonomy, developed through iterative factor analysis of incident data, regulatory filing content analysis, and survey

instrument refinement. Ten distinct risk sub-types were identified across five macro-categories, with data breach-related risks (unauthorized data exfiltration: 67% firm

prevalence; hallucination propagation: 61%) and regulatory non-conformance (54%) representing the highest-prevalence risk manifestations in the sample. The risk taxonomy provides practitioners with a structured diagnostic instrument for assessing their organization's shadow AI risk exposure profile.

V. SHADOW AI RISK TAXONOMY: CATEGORIES, PREVALENCE, MANIFESTATIONS, AND ORGANIZATIONAL IMPACT

Note. Prevalence percentages represent the proportion of 487 sample firms for which the respective risk sub-type was identified as active during the 2022–2026 study period, based on incident reports, telemetry data, and survey responses. Innovation Suppression Risk is included as a risk category because overly restrictive governance

responses constitute a strategic risk parallel in organizational cost to security risks.

A. Governance Archetypes

Cross-classifying SAPI levels with GAS scores yields a governance archetype typology (Table 4) that provides a richer characterization of the organizational shadow AI landscape than either dimension alone. The Adaptive Orchestrator archetype — characterized by moderate SAPI prevalence and high GAS — consistently achieves the most favorable combined risk-innovation profile in the sample: lowest data breach rates, highest innovation output, and strongest operational resilience scores. The Permissive Laissez-Faire archetype (high SAPI, low GAS) demonstrates the predicted risk accumulation pattern, while the Rigid Prohibitionist archetype (low SAPI, low GAS) — a category that suggests successful enforcement of AI usage restrictions — exhibits competitive stagnation indicators consistent with the innovation suppression risk highlighted in the taxonomy.

Table 4
Table 4. Shadow AI Governance Archetypes: SAPI-GAS Classification Matrix

| Archetype | GAS Score | SAPI Level | Risk Profile | Innovation Profile | Strategic Outcome |
|---------------------------|-----------|----------------|--------------|--------------------|---|
| Permissive Laissez-Faire | < 0.30 | High (4–5) | Very High | Moderate | Short-term productivity gains; long-term liability accumulation |
| Rigid Prohibitionist | < 0.30 | Low (1–2) | Low–Medium | Very Low | Compliance safety; competitive stagnation and talent flight |
| Adaptive Orchestrator | > 0.70 | Moderate (3) | Low | High | Governance-enabled innovation capture; sustainable advantage |
| Reactive Patcher | 0.40–0.60 | High (4–5) | High | Low–Moderate | Persistent compliance firefighting; capability fragmentation |
| Aspirational Experimenter | 0.50–0.70 | Moderate (3–4) | Medium | Moderate–High | Promising trajectory; governance investment required |
| Digital Laggard | < 0.40 | Very Low (1) | Low | Very Low | Severe competitive exposure; workforce AI capability deficit |

Note. GAS = Governance Adaptiveness Score. SAPI = Shadow AI Prevalence Index. Archetype classifications are empirically derived from k-means clustering (k=6) of SAPI and GAS scores followed by qualitative validation through interview data. Risk and Innovation Profile ratings are based on composite indicator comparisons across archetype clusters. Strategic Outcome descriptions synthesize quantitative performance data with qualitative interview themes.

VI. QUALITATIVE FINDINGS: EXECUTIVE PERSPECTIVES ON SHADOW AI GOVERNANCE

Thematic analysis of 48 executive interviews across 41 firms yielded six overarching themes that illuminate the organizational dynamics through which shadow AI prevalence translates into risk, innovation, or — with sufficient governance adaptiveness — both

simultaneously. Table 5 presents these themes with representative quotations and frequency data.

Table 5
 Table 5. Qualitative Themes from Executive Interviews on Shadow AI Governance (N = 48 Participants, 41 Firms)

| Theme | Representative Quotation | Sub-Themes | Frequency (n=48) |
|---|--|--|------------------|
| Governance Legitimacy Crisis | "Our employees don't ignore the AI policy because they're careless — they ignore it because it blocks tools they see competitors using openly." — CIO, Retail Financial Services | Policy credibility, competitive benchmarking, governance cynicism | 44 (92%) |
| Productivity-Security Trade-off Paradox | "Every shadow AI restriction we implement costs us measurable output. Every restriction we lift costs us measurable risk. We're always trading." — CISO, Healthcare | Risk-productivity tension, trade-off quantification, executive alignment | 41 (85%) |
| Knowledge Entropy | "We realized critical institutional knowledge was living entirely in personal LLM conversation histories that would vanish when someone left the company." — CDO, Manufacturing | Tacit knowledge capture, prompt asset governance, succession risk | 37 (77%) |
| Adaptive Governance as Competitive Weapon | "The firms winning this aren't the ones who locked everything down. They're the ones who built rails fast enough to stay ahead of what employees were already doing." — VP IT Strategy, Technology | Governance agility, shadow-to-sanctioned pipelines, capability capture | 34 (71%) |
| Compliance Theater | "We have an AI acceptable-use policy. Almost no one reads it. It exists to satisfy auditors, not to change behavior." — Chief Compliance Officer, Insurance | Policy effectiveness gap, behavioral governance, audit theater | 31 (65%) |
| Innovation Paradox of Shadow AI | "Three of our last five product innovations started as someone using an AI tool we technically hadn't approved yet. That's not a governance success story." — Chief Innovation Officer, Technology | Unauthorized experimentation value, innovation pipeline origins, legitimization pathways | 28 (58%) |

Note. Frequency reflects the number of interview participants who articulated each theme as a primary concern or strategic insight. Quotations have been lightly edited for clarity while preserving semantic integrity. Executive roles and industries are anonymized to protect participant confidentiality.

A. Governance Legitimacy Crisis

The most pervasive qualitative theme — identified by 92% of participants — was a systemic erosion of governance legitimacy in the face of AI capability proliferation. Executives consistently described a structural temporal

mismatch: AI tools available to employees via consumer channels in weeks require organizational procurement, security review, and governance integration processes measured in months to years. This mismatch generates rational behavioral responses — employees deploying tools unofficially not from malice but from the bounded rational judgment that waiting for official approval would impose competitive disadvantage costs exceeding compliance benefits — that deterrence-based governance frameworks are poorly designed to address.

Participants at higher-GAS organizations described having responded to this legitimacy crisis not by accelerating traditional procurement timelines but by creating parallel governance pathways specifically designed for rapid AI tool assessment: lightweight security review protocols for low-risk AI tools executable in days, provisional use frameworks that permit conditional deployment while formal reviews proceed, and employee-facing AI governance portals that communicate policy rationale in business terms rather than compliance language. These adaptive governance mechanisms directly address the behavioral root causes of shadow AI by reducing the governance friction that drives unauthorized adoption without sacrificing the risk assessment rigor necessary for responsible AI deployment.

B. The Productivity-Security Trade-Off Paradox

Eighty-five percent of participants articulated a persistent executive-level tension between AI governance restrictiveness and organizational productivity — a trade-off that conventional security frameworks, designed for binary allow/block access control decisions, provide insufficient conceptual apparatus to navigate. Chief Information Security Officers and Chief Information Officers frequently described being positioned on opposing sides of this tension: CISOs oriented toward risk minimization, CIOs oriented toward capability enablement, with the Chief Executive Officer positioned as an uncomfortable arbiter lacking the technical depth to evaluate either claim rigorously.

The quantitative findings partially illuminate this tension: the negative SAPI \times GAS interaction on risk outcomes, combined with the positive SAPI \times GAS interaction on innovation outcomes, suggests that governance adaptiveness is the mechanism through which the productivity-security trade-off can be partially dissolved rather than merely navigated. Organizations that develop genuine governance adaptiveness do not face a fixed trade-off frontier but can shift that frontier outward — achieving both lower risk exposure and higher innovation output simultaneously. This finding challenges the implicit assumption in conventional IS security governance that security and productivity exist in inherent tension, suggesting instead that governance quality is the moderating variable that determines whether this tension is intrinsic or contingent.

C. Knowledge Entropy

Seventy-seven percent of participants identified a form of organizational knowledge risk that existing IS security frameworks do not adequately conceptualize: the gradual displacement of institutional knowledge from governed organizational repositories into ungoverned personal AI environments. As employees develop sophisticated prompt libraries, workflow automations, and AI-assisted analytical frameworks for tasks previously managed through documented organizational processes, critical institutional knowledge becomes encoded in personal LLM conversation histories, private AI workspace environments, and informal prompt-sharing networks rather than formally governed knowledge management systems (Alavi & Leidner, 2001; Zack, 1999).

Multiple participants described discovering this knowledge entropy risk acutely during employee departures: upon the resignation of employees who had developed extensive personal AI workflows, organizations found that critical operational processes could not be reproduced because the AI-mediated knowledge assets through which they operated had been entirely personal and were therefore inaccessible or irrecoverable. This finding highlights a governance dimension of shadow AI — the knowledge integrity risk — that security-focused governance frameworks systematically underweight relative to data breach and compliance risks, despite its potentially greater long-term competitive significance.

D. Adaptive Governance As Competitive Weapon

Seventy-one percent of participants, predominantly from higher-GAS organizations, articulated an emergent strategic insight that reframes adaptive governance from a risk management function to a source of competitive advantage. The core argument: in an era when AI capability proliferation occurs faster than any governance framework can track, organizations that develop superior governance adaptiveness gain a structural capability advantage over peers that remain locked in reactive governance postures.

Superior governance adaptiveness enables faster sanctioning of valuable AI innovations, more effective channeling of employee AI creativity toward organizational goals, and more credible AI governance communications to regulators, clients, and investors — each of which constitutes a distinct competitive asset.

This reframing is consistent with dynamic capabilities theory (Teece et al., 1997; Pavlou & El Sawy, 2011): governance adaptiveness, conceptualized as a dynamic capability rather than a static control mechanism, generates competitive returns through the continuous sense-assess-integrate cycle it enables organizations to apply to the emerging AI tool landscape. Organizations that institutionalize this cycle develop a durable governance capability that compounds in value as AI tool proliferation accelerates, while organizations that respond reactively to each shadow AI incident remain permanently in deficit relative to the rate of change.

E. Compliance Theater

Sixty-five percent of participants, with notable concentration among compliance and risk function leaders, identified a phenomenon they variously termed compliance theater, governance performance, or policy hygiene: the organizational practice of maintaining formally documented AI governance policies whose primary function is audit compliance rather than behavioral influence. These policies — often comprising lengthy acceptable-use documents, AI risk frameworks, and governance committee charters — were described as satisfying external stakeholder expectations for governance sophistication while exercising minimal actual influence on employee AI behavior.

This compliance theater dynamic creates a particularly insidious governance risk: organizations that believe their shadow AI risk exposure is managed by formal policy documentation may be substantially underestimating their actual risk because their governance infrastructure operates as a signaling mechanism rather than a behavioral control system. Cram et al.'s (2017) review of IS security policy effectiveness provides theoretical grounding for this concern, documenting persistent gaps between policy existence and policy compliance in organizational IS security contexts.

The generative AI context amplifies this gap because the behavioral salience of AI tool productivity benefits — experienced daily and immediately — dramatically outweighs the abstract and probabilistic nature of policy compliance risks.

F. The Innovation Paradox Of Shadow Ai

Fifty-eight percent of participants articulated what several independently termed the shadow AI innovation paradox: the empirical observation that a significant proportion of their organizations' most valuable recent AI innovations had originated as shadow AI deployments — unauthorized experimentation that, had it been immediately discovered and blocked, would have eliminated the innovation before its value could be recognized. This paradox is not merely a historical artifact but an ongoing organizational dynamic: the exploratory experimentation that generates genuine innovation frequently precedes formal governance approval, meaning that organizations with highly effective shadow AI prevention may inadvertently be preventing the experimentation through which future competitive capabilities gestate.

Higher-GAS organizations had developed formal Shadow-to-Sanctioned conversion processes (presented in Figure 3) specifically designed to capture value from this innovation paradox by creating structured pathways through which promising shadow AI deployments could be rapidly assessed, legitimized, and integrated into the formal AI capability portfolio. These conversion processes differ from conventional AI procurement in both speed (targeted completion in days to weeks rather than months) and starting point (beginning from demonstrated employee value rather than theoretical business case), creating a fundamentally different risk-return profile for AI capability development investment.

VII. THE SHADOW-TO-SANCTIONED AI CONVERSION FRAMEWORK

Synthesizing quantitative findings on governance adaptiveness, qualitative themes from executive interviews, and adaptive governance practices documented across high-GAS organizations, the study develops the Shadow-to-Sanctioned AI Conversion Framework (S2SAI-CF) presented in Figure 3. The framework operationalizes governance adaptiveness as a four-stage organizational process through which shadow AI deployments are systematically detected, assessed, triaged, and — where appropriate — converted into sanctioned organizational capabilities.

Figure 3

| Stage 1 Detection | Stage 2 Assessment | Stage 3 Triage | Stage 4 Sanctioning |
|---|---|---|---|
| <ul style="list-style-type: none"> • AI usage monitoring • Anomaly detection • Employee disclosure • Vendor discovery • Network traffic analysis | <ul style="list-style-type: none"> • Risk classification • Value potential scoring • Compliance gap mapping • Data sensitivity audit • User community size | <ul style="list-style-type: none"> • Immediate block (high risk) • Conditional sandbox • Fast-track approval • Policy waiver pathway • Escalation protocol | <ul style="list-style-type: none"> • Formal procurement • Security hardening • DLP integration • Governance embedding • Center of Excellence |

Figure 3. The Shadow-to-Sanctioned AI Conversion Framework (S2SAI-CF): Stage-Gate Process

Note. The framework represents the normative process observed across Adaptive Orchestrator archetype firms in the qualitative sample. Stage durations in high-GAS organizations: Detection (continuous); Assessment (24–72 hours for Tier 1 tools); Triage (48 hours); Sanctioning (2–4 weeks for low-risk tools; 2–4 months for high-risk enterprise tools). The framework does not prescribe universal sanctioning — high-risk tools identified in Stage 3 are blocked promptly with documented rationale communicated to the originating employee or team.

Stage 1 (Detection) encompasses the organizational monitoring capabilities through which shadow AI deployments are identified. High-GAS organizations employed a portfolio of detection mechanisms: network traffic analysis for unauthorized API calls to LLM providers, endpoint detection tools identifying AI tool installation events, voluntary employee disclosure programs incentivized through a no-fault self-reporting policy, and periodic vendor discovery audits cross-referencing SaaS spend data against official procurement records. Critically, high-GAS organizations positioned detection as a value discovery process rather than a surveillance mechanism, communicating to employees that self-disclosure of unauthorized AI tools would be treated as an innovation contribution warranting assessment and potential integration rather than a compliance violation warranting sanction.

Stage 2 (Assessment) encompasses the rapid evaluation of detected shadow AI deployments across four dimensions: security and data governance risk (data classification levels accessed; third-party data processing provisions; authentication and access control adequacy); regulatory compliance risk (sector-specific regulatory requirements; explainability and audit trail adequacy; cross-border data transfer implications); business value potential

(productivity improvement evidence; innovation novelty; employee adoption breadth); and strategic fit (alignment with AI strategy roadmap; integration complexity with existing enterprise systems). High-GAS organizations had developed standardized assessment templates that enabled non-specialist business unit managers to complete preliminary assessments in hours, with specialist review reserved for tools identified as either high-risk or high-value.

Stage 3 (Triage) encompasses the governance decision process through which assessed tools are classified and routed. High-GAS organizations employed four triage pathways: immediate block with documented rationale (reserved for tools with material data security, compliance, or safety risks that cannot be mitigated through configuration or contractual controls); conditional sandbox (provisional deployment permitted within defined data scope and user population constraints while full assessment proceeds); fast-track approval (streamlined procurement and security integration for low-risk tools with demonstrated value); and escalation to enterprise AI strategy committee (for high-value, high-complexity tools with strategic implications beyond departmental scope). The existence of multiple triage pathways — rather than a binary allow/block decision — is a defining characteristic of high-GAS governance, enabling proportionate responses that preserve value while managing risk.

Stage 4 (Sanctioning) encompasses the integration of approved shadow AI tools into the formal enterprise AI governance architecture through formal procurement contracts, DLP (Data Loss Prevention) configuration, audit trail implementation, and Center of Excellence knowledge transfer. High-GAS organizations had developed streamlined sanctioning processes specifically for shadow AI conversions that were substantially faster than standard

enterprise software procurement, reflecting the recognized competitive cost of prolonged shadow-to-sanctioned transition periods during which value-generating tools remain in an ungoverned intermediate state. Blockchain-based frameworks offer a complementary technical architecture for implementing tamper-evident audit trails that can document the provenance and data handling history of sanctioned AI tools throughout their organizational lifecycle (Sammangi, Jagatha, & Liu, 2025c).

VIII. DISCUSSION

A. Theoretical Contributions

This study makes four primary theoretical contributions. First, the Shadow AI Duality Model advances IS security governance theory by reframing shadow AI from a unidimensional compliance problem to a bidimensional organizational phenomenon with simultaneous risk and innovation potentials. This reframing has significant implications for how governance researchers conceptualize the organizational functions of unauthorized technology adoption: rather than positioning unauthorized use as uniformly dysfunctional — a governance failure requiring correction — the SADM positions it as an organizational signal of governance friction and exploratory capability investment that governance systems should detect, assess, and strategically route rather than simply suppress. This extends Haag et al.'s (2015) empirical observation of shadow IT performance advantages into a theoretical account of the mechanisms through which unauthorized technology use generates value.

Second, the concept of Governance Adaptiveness as a distinct organizational capability extends IS governance theory beyond its traditional focus on governance structure (Tallon et al., 2013; Tiwana & Konsynski, 2010) to emphasize governance process agility as a competitive differentiator. The quantitative evidence that GAS predicts innovation output ($\beta = 0.41$) and operational resilience ($\beta = 0.48$) independently of SAPI suggests that governance adaptiveness generates organizational value that extends beyond shadow AI management to broader AI capability development and organizational resilience dimensions.

Third, the S2SAI-CF framework provides the IS governance literature with a normative process model specifically designed for the shadow AI conversion problem — a governance challenge with no direct

precedent in the shadow IT literature because of the qualitative differences in capability generation, data governance implications, and institutional knowledge dimensions that distinguish shadow AI from its predecessors. Fourth, the knowledge entropy construct contributes to knowledge management theory (Alavi & Leidner, 2001; Zack, 1999) by identifying AI-mediated knowledge displacement as a novel knowledge governance risk category requiring dedicated management attention. The broader challenge of governing AI deployments in high-stakes organizational contexts — ensuring that AI systems for critical prediction and decision-support tasks are subject to appropriate oversight, validation, and accountability mechanisms — extends across domains from enterprise shadow AI to specialized predictive applications (Sharma, Singh, Sammangi, Sharma, Pandey, Srivastava, Agarwal, & Singh, 2025a).

B. Practical Implications

The central practical implication of this study is that shadow AI governance requires a strategic reorientation from prevention to detection, assessment, and structured conversion. Organizations that invest governance resources predominantly in preventing shadow AI adoption will face a progressively losing battle against increasingly accessible AI capabilities, while simultaneously suppressing the organizational AI creativity that constitutes a genuine competitive resource. The governance adaptiveness evidence — and particularly the Innovation Paradox theme from the qualitative findings — suggests that many organizations are inadvertently destroying competitive value through excessively restrictive AI governance.

The governance archetype analysis provides practical diagnostic value: executives can assess their organization's archetype position and identify governance investment priorities accordingly. Rigid Prohibitionists require urgently developed legitimization pathways and provisional use frameworks before competitive stagnation becomes acute. Permissive Laissez-Faire organizations require immediate GAS investment — particularly in detection and triage capabilities — before accumulated risk exposure crystallizes into material incidents. Adaptive Orchestrators require sustained governance investment to maintain adaptiveness as AI tool proliferation continues to accelerate.

The knowledge entropy finding suggests a specific governance investment priority that current enterprise AI

strategies systematically underemphasize: AI knowledge asset governance. Organizations should establish formal programs for identifying, documenting, and institutionalizing AI workflows, prompt libraries, and automation architectures currently residing in personal AI environments — treating these assets with the same governance attention historically applied to source code, intellectual property, and institutional data. Employee transition protocols should be updated to include AI asset transfer provisions equivalent to traditional IT system access management procedures.

C. Limitations And Future Research

Several limitations constrain the generalizability of this study's findings and suggest productive avenues for future investigation. The measurement of shadow AI prevalence through a combination of network monitoring telemetry, incident reports, and survey self-reports likely underestimates true prevalence: by definition, the most carefully concealed shadow AI deployments are least visible to detection methods. Future research could explore ethnographic or insider methods that provide more complete prevalence estimates, though these raise significant research ethics considerations regarding employee privacy. Second, the study's sample, while substantially broader than most IS security research samples, is limited to large and mid-market public firms; shadow AI dynamics in smaller organizations, public sector entities, and non-profit organizations may differ substantially and warrant dedicated investigation.

Third, the rapidly evolving AI tool landscape means that the specific risk and innovation manifestations of shadow AI documented in this study — based on the 2022–2026 generative AI era — will evolve as agentic AI systems with greater operational autonomy become accessible as consumer tools. Longitudinal extension of this research is planned, with particular attention to the governance implications of shadow agentic AI: unauthorized deployments of autonomous multi-step AI agents create governance challenges qualitatively more complex than unauthorized LLM tool use, including autonomous data access, external API calls, and action execution with limited human oversight. Future research should develop governance frameworks specifically calibrated to these emerging shadow AI risk profiles.

IX. CONCLUSION

Shadow AI represents one of the most pressing and theoretically underexplored challenges in contemporary information systems governance. Its organizational character — simultaneously a significant security and compliance risk and a genuine source of innovation capability — demands governance frameworks that transcend the binary compliance logic of conventional IS security approaches. This study has provided the first large-scale longitudinal empirical examination of the shadow AI prevalence-performance relationship, developed two validated measurement instruments (SAPI and GAS), and advanced a theoretical model and practical framework for governance practitioners navigating this complex terrain.

The study's most significant empirical finding — that governance adaptiveness moderates the shadow AI-risk relationship while independently predicting innovation output and organizational resilience — carries a clear and actionable message for senior executives: the competitive question in the shadow AI era is not whether to govern AI, but how to govern it with sufficient adaptiveness to capture the innovation value that employee AI experimentation generates while containing the risk pathways through which ungoverned AI deployment creates organizational harm. Organizations that answer this question with rigid prohibition will find themselves governing a progressively obsolete AI landscape while their more adaptive competitors build compounding governance capabilities that enable them to exploit the full strategic potential of enterprise AI adoption.

The Shadow-to-Sanctioned AI Conversion Framework provides practitioners with a concrete operational architecture for building the governance adaptiveness that this study demonstrates is central to superior organizational outcomes. As AI capabilities continue to accelerate — and as the boundary between sanctioned and shadow AI becomes progressively more difficult to maintain — the organizations that invest seriously in governance adaptiveness will emerge as the governance leaders of the agentic AI era, transforming the perpetual challenge of unauthorized AI adoption from a compliance burden into a competitive advantage.

REFERENCES

1. Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, 54(2), 442–492. <https://doi.org/10.1257/jel.54.2.442>
2. Alavi, M., & Leidner, D. E. (2001). Review: Knowledge management and knowledge management systems: Conceptual foundations and research issues. *MIS Quarterly*, 25(1), 107–136.
3. Anderson, R., & Moore, T. (2006). The economics of information security.
4. *Science*, 314(5799), 610–613.
4. Argyris, C., & Schön, D. A. (1978). *Organizational learning: A theory of action perspective*. Addison-Wesley.
5. Barlow, J. B., Warkentin, M., Ormond, D., & Dennis, A. R. (2018). Don't even think about it! The effects of antineutralization, informational, and normative communication on information security compliance. *Journal of the Association for Information Systems*, 19(8), 689–715.
6. 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.
7. Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing artificial intelligence. *MIS Quarterly*, 45(3), 1433–1450.
8. Bulgurcu, B., Cavusoglu, H., & Benbasat, I. (2010). Information security policy compliance: An empirical study of rationality-based beliefs and information security awareness. *MIS Quarterly*, 34(3), 523–548.
9. 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.
10. D'Arcy, J., Hovav, A., & Galletta, D. (2009). User awareness of security countermeasures and its impact on information systems misuse: A deterrence approach. *Information Systems Research*, 20(1), 79–98.
11. Davenport, T. H., & Mittal, N. (2023). How generative AI is changing creative work. *Harvard Business Review Digital Articles*. <https://hbr.org/2022/11/how-generative-ai-is-changing-creative-work>
12. Di Gregorio, D., Musteen, M., & Thomas, D. E. (2008). Offshore outsourcing as a source of international competitiveness for SMEs. *Journal of International Business Studies*, 39(3), 526–543.
13. Dhillon, G., & Backhouse, J. (2001). Current directions in IS security research: Towards socio-organizational perspectives. *Information Systems Journal*, 11(2), 127–153.
14. Fernandez-Vidal, J., Perotti, F. A., González, R., & Gasco, J. (2022). Managing shadow IT: A systematic literature review and research agenda. *Information Systems Frontiers*, 24(4), 1113–1128.
15. Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, 254–280.
16. Gartner. (2024). *Hype cycle for emerging technologies, 2024*. Gartner Research.
17. Gartner. (2025). *Shadow AI governance: Emerging enterprise risk patterns*. Gartner Research.
18. Haag, S., & Eckhardt, A. (2017). Normalizing the shadows — the role of symbolic models for individuals' shadow IT usage. *Proceedings of the 38th International Conference on Information Systems*.
19. Haag, S., Eckhardt, A., & Bozoyan, C. (2015). Are shadow system users the better IS users? Insights of a lab experiment. *Proceedings of the 36th International Conference on Information Systems*.
20. Huber, G. P. (1991). Organizational learning: The contributing processes and the literatures. *Organization Science*, 2(1), 88–115.
21. IBM Institute for Business Value. (2024). *AI governance in the enterprise: Managing risk in an autonomous AI world*. IBM Corporation.
22. Iansiti, M., & Lakhani, K. R. (2023). *Competing in the age of AI: Strategy and leadership when algorithms and networks run the world*. Harvard Business Review Press.
23. Jarvenpaa, S. L., & Majchrzak, A. (2010). Research commentary: Vigilant interaction in knowledge collaboration: Challenges of online user participation under ambivalence. *Information Systems Research*, 21(4), 773–784.
24. 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. *Engineering: Open Access*, 3(6), 1–8.
25. Johnson, A. M., & Lederer, A. L. (2010). CEO/CIO mutual understanding, strategic alignment, and the

contribution of IS to organizational goals. *Information & Management*, 47(3), 138–149.

26. Kane, G. C., Phillips, A. N., Copulsky, J. R., & Andrus, G. R. (2019). *The technology fallacy: How people are the real key to digital transformation*. MIT Press.

27. Kark, K., & Shaikh, A. (2024). The shadow AI problem: What CISOs need to know now. *MIT Sloan Management Review Digital Articles*.

28. Kettinger, W. J., Marchand, D. A., & Davis, J. M. (2010). Designing enterprise IT architectures to optimize flexibility and standardization in global business. *MIS Quarterly Executive*, 9(2), 95–113.

29. Kwon, J., & Johnson, M. E. (2014). Proactive versus reactive security investments in the healthcare sector. *MIS Quarterly*, 38(2), 451–471.

30. Lacity, M. C., & Willcocks, L. P. (2016). A new approach to automating services. *MIT Sloan Management Review*, 58(1), 41–49.

31. Li, X., Hess, T. J., & Valacich, J. S. (2008). Why do we trust new technology? A study of initial trust formation with organizational information systems. *Journal of Strategic Information Systems*, 17(1), 39–71.

32. March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1), 71–87.

33. McKinsey Global Institute. (2023). *The economic potential of generative AI: The next productivity frontier*. McKinsey & Company.

34. McKinsey Global Institute. (2025). *The agentic AI enterprise: From co-pilots to autonomous workflows*. McKinsey & Company.

35. Nambisan, S., Lyytinen, K., Majchrzak, A., & Song, M. (2017). Digital innovation management. *MIS Quarterly*, 41(1), 223–238.

36. O'Brien, L., & Toms, E. G. (2008). What is user engagement? A conceptual framework for defining user engagement with technology. *Journal of the American Society for Information Science and Technology*, 59(6), 938–955.

37. Orlikowski, W. J. (1992). The duality of technology: Rethinking the concept of technology in organizations. *Organization Science*, 3(3), 398–427.

38. Pavlou, P. A., & El Sawy, O. A. (2011). Understanding the elusive black box of dynamic capabilities. *Decision Sciences*, 42(1), 239–273.

39. Ponemon Institute. (2024). *Cost of insider threats global report 2024*. Proofpoint/Ponemon Institute.

40. Sarbanes-Oxley Act. (2002). Pub. L. No. 107-204, 116 Stat. 745. U.S. Government Publishing Office.

41. Sammangi, H., Ambati, L. S., Liu, J., & Jagatha, A. (2025). AI-driven decentralized IoT for secure and scalable healthcare. *AMCIS 2025 Proceedings*. https://aisel.aisnet.org/amcis2025/health_it/sig_health/3

42. Sammangi, H., Jagatha, A., & Liu, J. (2025b). Harnessing generative AI and large language models for revolutionizing cybersecurity in the Internet of Things: Ethical and privacy implications. *Engineering: Open Access*, 3(6), 1–12.

43. Sammangi, H., Jagatha, A., & Liu, J. (2025c). Integrating blockchain technology into telemedicine: A framework for enhancing data privacy and security. *Engineering: Open Access*, 3(6), 1–7.

44. Siponen, M., & Vance, A. (2010). Neutralization: New insights into the problem of employee information systems security policy violations. *MIS Quarterly*, 34(3), 487–502.

45. Sharma, G., Singh, J., Sammangi, H., Sharma, M., Pandey, R., Srivastava, S., Agarwal, G., & Singh, I. (2025a). 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

46. Stahl, B. C., Antoniou, J., Ryan, M., Macnish, K., & Jiya, T. (2022). Organisational responses to the ethical issues of artificial intelligence. *AI & Society*, 37(1), 23–37.

47. Straub, D. W. (1990). Effective IS security: An empirical study. *Information Systems Research*, 1(3), 255–276.

48. 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.

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

50. Tiwana, A., & Konsynski, B. (2010). Complementarities between organizational IT architecture and governance structure. *Information Systems Research*, 21(2), 288–304.

51. Vance, A., Siponen, M., & Pahlila, S. (2012). Motivating IS security compliance: Insights from habit and



protection motivation theory. *Information & Management*, 49(3-4), 190-198.

52. Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425-478.

53. World Economic Forum. (2025). The future of jobs report 2025: AI, automation and workforce transformation. WEF.

54. Zack, M. H. (1999). Developing a knowledge strategy. *California Management Review*, 41(3), 125-145.

55. Zahra, S. A., & George, G. (2002). Absorptive capacity: A review, reconceptualization, and extension. *Academy of Management Review*, 27(2), 185-203.

56. Zmud, R. W. (1983). The effectiveness of external information channels in facilitating innovation within software development groups. *MIS Quarterly*, 7(2), 43-58.