
| RESEARCH ARTICLE

Risk-Calibrated Generative AI Governance for Responsible AI Adoption in U.S. Organizations: Evidence from Marketing and Corporate Disclosures

Umme Aimen Khan

MBA Program, College of Business Administration, Central Michigan University, Mount Pleasant, Michigan, United States

Corresponding Author: Umme Aimen Khan, E-mail: ummeaimenkhan69@gmail.com

| ABSTRACT

Generative artificial intelligence (GenAI) is increasingly embedded in marketing content creation, personalization, customer communication, advertising, and multi-channel campaign execution. Although GenAI can improve speed and scalability, its use creates risks involving accuracy, privacy, intellectual property, bias, compliance, brand consistency, and consumer protection. This study develops and empirically examines a risk-calibrated governance framework for cross-functional marketing projects. The research combines an integrative review and thematic synthesis of relevant scholarly literature with an archival content analysis of official AI-governance disclosures from a purposive sample of large U.S. public companies across technology, finance, and consumer/media sectors. Six governance gates were coded as publicly disclosed or not observed: risk classification, input and data governance, output quality review, legal and ethical compliance, approval and human oversight, and post-deployment monitoring and escalation. The mean governance disclosure score was 5.27 out of 6. Legal and ethical compliance was universally disclosed within the sample, while data governance, risk classification, output quality review, and approval or human oversight were also widely observed. Post-deployment monitoring and escalation was the least consistently disclosed governance stage. Fisher's exact test showed that firms disclosing risk classification were more likely to disclose approval or human-oversight controls (odds ratio = 36.00, $p = .038$). The corresponding association between risk classification and post-deployment monitoring was positive but not statistically significant (odds ratio = 10.67, $p = .117$). Governance disclosure scores did not differ significantly across sectors (Kruskal-Wallis $H = 0.067$, $p = .967$). The findings support risk-calibrated pre-deployment oversight while highlighting comparatively less consistent public disclosure of lifecycle monitoring after deployment.

| KEYWORDS

Generative AI; AI Governance; Risk Management; Marketing Project Management; Cross-Functional Teams; Human Oversight; Compliance; Multi-Channel Marketing

| ARTICLE INFORMATION

ACCEPTED: 02 June 2026

PUBLISHED: 26 August 2026

DOI: 10.32996/jefas.2026.8.9.3

1. Introduction

Generative artificial intelligence (GenAI) is increasingly used in marketing for content creation, personalization, customer communication, advertising, and campaign execution. AI already supports customer analysis, targeting, positioning, and marketing actions, while GenAI extends these capabilities by allowing firms to create and adapt marketing content quickly and at scale (Grewal et al., 2025; Huang & Rust, 2021; Wahid et al., 2025).

However, GenAI also creates risks involving inaccurate information, hallucinations, privacy, intellectual property, bias, brand consistency, and regulatory compliance. Marketing-focused research emphasizes consumer-protection, privacy, ethical, and legal concerns, while higher-risk applications may require stronger human review even when additional controls increase time and cost (Duivenvoorde, 2025; Grewal et al., 2025; Hari et al., 2025). NIST guidance likewise emphasizes that responsible AI adoption requires lifecycle risk governance, safeguards, and integration into organizational workflows (Autio et al., 2024).

Despite growing research on GenAI marketing and AI governance, limited attention has been given to how these controls are coordinated within cross-functional marketing project workflows or how closely real organizations publicly describe controls that

resemble such gates. This study therefore combines literature synthesis with an archival analysis of official corporate AI-governance disclosures to develop and examine a risk-calibrated framework covering risk classification, data use, output review, legal and ethical checks, cross-channel approval, and post-deployment monitoring.

2. Literature Review

2.1 Generative AI in Marketing

Generative artificial intelligence is expanding the role of AI in marketing by supporting content creation, personalization, advertising, customer communication, and other customer-facing activities. Earlier marketing research showed that AI can support segmentation, targeting, positioning, automation, and strategic decision-making, while recent GenAI research highlights its ability to create and adapt marketing content quickly and at scale (Grewal et al., 2025; Huang & Rust, 2021; Wahid et al., 2025).

Despite these advantages, GenAI introduces risks related to inaccurate information, hallucinations, privacy, bias, intellectual property, consumer protection, and inappropriate customer-facing outputs. Recent marketing studies therefore emphasize privacy-by-design, legal and ethical safeguards, and risk-sensitive human review, especially when content is externally visible or customer facing (Duivenvoorde, 2025; Grewal et al., 2025; Hari et al., 2025; Kumar & Suthar, 2024; Saura et al., 2024).

2.2 Responsible AI Governance and Human Oversight

Responsible AI governance focuses on ensuring that AI systems are used with clear accountability, appropriate controls, monitoring, and escalation procedures. Research shows that governance should combine structural mechanisms such as defined roles and decision rights, procedural mechanisms such as review and compliance monitoring, and relational mechanisms such as communication and cross-functional collaboration (Janssen, 2025; Papagiannidis et al., 2025).

Human oversight is important, but the literature suggests that oversight should be calibrated to risk and supported by broader organizational controls. Studies document approval boards, human-in-the-loop review, escalation paths, audits, technical guardrails, and post-deployment monitoring, while also warning that excessive or poorly designed controls can create friction, resistance, or overreliance on reviewers (Corrêa et al., 2025; Hadley et al., 2025; Jorgensen et al., 2025; Leon, 2026; Salgado-Criado, 2025). Cross-functional coordination and role clarity are therefore central to implementation (Jafari Meimandi et al., 2025; Sayan & Bueechl, 2025).

2.3 Research Gap and Conceptual Framework

Although previous studies discuss GenAI-related marketing risks, responsible AI principles, human oversight, organizational governance, and AI-supported project management, these research streams remain fragmented. Existing work shows that governance mechanisms already exist, but limited research integrates them into a coherent, risk-calibrated architecture for cross-functional, multi-channel marketing project workflows and compares that architecture with current corporate governance disclosures (Autio et al., 2024; Leon, 2026; Papagiannidis et al., 2025).

To address this gap, the study asks three questions:

RQ1: What GenAI-related risks and governance mechanisms are identified across marketing, responsible-AI, and project-management literature?

RQ2: How can these mechanisms be integrated into risk-calibrated quality and compliance gates for cross-functional marketing projects?

RQ3: To what extent do large U.S. companies publicly disclose controls corresponding to these gates, and are risk-classification disclosures associated with approval and post-deployment monitoring?

3. Methodology

This study uses a two-stage design combining a structured integrative literature review with an exploratory archival content analysis of corporate AI-governance disclosures. The literature stage develops the six-gate framework; the archival stage examines whether comparable controls are visible in current real-world organizational disclosures.

The two evidence layers are analyzed separately and then interpreted together. This design avoids treating public disclosure as proof of internal control effectiveness while still providing an observable real-world test of whether the framework corresponds to governance practices described by major organizations.

3.1 Review Design and Thematic Synthesis

The literature review focused on studies related to generative AI in marketing, responsible AI governance, human oversight, organizational controls, and project management. Relevance-based searches used combinations of terms including “generative AI marketing,” “AI governance,” “responsible AI,” “human oversight,” “AI risk management,” and “AI project management.” Relevant scholarly studies were selected based on their direct contribution to marketing-related GenAI risks, governance mechanisms, cross-functional coordination, and operational oversight and compliance. The selected literature was coded thematically for risk categories, accountability, human oversight, risk calibration, approval processes, monitoring, and escalation. The review is a focused integrative synthesis and does not claim exhaustive coverage of all publications.

3.2 Archival Corporate Sample and Coding

The empirical layer uses a purposive archival sample of large U.S. public companies with accessible official AI-governance materials as of August 22, 2026, spanning the technology, financial, and consumer/media sectors. Evidence was taken only from official company webpages, governance reports, policy materials, technology articles, or current operational role descriptions. The full company-level coding audit trail and source URLs are provided in Supplementary Table S1. Each company was coded on six binary disclosure variables corresponding to the proposed gates: risk classification, data governance, output-quality review, legal/ethical compliance, approval or human oversight, and post-deployment monitoring or escalation. A governance disclosure score from 0 to 6 was calculated as the sum of these indicators. A value of 0 means the control was not observed in the selected public source; it does not establish that the company lacks the control internally.

3.3 Hypotheses and Statistical Analysis

Three exploratory hypotheses were evaluated. H1: Companies disclosing risk classification are more likely to disclose approval or human-oversight controls. H2: Companies disclosing risk classification are more likely to disclose post-deployment monitoring or escalation. H3: Overall governance disclosure scores differ across technology, finance, and consumer/media sectors. Fisher's exact tests were used for H1 and H2 because of the small binary contingency tables. H3 was evaluated with the Kruskal-Wallis test because the disclosure score is bounded from 0 to 6 and the sector groups are small. Statistical tests are interpreted as exploratory rather than confirmatory.

4. Findings and Empirical Results

4.1 Key GenAI Risk Themes

The reviewed literature identifies several recurring risks associated with GenAI use in marketing, including inaccurate or fabricated outputs, privacy concerns, intellectual-property issues, bias, transparency problems, regulatory non-compliance, and potential damage to brand reputation. These risks become more significant when GenAI-generated material is used in customer-facing or externally published marketing activities (Duivenvoorde, 2025; Grewal et al., 2025; Hari et al., 2025; Kumar & Suthar, 2024; Saura et al., 2024).

The literature also indicates that GenAI risks vary according to the nature and potential impact of the task. Higher-risk applications involving sensitive data, personalized communications, public claims, or customer interactions generally require stronger review and control than low-risk internal activities. This supports a risk-calibrated rather than uniform approach to GenAI governance (Grewal et al., 2025; Janssen, 2025; Wahid et al., 2025).

4.2 Governance Mechanisms and Cross-Functional Coordination

The reviewed studies show that effective AI governance depends on more than formal policies. Important mechanisms include clearly assigned responsibilities, approval procedures, documentation, monitoring, audits, escalation paths, and defined decision rights. Responsible AI governance therefore requires structural, procedural, and relational controls that connect policy with day-to-day organizational practice (Hadley et al., 2025; Papagiannidis et al., 2025).

Cross-functional coordination is especially important because GenAI-related decisions may involve marketing, analytics, technology, brand, legal, and compliance functions. Research on organizational AI governance shows that unclear roles and fragmented responsibility can weaken implementation, while better coordination improves consistency in review, escalation, and accountability (Jafari Meimandi et al., 2025; Sayan & Bueechl, 2025).

4.3 Risk-Calibrated Oversight and Approval Gates

The reviewed literature suggests that GenAI governance should be risk-calibrated rather than applied uniformly. Low-risk internal tasks may require limited oversight, whereas customer-facing content, sensitive data use, personalized communications, or high-impact claims may require stronger review, approval, and documentation. This approach helps organizations balance responsible AI use with efficiency and avoids unnecessary governance burdens (Grewal et al., 2025; Hadley et al., 2025; Janssen, 2025).

The literature also supports the use of practical controls such as risk triage, approval gates, output verification, escalation procedures, audit trails, and post-deployment monitoring. These mechanisms can translate AI governance principles into

operational decision points throughout a marketing project and provide the basis for the risk-calibrated governance framework proposed in this study (Jorgensen et al., 2025; Leon, 2026; Salgado-Criado, 2025).

4.4 Archival Validation of the Six-Gate Framework

Across the archival sample, the mean governance disclosure score was 5.27 out of 6 and the median was 6. Legal and ethical compliance was disclosed by all firms, while data governance, risk classification, output-quality review, and approval or human oversight were also widely disclosed. Post-deployment monitoring or escalation was the least consistently disclosed governance gate. Table 1 summarizes the gate-level results.

Table 1. Public disclosure coverage of the six governance gates (N = 22)

Governance Gate	Companies (n)	Percent
G1 Risk classification	19	86.4%
G2 Data governance	20	90.9%
G3 Output quality review	19	86.4%
G4 Legal/ethical compliance	22	100.0%
G5 Approval/human oversight	19	86.4%
G6 Monitoring/escalation	17	77.3%

Source: Author's analysis based on publicly available corporate AI-governance disclosures.

H1 was supported: risk-classification disclosure was associated with approval or human-oversight disclosure (Fisher's exact test: OR = 36.00, $p = .038$). H2 was not statistically supported: the association between risk classification and monitoring or escalation was positive but not significant (OR = 10.67, $p = .117$). H3 was also not supported; disclosure scores did not differ across sectors (Kruskal-Wallis $H = 0.067$, $p = .967$). Mean scores were 5.60 for technology, 5.17 for finance, and 4.83 for consumer/media firms.

Table 2. Exploratory hypothesis tests

Hypothesis	Test	Statistic	p	Result
H1 Risk classification -> approval / oversight	Fisher's exact	OR = 36.00	.038	Supported
H2 Risk classification -> monitoring / escalation	Fisher's exact	OR = 10.67	.117	Not supported
H3 Sector differences in total score	Kruskal-Wallis	$H = 0.067$	.967	Not supported

Source: Author's analysis based on publicly available corporate AI-governance disclosures.

5. Proposed Risk-Calibrated GenAI Governance Framework

Based on the literature synthesis and archival evidence, this study proposes a risk-calibrated governance framework for managing GenAI within cross-functional marketing projects. The archival findings provide empirical support for the linkage between risk classification and pre-deployment approval or human oversight, while the lower prevalence of monitoring disclosures indicates that lifecycle governance should not be assumed to follow automatically from front-end risk assessment. The framework consists of six governance gates:

(1) use-case and risk classification, to determine exposure and review intensity; (2) input and data governance, to control sensitive or proprietary information; (3) output quality review, to verify accuracy, factual claims, and brand alignment; (4) legal and ethical compliance review, to address privacy, intellectual property, bias, disclosure, and consumer-protection concerns; (5) cross-channel release approval, to confirm consistency and authorization across campaign channels; and (6) post-deployment monitoring and escalation, to detect incidents, complaints, drift, or required corrections. Project managers coordinate these gates by connecting marketing, analytics, technology, brand, legal, and compliance specialists rather than acting as the sole decision-maker.

6. Discussion

The combined findings suggest that effective GenAI governance in marketing should not rely on policies or human review alone. The archival analysis shows that public disclosure of risk classification is strongly associated with approval or human-oversight disclosure, which supports the central risk-calibrated logic of the framework. However, the corresponding relationship with post-deployment monitoring was not statistically significant, indicating that front-end risk assessment does not necessarily imply equally visible lifecycle monitoring.

The contribution of the proposed framework is not the invention of approval gates themselves, which are already documented in organizational AI governance research. Rather, it synthesizes existing controls into a six-stage architecture tailored to cross-functional marketing projects and tests whether comparable controls are visible in current corporate disclosures. The absence of significant sector differences further suggests that the governance logic may be cross-sectoral, consistent with NIST guidance framing GenAI risk management as a lifecycle governance challenge across organizational contexts (Autio et al., 2024).

7. Theoretical and Managerial Implications

Theoretically, this study connects GenAI marketing, responsible AI governance, and project management while adding an archival validation layer. The findings suggest that risk classification and pre-deployment oversight frequently appear together in corporate disclosures, but monitoring is less consistently visible. This distinction reinforces the need to conceptualize governance as a lifecycle architecture rather than a single approval event.

From a managerial perspective, the six-gate framework can be used as a project checklist for determining when GenAI-generated marketing work requires additional review, approval, or escalation. The empirical results highlight post-deployment monitoring as the weakest publicly disclosed gate, making complaint handling, drift detection, incident response, and corrective-action ownership especially important. Project managers can coordinate these controls across marketing, analytics, technology, brand, legal, and compliance teams while applying lighter oversight to genuinely lower-risk activities.

8. Limitations and Future Research

This study has several limitations. The literature review is focused rather than exhaustive, and the archival sample is purposive rather than statistically representative of all U.S. firms. Corporate coding is based on selected public official sources, so a zero means that a control was not observed in the reviewed disclosure and should not be interpreted as evidence that the control is absent internally. Source types also vary in detail, and the coding was conducted by a single researcher without an independent inter-rater reliability assessment.

Future research should expand the archival sample, use multiple independent coders, and test whether disclosed governance controls correspond to measurable outcomes such as AI incidents, rework, regulatory findings, customer complaints, or project delays. Longitudinal studies could also examine whether organizations strengthen post-deployment monitoring as GenAI moves from experimentation into higher-autonomy customer-facing workflows.

9. Conclusion

Generative AI offers substantial opportunities for marketing, but its use also introduces risks related to accuracy, privacy, intellectual property, bias, compliance, and brand reputation. This study integrates insights from relevant scholarly literature with an archival analysis of official AI-governance disclosures from large U.S. public companies to develop and examine a risk-calibrated governance framework for cross-functional marketing projects.

The results show broad public disclosure of legal and ethical compliance, data governance, risk classification, output review, and approval controls, but comparatively weaker disclosure of post-deployment monitoring. Risk classification was significantly associated with approval or human oversight, supporting risk-calibrated pre-deployment governance, while sector differences were not significant. The six-gate framework therefore provides a practical structure for balancing GenAI-enabled innovation with accountability, compliance, quality, and lifecycle risk management.

Data Availability: The company-level coding audit trail, source URLs, and statistical summary are provided in Supplementary Table S1.

Funding: This research received no external funding.

Conflicts of Interest: The author declares no conflict of interest.

ORCID iD: <https://orcid.org/0009-0005-8331-0120>

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

- [1]. Autio, C., Schwartz, R., Dunietz, J., Jain, S., Stanley, M., Tabassi, E., Hall, P., & Roberts, K. (2024). *Artificial intelligence risk management framework: Generative artificial intelligence profile* (NIST AI 600-1). National Institute of Standards and Technology (U.S.). <https://doi.org/10.6028/NIST.AI.600-1>
- [2]. Corrêa, A. M., Garsia, S., & Elbi, A. (2025). Better together? Human oversight as means to achieve fairness in the European AI Act governance. *Cambridge Forum on AI: Law and Governance*, 1, e29. <https://doi.org/10.1017/cfl.2025.10010>
- [3]. Duivenvoorde, B. (2025). Generative AI and the future of marketing: A consumer protection perspective. *Computer Law & Security Review*, 57, 106141. <https://doi.org/10.1016/j.clsr.2025.106141>
- [4]. Grewal, D., Satornino, C. B., Davenport, T., & Guha, A. (2025). How generative AI is shaping the future of marketing. *Journal of the Academy of Marketing Science*, 53(3), 702–722. <https://doi.org/10.1007/s11747-024-01064-3>
- [5]. Hadley, E., Blatecky, A., & Comfort, M. (2025). Investigating algorithm review boards for organizational responsible artificial intelligence governance. *AI and Ethics*, 5(3), 2485–2495. <https://doi.org/10.1007/s43681-024-00574-8>

- [6]. Hari, H., Sharma, A., Verma, S., & Chaturvedi, R. (2025). Exploring ethical frontiers of artificial intelligence in marketing. *Journal of Responsible Technology*, 21, 100103. <https://doi.org/10.1016/j.jrt.2024.100103>
- [7]. Huang, M.-H., & Rust, R. T. (2021). A strategic framework for artificial intelligence in marketing. *Journal of the Academy of Marketing Science*, 49(1), 30–50. <https://doi.org/10.1007/s11747-020-00749-9>
- [8]. Jafari Meimandi, K., Reuel, A., Aranguiz-Dias, G., Rahama, H., Ayadi, A.-E., Boullier, X., Verdo, J., Montanie, L., & Kochenderfer, M. (2025). An Adaptive Responsible AI Governance Framework for Decentralized Organizations. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 8(2), 1324–1334. <https://doi.org/10.1609/aies.v8i2.36633>
- [9]. Janssen, M. (2025). Responsible governance of generative AI: conceptualizing GenAI as complex adaptive systems. *Policy and Society*, 44(1), 38–51. <https://doi.org/10.1093/polsoc/puae040>
- [10]. Jorgensen, M., Brogle, K., Collins, K. M., Ibrahim, L., Shah, A., Ivanovic, P., Broestl, N., Piles, G., Dongha, P., Abdulhussein, H., Weller, A., Powers, J., & Bhatt, U. (2025). Documenting Deployment with Fabric: A Repository of Real-World AI Governance. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 8(2), 1350–1362. <https://doi.org/10.1609/aies.v8i2.36636>
- [11]. Kumar, D., & Suthar, N. (2024). Ethical and legal challenges of AI in marketing: An exploration of solutions. *Journal of Information, Communication and Ethics in Society*, 22(1), 124–144. <https://doi.org/10.1108/JICES-05-2023-0068>
- [12]. Leon, M. (2026). Governing generative AI in organizations: A design theory and quasi-experimental field study of sociotechnical guardrails. *The Journal of Supercomputing*, 82(12), 641. <https://doi.org/10.1007/s11227-026-08760-7>
- [13]. Papagiannidis, E., Mikalef, P., & Conboy, K. (2025). Responsible artificial intelligence governance: A review and research framework. *The Journal of Strategic Information Systems*, 34(2), 101885. <https://doi.org/10.1016/j.jsis.2024.101885>
- [14]. Salgado-Criado, J. (2025). Human oversight of artificial intelligence: An operations management perspective. *Journal of Industrial Engineering and Management*, 18(2), 285–304. <https://doi.org/10.3926/jiem.8567>
- [15]. Saura, J. R., Škare, V., & Dosen, D. O. (2024). Is AI-based digital marketing ethical? Assessing a new data privacy paradox. *Journal of Innovation & Knowledge*, 9(4), 100597. <https://doi.org/10.1016/j.jik.2024.100597>
- [16]. Sayan, M. M., & Bueechl, J. (2025). Embedding AI Governance in Organizations: A Multi-Level Approach to Responsible AI Use. *Procedia Computer Science*, 270, 4144–4154. <https://doi.org/10.1016/j.procs.2025.09.539>
- [17]. Wahid, R., Mero, J., & Ritala, P. (2025). Technology-enabled democratization: Impact of generative AI on content marketing agencies. *Industrial Marketing Management*, 131, 1–16. <https://doi.org/10.1016/j.indmarman.2025.09.007>