



(RESEARCH ARTICLE)



A TIERED LIFECYCLE GOVERNANCE FRAMEWORK FOR AGENTIC ARTIFICIAL INTELLIGENCE IN FINANCIAL SERVICES

Jimmy Kinyonyi Bagonza ^{1,*} and Mathias Ndungu ²

¹ School of Computer and Information Sciences, University of the Cumberlands, Kentucky, USA.

² Kogod School of Business, American University.

* Corresponding Author

World Journal of Advanced Research and Reviews, 2026, 32(01), 1283–1294

Publication history: Received on 06 September 2026; revised on 18 October 2026; accepted on 20 October 2026

Article DOI: <https://doi.org/10.30574/wjarr.2026.32.1.2684>

Copyright © 2026. Author(s) retain the copyright of this article. This article is published under the terms of the Creative Commons Attribution License 4.0

ABSTRACT

The rapid emergence of agentic artificial intelligence (AI) presents significant governance challenges for financial services, where autonomous systems increasingly support high-impact financial activities. Existing regulatory approaches remain fragmented and continue to rely primarily on non-binding supervisory guidance, leaving important gaps in lifecycle oversight, accountability, and systemic risk management. This paper proposes the Tiered Lifecycle Governance (TLG) Framework, a strategic governance model that integrates lifecycle-based AI oversight with responsive, risk-based regulation. The framework combines three transaction-risk tiers with four AI lifecycle stages to establish proportionate oversight obligations that balance innovation, regulatory accountability, and financial stability. Drawing on technological, economic, legal, and institutional perspectives, the framework is applied to the 2026 U.S. interagency review of model risk management guidance (SR 26-2) to demonstrate how lifecycle governance can strengthen oversight of agentic AI in financial services. The analysis shows that the TLG Framework provides a structured alternative to the existing guidance-versus-rulemaking approach by aligning regulatory obligations with transaction risk rather than institutional size or technology labels. The paper concludes that binding, lifecycle-based oversight offers a practical pathway for improving resilience, transparency, and accountability in financial AI governance while providing a transferable governance model for other emerging AI policy domains.

Keywords: Agentic Artificial Intelligence; Financial Services; AI Governance; Lifecycle Governance; Risk-Based Regulation; Model Risk Management; Financial Regulation; Responsible AI.

1 INTRODUCTION

The rapid advancement of agentic artificial intelligence (AI) is transforming financial services by enabling autonomous systems to support high-impact activities such as credit assessment, investment management, and financial decision-making. While these capabilities offer significant opportunities for innovation and operational efficiency, they also introduce new governance challenges related to accountability, transparency, systemic risk, and regulatory oversight. Existing supervisory approaches have not kept pace with these developments, leaving important questions regarding the appropriate governance of increasingly autonomous AI systems.

In the United States, these challenges have become particularly evident following the release of the revised interagency model risk management guidance (SR 26-2 and OCC Bulletin 2026-13) by the Federal Reserve, the Office of the Comptroller of the Currency (OCC), and the Federal Deposit Insurance Corporation (FDIC). The revised guidance explicitly excludes generative and agentic AI from binding supervisory requirements, instead deferring consideration of these technologies to a future Request for Information (RFI) (Federal Reserve, 2026). This decision presents a significant strategic policy challenge by leaving oversight dependent on non-binding guidance despite growing evidence of regulatory fragmentation and increasing systemic risk associated with autonomous AI systems.

Previous research identifies three interconnected concerns. First, scenario-based forecasting suggests that regulatory fragmentation is the most likely trajectory for generative AI governance in the absence of deliberate institutional investment (Bagonza & Mathias, 2026). Second, sector-specific evidence indicates that fragmentation is already emerging within financial services, with many institutions lacking the capability to effectively control malfunctioning AI agents (Financial Stability Board, 2025; Kapron, 2026). Third, existing legal and policy analyses indicate that federal banking agencies possess sufficient statutory authority to establish binding lifecycle oversight under existing safety-and-soundness legislation, although that authority has not yet been fully exercised (Consumer Financial Protection Bureau, 2022; Federal Reserve, 2026; Judge et al., 2024). Together, these findings highlight the growing disconnect between technological capability and existing regulatory practice.

Responding to this challenge requires a governance approach that extends beyond the traditional choice between voluntary guidance and comprehensive regulation. Rather than treating all AI applications uniformly, effective governance should recognize differences in transaction risk while addressing oversight throughout the entire AI lifecycle. This paper addresses that need by proposing the **Tiered Lifecycle Governance (TLG) Framework**, which integrates lifecycle-based AI governance with responsive, risk-based regulation to establish proportionate oversight obligations according to both transaction risk and lifecycle stage (Ayres & Braithwaite, 1992; Baldwin et al., 2012; Judge et al., 2024).

The paper contributes to the emerging literature on AI governance by synthesizing technological, economic, legal, and institutional perspectives into a unified governance framework and demonstrating its application to the 2026 U.S. interagency review of model risk management guidance (SR 26-2). Through this application, the study illustrates how a tiered lifecycle approach can strengthen resilience, accountability, and transparency while balancing innovation with financial stability. The framework further provides a strategic governance model to inform future policy development for agentic AI in financial services and other emerging technology domains.

2 LITERATURE REVIEW

2.1 Agentic Artificial Intelligence in Financial Services

The emergence of agentic artificial intelligence (AI) represents a significant evolution beyond conventional AI systems by enabling autonomous agents to perceive, reason, plan, and execute complex tasks with limited human intervention. Unlike traditional predictive models or generative AI systems that primarily respond to user prompts, agentic AI systems can initiate and coordinate multi-step actions, interact with external environments, and adapt their behavior in response to dynamic objectives (Judge et al., 2024; Gong, 2026). These capabilities have accelerated interest in deploying agentic AI across financial services, where autonomous systems increasingly support activities such as credit underwriting, fraud detection, portfolio management, customer service, and payment processing.

The financial sector presents a particularly challenging environment for deploying agentic AI because autonomous decisions can directly affect financial stability, consumer protection, and market integrity. While AI-driven automation offers substantial opportunities for operational efficiency, improved risk assessment, and enhanced customer experience, it also introduces governance concerns that extend beyond those associated with conventional machine learning models. Financial Stability Board (2024) observes that increasing reliance on AI creates new channels through which operational failures, model errors, third-party dependencies, and correlated decision-making may amplify systemic risk across financial institutions. These risks become more pronounced as AI systems gain greater autonomy in executing financial decisions without continuous human oversight.

Recent scholarship argues that existing regulatory frameworks have not kept pace with advances in agentic AI. Judge et al. (2024) contend that traditional model governance approaches, which primarily evaluate system outputs after deployment, are inadequate for governing emergent AI systems whose behaviors evolve throughout their operational lifecycle. Similarly, Gong (2026) highlights that autonomous AI agents operating within financial markets introduce novel governance challenges, including explainability, accountability, controllability, and coordinated market behavior that are not sufficiently addressed by existing supervisory mechanisms. These concerns have prompted regulators and

international organizations to reconsider the adequacy of current model risk management practices for increasingly autonomous AI applications.

Growing regulatory attention reflects recognition that governance challenges extend beyond individual institutions to the broader financial ecosystem. The Financial Stability Board (2025) reports growing concern about the concentration of AI service providers, interconnected technological infrastructure, and limited institutional capacity to monitor or intervene when autonomous agents behave unexpectedly. Such developments suggest that effective governance must move beyond institution-specific compliance towards regulatory approaches that manage lifecycle risks associated with autonomous AI systems operating across complex financial environments.

Although the literature provides important insights into the risks posed by agentic AI in financial services, it largely focuses on identifying emerging governance challenges rather than on proposing integrated governance mechanisms to address them. Existing studies emphasize systemic risk, model oversight, and regulatory preparedness, yet provide limited guidance on how lifecycle governance and proportionate regulatory supervision can be combined into a coherent governance framework. This limitation provides the foundation for examining lifecycle governance approaches and responsive regulatory models in the following sections.

2.2 Lifecycle Governance of Artificial Intelligence

The increasing autonomy of agentic AI has prompted growing recognition that governance must extend beyond evaluating system outputs to encompass the entire AI lifecycle. Traditional model risk management approaches have largely focused on validating models before deployment and monitoring performance after implementation. However, these approaches are less effective for agentic AI systems whose behavior evolves through continuous interaction with data, users, and external environments (Judge et al., 2024). Consequently, lifecycle governance has emerged as a more comprehensive paradigm that incorporates oversight from system design and data governance through deployment, monitoring, and incident response.

Judge et al. (2024) argue that governance of general-purpose AI should mirror the regulatory approaches employed in highly regulated industries such as aviation and nuclear energy, where oversight extends throughout the complete system lifecycle rather than concentrating solely on operational outcomes. Such an approach recognizes that many AI-related failures originate during system design, training data preparation, model development, or deployment rather than at the point of use. Accordingly, effective governance requires continuous verification of model behavior, documentation of development processes, independent auditing, and ongoing monitoring after deployment.

Recent regulatory developments demonstrate a growing shift toward lifecycle-based governance. The National Institute of Standards and Technology's Artificial Intelligence Risk Management Framework (NIST AI RMF 1.0) promotes continuous risk management throughout AI system development and operation by integrating governance, mapping, measurement, and management functions (National Institute of Standards and Technology, 2023). Similarly, the European Union's Artificial Intelligence Act adopts a lifecycle perspective by imposing obligations relating to data governance, transparency, human oversight, technical documentation, post-market monitoring, and incident reporting for high-risk AI systems (European Parliament, 2024). Although the revised U.S. interagency model risk management guidance (SR 26-2) continues to rely primarily on existing supervisory principles, the increasing complexity of agentic AI highlights the need for more explicit lifecycle-oriented regulatory approaches (Federal Reserve, 2026).

Despite growing consensus on the importance of lifecycle governance, existing literature offers limited guidance on how lifecycle oversight should be calibrated to differences in AI risk exposure. Most proposed frameworks assume relatively uniform governance obligations across AI applications, creating potential compliance burdens for lower-risk systems while failing to distinguish among varying levels of financial and systemic risk adequately. This limitation suggests the need to complement lifecycle governance with proportionate regulatory approaches.

2.3 Responsive and Risk-Based Regulatory Approaches

Responsive regulation has become one of the most influential theories in modern regulatory governance by recognizing that regulatory interventions should be proportionate to both the level of risk and the behavior of regulated entities. Rather than relying exclusively on prescriptive rules or voluntary compliance, responsive regulation advocates graduated enforcement that escalates according to demonstrated risk and regulatory performance (Ayres & Braithwaite, 1992). Baldwin et al. (2012) further argue that effective regulation should balance flexibility, accountability, and proportionality by matching supervisory intensity to the characteristics of regulated activities.

Risk-based regulation has been widely adopted across financial supervision because it enables regulators to allocate supervisory resources according to the potential consequences of institutional failures. Instead of imposing identical compliance obligations across all regulated entities, risk-based supervision prioritizes activities that present the

greatest threats to financial stability, consumer protection, or market integrity. This approach has become increasingly relevant as financial institutions deploy AI systems with varying levels of autonomy and decision-making authority.

The emergence of agentic AI further reinforces the need for proportionate regulatory oversight. Autonomous AI systems differ substantially in their operational scope, ranging from low-risk informational chatbots to high-risk autonomous trading systems capable of executing complex financial transactions. Applying identical governance requirements across these heterogeneous applications may either impose unnecessary compliance costs or fail to adequately mitigate systemic risks. Consequently, several regulatory authorities have increasingly emphasized proportionality in developing AI governance strategies, recognizing that supervisory requirements should reflect differences in operational risk rather than technology alone (Financial Conduct Authority, 2026).

Although responsive regulation provides a valuable theoretical foundation for proportionate supervision, it offers relatively limited guidance regarding governance responsibilities throughout the AI lifecycle. Conversely, lifecycle governance provides comprehensive oversight across the stages of AI development but generally assumes relatively uniform governance requirements. Integrating these complementary approaches therefore offers the potential to develop governance mechanisms that are both lifecycle-oriented and proportionate to transaction risk.

2.4 Cross-Disciplinary Perspectives on AI Governance

Governance of agentic AI extends beyond technological considerations and increasingly requires integration of economic, legal, and institutional perspectives. Recent scholarship suggests that effective AI governance cannot rely on technical safeguards alone but must address broader questions concerning market incentives, regulatory authority, organizational capability, and public accountability (Asensio et al., 2025; Hanisch et al., 2023).

From a technological perspective, agentic AI presents governance challenges because autonomous systems exhibit emergent behaviors that cannot always be predicted during development. Judge et al. (2024) argue that such characteristics render traditional output-based regulation insufficient and instead support governance mechanisms that emphasize independent testing, continuous monitoring, and lifecycle oversight. These approaches acknowledge that AI failures often stem from interactions among system design, training data, operational environments, and adaptive learning processes.

Economic perspectives similarly emphasize market failures associated with AI deployment. Information asymmetry limits consumers' and regulators' ability to evaluate AI quality and governance practices (Akerlof, 1970), while interconnected AI infrastructures may generate negative externalities that extend beyond individual financial institutions (Financial Stability Board, 2024). Gordon and Loeb (2002) further argue that information security investments often remain below socially optimal levels without appropriate regulatory incentives, suggesting that governance frameworks should align accountability with the entities generating systemic risks.

The legal dimension of AI governance focuses on establishing clear statutory authority and regulatory accountability. Existing U.S. banking legislation provides mechanisms for addressing safety-and-soundness risks associated with financial institutions, while consumer protection frameworks increasingly emphasize explainability and fairness in automated decision-making (Consumer Financial Protection Bureau, 2022). At the same time, evolving judicial standards have increased the importance of evidence-based rulemaking when introducing new regulatory obligations, reinforcing the need for governance frameworks that are both legally defensible and operationally practical.

Institutional perspectives highlight the evolution of governance from traditional supervisory models towards digitally enabled regulatory systems. Hanisch et al. (2023) distinguish between analog, augmented, and automated governance, arguing that effective regulation increasingly requires combining human oversight with technology-enabled monitoring. Asensio et al. (2025) similarly contend that integrating technological capability with institutional capacity represents a critical prerequisite for effective AI governance, particularly in complex policy environments characterized by rapid technological change.

Collectively, these perspectives demonstrate that the governance of agentic AI requires an integrated approach that addresses technological uncertainty, economic incentives, legal authority, and institutional capacity simultaneously. Such multidisciplinary foundations provide the theoretical basis for developing a governance framework that balances innovation with accountability and systemic resilience.

2.5 Research Gap and Conceptual Foundation

The literature demonstrates a growing consensus that agentic AI poses governance challenges that cannot be adequately addressed by conventional model risk management approaches. Existing scholarship has advanced understanding of lifecycle AI governance, responsive regulation, financial AI risks, and digital governance. Collectively, these studies recognize the need for governance approaches that enhance accountability, transparency, and resilience

as AI systems become increasingly autonomous (Asensio et al., 2025; Financial Stability Board, 2024; Hanisch et al., 2023; Judge et al., 2024).

Despite these advances, important gaps remain. Studies advocating lifecycle governance emphasize continuous oversight throughout the AI development and deployment process but provide limited guidance on calibrating governance obligations according to differences in transaction risk (Judge et al., 2024). Conversely, responsive and risk-based regulatory approaches offer well-established principles for proportionate supervision but do not explicitly address governance responsibilities across the AI lifecycle (Ayres & Braithwaite, 1992; Baldwin et al., 2012). Similarly, technological, economic, legal, and institutional perspectives identify complementary governance challenges but are generally examined in isolation rather than integrated into a unified approach to financial services governance.

Consequently, the current literature provides limited guidance on how lifecycle oversight, responsive regulation, and multidisciplinary governance perspectives can be combined into a coherent framework for regulating agentic AI in financial services. Addressing this gap requires an integrated governance approach that aligns oversight responsibilities with both transaction risk and the AI lifecycle while balancing innovation, accountability, and systemic resilience. The following section describes the research design adopted to develop such a governance framework.

2.6 Research Design

This study adopts a conceptual research design based on strategic policy synthesis to develop a governance framework for agentic artificial intelligence (AI) in financial services. Rather than generating new empirical data, the study synthesizes established theoretical perspectives and contemporary policy developments to address an emerging regulatory challenge for financial-sector AI governance. Conceptual research is particularly appropriate for rapidly evolving technological domains where regulatory institutions and governance frameworks are still developing (Asensio et al., 2025).

The framework was developed through an integrative review of four complementary bodies of literature: lifecycle governance of artificial intelligence, responsive and risk-based regulation, financial AI governance, and digital governance. These perspectives were selected because they collectively address the technological, economic, legal, and institutional dimensions of governing increasingly autonomous AI systems. Lifecycle governance provides the foundation for continuous oversight throughout the AI development and deployment process (Judge et al., 2024), while responsive regulation contributes the principle of proportionality, aligning regulatory obligations with levels of transaction risk (Ayres & Braithwaite, 1992; Baldwin et al., 2012). These theoretical perspectives were complemented by contemporary regulatory guidance and policy documents addressing AI governance within financial services, including publications from the Federal Reserve, the Financial Stability Board, the National Institute of Standards and Technology, and the European Union.

Building on these foundations, the study develops the Tiered Lifecycle Governance (TLG) Framework by integrating transaction-risk classification with lifecycle-based oversight responsibilities. The framework combines three transaction-risk tiers with four AI lifecycle stages to establish proportionate governance obligations that balance innovation, accountability, and systemic resilience. Rather than proposing an entirely new regulatory philosophy, the framework synthesizes existing governance principles into a unified structure applicable to agentic AI within financial services.

To demonstrate its practical relevance, the proposed framework is applied to the 2026 U.S. interagency review of model risk management guidance (SR 26-2). This policy context provides an appropriate case for illustrating how the framework can inform regulatory decision-making concerning lifecycle oversight, proportionate supervision, and governance of increasingly autonomous AI systems. The case application serves as an analytical demonstration of the framework's utility rather than an empirical validation, providing a structured basis for evaluating its policy implications and broader applicability to emerging AI governance challenges.

2.7 Framework Design

Building upon these complementary theoretical foundations, this study proposes the Tiered Lifecycle Governance (TLG) Framework, which integrates transaction-risk classification with lifecycle-based AI oversight. The framework combines three transaction-risk tiers—low, moderate, and high/systemic risk—with four lifecycle stages comprising design and data governance, pre-deployment testing, deployment and monitoring, and incident response. The interaction of these dimensions establishes proportionate governance obligations that correspond to both the operational risk posed by an AI system and its lifecycle stage.

Unlike governance approaches that apply uniform compliance requirements across diverse AI applications, the TLG Framework recognizes that regulatory obligations should increase progressively with transaction risk and potential systemic consequences. Consequently, lower-risk applications are subject to lighter supervisory requirements. At the

same time, autonomous AI systems capable of generating significant financial or systemic impacts are expected to satisfy more comprehensive governance obligations throughout their operational lifecycle.

The framework is organized around a central strategic objective of promoting resilient, accountable, and inclusive governance of agentic AI within financial services. Four analytical perspectives- technological, economic, legal, and institutional- inform the governance model and collectively support the framework's lifecycle oversight engine. Two policy decision levers further guide regulatory implementation. The first concerns the transition from voluntary supervisory guidance to binding regulatory requirements, while the second determines the degree of coordination required across regulatory jurisdictions. Together, these components enable policymakers to balance innovation with accountability while managing the trade-off between regulatory burden and systemic risk.




Figure 1: Tiered Lifecycle Governance (TLG) Framework for Agentic Artificial Intelligence in Financial Services

2.8 Multidisciplinary Foundations of the Framework

The TLG Framework is supported by four complementary analytical perspectives that collectively strengthen its governance approach.

From a technological perspective, agentic AI systems exhibit emergent behaviors that cannot always be anticipated during development, rendering static regulatory approaches insufficient to manage evolving operational risks (Judge et al., 2024). Consequently, the framework emphasizes independent testing, continuous monitoring, and lifecycle oversight as essential governance mechanisms for increasingly autonomous AI systems.

From an economic perspective, the framework addresses information asymmetry, systemic externalities, and the unequal distribution of regulatory costs associated with AI deployment. Disclosure requirements, action-level logging, and proportionate governance obligations aim to improve transparency while ensuring that organizations that generate higher systemic risks bear correspondingly greater governance responsibilities (Akerlof, 1970; Financial Stability Board, 2024; Gordon & Loeb, 2002).

The legal perspective recognizes that effective governance must operate within existing statutory and regulatory authority while remaining adaptable to evolving legal standards. Existing banking legislation and consumer protection requirements provide an important foundation for lifecycle governance, although emerging judicial expectations increasingly favor evidence-based and procedurally robust regulatory approaches (Consumer Financial Protection Bureau, 2022; Federal Reserve, 2026).

Finally, the institutional perspective acknowledges that governance capacity must evolve alongside technological capability. Contemporary scholarship increasingly characterizes AI governance as progressing from traditional supervisory models towards digitally enabled governance systems that integrate human oversight with automated monitoring and regulatory intelligence (Asensio et al., 2025; Hanisch et al., 2023). The TLG Framework reflects this transition by combining lifecycle oversight with proportionate regulatory supervision to support adaptive governance of increasingly autonomous financial AI systems.

Collectively, these technological, economic, legal, and institutional perspectives establish the multidisciplinary foundation upon which the Tiered Lifecycle Governance Framework is constructed. By integrating these complementary dimensions into a single governance model, the framework provides the conceptual basis for its application to contemporary regulatory decision-making in financial services.

2.9 Design Principles of the Tiered Lifecycle Governance Framework

The Tiered Lifecycle Governance (TLG) Framework is guided by four interrelated design principles that collectively support proportionate and adaptive governance of agentic AI in financial services.

The first principle is lifecycle oversight, which recognizes that governance should extend across the entire AI lifecycle, including system design, data governance, pre-deployment testing, operational monitoring, and incident response. Continuous oversight enables regulators and institutions to identify and mitigate risks before they escalate into operational or systemic failures (Judge et al., 2024).

The second principle is proportionality, whereby regulatory obligations are aligned with the level of transaction risk presented by an AI system. Rather than applying uniform compliance requirements, the framework adopts a graduated approach in which supervisory expectations increase according to the potential impact of autonomous decision-making on consumers, financial institutions, and market stability (Ayres & Braithwaite, 1992; Baldwin et al., 2012).

The third principle is accountability, which emphasizes transparency, auditability, and clearly defined governance responsibilities throughout the AI lifecycle. Effective accountability requires mechanisms such as documentation, action-level logging, independent evaluation, and clearly assigned organizational responsibilities to ensure that autonomous systems remain subject to meaningful human oversight (Consumer Financial Protection Bureau, 2022; Financial Stability Board, 2024).

The fourth principle is adaptive governance, recognizing that governance frameworks must evolve alongside technological innovation. Agentic AI systems continuously develop in capability and complexity, requiring regulatory approaches that support ongoing monitoring, periodic reassessment, and institutional learning rather than static compliance requirements (Asensio et al., 2025; Hanisch et al., 2023).

Together, these design principles provide the normative foundation of the TLG Framework and guide the implementation of governance obligations across transaction-risk tiers and AI lifecycle stages. By integrating lifecycle oversight, proportionality, accountability, and adaptive governance into a single conceptual model, the framework provides a structured basis for strengthening resilience, transparency, and regulatory effectiveness in the governance of agentic AI within financial services.

3 APPLICATION OF THE TIERED LIFECYCLE GOVERNANCE FRAMEWORK TO THE U.S. SR 26-2 REGULATORY REVIEW

The Tiered Lifecycle Governance (TLG) Framework is applied to the 2026 U.S. interagency review of model risk management guidance (SR 26-2) to illustrate how proportionate lifecycle governance can inform regulatory decision-making for agentic AI within financial services. The revised guidance issued by the Federal Reserve, Office of the Comptroller of the Currency (OCC), and Federal Deposit Insurance Corporation (FDIC) explicitly excludes generative and agentic AI from binding supervisory requirements, instead deferring consideration of these technologies to a future Request for Information (RFI) (Federal Reserve, 2026). This policy context provides an appropriate case for demonstrating how the proposed framework addresses the limitations of existing supervisory approaches.

Using the TLG Framework, the governance requirements applicable to an AI system are determined by its transaction-risk exposure rather than institutional size or technology classification. This approach enables supervisory obligations to increase in proportion to operational and systemic risk while maintaining lifecycle oversight throughout system development and deployment.

- At **Tier 1**, AI applications associated with low transaction risk, such as informational customer-service chatbots or advisory systems that do not independently execute financial decisions, are subject to proportionate

governance requirements. These obligations include organizational self-assessment, basic documentation, and periodic supervisory reporting, thereby maintaining regulatory oversight without imposing unnecessary compliance burdens on lower-risk applications. This approach is broadly consistent with principles-based supervisory models adopted by regulatory authorities, such as the Financial Conduct Authority (2026).

- At **Tier 2**, AI systems supporting higher-impact financial decisions, including automated credit underwriting and risk assessment, require enhanced lifecycle governance. These systems are subject to structured pre-deployment testing, mapping against recognized AI risk management standards, compulsory incident reporting, and strengthened governance documentation. Such measures align with emerging supervisory expectations for autonomous financial AI while recognizing the growing operational consequences of partially autonomous decision-making.
- At **Tier 3**, autonomous AI systems capable of initiating or executing high-risk financial transactions, including autonomous trading platforms and fund-transfer agents, are subject to the most comprehensive governance requirements. These include independent pre-deployment audits, continuous operational monitoring, mandatory action-level logging, and verified, non-removable shutdown capabilities to ensure meaningful human intervention in the event of abnormal system behavior. These governance obligations directly address concerns about institutional preparedness to control autonomous AI systems operating at a systemic scale (Financial Stability Board, 2025; Kapron, 2026).

Beyond defining governance obligations across transaction-risk tiers, the TLG Framework also informs strategic regulatory decision-making through two policy implementation levers. The first concerns the transition from supervisory guidance to binding regulation. Rather than relying solely on interpretive guidance, the framework supports the development of evidence-based, notice-and-comment rulemaking that can establish durable lifecycle governance requirements consistent with evolving legal standards following *Loper Bright Enterprises v. Raimondo* (2024). The second lever concerns regulatory coordination across jurisdictions. Recognizing increasing international divergence in AI governance approaches, the framework supports targeted coordination on minimum governance standards, incident reporting, and AI registries, while allowing flexibility for jurisdiction-specific implementation.

The application of the framework demonstrates that governance decisions should be driven primarily by transaction risk and lifecycle stage, rather than by technology labels or organizational size alone. This approach enables regulators to balance innovation with accountability while reducing systemic vulnerabilities associated with increasingly autonomous financial AI systems. By integrating lifecycle oversight with proportionate supervision, the TLG Framework provides a structured mechanism to strengthen resilience, transparency, and regulatory consistency in emerging financial AI governance.

4 DISCUSSION

4.1 Strengths of the Tiered Lifecycle Governance Framework

The primary strength of the Tiered Lifecycle Governance (TLG) Framework lies in its integration of lifecycle governance and responsive regulation into a single governance model for agentic AI within financial services. Rather than framing regulatory choices as a binary decision between voluntary guidance and comprehensive rulemaking, the framework establishes proportionate governance obligations that vary according to transaction risk and AI lifecycle stage. This approach enables regulators to balance innovation with accountability while reducing unnecessary compliance burdens for lower-risk AI applications.

A further strength of the framework is its multidisciplinary foundation. By integrating technological, economic, legal, and institutional perspectives, the framework provides a comprehensive governance approach that addresses the diverse challenges associated with autonomous AI systems. This integration reduces fragmentation in policy analysis and supports more coherent regulatory decision-making, consistent with Asensio et al. (2025), who observe that effective AI governance requires coordination across multiple analytical domains.

The framework also enhances regulatory adaptability by recognizing that governance obligations should evolve alongside technological capability. Its tiered structure allows supervisory intensity to increase progressively with transaction risk while maintaining continuous oversight throughout the AI lifecycle. This flexibility provides regulators with a structured mechanism for responding to emerging AI capabilities without requiring fundamental redesign of the governance architecture.

Limitations of the Framework

Despite its conceptual strengths, the proposed framework has several limitations. First, assigning AI systems to appropriate transaction-risk tiers requires robust methods for assessing operational and systemic risk. As autonomous AI capabilities continue to evolve, determining appropriate risk classifications may become increasingly complex, reflecting the broader measurement challenges identified by Judge et al. (2024).

Second, effective implementation depends on the technical capacity of regulatory institutions. Lifecycle governance requires expertise in AI systems, model validation, continuous monitoring, and digital supervision. Asensio et al. (2025) observe that institutional capabilities vary considerably across jurisdictions, potentially affecting the consistency with which lifecycle governance can be implemented.

Third, the framework has been developed within the context of the United States financial regulatory system and therefore reflects existing statutory authorities, supervisory practices, and institutional arrangements. While the conceptual architecture may be transferable, implementation within other jurisdictions would require adaptation to local legal frameworks, regulatory mandates, and governance structures.

Finally, the framework has been demonstrated through conceptual application rather than empirical validation. Although application to the SR 26-2 regulatory review illustrates its practical utility, further empirical research is required to evaluate its operational effectiveness across diverse financial institutions and regulatory environments.

4.2 Policy Implications

The findings of this study suggest that future AI governance should move beyond technology-specific regulation towards governance models that combine lifecycle oversight with proportionate supervision. Such an approach enables regulatory requirements to reflect differences in transaction risk while ensuring continuous governance throughout AI development, deployment, and operation.

For financial regulators, the framework provides a structured basis for strengthening model risk governance as AI systems become increasingly autonomous. By aligning governance obligations with transaction risk rather than institutional size or technology labels, regulators can improve accountability while maintaining flexibility for lower-risk AI applications. This proportionate approach may also reduce unnecessary compliance burdens while directing supervisory resources toward systems with the greatest potential to generate systemic consequences.

The framework further highlights the importance of multidisciplinary governance. Effective oversight of agentic AI requires a coordinated consideration of technological capability, economic incentives, legal authority, and institutional capacity, rather than addressing these dimensions in isolation. Such integration is essential for developing governance approaches that can adapt to ongoing advances in autonomous AI.

4.3 Strategic Recommendations

Based on the findings of this study, four strategic recommendations emerge for policymakers and financial regulators.

First, regulatory authorities should transition from reliance on supervisory guidance towards evidence-based lifecycle governance for higher-risk agentic AI applications. Where transaction risk poses risks to financial stability or consumer protection, governance obligations should be established through transparent regulatory processes that provide durable and enforceable oversight.

Second, governance requirements for high-risk autonomous AI systems should incorporate independent pre-deployment evaluation, continuous operational monitoring, action-level logging, and verified shutdown mechanisms. These measures strengthen accountability while ensuring meaningful human oversight of increasingly autonomous financial AI systems.

Third, regulatory frameworks should support proportionate governance by aligning supervisory obligations with transaction risk rather than organizational size or technology classification. Such an approach promotes innovation while ensuring that governance expectations remain commensurate with potential operational and systemic consequences.

Finally, regulators should strengthen international cooperation on AI governance by developing common minimum standards for incident reporting, AI registries, and lifecycle governance principles. While comprehensive regulatory harmonization may remain challenging, targeted coordination can improve regulatory consistency and reduce governance fragmentation across jurisdictions.

The implementation of these recommendations would strengthen resilience, transparency, and accountability in financial AI governance while supporting responsible innovation as agentic AI continues to evolve.

5 GENERALISABILITY OF THE TIERED LIFECYCLE GOVERNANCE FRAMEWORK

The Tiered Lifecycle Governance (TLG) Framework was developed in the context of financial services; however, its underlying conceptual architecture is sufficiently flexible to support its application across other policy domains characterized by autonomous AI systems, fragmented regulatory environments, and evolving risk profiles. By integrating lifecycle governance with responsive, risk-based regulation, the framework provides a structured approach to governance that is not limited to a single technological application or regulatory sector.

The core strength of the framework lies in its integration of transaction-risk classification with lifecycle oversight. This governance logic applies wherever autonomous AI systems perform functions that can affect public safety, economic stability, or individual rights. Potential areas of application include healthcare, autonomous transportation, critical infrastructure, cybersecurity, and algorithmic content moderation, where AI systems increasingly operate with varying degrees of autonomy and where governance must balance innovation with accountability.

The technological and institutional foundations of the framework are particularly transferable across policy domains. Autonomous AI systems share common characteristics, including adaptive behavior, operational uncertainty, and continuous interaction with dynamic environments, making lifecycle oversight an important governance requirement across sectors (Judge et al., 2024). Likewise, the institutional transition from traditional supervisory approaches towards digitally enabled governance systems is increasingly recognized across diverse regulatory contexts (Hanisch et al., 2023). These shared characteristics support broader application of the framework's principles of lifecycle oversight, proportionality, accountability, and adaptive governance.

The economic rationale underlying the framework is similarly applicable beyond financial services. Information asymmetry, systemic externalities, and uneven distribution of governance costs are common features of many AI-enabled systems, suggesting that governance mechanisms should continue to align regulatory obligations with the risk level generated by autonomous technologies (Akerlof, 1970; Gordon & Loeb, 2002). Consequently, proportionate governance based on transaction or operational risk may provide an effective foundation for regulating AI systems across multiple sectors.

Nevertheless, successful transfer of the framework requires careful adaptation to sector-specific legal and institutional contexts. The legal foundations supporting governance within U.S. financial services, including safety-and-soundness authority and existing model risk management practices, may not exist in other jurisdictions or policy domains. Application within the European Union, for example, would require alignment with the Artificial Intelligence Act and its risk-based conformity assessment requirements, while implementation within the United Kingdom would need to reflect the Financial Conduct Authority's principles-based regulatory approach (European Parliament, 2024; Financial Conduct Authority, 2026). Similar adaptation would be necessary for other sectors where regulatory mandates, institutional responsibilities, and acceptable levels of operational risk differ substantially from those of financial services.

Accordingly, the principal contribution of the TLG Framework is not the specification of fixed regulatory requirements but the provision of a transferable governance architecture adaptable to different regulatory environments. While the specific transaction-risk thresholds, lifecycle obligations, and supervisory mechanisms should be calibrated according to sectoral and jurisdictional requirements, the underlying governance logic provides a coherent conceptual foundation for regulating increasingly autonomous AI systems. Future research should explore the framework's applicability across additional policy domains and evaluate its performance under diverse regulatory and institutional settings.

6 CONCLUSION

The rapid advancement of agentic artificial intelligence is transforming financial services by introducing increasingly autonomous systems that support high-impact financial decision-making. While these technologies offer substantial opportunities for innovation, efficiency, and improved service delivery, they also create governance challenges that extend beyond the scope of existing model risk management approaches. Current regulatory frameworks remain fragmented and continue to rely predominantly on supervisory guidance, providing limited lifecycle oversight for autonomous AI systems operating in increasingly complex financial environments.

This study addressed this governance challenge by developing the Tiered Lifecycle Governance (TLG) Framework, a conceptual governance model that integrates lifecycle AI governance with responsive, risk-based regulation. By combining transaction-risk classification with lifecycle oversight, the framework provides a structured approach for

aligning governance obligations with both the operational risk and developmental stage of agentic AI systems. The framework further synthesizes technological, economic, legal, and institutional perspectives to establish a multidisciplinary governance model that strengthens accountability, transparency, and systemic resilience.

Application of the framework to the 2026 U.S. interagency review of model risk management guidance (SR 26-2) demonstrates how proportionate lifecycle governance can support more effective regulatory decision-making for financial AI. Rather than applying uniform governance obligations across diverse AI applications, the framework illustrates how supervisory expectations can be calibrated according to transaction risk while maintaining continuous oversight throughout the AI lifecycle. This proportionate approach provides a practical alternative to the current guidance-versus-rulemaking debate by balancing innovation with regulatory accountability.

The study also contributes to the emerging literature on AI governance by integrating two complementary regulatory traditions lifecycle governance and responsive regulation into a unified conceptual framework tailored to the governance of agentic AI in financial services. In doing so, it provides policymakers and regulators with a structured governance approach that is adaptable to evolving technological capabilities while remaining grounded in established regulatory principles.

Although developed in the context of financial services, the conceptual architecture of the TLG Framework is broadly relevant to other policy domains where increasingly autonomous AI systems pose comparable governance challenges. Future research should evaluate the framework through empirical implementation, comparative regulatory analysis, and cross-sector applications to assess its effectiveness, adaptability, and policy impact across diverse institutional and jurisdictional contexts.

In conclusion, effective governance of agentic AI requires regulatory approaches that move beyond static compliance towards adaptive, lifecycle-oriented supervision that reflects differences in transaction risk and evolving technological capability. The Tiered Lifecycle Governance Framework offers a structured contribution to this emerging field by providing a coherent governance model to strengthen resilience, transparency, accountability, and responsible innovation in the regulation of autonomous AI systems in financial services.

REFERENCES

- [1] Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3), 488–500. <https://doi.org/10.2307/1879431>
- [2] Asensio, I. O., Moore, E. C., Ulibarri, N., Simsekler, C. M., Lan, T., & Rivero, G. (2025). *Data technologies and analytics for policy and governance: A landscape review*. Cambridge University Press.
- [3] Ayres, I., & Braithwaite, J. (1992). *Responsive regulation: Transcending the deregulation debate*. Oxford University Press.
- [4] Bagonza, J. K., & Ndungu, M. (2026). A definitive forecast on the regulation of generative artificial intelligence for the years 2025 to 2035. *World Journal of Advanced Research and Reviews*, 31(1), 523–532. <https://doi.org/10.30574/wjarr.2026.31.1.1855>
- [5] Baldwin, R., Cave, M., & Lodge, M. (2012). *Understanding regulation: Theory, strategy, and practice* (2nd ed.). Oxford University Press.
- [6] Consumer Financial Protection Bureau. (2022). *Consumer Financial Protection Circular 2022-03: Adverse action notification requirements in connection with credit decisions based on complex algorithms*. <https://www.consumerfinance.gov/compliance/circulars/circular-2022-03-adverse-action-notification-requirements-in-connection-with-credit-decisions-based-on-complex-algorithms/>
- [7] European Parliament. (2024). *Regulation (EU) 2024/1689 laying down harmonized rules on artificial intelligence (the Artificial Intelligence Act)*. Official Journal of the European Union.
- [8] Federal Reserve. (2026). *Supervisory Letter SR 26-2: Revised guidance on model risk management*. Board of Governors of the Federal Reserve System. <https://www.federalreserve.gov/supervisionreg/srletters/SR2602.htm>
- [9] Financial Conduct Authority. (2026). *AI and the FCA: Our approach*. <https://www.fca.org.uk/firms/innovation/ai-approach>
- [10] Financial Stability Board. (2024). *The financial stability implications of artificial intelligence*. <https://www.fsb.org/uploads/P14112024.pdf>

- [11] Financial Stability Board. (2025). *Monitoring the adoption of artificial intelligence and related vulnerabilities in the financial sector*. <https://www.fsb.org/2025/10/monitoring-adoption-of-artificial-intelligence-and-related-vulnerabilities-in-the-financial-sector/>
- [12] Gong, H. (2026). AI agents in financial markets: Architecture, applications, and systemic implications. *FinTech*, 5(2), Article 34. <https://doi.org/10.3390/fintech5020034>
- [13] Gordon, L. A., & Loeb, M. P. (2002). The economics of information security investment. *ACM Transactions on Information and System Security*, 5(4), 438–457. <https://doi.org/10.1145/581271.581274>
- [14] Hanisch, M., Goldsby, C. M., Fabian, N. E., & Oehmichen, J. (2023). Digital governance: A conceptual framework and research agenda. *Journal of Business Research*, 162, Article 113777. <https://doi.org/10.1016/j.jbusres.2023.113777>
- [15] Judge, B., Nitzberg, M., & Russell, S. (2024). When code is not law: Rethinking regulation for artificial intelligence. *Policy and Society*, 44(1), 85–97. <https://doi.org/10.1093/polsoc/puae005>
- [16] Kapron, Z. (2026, April 22). *The governance gap that could undermine financial markets*. *Forbes*. <https://www.forbes.com/sites/zennonkapron/2026/04/22/the-governance-gap-that-could-break-financial-markets/>
- [17] National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (AI 100-1). U.S. Department of Commerce.