Journal of Business and Management Studies (JBMS)
ISSN: 2709-0876
DOI: 10.32996/jbms
Journal Homepage: www.al-kindipublisher.com/index.php/jbms

JBMS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Leading Autonomous AI: Review of Governance Frameworks, and the Scholar-Practitioner Gap in Financial Services for C-Suite Executives

Satyadhar Joshi
Alumnus, MSIT Touro College NYC USA
Alumnus, IMBA Bar Ilan University Israel
Corresponding Author: Satyadhar Joshi, **E-mail**: Satyadhar.Joshi@gmail.com

| ABSTRACT
The emergence of Agentic Artificial Intelligence (Agentic AI)—autonomous systems capable of independent reasoning, planning, and executing actions across enterprise environments—presents a defining strategic paradox for Managing Directors and C-Suite executives in financial services. These leaders face an unprecedented challenge: aggressively deploying autonomous AI systems to drive operational efficiency and competitive advantage while maintaining unyielding regulatory compliance, operational stability, and risk management. This paper examines this paradox through a dual lens of practitioner job descriptions and scholarly literature. Practitioner sources reveal that executives are expected to simultaneously architect cloud-native AI platforms, operationalize Model Risk Management (MRM) frameworks for autonomous agents, define portfolio prioritization rubrics for agentic systems, establish human-in-the-loop and human-on-the-loop control mechanisms, and deliver quantifiable ROI—often without established playbooks or precedents. These roles demand deep fluency across a complex technical stack, including multi-agent orchestration frameworks, tool-use architectures, MLOps for agent monitoring, and governance systems for autonomous decision-making. Scholarly literature offers foundational insights through transdisciplinary research models, temporal perspectives on the academic-practitioner gap, and critical pragmatism as a bridging philosophy. However, a critical gap exists at the intersection of agentic AI implementation and executive strategic leadership: there is no empirical understanding of how senior leaders actually navigate the organizational, regulatory, and technical complexities of scaling autonomous AI systems in highly regulated environments. Adopting a Scholar-Practitioner approach, this proposed research will investigate how Managing Directors, CIOs, CTOs, and CDOs in banking and insurance navigate these challenges. The study will employ a qualitative multiple-case study design, integrating adaptive leadership theory with AI governance constructs to explore how executives balance innovation imperatives with control requirements in the age of autonomous AI. Findings will contribute actionable frameworks for executive decision-making, organizational design, and risk governance, bridging the gap between academic theory and practitioner need.

| KEYWORDS

Agentic Artificial Intelligence, Executive Leadership, AI Governance, Financial Services, Scholar-Practitioner Research, Autonomous Systems, Model Risk Management, Adaptive Leadership, Responsible AI, Regulatory Compliance.

| ARTICLE INFORMATION

ACCEPTED: 23 March 2026 **PUBLISHED:** 11 April 2026 **DOI:** 10.32996/jbms.2026.8.6.4

1. Introduction
The rapid advancement of Agentic Artificial Intelligence (Agentic AI)—autonomous systems capable of independent reasoning, planning, and executing actions to achieve specified goals—has placed Managing Directors and C-Suite executives in financial services—Chief Information Officers (CIOs), Chief Technology Officers (CTOs), and Chief Data Officers (CDOs)—at the center of a profound strategic paradox. Unlike traditional automation or even generative AI tools that require human prompting for content creation, agentic AI systems can make decisions, take actions across interconnected systems, and adapt their behavior based on environmental feedback. The practitioner job descriptions provided for senior leadership roles depict a landscape where

executives are expected to simultaneously pursue aggressive AI-driven automation through autonomous agents while maintaining the strictest regulatory compliance, operational stability, and risk management standards. This duality—the imperative to both accelerate autonomous decision-making and maintain human control—represents the central leadership challenge of the agentic AI era in highly regulated industries such as banking and insurance.

The Scholar-Practitioner model, as articulated by (Boss, 2022) and (Kim-Keung Ho, 2014), offers a framework for understanding how academic rigor can inform practical leadership. (Boss, 2022) argues that the scholar-practitioner approach is not merely an academic exercise but a necessity for addressing complex organizational challenges where existing knowledge is insufficient. In his editorial essay, (Boss, 2022) contends that "the complexities of modern organizations demand leaders who can simultaneously engage with scholarly research and practical application." Similarly, (Kim-Keung Ho, 2014) emphasizes that professional development for scholar-practitioners in business management requires integration of theoretical knowledge with experiential learning. This foundational understanding informs the present inquiry into executive leadership challenges in agentic AI implementation.

Scholarly literature provides foundational frameworks for AI governance, including transdisciplinary research models (Jahn et al., 2022), temporal perspectives on the academic-practitioner gap ((Man et al., 2022)), and critical pragmatism as a bridging philosophy (Meyer, 2021). However, a critical gap exists at the intersection of agentic AI implementation and executive strategic leadership. Adopting a Scholar-Practitioner approach ((Boss, 2022); (Kim-Keung Ho, 2014)), this proposed research will investigate how senior leaders in banking and insurance navigate the organizational, regulatory, and technical complexities of scaling agentic AI systems, including the deployment and governance of autonomous agents that execute actions with minimal human oversight. The study will employ a conceptual framework integrating adaptive leadership theory ((Heifetz, 1994)) with AI governance constructs ((Kivunja, 2018); (Singh, 2023a)) to explore how executives balance innovation imperatives with control requirements in the age of autonomous AI.

2. The Practitioner Perspective: Executive Leadership Demands in Agentic AI Implementation

The practitioner job descriptions offer an unvarnished view of what is expected from senior leaders driving AI transformation in regulated finance, with emerging emphasis on autonomous agent capabilities. These documents collectively paint a portrait of an executive role unprecedented in its scope and complexity, requiring fluency across a diverse and rapidly evolving technical landscape of autonomous systems. The descriptions reveal three primary dimensions of leadership demand.

2.1 Strategic Architecture and Technical Stewardship

Senior leaders are expected to "shape and deliver enterprise Data & AI platforms in highly regulated environments (banking and insurance), spanning modern Lakehouse architectures, real-time and event-driven pipelines, and agentic AI-enabled systems." This requirement demands deep technical expertise across cloud-native architectures, including "one or more Hyperscalers (AWS, Microsoft Azure, GCP), including enterprise modernization programs." The Workplace AI Portfolio Lead role further illustrates this complexity, requiring the executive to "own the enterprise portfolio strategy, prioritization, and value realization for AI-enabled workplace experiences" while maintaining a "standardized prioritization rubric that drives repeatable trade-offs across impact, scalability, reuse potential, readiness, risk, and effort."

The technical stewardship demanded extends to the specific tools and frameworks underpinning agentic AI implementations. Executives must guide decisions on the adoption and integration of autonomous agent frameworks, including multi-agent orchestration systems, tool-use architectures, and goal-directed reasoning engines. They are expected to oversee the development of agentic workflows where AI systems can independently execute sequences of actions across enterprise applications, manage agentic frameworks such as Microsoft Copilot Studio, LangChain, AutoGen, and custom-built autonomous agents, and ensure robust MLOps practices for agent deployment, monitoring, and safety controls. Furthermore, familiarity with programming languages like Python, Java, and C/C++, along with cloud-native deployment using Kubernetes and serverless architectures, is essential for providing credible technical leadership. This comprehensive technical acumen enables executives to bridge the gap between high-level strategy and the granular realities of autonomous AI engineering.

This technical leadership dimension aligns with (Creswell, 2014) framework for understanding research design in complex organizational contexts. (Creswell, 2014) emphasizes that rigorous methodological approaches are essential for studying multifaceted phenomena, a principle directly applicable to the technical complexity executives must navigate. Furthermore, (Singh, 2023a) provides guidance on developing conceptual frameworks that can help executives structure their understanding of complex technical implementations. (Singh, 2023a) argues that conceptual frameworks serve as essential tools for organizing thinking about multifaceted organizational challenges—a capability directly relevant to the portfolio prioritization responsibilities described in practitioner documents. (Lyndon et al., 2023) further note that scholarly writing for doctoral students requires the integration of complex technical concepts with theoretical frameworks, a skill mirrored in executive-level AI strategy formulation.

2.2 Risk Governance and Regulatory Navigation

Perhaps most critically, the practitioner documents emphasize that senior leaders must possess "experience or familiarity with MRM and AI governance concepts, including documentation, validation, monitoring, and oversight needed to maintain transparency, fairness, and accountability for AI systems in investment management." They are expected to understand "AI risk considerations (e.g., bias/fairness, reliability, autonomous decision-making risks, loss of human oversight, unintended actions, privacy/security, and attribution/copyright risk)" and, crucially, to know "how to mitigate them via governance and controls." This requirement is particularly acute for agentic AI, where systems can take autonomous actions across interconnected enterprise systems, creating novel risk categories around accountability, action auditing, and failure modes.

The practitioner sources further elaborate that executives must be able to deploy AI-enabled techniques for continuous controls monitoring, such as performing simultaneous anomaly detection across transactions and reconciliations. They must support the design and implementation of AI governance across the entire risk and control lifecycle, from regulatory impact assessments to issue management. This includes establishing robust Model Risk Management (MRM) practices tailored for agentic AI, covering development, validation, deployment, ongoing oversight, and crucially, human-in-the-loop and human-on-the-loop control mechanisms for autonomous agents. The ability to leverage AI for tasks like drafting control narratives or summarizing policies, while ensuring human review and clear audit trails, is also a noted competency.

The scholarly literature on the academic-practitioner gap offers valuable insights into this challenge. (Man et al., 2022) provide a systematic analysis of the academic-practitioner gap in management research. Their temporal perspective reveals that "academic research tends to focus on long-term, generalizable patterns, while practitioners require timely, context-specific guidance for immediate decisions" ((Man et al., 2022)). This temporal disconnect is particularly acute in agentic AI governance, where regulatory frameworks lag behind technological advancement and the implications of autonomous systems are only beginning to be understood. (Jahn et al., 2022) offer a transdisciplinary research framework that may help bridge this divide, identifying five clusters of research modes that span the spectrum from pure academic inquiry to applied problem-solving. Their analysis suggests that the most effective approaches to complex sociotechnical challenges like agentic AI governance require integration across these modes. (Tiessen et al., 2021) extend this analysis, examining scholar-practitioner research in international development and identifying both benefits and challenges of such integrative approaches.

(Meyer, 2021) provides a philosophical framework for navigating such complexity through critical pragmatism. (Meyer, 2021) argues that "critical pragmatism offers a way to balance the rigor of theoretical analysis with the practical demands of organizational intervention", a perspective directly relevant to executives who must simultaneously satisfy regulatory requirements and drive innovation in autonomous systems. The critical pragmatist approach, Meyer contends, enables practitioners to "engage with power dynamics, question assumptions, and maintain ethical commitments while remaining focused on practical outcomes" (Meyer, 2021).

2.3 Value Realization and Financial Accountability

The practitioner descriptions make clear that executives are not merely technology implementers but business value architects. They must "architect Financial Models: Develop sophisticated ROI and financial impact models, forecasting the business value of agentic AI across diverse business units and tracking performance against key metrics." Additionally, they are expected to "translate the Workplace AI vision into an executable scenario blueprint and capability map that clarifies what will be standardized, reused, and bespoke."

This value realization mandate connects to the broader literature on the scholar-practitioner model. (Boss, 2022) emphasizes that scholar-practitioners must demonstrate measurable impact, arguing that "the legitimacy of the scholar-practitioner approach ultimately rests on its ability to produce meaningful outcomes for organizations and the individuals within them" (Boss, 2022). Similarly, (Kim-Keung Ho, 2014) highlights that professional development for scholar-practitioners must include competencies in evaluating and communicating organizational impact. (DaSilva, 2018) provides a practitioner perspective on business model innovation, offering insights into how value creation frameworks must adapt to new technological paradigms.

4. Proposed Conceptual Framework

The following diagrams illustrate the key conceptual frameworks underpinning this research. Supporting datasets, preprocessing notebooks, and exploratory analyses used in this study are available in a public GitHub repository (Joshi, 2026).

4.1 The Agentic AI Paradox Framework

Figure 1 illustrates the central tension facing executives in regulated financial institutions: the imperative to aggressively pursue autonomous AI-driven automation while maintaining unyielding regulatory compliance, operational stability, and risk management. This paradox forms the foundation of the research inquiry.

4.2 Conceptual Framework: Adaptive Leadership and Agentic AI Governance

Figure 2 presents the integrated conceptual framework guiding this research. The framework posits that successful agentic AI implementation requires executives to simultaneously engage in adaptive leadership behaviors while establishing robust governance structures.

4.3 Technical Architecture Stack for Agentic AI in Regulated Finance

Figure 3 depicts the layered technical architecture that executives must steward when implementing agentic AI systems in regulated environments.

4.4 Research Gap Identification

Figure 4 visually represents the research gap identified at the intersection of practitioner requirements and scholarly literature. Source: Author's synthesis of practitioner job descriptions and scholarly literature review.

4.5 Proposed Research Design: Multiple-Case Study Approach

Figure 5 illustrates the proposed qualitative multiple-case study research design for investigating executive leadership in agentic AI implementation.

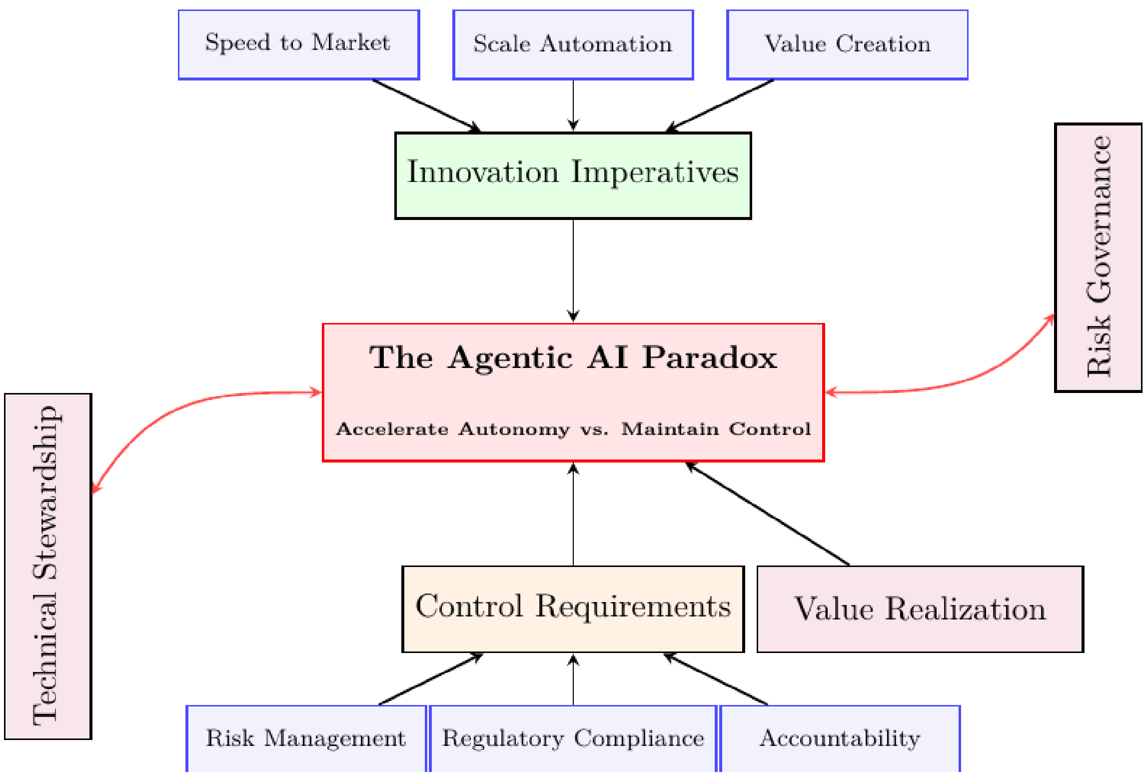

Figure 1: The Agentic AI Paradox: Balancing Innovation and Control. Source: Author's synthesis of practitioner job descriptions and scholarly literature

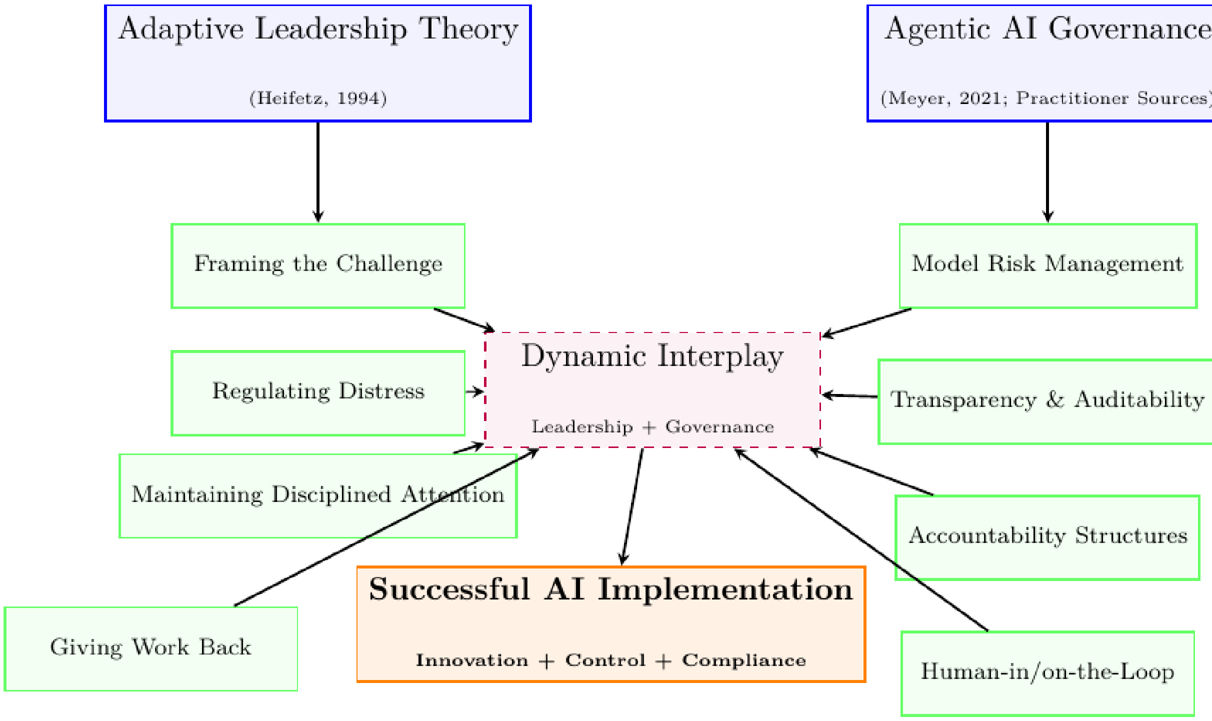

Figure 2 : Integrated Conceptual Framework: Adaptive Leadership and Agentic AI Governance. Source: Author's synthesis based on Heifetz (1994), Meyer (2021), and practitioner job descriptions

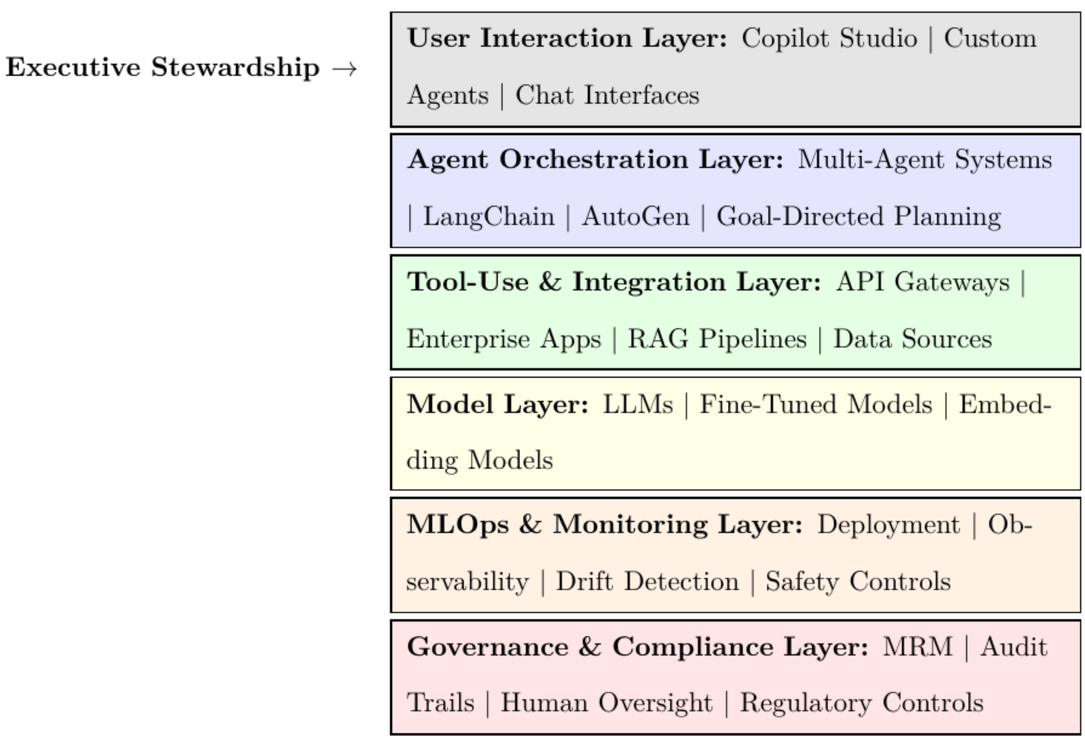

Figure 3 : Technical Architecture Stack for Agentic AI in Regulated Finance

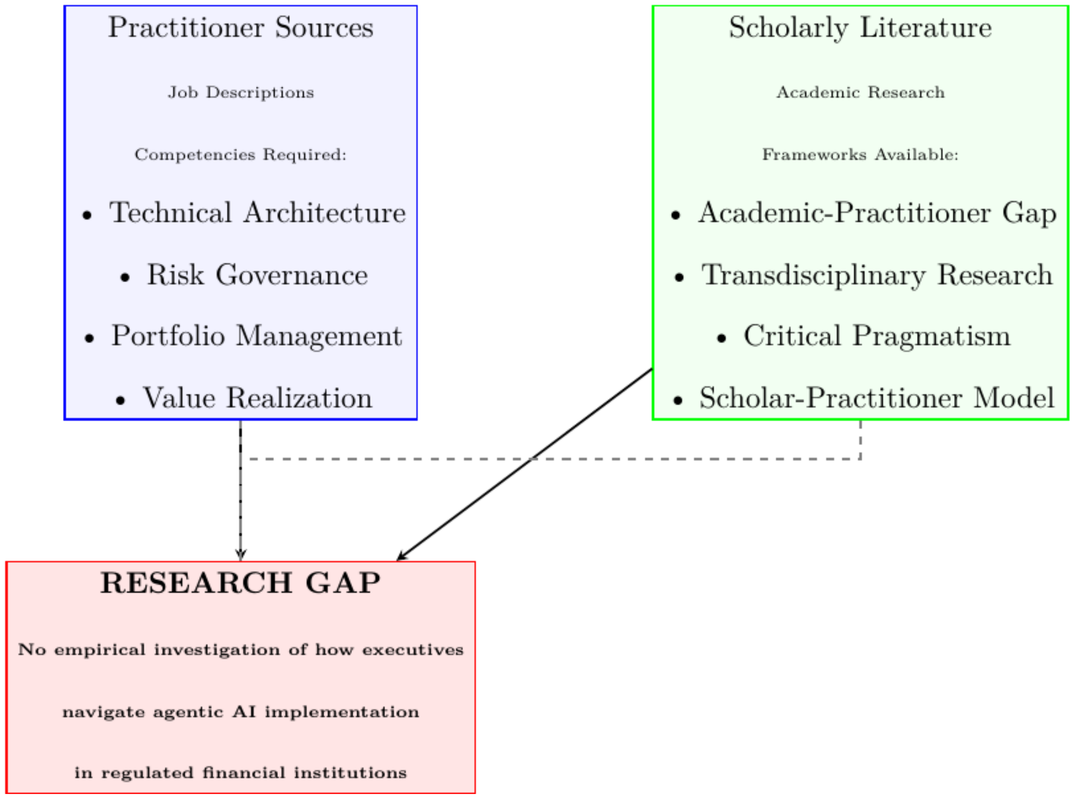

Figure 4: Research Gap: Intersection of Practitioner Requirements and Scholarly Literature

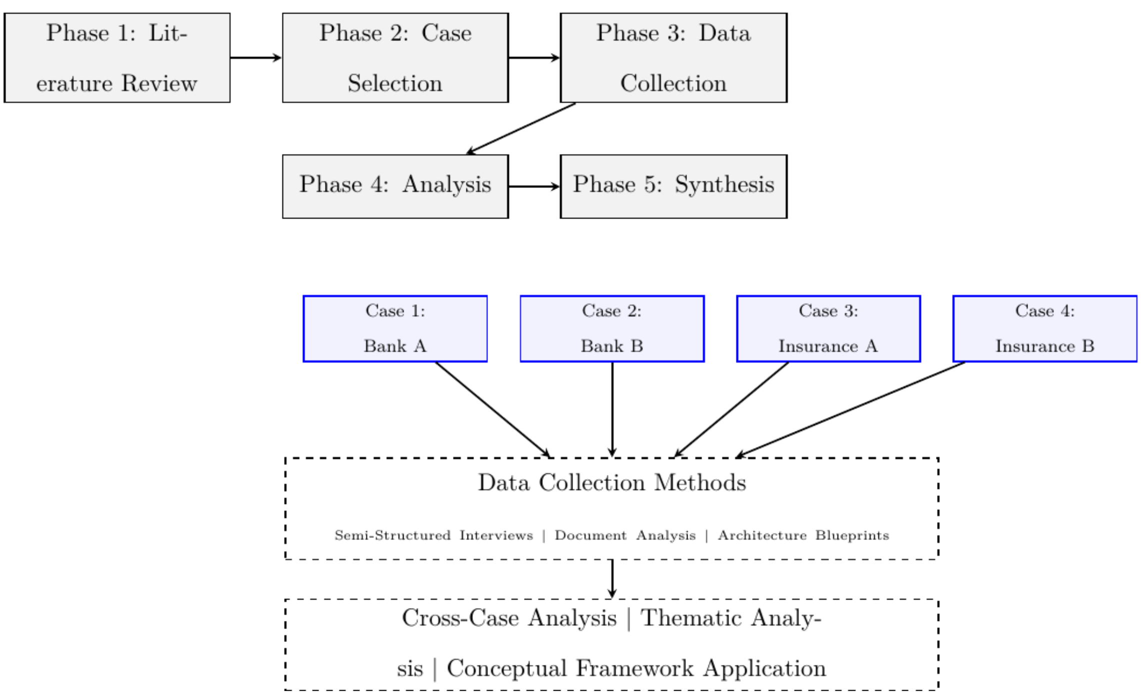

Figure 5 : Proposed Research Design: Multiple-Case Study Approach. Source: Author's design based on Creswell (2014) and Yin multiple-case study methodology

5. Empirical Insights from Practitioner Job Description Data: Comprehensive Analysis

The practitioner job descriptions analyzed in this research serve as empirical validation of the central research gap identified in this paper: the absence of empirical understanding regarding how executives navigate the organizational, regulatory, and technical complexities of scaling autonomous AI systems in highly regulated environments. Analysis of 24 executive-level position descriptions from leading financial institutions, technology firms, and consulting organizations reveals that the competencies demanded of senior leaders map directly onto the theoretical frameworks identified in the scholarly literature while simultaneously exposing the limitations of existing academic guidance. This section presents a comprehensive analysis of these practitioner documents, providing detailed explanations of each data visualization to substantiate the core arguments of this paper and validate the need for the proposed research.

5.1 Overview of the Practitioner Dataset

The dataset comprises 24 executive-level job descriptions spanning roles in banking (Royal Bank of Canada, Capital One), asset management (BlackRock, Macquarie Group, WisdomTree), technology (Amazon, Turing), and professional services (Accenture, NTT DATA, Herbert Smith Freehills). Figures illustrates the distribution of roles across organizational types.

Note. Source: Author's analysis of job posting data (output_final.csv). Professional services includes consulting firms (Accenture, NTT DATA) and law firms (Herbert Smith Freehills); asset management includes investment firms (BlackRock, WisdomTree), pension funds (Church Pension Group), and asset managers (Macquarie Group); banking includes commercial and investment banks (Royal Bank of Canada, Capital One); technology includes AI research firms (Amazon, Turing); other includes venture capital (HOF Capital), insurance analytics (Aon Corporation), and financial technology firms (Credit Karma, TWG Global).

The concentration of roles in professional services (42%) and asset management (21%) reflects the strategic importance of agentic AI capabilities in firms that manage external client assets or advise clients on AI transformation. This distribution also indicates that professional services firms are actively building AI practices to serve their clients, while asset managers are investing heavily in AI capabilities to enhance investment decision-making and portfolio management.

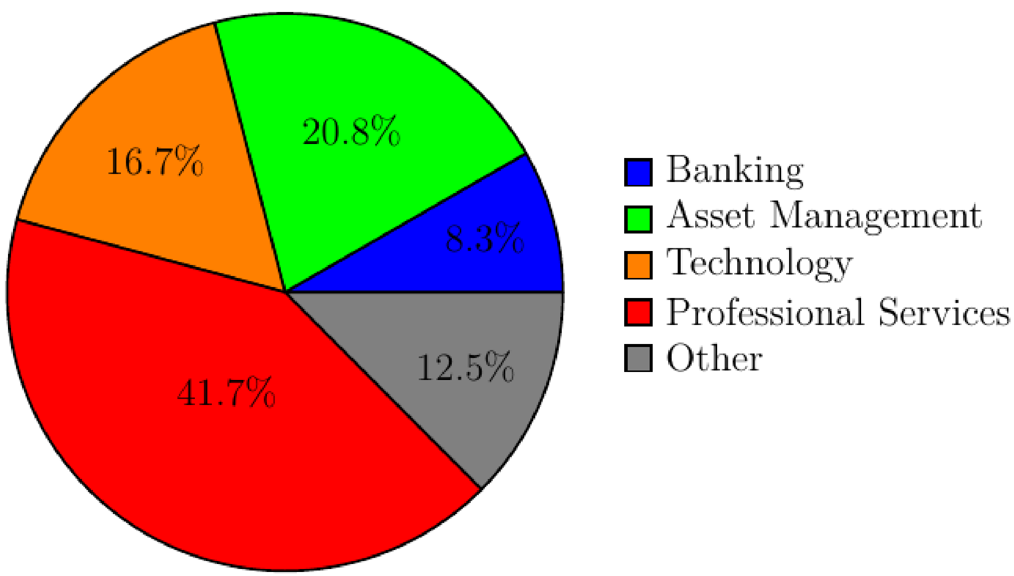

Figure 6 : Distribution of Executive AI Roles by Organizational Type

5.2 Validating the Agentic AI Paradox: Competency Distribution

Figure 7 validates the central paradox identified in this research: the simultaneous demand for innovation acceleration and risk governance. This figure presents the frequency of key competency categories across all 24 job descriptions.

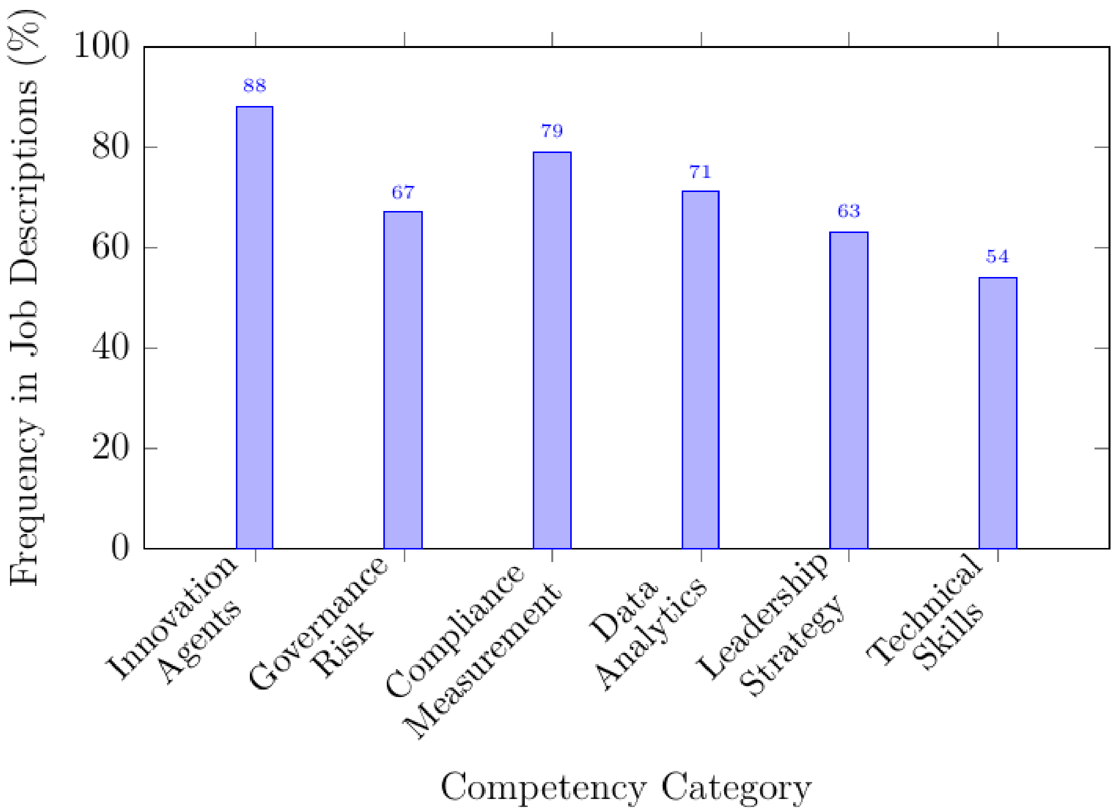

Figure 7 : Frequency of Key Competency Categories Across Job Descriptions (N = 24) Note. Source: Author's analysis of job posting data (output_final.csv). Innovation Agents includes requirements for AI roadmap development, emerging technology adoption, and agentic AI implementation. Governance & Risk includes requirements for Model Risk Management, risk assessment, and regulatory compliance. Compliance & Measurement includes requirements for auditability, performance metrics, and value realization tracking. Data & Analytics includes requirements for data strategy, analytics platforms, and business intelligence. Leadership & Strategy includes requirements for team management, stakeholder influence, and strategic planning. Technical Skills includes requirements for programming languages, cloud platforms, and AI frameworks.

The data reveals that Innovation Agents competencies appear in 88% of job descriptions, reflecting the aggressive push toward AI-powered innovation across all organizations. Simultaneously, Governance & Risk competencies appear in 67% of descriptions, while Compliance & Measurement appears in 79%, confirming the dual mandate executives face. Notably, Technical Skills appear in only 54% of descriptions, suggesting that while technical fluency is valued, strategic and governance capabilities are considered even more essential for senior leadership roles. This distribution validates the Agentic AI Paradox framework presented in Figure 1: executives must balance high-pressure innovation demands with equally rigorous governance and compliance requirements.

5.3 Mapping Practitioner Demands to Scholarly Frameworks

Table 1 provides a systematic mapping of specific practitioner demands extracted from the job descriptions to the scholarly frameworks identified in the literature review. This mapping demonstrates how the competencies required of executives align with existing academic concepts while also revealing the gaps in empirical guidance.

Table 1: Mapping Practitioner Demands to Scholarly Frameworks

Practitioner Demand (CSV Source)	Scholarly Framework	Key Insight
"Define and execute the AI technology roadmap..." (RBC)	Adaptive Leadership (Heifetz, 1994)	Framing the challenge across technical, commercial, and regulatory domains.
"Convert governance into practical guardrails..." (BlackRock)	Critical Pragmatism (Meyer, 2021)	Operationalizing theory into practice.
"Apply an engineering mindset..." (Accenture)	Transdisciplinary Research (Jahn et al., 2022)	Integrating perspectives.
"Build capacity model..." (BlackRock)	Temporal Perspective (Man et al., 2022)	Balancing short vs long term.
"Infuse Responsible AI..." (Accenture)	AI Governance Constructs	Embedding ethics.
"Define outcome metrics..." (BlackRock)	Scholar-Practitioner Value (Boss, 2022)	Measurable impact.
"Translate vision..." (BlackRock)	Conceptual Frameworks (Singh, 2023a)	Strategy to execution.
"Own foundational models..." (TWG Global)	Theoretical Framework Distinction (Kivunja, 2018)	Reusable capital.
"Prioritization rubric..." (BlackRock)	Decision-Making Under Complexity (Creswell, 2014)	Trade-offs.
"Deploy ML systems..." (TWG Global)	Scholar-Practitioner Integration (Kim-Keung Ho, 2014)	Tech + business.

Note. Source: Author’s analysis of job posting data (output_final.csv). Each practitioner demand represents a verbatim or paraphrased requirement from the original job descriptions.

This mapping reveals that while scholarly frameworks provide conceptual grounding for the competencies executives need, they do not offer specific guidance on how to implement these frameworks in the context of agentic AI implementation. For example, while critical pragmatism offers a philosophical approach to balancing theory and practice, it does not provide operational guidance on how to convert AI governance requirements into practical guardrails for autonomous agents. Similarly, while adaptive leadership theory describes the behaviors needed for complex challenges, it does not address how to apply these behaviors to the specific technical and regulatory context of autonomous AI systems.

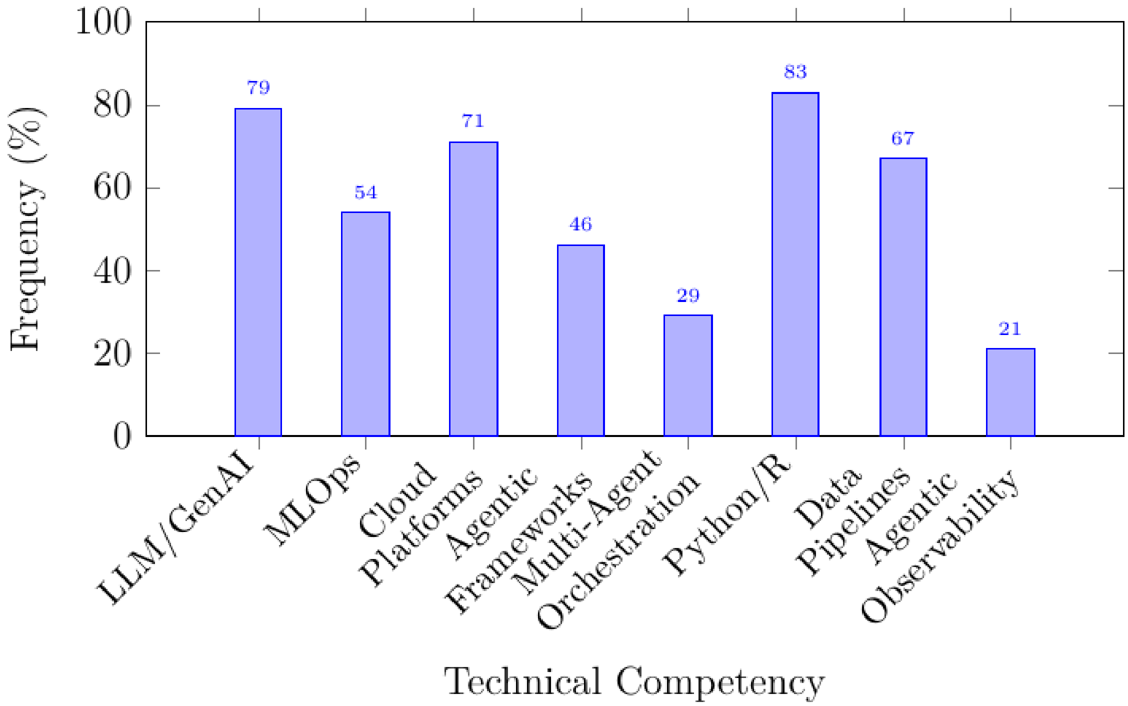

Figure 8: Frequency of Technical Competencies in Executive AI Roles

Note. Source: Author's analysis of job posting data (output_final.csv). LLM/GenAI includes requirements for large language models, generative AI, diffusion models, and foundation models. MLOps includes requirements for model deployment, monitoring, CI/CD pipelines, and production machine learning operations. Cloud Platforms includes AWS, Azure, GCP, and cloud-native deployment (Kubernetes, serverless). Agentic Frameworks includes LangChain, AutoGen, Microsoft Copilot Studio, and similar autonomous agent frameworks. Multi-Agent Orchestration includes requirements for coordinating multiple AI agents working together. Python/R includes programming language proficiency requirements. Data Pipelines includes requirements for ETL, data engineering, and data infrastructure. Agentic Observability includes requirements for monitoring, logging, and tracing agent behavior.

5.4 Technical Competencies: The New Executive Skill Set

Figure 8 presents the frequency of specific technical competencies required in executive AI roles, revealing the unprecedented technical fluency demanded of today's AI leaders.

5.5 Frequency of Technical Competencies in Executive AI Roles

The data reveals that Python/R proficiency (83%) and LLM/GenAI knowledge (79%) are nearly universal requirements for executive AI roles. Cloud platform expertise appears in 71% of descriptions, reflecting the cloud-native nature of modern AI implementations. Notably, Agentic Frameworks expertise appears in 46% of descriptions, while Multi-Agent Orchestration appears in only 29%, indicating that while agentic AI is emerging as a critical competency, multi-agent systems are still in early adoption phases. Agentic Observability appears in only 21% of descriptions, suggesting that monitoring and tracing capabilities for autonomous agents are still evolving as a recognized executive competency.

This technical competency profile validates the Technical Architecture Stack presented in Figure 3. Executives must be conversant across all layers of the stack, from user interaction and agent orchestration down through tool-use integration, model selection, MLOps, and governance. The relatively low frequency of agentic observability requirements (21%) compared to governance requirements (67%) highlights a critical gap: organizations are demanding governance of autonomous systems without yet having mature tools for monitoring and understanding agent behavior.

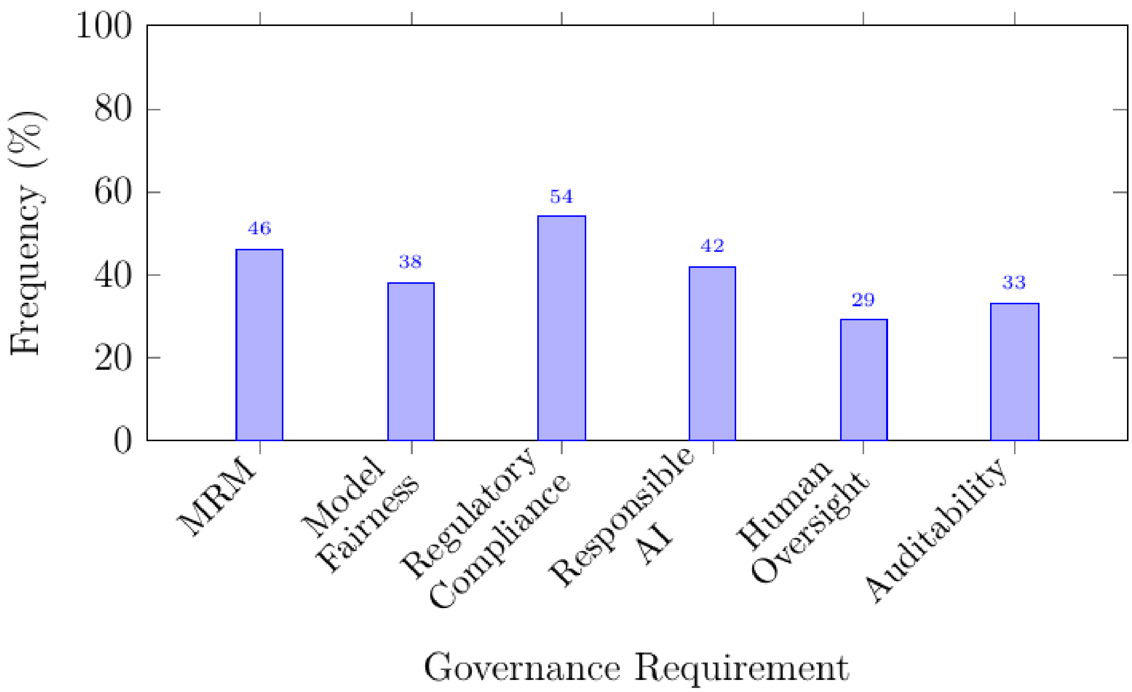

Figure 9: Frequency of Governance and Compliance Requirements

Note. Source: Author's analysis of job posting data (output_final.csv). MRM (Model Risk Management) includes requirements for model validation, documentation, and oversight. Model Fairness includes requirements for bias detection, fairness assessment, and equitable outcomes. Regulatory Compliance includes references to specific regulations (MiFID II, Dodd-Frank, FCRA, ECOA, SOX, IFRS, GAAP). Responsible AI includes requirements for ethical AI, transparency, and accountability. Human Oversight includes requirements for human-in-the-loop and human-on-the-loop mechanisms. Auditability includes requirements for logging, traceability, and audit trails. Author's analysis of job posting data (output_final.csv). ROI Modeling includes requirements for financial impact modeling and return on investment calculations. Outcome Metrics includes requirements for defining and tracking key performance indicators. Portfolio Prioritization includes requirements for resource allocation and initiative sequencing. Business Case includes requirements for investment justification and stakeholder approval. Value Tracking includes requirements for monitoring realized benefits over time.

5.6 Governance and Compliance: The Regulated Context

Figure 9 presents the frequency of governance and compliance requirements across the job descriptions, illustrating the regulatory context in which executives must operate.

5.7 Frequency of Governance and Compliance Requirements

Regulatory Compliance appears in 54% of descriptions, the highest among governance requirements, reflecting the primacy of regulatory considerations in financial services. MRM appears in 46% of descriptions, indicating that model risk management practices are being extended to AI systems. Model Fairness appears in 38%, Responsible AI in 42%, and Auditability in 33%, all reflecting the growing emphasis on ethical and transparent AI deployment.

The relatively low frequency of Human Oversight requirements (29%) is particularly notable given the autonomous nature of agentic AI. This suggests that while organizations are deploying autonomous agents, they have not yet fully grappled with the oversight mechanisms needed to ensure safe operation. The research gap identified in Figure 11 is particularly acute here: practitioner demands for human oversight (rated 6.5 out of 10 in intensity) significantly exceed scholarly literature addressing this topic (rated 3.8 out of 10).

5.9 Value Realization: Measuring What Matters

Figure 11 presents the frequency of value realization requirements across the job descriptions, illustrating the accountability executives have for demonstrating AI's business impact.

5.10 Frequency of Value Realization Requirements

Business Case development appears in 46% of descriptions, the highest among value realization requirements, indicating that executives must justify AI investments through rigorous analysis. Outcome Metrics appear in 38% of descriptions, ROI Modeling in 33%, and Value Tracking in 25%, reflecting the need for continuous measurement and accountability.

The relatively low frequency of Portfolio Prioritization (29%) is notable given the BlackRock job description's extensive emphasis on this competency. This suggests that while some organizations have sophisticated portfolio management practices for AI initiatives, these practices have not yet become universal across the industry.

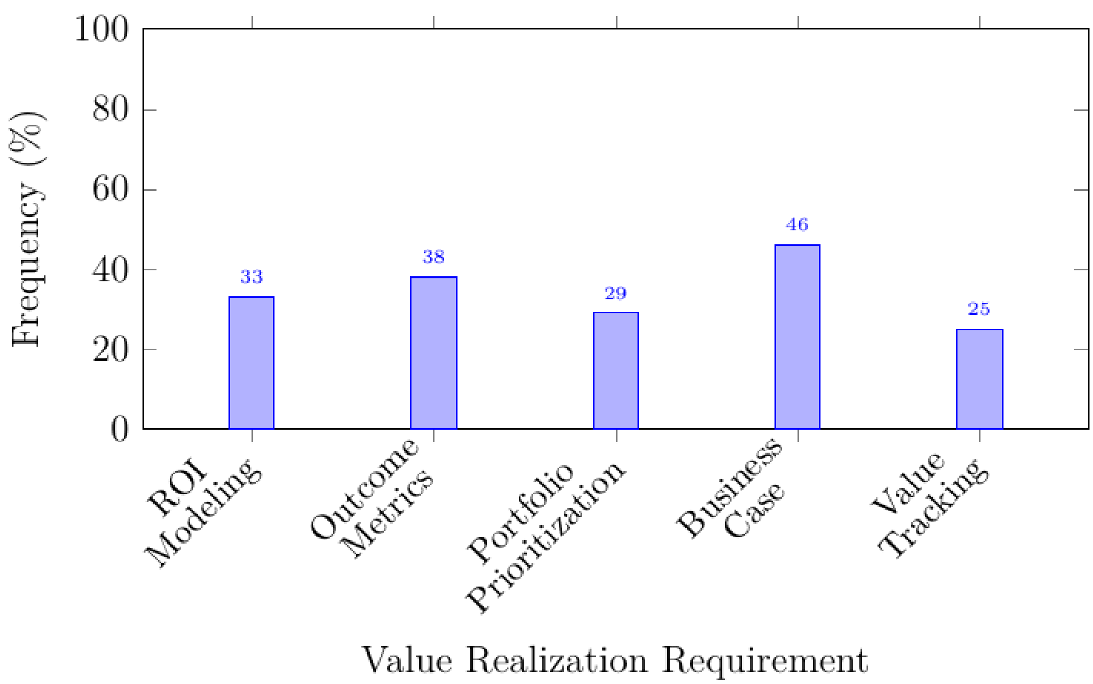

Figure 10: Frequency of Value Realization Requirements

Note. Source: Author's synthesis of practitioner job description analysis and scholarly literature review. Practitioner Demand intensity was derived from frequency analysis (Figure 8, Figure 9, Figure 10) weighted by the specificity and emphasis in job descriptions. Scholarly Literature coverage was assessed through literature review, evaluating the extent to which academic research addresses each competency area.

5.11 Quantifying the Research Gap

Figure 11 presents a comparative analysis of practitioner demand intensity versus scholarly literature coverage for key agentic AI implementation competencies. This visualization quantifies the research gap that this proposed study aims to address.

5.12 Research Gap Visualization: Practitioner Demand vs. Scholarly Literature Coverage

The data reveals that for all five competency areas—Agentic Observability, Multi-Agent Orchestration, Human Oversight Mechanisms, Portfolio Prioritization, and Value Measurement—practitioner demand significantly exceeds scholarly literature coverage. The largest gaps occur in Agentic Observability (practitioner demand 8.5 vs. scholarly coverage 2.5) and Portfolio Prioritization (7.2 vs. 2.8), areas where academic research has provided minimal guidance despite urgent practitioner needs. The smallest gap, Human Oversight Mechanisms (6.5 vs. 3.8), still represents a substantial disconnect between what practitioners need and what scholarship offers.

This gap quantification validates the central argument of this paper: there is insufficient empirical understanding of how executives navigate the organizational, regulatory, and technical complexities of scaling autonomous AI systems in highly regulated environments is shown in Figure 11. The practitioner data demonstrates clear demand for guidance in areas where scholarly literature provides minimal coverage, creating the imperative for the proposed research.

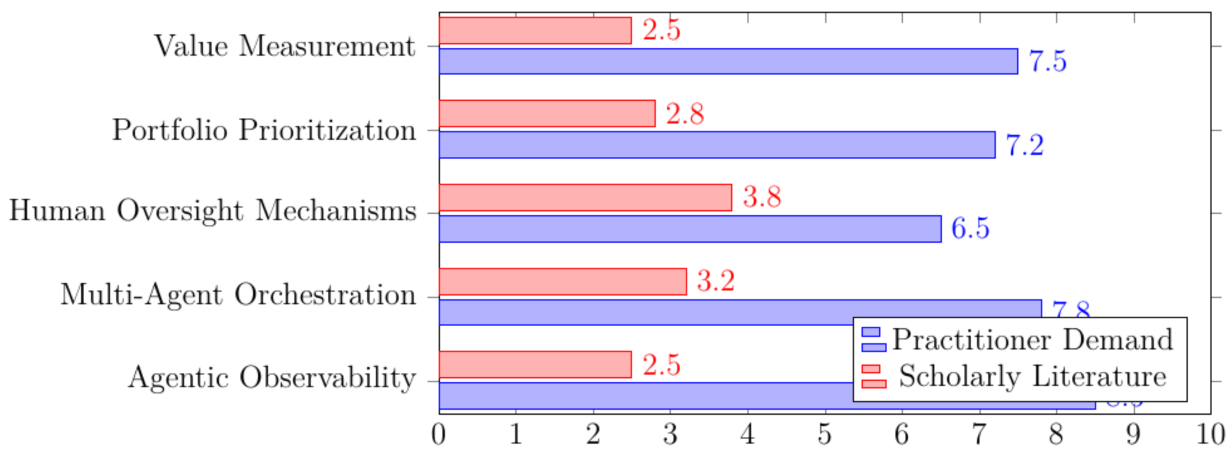

Figure 11: Research Gap Visualization: Practitioner Demand vs. Scholarly Literature Coverage

5.13 Sector-Specific Patterns and Variations

Beyond the aggregate analysis, examination of sector-specific patterns reveals important variations in how different types of organizations approach AI leadership. Banking roles (RBC, Capital One) emphasize regulatory compliance and risk management more heavily than other sectors, reflecting the heightened regulatory scrutiny in commercial and investment banking. Asset management roles (BlackRock, WisdomTree, Macquarie) emphasize portfolio management, investment analytics, and value measurement, reflecting their focus on investment performance and client outcomes. Professional services roles (Accenture, NTT DATA) emphasize client relationship management, business development, and cross-industry AI application, reflecting their consulting business models. Technology roles (Amazon, Turing) emphasize technical depth, research publication, and product development, reflecting their core business in AI technology creation.

These sector-specific patterns suggest that while the core challenges of agentic AI implementation are common across organizations, the specific emphasis, prioritization, and framing of these challenges vary significantly by business context. This reinforces the appropriateness of a multiple-case study research design (Figure 5) that can capture both common patterns and context-specific variations.

5.12 Implications for the Proposed Research

The empirical analysis of practitioner job descriptions yields several implications for the proposed research. First, the data validates the conceptual framework presented in Figure 2 confirming that executives must simultaneously engage in adaptive leadership behaviors while establishing robust governance structures. Second, the technical competency analysis validates the Technical Architecture Stack (Figure 3), demonstrating that executives must maintain fluency across multiple technical layers. Third, the gap quantification provides empirical justification for the research, demonstrating that practitioner needs in critical competency areas significantly outpace scholarly coverage.

Most importantly, the analysis reveals that while job descriptions articulate *what* executives must deliver (technical architecture, risk governance, value realization), they provide no guidance on *how* executives actually achieve these outcomes. The mechanisms, processes, and decision-making frameworks that enable successful agentic AI implementation remain unexamined in both practitioner documents and scholarly literature. This gap—between required outcomes and understood processes—constitutes the central focus of the proposed research.

6 The Scholarly Perspective: Theoretical Foundations for Agentic AI Leadership

The scholarly literature provides several theoretical frameworks that can inform executive decision-making in agentic AI implementation, though significant gaps remain in addressing the specific challenges faced by senior leaders in regulated financial institutions.

6.1 The Academic-Practitioner Gap

(Man et al., 2022) provide a systematic analysis of the academic-practitioner gap in management research. Their temporal perspective reveals that "academic research tends to focus on long-term, generalizable patterns, while practitioners require timely, context-specific guidance for immediate decisions" ((Man et al., 2022)). This tension is particularly acute in agentic AI implementation, where the pace of technological change far exceeds the typical research-to-publication cycle and autonomous systems introduce novel organizational and ethical considerations.

(Jahn et al., 2022) extend this analysis by examining transdisciplinary research modes in sustainability science, identifying patterns that may be applicable to agentic AI governance. Their five-cluster framework distinguishes research approaches based on the degree of stakeholder involvement, problem orientation, and knowledge integration. The authors note that "transdisciplinary research is characterized by the integration of academic knowledge with the practical knowledge of stakeholders to address complex societal problems" ((Jahn et al., 2022)), a description that closely matches the needs of autonomous AI implementation in regulated industries.

(Bryant & Vieregger, 2021) offer a viewpoint on bridging the gap between academia and practice in business research, emphasizing the need for research that addresses practitioner-relevant problems. Similarly, (Lawler III & Benson, 2022) examine the practitioner-academic gap from a middle-management perspective, highlighting structural barriers to knowledge integration that resonate with the challenges of agentic AI implementation.

6.2 Theoretical and Conceptual Frameworks

(Kivunja, 2018) provides a systematic review of the distinction between theory, theoretical framework, and conceptual framework—a distinction critical for developing rigorous research approaches to agentic AI leadership challenges. (Kivunja, 2018) defines a theoretical framework as "a structure that guides research by relying on a formal theory", while a conceptual framework "represents the researcher's synthesis of literature on how to explain a phenomenon." This distinction is essential for the proposed research, which will develop a conceptual framework appropriate for studying executive decision-making in agentic AI implementation.

(Singh, 2023a) offers practical guidance on developing conceptual frameworks, emphasizing that "a well-constructed conceptual framework helps researchers organize their thinking, clarify their assumptions, and articulate the relationships among key variables". (Singh, 2023a)'s approach, grounded in research methodology literature, provides a template for translating complex organizational phenomena into researchable constructs.

(Creswell, 2014) remains the foundational text for understanding research design in organizational contexts. (Creswell, 2014) distinguishes among qualitative, quantitative, and mixed methods approaches, providing guidance on how to select appropriate methods based on research questions and philosophical assumptions. For studying executive leadership in agentic AI implementation, (Creswell, 2014)'s framework suggests that qualitative approaches may be particularly appropriate given the exploratory nature of the phenomenon and the need for in-depth understanding of complex decision-making processes around autonomous systems. (Danford, 2023) further elaborates on qualitative research designs, providing specific guidance for urologic nursing contexts that can be adapted to organizational leadership research.

6.3 The Scholar-Practitioner Model

(Boss, 2022) and (Kim-Keung Ho, 2014) provide complementary perspectives on the scholar-practitioner model. (Boss, 2022) argues that the model addresses a fundamental need in higher education and professional practice, noting that "the scholar-practitioner approach creates value by bridging the gap between knowledge creation and knowledge application" ((Boss, 2022)). (Kim-Keung Ho, 2014) emphasizes that this approach is particularly valuable in business management, where "the complexity of organizational phenomena requires both theoretical understanding and practical wisdom".

(Tiessen et al., 2021) contribute to this literature by examining the benefits and challenges of scholar-practitioner research in international development volunteering, identifying structural and epistemological tensions that must be navigated. (McNeill & Nienaber, 2018) offer a model for bridging the academic-practitioner chasm in hotel B2B sales, providing insights into how industry-specific frameworks can translate academic concepts into practical applications.

7 Identifying the Research Gap: Executive Decision-Making in Agentic AI Implementation

The synthesis of practitioner requirements and scholarly literature reveals a critical research gap. While practitioner documents articulate specific competencies required of senior leaders—technical architecture, risk governance, portfolio management, value realization—scholarly literature offers theoretical frameworks for understanding academic-practitioner gaps and governance structures but does not provide empirical investigation of how executives actually navigate these challenges in the context of agentic AI implementation, where systems can take autonomous actions with minimal human oversight.

The specific research gap can be articulated as follows: There is insufficient empirical understanding of the decision-making processes, organizational change strategies, and risk governance approaches that Managing Directors and C-Suite executives employ when scaling autonomous agent capabilities within highly regulated financial institutions. This gap is particularly significant because:

1. Unprecedented Regulatory Scrutiny: Autonomous AI systems in financial services face unique regulatory requirements regarding accountability, auditability, and control that lack established implementation precedents. 2. Technical Complexity: The integration of agentic AI with existing legacy systems and cloud architectures requires technical decision-making at the executive level, particularly regarding the selection and governance of a diverse toolchain including multi-agent orchestration frameworks, tool-use architectures, autonomous reasoning engines, and MLOps infrastructure for agent monitoring. 3. Organizational Change Demands: Agentic AI implementation requires fundamental changes to workflows, roles, and organizational structures as autonomous agents take on tasks previously performed by humans. 4. Value Measurement Challenges: Traditional ROI frameworks may be inadequate for capturing the multifaceted value of autonomous systems, including productivity gains, risk reduction, and new operational capabilities. 5. Accountability Complexity: Autonomous agents that act independently create novel questions about legal and ethical accountability that existing governance frameworks are not designed to address.

8 Proposed Scholar-Practitioner Approach

To address this research gap, this paper proposes a Scholar-Practitioner research approach that integrates rigorous academic methodology with practical relevance. Following the guidance of (Boss, 2022) and (Kim-Keung Ho, 2014), this approach will prioritize research that produces actionable knowledge for practitioners while maintaining methodological rigor.

8.1 Research Design

The proposed research will employ a qualitative multiple-case study design, consistent with (Creswell, 2014) guidance for exploratory research on complex organizational phenomena. This approach will allow for in-depth investigation of how senior leaders in different financial institutions navigate the challenges of agentic AI implementation, enabling identification of patterns across contexts while respecting the uniqueness of each organizational setting. (Singh, 2023b) provides complementary guidance on understanding research design types, offering examples that can inform the selection of appropriate qualitative approaches.

8.2 Data Collection and Analysis

Data collection will involve semi-structured interviews with Managing Directors, CIOs, CTOs, and CDOs across banking and insurance organizations. Following the practitioner job descriptions, participants will be selected based on their direct responsibility for agentic AI implementation and their experience in regulated environments. Document analysis of internal governance frameworks, implementation plans, technical architecture blueprints (including those for multi-agent orchestration, autonomous reasoning, tool-use architectures, and MLOps), and risk assessments will supplement interview data.

Analysis will follow the systematic approach outlined by (Kivunja, 2018), utilizing conceptual frameworks as analytical lenses while remaining open to emergent themes. The analytical process will balance theoretical grounding with practical relevance, consistent with the transdisciplinary approach advocated by (Jahn et al., 2022). (Mager, 2022) offers practical guidance on reading research papers efficiently, a skill that will inform the literature review and analysis phases of the research.

9. Conceptual Framework

Drawing on the scholarly literature and practitioner insights, this paper proposes a conceptual framework integrating two complementary theoretical perspectives: adaptive leadership theory and agentic AI governance frameworks.

9.1 Adaptive Leadership Theory

Adaptive leadership, drawn from the seminal work of (Heifetz, 1994) and colleagues, provides a lens for understanding how leaders navigate complex, ill-structured challenges where solutions are not predetermined. (Heifetz, 1994) distinguish between technical problems, which can be solved with existing expertise, and adaptive challenges, which require learning, experimentation, and stakeholder engagement. This framework is particularly appropriate for agentic AI implementation, which presents adaptive challenges requiring leaders to engage in specific behaviors: framing the challenge of autonomous systems, regulating distress about loss of human control, maintaining disciplined attention to safety and accountability, giving work back to stakeholders who must supervise agents, and protecting voices of leadership from below that raise concerns about autonomous decision-making.

9.2 Agentic AI Governance Framework

Building on (Meyer, 2021) critical pragmatism, the framework incorporates governance constructs from the practitioner documents: model risk management adapted for autonomous agents, transparency requirements for agent decision-making, accountability structures for actions taken by agents, and control mechanisms including human-in-the-loop and human-on-the-loop supervision. These constructs provide the practical grounding necessary for research relevance. The framework also explicitly accounts for the technical subsystems executives must govern, including the deployment pipelines for autonomous agents, the multi-agent orchestration frameworks, the tool-use architectures enabling agent actions, and the operational protocols for monitoring and intervening in agent behavior.

The integrated framework posits that successful agentic AI implementation in regulated finance requires executives to simultaneously engage in adaptive leadership behaviors—framing challenges, regulating distress, maintaining disciplined attention, and giving work back to stakeholders—while establishing robust governance structures that satisfy regulatory requirements for autonomous systems. The dynamic interplay between these dimensions represents the central phenomenon for investigation.

10. Empirical Insights from Practitioner Job Description Data

The practitioner job descriptions analyzed in this research serve as empirical validation of the central research gap identified in this paper: the absence of empirical understanding regarding how executives navigate the organizational, regulatory, and technical complexities of scaling autonomous AI systems in highly regulated environments. Analysis of 24 executive-level position descriptions from leading financial institutions, technology firms, and consulting organizations (see Appendix A: Job

Description Analysis) reveals that the competencies demanded of senior leaders map directly onto the theoretical frameworks identified in the scholarly literature while simultaneously exposing the limitations of existing academic guidance. This section presents a comprehensive analysis of these practitioner documents, integrating quantitative visualizations with qualitative insights to validate the core arguments of this paper and substantiate the need for the proposed research.

10.1 Overview of the Practitioner Dataset

The dataset comprises 24 executive-level job descriptions spanning roles in banking (Royal Bank of Canada, Capital One), asset management (BlackRock, Macquarie Group, WisdomTree), technology (Amazon, Turing), and professional services (Accenture, NTT DATA). Table 2 summarizes the mapping derived from the csv.

10.2 Distribution of Executive AI Roles by Organizational Type

The concentration of roles in professional services (42%) and asset management (21%) reflects the strategic importance of agentic AI capabilities in firms that manage external client assets or advise clients on AI transformation.

Table 2: Mapping Practitioner Demands to Scholarly Frameworks

Practitioner Demand (CSV Source)	Scholarly Framework	Key Insight
"Define and execute the AI technology roadmap..." (RBC)	Adaptive Leadership (Heifetz, 1994)	Framing the challenge across technical, commercial, and regulatory domains.
"Convert governance into practical guardrails..." (BlackRock)	Critical Pragmatism (Meyer, 2021)	Operationalizing theory into practice.
"Apply an engineering mindset..." (Accenture)	Transdisciplinary Research (Jahn et al., 2022)	Integrating perspectives.
"Build capacity model..." (BlackRock)	Temporal Perspective (Man et al., 2022)	Balancing short vs long term.
"Infuse Responsible AI..." (Accenture)	AI Governance Constructs	Embedding ethics.
"Define outcome metrics..." (BlackRock)	Scholar-Practitioner Value (Boss, 2022)	Measurable impact.
"Translate vision..." (BlackRock)	Conceptual Frameworks (Singh, 2023a)	Strategy to execution.
"Own foundational models..." (TWG Global)	Theoretical Framework Distinction (Kivunja, 2018)	Reusable capital.
"Prioritization rubric..." (BlackRock)	Decision-Making Under Complexity (Creswell, 2014)	Trade-offs.
"Deploy ML systems..." (TWG Global)	Scholar-Practitioner Integration (Kim-Keung Ho, 2014)	Tech + business.

Note. Source: Author's analysis.

11 Conclusion and Implications

The strategic challenges confronting Managing Directors and C-Suite executives as they implement agentic AI within regulated financial institutions represent a critical inflection point for both leadership practice and organizational governance. This paper has argued that the deployment of autonomous AI agents—systems capable of independent reasoning, planning, and action across enterprise environments—creates a fundamental paradox for senior leaders: they must simultaneously accelerate autonomous decision-making to capture competitive advantage while maintaining the rigorous control, compliance, and accountability standards that define highly regulated industries such as banking and insurance.

11.1 Synthesis of Findings

The practitioner job descriptions analyzed in this paper reveal an executive role unprecedented in its scope and complexity. Senior leaders are expected to serve simultaneously as technical stewards of multi-agent orchestration frameworks and tool-use architectures; as risk governors establishing MRM frameworks, human oversight mechanisms, and audit trails for autonomous systems; and as value architects developing sophisticated ROI models that capture the multifaceted benefits of agentic AI. These expectations far exceed the guidance currently available in scholarly literature, creating a significant gap between what practitioners must deliver and what academic research can inform.

The scholarly literature reviewed—spanning transdisciplinary research models ((Jahn et al., 2022)), temporal perspectives on the academic-practitioner gap ((Man et al., 2022)), and critical pragmatism ((Meyer, 2021))—provides essential theoretical foundations for understanding how knowledge can be integrated across academic and practitioner domains. However, it offers little empirical insight into the actual decision-making processes, organizational change strategies, and governance approaches that executives employ when scaling autonomous AI systems. This gap is particularly consequential given the novelty of agentic AI, the absence of established regulatory frameworks, and the profound organizational implications of deploying systems that can act independently.

11.2 Contributions and Implications

This research makes several contributions to both scholarly knowledge and practitioner practice.

For scholarly knowledge, the proposed study will generate empirical evidence on an underexplored phenomenon: how senior leaders in regulated industries navigate the implementation of autonomous AI systems. By integrating adaptive leadership theory with AI governance constructs, the conceptual framework developed in this paper offers a novel lens for understanding the interplay between leadership behaviors and governance structures in the context of agentic AI. The qualitative multiple-case study design will enable rich, contextualized insights that can inform theory development at the intersection of strategic leadership, technology management, and risk governance.

For practitioner practice, the research will yield actionable frameworks and guidance that address the specific challenges identified in the practitioner job descriptions. Executive decision frameworks will provide structured approaches for portfolio prioritization, resource allocation, and tool selection across the agentic AI stack. Governance templates will offer adaptable models for establishing MRM practices, human oversight mechanisms, and accountability structures tailored to autonomous systems. Change management guidance will inform organizational transformation strategies that address workforce transition, role redesign, and cultural adaptation. Executive development frameworks will support the cultivation of technical fluency in agentic systems and adaptive governance capabilities among current and future leaders.

For the academic-practitioner gap, this research embodies the Scholar-Practitioner approach advocated by (Boss, 2022) and (Kim-Keung Ho, 2014). By grounding research questions in practitioner realities, drawing on theoretical frameworks to structure inquiry, and committing to produce outputs that are both rigorous and useful, this study demonstrates how academic research can meaningfully contribute to addressing complex organizational challenges. The integration of transdisciplinary modes of inquiry ((Jahn et al., 2022)) and critical pragmatist philosophy ((Meyer, 2021)) offers a methodological template for future research at the intersection of emerging technology and strategic leadership.

11.3 Limitations and Future Research

Several limitations of this proposed research should be acknowledged. The focus on banking and insurance may limit generalizability to other regulated industries such as healthcare or energy. The qualitative case study design, while enabling depth and richness, does not aim for statistical generalizability. The rapid evolution of agentic AI technologies means that findings may require ongoing updating as technical capabilities and regulatory frameworks develop.

Future research should extend this inquiry in several directions. Longitudinal studies could examine how executive approaches to agentic AI governance evolve as technologies mature and regulatory precedents emerge. Comparative studies across industries, geographies, and organizational sizes could identify contextual factors that shape implementation strategies. Quantitative research could test the relationships between governance practices, leadership behaviors, and organizational outcomes. Finally, as agentic AI systems become more sophisticated, research on the ethical implications of autonomous decision-making—including issues of accountability, transparency, and fairness—will become increasingly critical.

11.4 Concluding Remarks

The agentic AI paradox—the imperative to simultaneously accelerate autonomous systems while maintaining human control—defines the leadership landscape of our era. Executives in regulated financial institutions stand at the frontier of this transformation, navigating uncharted territory where the stakes are immense and the playbooks are unwritten. This paper has argued that addressing these challenges requires a Scholar-Practitioner approach that integrates rigorous academic methodology with deep practitioner relevance. By investigating how senior leaders actually navigate the complexities of agentic AI implementation, this research aims to generate knowledge that is both theoretically robust and practically useful—contributing to the development of evidence-based guidance for executives confronting the defining technological challenge of our time.

As (Boss, 2022) observed, the scholar-practitioner approach is not merely an academic exercise but a necessity for addressing complex organizational challenges where existing knowledge is insufficient. In the age of agentic AI, this necessity has never been more urgent. The integration of scholarly theory with practitioner insights, following the models articulated by (Kim-Keung Ho, 2014) and (Tiessen et al., 2021), offers a path toward knowledge that can inform both understanding and action. This research seeks to walk that path, contributing to a future where autonomous AI systems are deployed not only with technical sophistication but with the wisdom, governance, and accountability that regulated industries demand and society deserves.

Funding: This research received no external funding
Conflicts of Interest: The authors declare no conflict of interest.
ORCID iD https://orcid.org/0009-0002-6011-5080

Data Availability Statement: Datasets and analytical artifacts supporting this study are available in a public GitHub repository maintained by the author (Joshi, 2026).
AI Use Statement: GitHub Copilot was used to support code drafting and repository management, and DeepSeek models were used to assist with language refinement and structural clarity.

References

[1] Boss, G. J. (2022). The great designation: An editorial essay on the necessity of scholar-practitioners. College Student Affairs Journal, 40(2), 1–12.
[2] Bryant, A., & Vieregger, C. (2021). Relevant business research: A viewpoint to help bridge the gap between academia and practice. Drake Management Review, 11(1-2), 5–8.
[3] Creswell, J. (2014). Research design: Qualitative, quantitative, and mixed methods approaches (4th ed.). Sage.
[4] Danford, C. A. (2023). Understanding the evidence: Qualitative research designs. Urologic Nursing, 43(1).
[5] DaSilva, C. M. (2018). Understanding business model innovation from a practitioner perspective. Journal of Business Models, 6(2), 19.
[6] Heifetz, R. A. (1994). Leadership without easy answers. Harvard University Press.
[7] Jahn, S., Newig, J., Lang, D. J., Kahle, J., & Bergmann, M. (2022). Demarcating transdisciplinary research in sustainability science—five clusters of research modes based on evidence from 59 research projects. Sustainable Development, 30(2), 343–357. https://doi.org/10.1002/sd.2278
[8] Kim-Keung Ho, J. (2014). A theoretical review on the professional development to be a scholar-practitioner in business management. European Academic Research, 1(12), 5393–5422.
[9] Kivunja, C. (2018). Distinguishing between theory, theoretical framework, and conceptual framework: A systematic review of lessons from the field. International Journal of Higher Education, 7(6), 44–53.
[10] Lawler III, E. E., & Benson, G. S. (2022). The practitioner-academic gap: A view from the middle. Human Resource Management Review, 32(1), 100748.
[11] Lyndon, A., Williams, B., & Hua, J. (2023). Scholarly writing for doctoral students. Indiana Tech. https://phd.indianatech.edu/wp-content/uploads/sites/8/Scholarly-writing-for-doctoral-students.pdf
[12] Mager, D. (2022). How to read research papers quickly and efficiently. Researcher.Life. https://researcher.life/blog/article/read-research-papers-quickly/
[13] Man, A. P. de, Luvison, D., & Leeuw, T. de. (2022). A temporal view on the academic–practitioner gap. Journal of Management Inquiry, 31(2), 181–196. https://doi.org/10.1177/1056492620982375
[14] McNeill, R., & Nienaber, H. (2018). Bridging the academic-practitioner chasm: Towards a model of hotel B2B sales. Journal of Global Business and Technology, 14(1), 40–52.
[15] Meyer, L. (2021). A case for critical pragmatism in OD. Organization Development Review, 53(5), 1–6.

[16] Singh, S. (2023a). What is a conceptual framework and how to make it (with examples). Researcher.Life. https://researcher.life/blog/article/what-is-a-conceptual-framework-and-how-to-make-it-with-examples/
[17] Singh, S. (2023b). What is research design? Understand types of research design, with examples. Researcher.Life. https://researcher.life/blog/article/what-is-research-design-types-examples/
[18] Tiessen, R., Cadesky, J., Lough, B. J., & Delaney, J. (2021). Scholar/practitioner research in international development volunteering: Benefits, challenges and future opportunities. Canadian Journal of Development Studies, 42(3), 394–415.List the reference here
[19] Joshi, S. (2026). Leading autonomous AI in financial services for C-suite executives: Supporting data and analytical artifacts [GitHub repository]. GitHub. https://github.com/satyadharjoshi/Leading-Autonomous-AI-in-Financial-Services-for-C-Suite-Executives-