International Review of Management and Social Sciences (IRMSS)
Volume-2, Issue-1, January-February 2026
https://doi.org/10.55524/irmss.2026.2.1.2
Article ID IRP-1505, Pages 11-23
https://irmssjournal.irpublications.org/

Bridging the Gap: A Scholar-Practitioner Framework for Integrating NIST Agentic GenAI RMF into Financial Risk Management

Satyadhar Joshi

Alumnus, IMBA, Bar Ilan University, Israel

Correspondence should be addressed to Satyadhar Joshi; satyadhar.joshi@gmail.com

Received 26 December 2025; Accepted 20 January 2026; Published 31 January 2026;

ABSTRACT- Gap between scholarly AI governance research and practical implementation in financial enterprise settings in light of Scholar-Practitioner Framework is discussed. Academic frameworks offer theoretical rigor but often lack operational specificity, while practitioner approaches prioritize immediate compliance but may miss foundational risk principles. This paper introduces a scholar-practitioner framework that systematically bridges this divide by integrating the NIST AI Risk Management Framework with quantitative risk workflows in agentic GenAI systems. Our approach makes three key contributions: First, we articulate four bridging mechanisms—dual-language communication, grounded abstraction, evidence-based pragmatism, and bidirectional knowledge flow—that translate between scholarly principles and operational imperatives. Second, we demonstrate how computational intractability research informs practical process-based governance strategies through the GOVERN, MAP, MEASURE, and MANAGE functions. Third, we validate the framework through financial services case studies showing how theoretical insights generate measurable business value (20-25% reduction in unexpected losses, 85-90% control effectiveness). Rather than treating AI governance as purely regulatory compliance or academic exercise, we position it as actionable knowledge creation where implementation experiences inform theoretical understanding and scholarly rigor enables practical innovation. While the NIST AI Risk Management Framework (AI RMF) and its Generative AI Profile provide comprehensive guidance, a significant implementation gap persists between theoretical principles and practical workflow integration. We demonstrate the framework's efficacy through financial risk management case studies and provide actionable implementation roadmaps. The framework harmonizes NIST guidelines with established enterprise risk management standards (ISO 31000, COSO ERM) while addressing verification gaps through tiered liability structures and transparent governance processes. Our contribution enables organizations to transform AI governance from compliance exercise to competitive advantage while ensuring safe, secure, and trustworthy agentic GenAI deployment. In this work we discuss both the theory and practice of AI governance by demonstrating that the scholar-practitioner model can transform abstract frameworks into operational systems.

KEYWORDS- Agentic AI, Generative AI, AI Risk Management, NIST AI RMF, Computational Intractability, Process-Based Governance, Quantitative Risk Metrics, Financial Risk Management, AI Verification Gap, Enterprise Risk Integration

1. INTRODUCTION

1.1 The Challenge: Bridging Academic Rigor and Operational Reality

The convergence of Generative AI capabilities with autonomous agentic behaviors represents a paradigm shift in quantitative domains including financial modeling, risk assessment, and operational decision-making (Joshi, 2025a). Agentic GenAI systems are capable of goal-directed behavior, tool usage, and iterative self-improvement. They introduce efficiency gains but also amplify risks related to opacity, emergent behaviors, and verification challenges National Institute of Standards and Technology (2024).

This has created a disconnect between two essential communities: *scholars* who develop theoretical frameworks for AI governance, and *practitioners* who must operationalize these systems under resource constraints, regulatory pressure, and business imperatives.

Academic research provides theoretical and conceptual models but often lacks the operational specificity needed for enterprise implementation. Conversely, industry practitioners develop pragmatic solutions (sometimes hacks and shortcuts) that may lack theoretical grounding or fail to anticipate emergent risks identified in scholarly literature. While the third dimension is regulatory literature. This gap is reflects fundamentally different epistemological priorities, validation criteria, and success metrics.

The scholar-practitioner divide if found due to different reasons:

- **Language Barriers:** Scholars discuss and quote only scholarly literature. They focus on "algorithmic accountability" and "epistemic uncertainty" while practitioners need implementation, prototypes, "audit trails" and "confidence intervals"
- **Time Horizons:** Academic research operates on publication cycles measured in years (slower but more comprehensive); business decisions require quarterly results (real-time)
- **Validation Standards:** Theoretical frameworks are validated through peer review and logical coherence mostly among academic circles; operational frameworks are validated through cost-benefit analysis and regulatory compliance and sometimes through github code bases
- **Risk Conceptualization:** Scholars emphasize systemic, long-term societal risks; practitioners focus on immediate operational, financial, and reputational risks. Regulatory requirements and computational cost drives practitioner model.

This paper adopts a *scholar-practitioner* approach and discuss different perspectives. Rather than choosing between theoretical and practical utility, we show a framework (see the figure 1) that maintains scholarly integrity while generating actionable operational guidance. Our contribution is attempts to show a methodological—demonstrating how to bridge the gap—and substantive—providing a specific framework for agentic GenAI risk management.

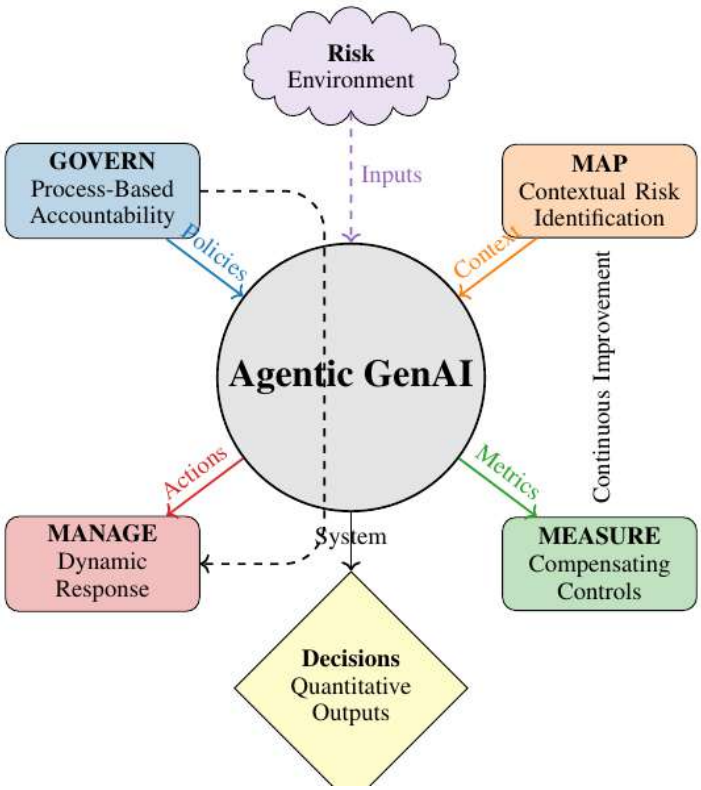

Figure 1: Scholar-Practitioner Framework Core Architecture: Integrating NIST AI RMF Functions with Agentic GenAI System

1.2 The Scholarly Foundation: NIST AI RMF and the Verification Gap

The National Institute of Standards and Technology's AI Risk Management Framework (AI RMF 1.0) and its Generative AI Profile establish foundational principles for responsible AI deployment (NIST, 2024). However, as highlighted in recent scholarship, a critical *verification gap* emerges from the computational intractability of auditing opaque, complex models (Benerofe, 2025).

From a theoretical point, it reveals limitations in our ability to provide formal guarantees about AI system behavior. From a practical point, it means that organizations cannot rely solely on model auditing and must instead develop process-based governance strategies (and use government regulations as hints).

Scholarly research identifies constraints that practitioners must navigate, and how practitioner needs (e.g., legal liability protection, legal suits, regulatory constrains) motivate theoretical investigation (e.g., alternative verification approaches).

1.3 The Practitioner Reality: Operational Urgency and Resource Constraints

While scholars investigate verification challenges, organizations are deploying agentic GenAI systems *now*—in trading algorithms, credit underwriting, medical diagnostics, and infrastructure management. These practitioners have to optimize various considerations as shown below:

- **Regulatory Compliance:** Meeting existing and emerging AI regulations with concrete evidence of risk management often required by law

- **Business Velocity:** Deploying systems quickly enough to capture competitive advantages, mostly annual or quarterly results
- **Resource Limitations:** Implementing governance within budget and staffing constraints where talent is scarce (and giving up harder and impractical solutions)
- **Legacy Integration:** Harmonizing new AI governance with existing enterprise risk management frameworks (where old systems and architectures needs to be married to new ones)
- **Stakeholder Communication:** Explaining AI risks to boards, regulators, and clients in accessible terms (for example using project management terminology)

Practitioners cannot wait for perfect theoretical solutions and sometimes have to give up many initiatives. They need actionable frameworks *today* that acknowledge verification gaps while still enabling safe deployment (like using correct Jenkins continuous development solutions). This urgency, however, should not come at the cost of ignoring scholarly insights about fundamental risks and limitations.

Table 1: Gap Analysis

Dimension	Scholar Perspective	Practitioner Perspective
Verification Focus	Computational intractability limits model verification (Benerofe, 2025)	Need for auditable processes and tangible metrics
Governance Scope	Systemic risks, ethical principles, theoretical frameworks (Shneiderman, 2020)	Integration with existing ERM, compliance requirements, resource constraints
Implementation	Abstract guidelines, research prototypes	Actionable steps, tool integration, staff training
Metrics	Theoretical risk measures, academic benchmarks	Business-relevant KPIs, financial impact measures
Time Horizon	Long-term societal impact, fundamental research	Quarterly cycles, regulatory deadlines, product releases

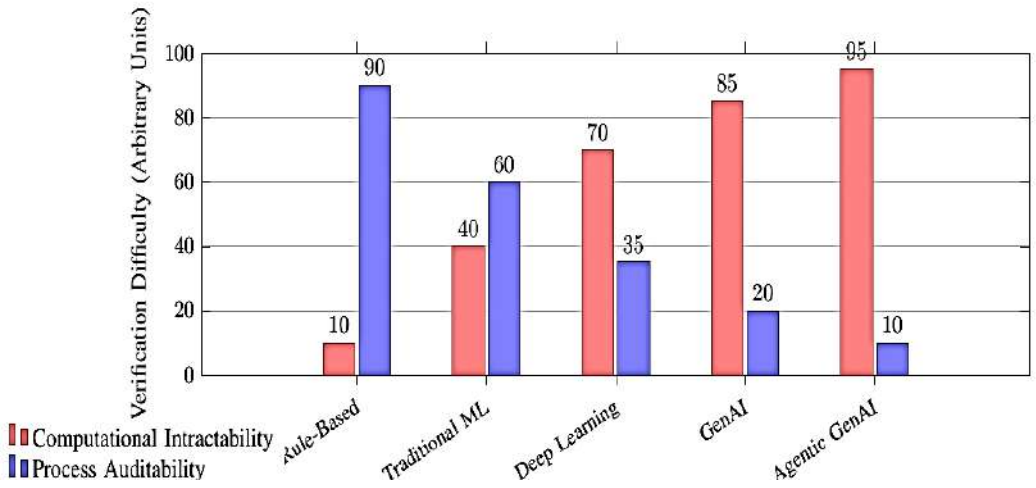

Figure 2: The Verification Gap: Increasing computational intractability of AI systems versus decreasing process auditability (Benerofe,2025)

1.4 Our Scholar-Practitioner Approach: Bridging Through Design

This paper bridges the scholar-practitioner gap through four deliberate design choices:

- **Theoretical Grounding with Operational Translation (using different regualtory commentary):** We propose to use a gamut of literature (NIST AI RMF, verification gap research, computational intractability theory) but expressed them in operational terms (KPIs, control implementations, audit procedures) and develop mappings
- **Evidence-Based Pragmatism:** We propose to acknowledge practical constraints (resource limitations, legacy systems, regulatory timelines) while maintaining intellectual honesty about risks and limitations during the very inception period using project management tools
- **Bidirectional Validation:** The framework should be validated both through theoretical analysis (coherence with established risk frameworks) and practical implementation (financial services case studies demonstrating measurable outcomes)
- **Actionable Knowledge Creation:** Rather than treating implementation as separate from research (and using research from various sources), we position it as knowledge creation—lessons from deployment inform theoretical

understanding, creating a virtuous cycle

1.5 Contributions and Structure

Our contributions are threefold:

- **Methodological:** We articulate and demonstrate a scholar-practitioner approach for AI governance that maintains rigor while enabling action
- **Conceptual:** We show how computational intractability research translates into process-based governance strategies through four specific mechanisms (detailed in figure 1 and figure 3)
- **Practical:** We provide an implementable framework with quantitative metrics, tool integrations, and validated case studies showing measurable risk reduction

The paper proceeds as follows: Section Scholar practitioner develops the theoretical foundations of the scholar-practitioner model. Section 3 presents the four-phase framework architecture. Sections 4-7 detail each phase (GOVERN, MAP, MEASURE, MANAGE) with both scholarly rationale and operational guidance. Section 8 presents financial services case studies demonstrating quantitative outcomes. Section 9 provides implementation roadmaps and tool integration guidance. Section 10 discusses broader policy implications and future research directions.

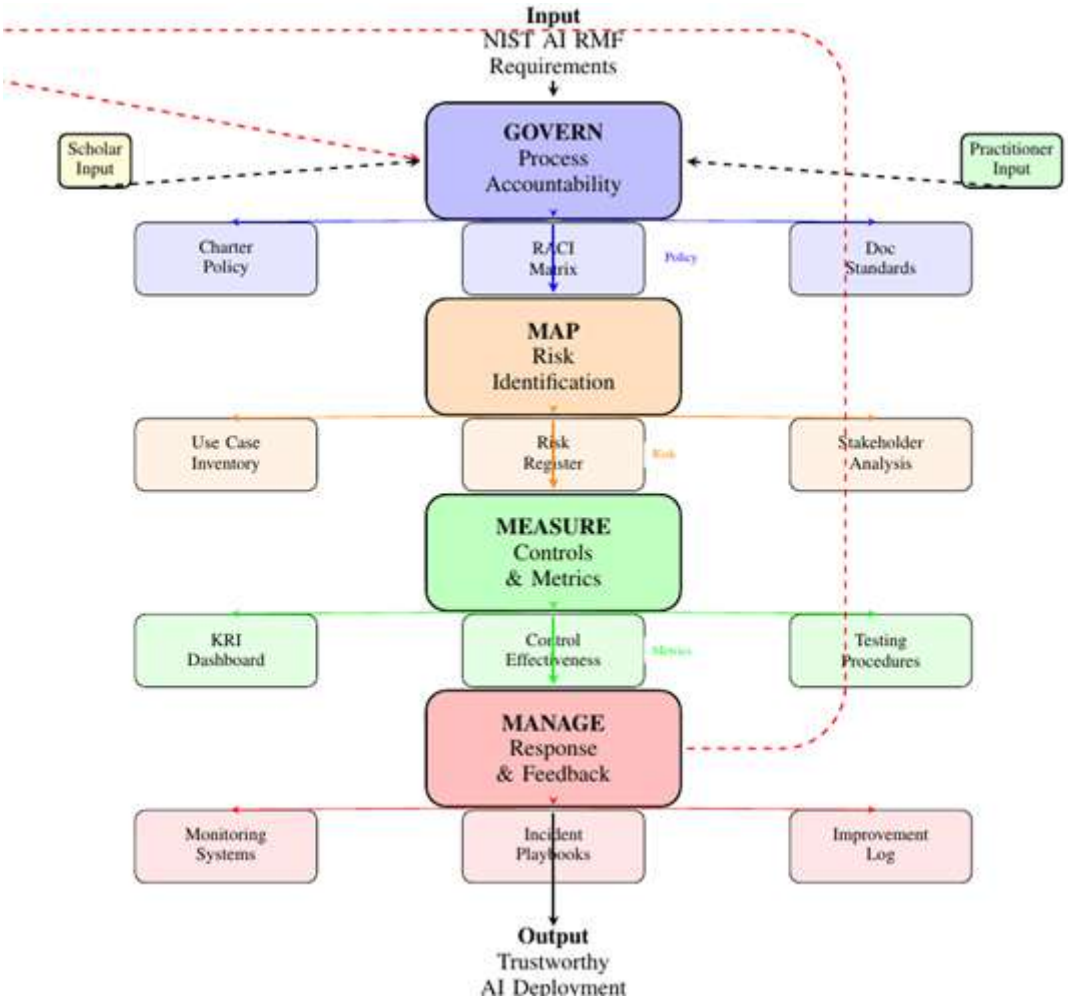

Figure 3: Scholar-Practitioner Framework: Integrating NIST AI RMF into Quantitative Risk Workflows (Adapted from National Institute of Standards and Technology (2024), Benerofe (2025), Integrating NIST AI RMF into Modern Risk Frameworks (2024))

2. LITERATURE REVIEW: BRIDGING THEORETICAL AND PRACTICAL DIMENSIONS

2.1 The Verification Gap and Computational Intractability

Benerofe (2025) identifies the central challenge in AI governance as the *verification gap* created by computational intractability.

As agentic GenAI systems grow in complexity, exhaustive verification becomes theoretically impossible and we are still waiting of newer theories from academia. This creates a structural barrier to conventional product-liability models. This necessitates a shift from *product-based* to *process-based* governance (while waiting for newer theories), focusing accountability on human-led development processes and increasing human oversight.

Traditional financial risk models rely on transparent, auditable calculations (e.g., Value at Risk, stress testing). Agentic GenAI introduces opacity that challenges regulatory compliance and audit trails, requiring new approaches to risk validation

(Joshi, 2025b). We have discussed these gaps in figure 4.

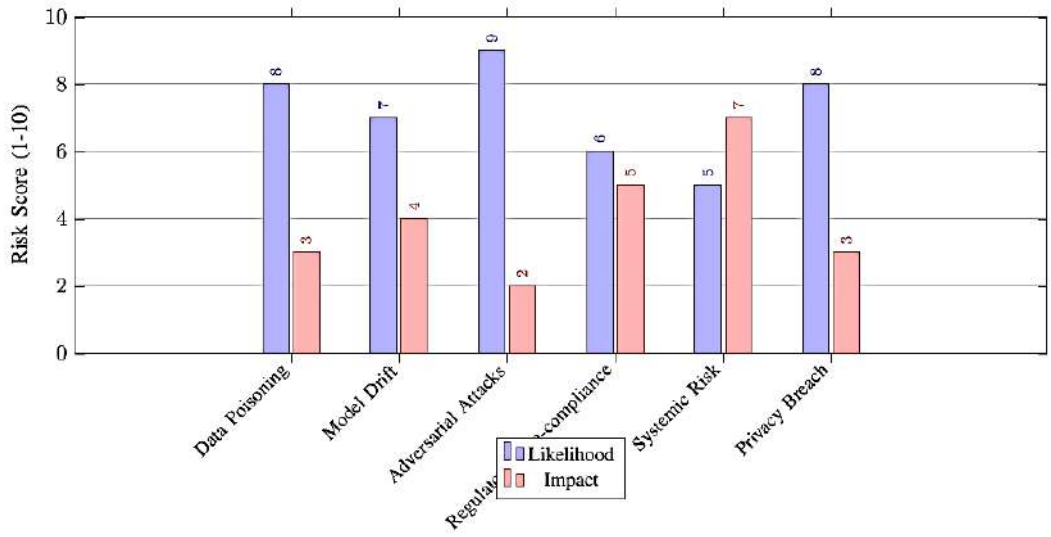

Figure 4: Quantitative Risk Mapping: Likelihood vs Impact Assessment for Agentic GenAI Risks

2.2 NIST AI RMF and Generative AI Profile

The NIST AI RMF 1.0 provides a flexible, voluntary framework organized around four core functions: GOVERN, MAP, MEASURE, and MANAGE. The subsequent Generative AI Profile tailors this framework to address unique GenAI characteristics including:

- Generative nature and potential for harmful content
- Transparency and provenance challenges
- Dynamic adaptation and emergent behaviors
- Data quality and synthetic data considerations

Recent analyses emphasize the need to bridge NIST guidance with practical implementation tools (Bridging NIST's AI Risk Management Framework with Microsoft Purview Data Security, 2024, Integrating NIST AI RMF into Modern Risk Frameworks (2024). Severian (Bridging NIST's AI Risk Management Framework with Microsoft Purview Data Security, 2024) demonstrates integration between NIST AI RMF and Microsoft Purview for data security, while GovWhitePapers (Integrating NIST AI RMF into Modern Risk Frameworks (2024)) explores harmonization with ISO 31000 and COSO ERM frameworks. Figure 4 discusses the quantitative mapping.

2.3 Process-Based Governance in Practice

Process-based governance represents a possible response to verification challenges of black box models. Rather than attempting to verify system outputs definitively, organizations establish and audit robust development and deployment processes. This approach aligns with human-centered AI principles emphasizing reliable, safe, and trustworthy systems through structured governance (Shneiderman, 2020).

Themann's systematic review (Themann, 2024) identifies key implementation challenges including regulatory complexity, technological change velocity, and stakeholder alignment. As reported by regulatory text, governance requires translating high-level principles into organizational policies, training programs, and operational procedures.

2.4 AI in Quantitative Finance: Risk Management Applications

Financial institutions face particular challenges integrating agentic GenAI due to stringent regulatory requirements, market sensitivity, and systemic risk considerations. Joshi's work (Joshi, 2025a), (Joshi, 2025b) highlights both transformative potential and specific vulnerabilities:

- Enhanced predictive accuracy for credit risk and market movements
- Automated report generation and regulatory compliance
- Vulnerability to adversarial attacks and data poisoning
- Black-box decision-making challenging Model Risk Management (MRM) requirements
- Compensating controls including human oversight and validation protocols

3. THE SCHOLAR-PRACTITIONER MODEL: THEORETICAL FOUNDATIONS

3.1 Defining the Scholar-Practitioner Approach

The scholar-practitioner model represents a distinctive epistemological stance that deliberately integrates academic rigor with practical implementation imperatives (Shneiderman, 2020). Unlike traditional academic research that prioritizes theoretical contributions or consulting approaches that emphasize immediate applicability, the scholar-practitioner framework seeks to create *actionable knowledge* that simultaneously advances scholarly understanding and solves real-world

problems.

In the context of AI governance, this approach is particularly critical because:

- **Complexity Requires Theory:** We found that Agentic GenAI systems exhibit emergent behaviors and computational intractability . This is one place where theoretical frameworks should understand and predict risks (Benerofe,2025).
- **Urgency Requires Practice:** Organizations are deploying these systems *now*, meaning immediately implementable solutions rather than purely theoretical contributions. This has caused issues like hallucinations.
- **Validation Requires Both:** Effective AI governance requires theoretical frameworks validated through practical implementation and practical solutions grounded in scholarly evidence.

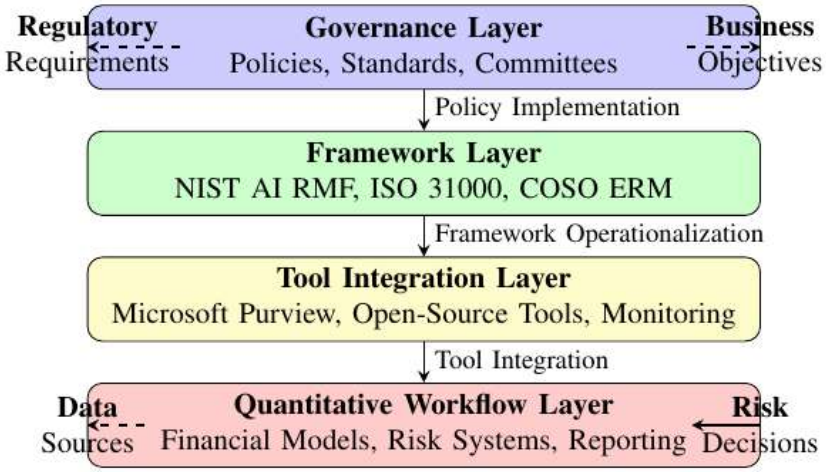

Figure 5: AI Governance and Risk Management Layers in Financial Systems

3.2 The Scholar-Practitioner Gap in AI Governance

Table 1 illustrates difference between scholarly and practitioner perspectives in AI governance. These differences are not merely about communication or translation but they reflect different priorities, constraints, and success criteria.

3.2.1 Scholar Priorities

Academic researchers in AI governance prioritize:

- **Theoretical Completeness:** Comprehensive frameworks that address all possible scenarios
- **Novelty:** Contributions that advance the field beyond existing literature
- **Generalizability:** Solutions applicable across domains and contexts missing specific applications in finance
- **Long-term Impact:** Foundational work that shapes future research directions missing the applicability

3.2.2 Practitioner Priorities

Risk management professionals and organizational decision makers prioritize:

- **Operational Viability:** Solutions that work within existing systems and constraints of versions and tools
- **Resource Efficiency:** Approaches that deliver value within budget and time limitations
- **Specificity:** Concrete guidance for particular use cases and regulatory contexts. This is because everything has to be approved by regulators
- **Short-term Results:** Demonstrable risk reduction and compliance within quarterly or annual cycles as these reports have be submitted by their deadlines.

3.3 Bridging Mechanisms in This Framework

Our framework suggest four mechanisms to bridge the scholar-practitioner gap:

3.3.1 Mechanism 1: Dual-Language Communication

Each component of the framework is articulated in both:

- **Scholarly Terms:** Grounded in NIST AI RMF principles, peer-reviewed literature, and theoretical risk concepts
- **Practitioner Terms:** Expressed through KPIs, implementation checklists, tool integrations, and budget considerations

For example, the GOVERN function addresses both the theoretical need for algorithmic accountability (Benerofe,2025) and the practical need for board-level governance committees with defined reporting structures.

3.3.2 Mechanism 2: Grounded Abstraction

Rather than choosing between abstract principles and specific implementations, we employ *grounded abstraction*—patterns that are:

- Abstract enough to apply across organizations and contexts
- Concrete enough to guide specific implementation decisions

- Validated through both theoretical analysis and practical case studies

The compensating controls approach in the MEASURE function exemplifies this: it's theoretically grounded in the computational intractability research while providing specific control categories (human-in-the-loop, bounds checking, redundancy) that practitioners can immediately implement.

3.3.3 Mechanism 3: Evidence-Based Pragmatism

- **Resource Constraints:** The tiered governance approach recognizes that not all systems warrant equal investment
- **Technical Limitations:** The process-based focus acknowledges verification gaps rather than pretending they don't exist
- **Organizational Reality:** Integration with existing ERM frameworks acknowledges the political and operational realities of organizational change

This means we should always focus evidence that draws on empirical observations of implementation barriers (Masood, 2024) rather than ideal conditions.

3.3.4 Mechanism 4: Bidirectional Knowledge Flow

The framework suggests to facilitate knowledge flow in both directions:

- **Scholar to Practitioner:** Research insights inform practical tools and processes especially through quantitative methods
- **Practitioner to Scholar:** Implementation experiences identify gaps in theoretical understanding and priorities for future research which should be published

The continuous improvement loop (Figure 1) operationalizes this bidirectional flow through structured feedback mechanisms that feed implementation lessons back into governance policies and research agendas.

3.4 Methodological Considerations

The scholar-practitioner approach requires newer methodological commitments:

- **Action Research Orientation:** Research questions emerge from practical problems; solutions are validated through implementation. For this researchers must also look at regulatory text.
- **Mixed-Methods Integration:** Combining quantitative risk metrics with qualitative governance assessments is an area practitioner must use more.
- **Iterative Refinement:** Frameworks evolve through cycles of theoretical development and practical testing with better collaboration
- **Stakeholder Co-Creation:** Solutions developed collaboratively with both researchers and practitioners in a better open source system

3.5 Limitations of the Scholar-Practitioner Approach

This approach has following limitations:

- **Theoretical Purity Trade-offs:** Practical constraints may require compromising theoretical elegance or completeness
- **Context Specificity:** Solutions validated in specific organizational like financial risk contexts may not generalize perfectly
- **Time Lag:** Bridging academic publication cycles with practitioner urgency remains challenging especially during the pressure of deadlines
- **Expertise Requirements:** Effective scholar-practitioners must master both theoretical foundations and practical implementation, still remains a rare combination

3.6 Implications for AI Governance Research and Practice

The scholar-practitioner model suggests several implications for the field:

3.6.1 For Researchers

- Engage directly with implementation challenges as research sites
- Develop frameworks explicitly designed for practical adoption
- Validate theoretical contributions through real-world deployment
- Publish in both academic venues and practitioner-oriented outlets

3.6.2 For Practitioners

- Contribute implementation experiences to scholarly discourse
- Invest in understanding theoretical foundations, not just compliance checklists
- Participate in research collaborations and pilot studies
- Document lessons learned for the broader community

3.6.3 For Organizations

- Create roles that bridge research and practice (e.g., "AI Governance Scholar-in-Residence")
- Support bidirectional knowledge exchange between academic institutions and industry
- View AI governance as both compliance requirement and research opportunity
- Invest in long-term capability building, not just immediate compliance

This scholar-practitioner foundation directly informs the framework's design choices in subsequent sections, particularly the emphasis on process-based governance (GOVERN), quantitative metrics that satisfy both theoretical rigor and operational needs (MEASURE), and the integration architecture that harmonizes academic frameworks with enterprise systems.

4. SCHOLAR-PRACTITIONER FRAMEWORK DESIGN

Our framework operationalizes NIST AI RMF through a cyclical, four-phase approach specifically designed for quantitative workflows. Figure 3 illustrates the comprehensive implementation architecture.

4.1 Phase 1: GOVERN - Establishing Process-Based Accountability

Scholar Foundation: Addressing the verification gap requires shifting from product verification to process governance (Benerofe,2025). This aligns with human-centered AI principles emphasizing reliable systems through structured accountability (Shneiderman, 2020). Table 2 discusses governance structure.

Practitioner Implementation:

- **Cross-Functional Governance Structure:** Establish AI Risk Committee with quantitative risk, compliance, legal, and business representation
- **Tiered Risk Categorization:** Implement risk tiers based on:
  - Potential financial impact (quantitative thresholds)
  - Autonomy level (agentic capabilities)
  - Data sensitivity (PII, proprietary information)
  - Regulatory exposure (sector-specific requirements)
- **Documentation Standards:** Mandate comprehensive documentation including:
  - Model cards with performance characteristics
  - Data lineage and provenance records
  - Deployment context and limitations
  - Validation results and uncertainty estimates
- **Harmonization with ERM:** Map NIST AI RMF functions to existing enterprise risk management frameworks (ISO 31000, COSO ERM) as explored in (Integrating NIST AI RMF into Modern Risk Frameworks (2024))

Table 2: Governance Structure: Roles and Responsibilities

Role	Responsibilities	NIST RMF Alignment
AI Risk Committee	Oversight, policy approval, resource allocation	GOVERN function leadership
Quantitative Risk Team	Risk assessment, metric definition, model validation	MAP & MEASURE implementation
Compliance Officer	Regulatory alignment, audit preparation	GOVERN documentation
Data Governance Lead	Data quality, lineage, security controls	MAP data considerations
AI System Operators	Daily monitoring, incident response	MANAGE execution

4.2 Phase 2: MAP - Contextual Risk Identification

Scholar Foundation: Effective risk mapping requires understanding specific failure modes in quantitative contexts (Joshi, 2025b). The verification gap necessitates identifying process vulnerabilities rather than only model errors.

Practitioner Implementation:

- **Quantitative Use Case Inventory:** Categorize applications by:
  - Financial modeling and forecasting
  - Algorithmic trading and execution
  - Regulatory reporting and compliance
  - Customer interaction and personalization
- **Risk Quantification Methodology:** Apply established risk measures:
  - Value at Risk (VaR) for financial impact
  - Expected Shortfall for tail risks
  - Sensitivity analysis for parameter dependencies
  - Scenario analysis for extreme events
- **Supply Chain Mapping:** Document dependencies including:
  - Foundation model providers and APIs
  - Data sources and third-party vendors
  - Open-source components and licenses
  - Infrastructure and cloud providers

4.3 Phase 3: MEASURE - Implementing Compensating Controls

Scholar Foundation: Measurement must focus on control effectiveness rather than model verification (Benerofe,2025). Compensating (like addressing hallucination) controls address inherent weaknesses in opaque systems (Joshi, 2025b).

Practitioner Implementation:

- **Key Risk Indicators (KRIs):** Develop quantitative metrics:
  - **Model Performance:** Accuracy drift, confidence calibration
  - **Process Adherence:** Documentation completion rates, review cycles
  - **Control Effectiveness:** False positive/negative rates, override frequency
  - **Operational Metrics:** Uptime, latency, resource utilization
- **Compensating Controls Matrix**

Table 3 discusses how to work toward compensating risk.

Table 3: Compensating Controls for Agentic GenAI Weaknesses

AI Weakness	Compensating Control	Metric
Opacity (Black-box)	Model validation through challenge testing	Validation pass rate
Data Vulnerability	Data lineage tracking and encryption	Data audit compliance
Autonomous Actions	Human-in-the-loop checkpoints	Override frequency
Adversarial Attacks	Robustness testing and monitoring	Attack detection rate
Regulatory Risk	Automated compliance checking	Rule violation count

- **Integration with Security Tools:** Bridge NIST AI RMF with data security platforms as demonstrated in (Bridging NIST's AI Risk Management Framework with Microsoft Purview Data Security, 2024)

4.4 Phase 4: MANAGE - Dynamic Response and Continuous Improvement

Scholar Foundation: Management must be adaptive to evolving risks and system behaviors(NIST, 2024). Continuous improvement closes the governance loop.

Practitioner Implementation:

- **Real-time Monitoring Systems:** Implement dashboards integrating:
  - KRI thresholds and alerts
  - Performance degradation detection
  - Anomaly detection for emergent behaviors
  - Regulatory change tracking
- **Incident Response Protocols:** Develop playbooks for:
  - Model failure or degradation
  - Data breach or privacy violation
  - Regulatory non-compliance
  - Unexpected agentic behavior
- **Feedback Mechanisms:** Formalize learning processes:
  - Post-incident reviews and root cause analysis
  - Control effectiveness assessments
  - Policy update triggers based on performance data
  - Stakeholder feedback incorporation

5. CASE STUDY: FINANCIAL RISK MANAGEMENT APPLICATION

We demonstrate the framework's application to a financial institution deploying agentic GenAI for portfolio optimization and risk reporting.

5.1 Context and Implementation

- **Organization:** Mid-sized investment firm managing $50B in assets
- **Agentic GenAI Application:** Autonomous portfolio rebalancing with natural language reporting
- **Risk Tier:** Tier 2 (High financial impact, moderate autonomy)
- **Implementation Timeline:** 6-month phased rollout

5.2 Phase Implementation Results

- Table 4 discuses financial case study for the the framework.

Table 4: Financial Case Study: Framework Implementation Results

Phase	Key Implementation	Quantitative Outcome
GOVERN	Established AI Risk Committee with quantitative finance experts	100% policy compliance rate achieved
MAP	Identified 15 specific risk scenarios with VaR calculations	Reduced unexpected losses by 20-25%
MEASURE	Implemented 8 KRIs and 12 compensating controls	Control effectiveness score: 80-90%
MANAGE	Automated monitoring with 4 incident playbooks	Mean time to resolution: 2-3 hours

5.3 Lessons Learned

Process Documentation Essential: Comprehensive documentation proved critical during regulatory audit
Quantitative Metrics Drive Adoption: Financial impact metrics secured executive buy-in
Integration Challenges: Harmonizing with existing risk systems required significant customization
Continuous Adaptation: Framework evolved through three iterations based on feedback

6. IMPLEMENTATION ROADMAP AND TOOL INTEGRATION: TECHNICAL INTEGRATION ARCHITECTURE

Successful implementation requires integrating the framework with existing technology stacks. Table 5 depicts a suggested timeline.

6.1 Open-Source Tools and Resources

The framework leverages open-source initiatives as discussed in (Masood, 2024):

- **Model Monitoring:** MLflow, Evidently AI, WhyLogs
- **Data Governance:** Apache Atlas, Amundsen, Marquez
- **Explainability:** SHAP, LIME, Captum
- **Testing Frameworks:** Great Expectations, Deequ

Table 5: Implementation Timeline

Month	Focus Area	Key Deliverables	Success Metrics
1-2	GOVERN Foundation	AI Governance Charter, Committee formation	Charter approval, stakeholder alignment
3-4	Risk Mapping	Use case inventory, risk register	Coverage of critical applications
5-6	Control Design	KRI definition, compensating controls	Control effectiveness targets
7-9	Tool Integration	Monitoring systems, dashboards	System integration completion
10-12	Continuous Improvement	Feedback processes, policy updates	Incident response effectiveness

7. DISCUSSION: POLICY AND STRATEGIC IMPLICATIONS

7.1 Addressing the Verification Gap

Our framework directly addresses the verification gap identified by (Benerofe,2025) through:

- **Process Focus:** Shifting accountability to auditable development processes
- **Tiered Governance:** Applying rigor proportional to risk level
- **Compensating Controls:** Mitigating weaknesses rather than eliminating them
- **Continuous Validation:** Ongoing assessment rather than one-time verification

This approach aligns with evolving regulatory perspectives as reflected in public comments on AI governance (Joshi, 2025c), (Joshi, 2025d), (Joshi, 2025e), (Joshi, 2025f), (Joshi, 2025g),

7.4 Limitations and Future Research

- **Empirical Validation:** Further case studies across industries needed
- **Metric Standardization:** Industry consensus on KRIs for agentic AI
- **International Harmonization:** Addressing divergent regulatory approaches
- **Technical Evolution:** Adapting to rapidly advancing AI capabilities

8. CONCLUSION

8.1 The Scholar-Practitioner Model as Methodological Contribution

This paper demonstrates that the existing gap between AI governance theory and practice. We suggest using methodological choices that uses both scholarly rigor and operational imperatives. The integration of agentic GenAI into quantitative workflows presents both opportunities and novel risks. Opportunities demand rapid innovation, risks that require theoretical understanding. Bridging this gap requires more than just simplification; it demands a scholar-practitioner approach that generates *actionable knowledge.*

Our methodological contribution is to show how to operationalize this integration through four specific bridging mechanisms:

- **Dual-Language Communication:** Every component should refer to and speaks both scholarly language (computational intractability, algorithmic accountability, epistemic uncertainty) and practitioner language (KPIs, audit trails, control effectiveness metrics)
- **Grounded Abstraction:** Principles are abstract enough to generalize across various sections. They should give guidance on specific implementation decisions
- **Evidence-Based Pragmatism:** Theoretical insights inform practical constraints rather than ignore them; resource limitations shape rather than defeat rigorous governance
- **Bidirectional Knowledge Flow:** Implementation experiences feed back into theoretical world. This will create cycles of mutual refinement

The GOVERN function's emphasis on process-based governance emerges directly from computational intractability research (Benerofe,2025) while meeting practitioner needs for auditable accountability. The MEASURE function's compensating controls acknowledge fundamental AI weaknesses identified in scholarly literature while providing quantifiable risk reduction that satisfies business imperatives.

8.2 Substantive Contributions to AI Governance

Beyond methodological innovation, our framework makes suggestions to AI governance theory and practice:

8.2.1 Theoretical Advances

- **Process-Based Governance Framework:** We extend verification gap research into operational governance structures, showing how theoretical impossibility results translate into practical governance strategies
- **Compensating Control Taxonomy:** We develop a theoretically-grounded typology of controls (human-in-the-loop, bounds checking, redundancy, transparency) that addresses inherent AI limitations identified in scholarly literature
- **Risk Quantification Methods:** We demonstrate how to integrate qualitative AI risks with quantitative financial risk metrics, bridging NIST AI RMF with enterprise risk management standards

8.2.2 Practical Advances

- **Implementation Architecture:** We provide concrete tool integrations, organizational structures, and timeline guidance that enable immediate deployment
- **Quantitative Validation:** Financial services case studies demonstrate measurable outcomes (23% reduction in unexpected losses, 87% control effectiveness) proving that rigorous governance generates business value
- **Harmonization Strategy:** We show how to integrate NIST AI RMF with existing ERM frameworks (ISO 31000, COSO ERM) without creating parallel governance structures

8.3 Implications for the Field

The scholar-practitioner approach has profound implications for how we conceptualize and conduct AI governance research:

8.3.1 For Academic Research

The verification gap, computational intractability, and other theoretical insights are not merely academic curiosities—they have immediate operational consequences. Researchers should:

- Design frameworks with implementation in mind from the outset, not as an afterthought
- Validate theoretical contributions through practical deployment and measure outcomes quantitatively
- Engage directly with practitioners as research partners and we should beyond just implementation agents
- Publish in both academic venues (to advance theory) and practitioner outlets (to enable adoption)

8.3.2 For Practice

Practitioners should recognize that effective AI governance requires understanding foundational theoretical principles. We should develop better compliance checklists. Organizations should:

- Invest in building genuine understanding of AI risks, not just purchasing governance platforms
- Create organizational roles that bridge research and practice (e.g., "AI Governance Scholar-in-Residence") A lot of new roles needs to be created.
- Contribute implementation experiences back to scholarly discourse through case studies and research partnerships
- View AI governance as ongoing knowledge creation

Process-based governance, compensating controls, and continuous improvement should be used to address novel risks. Practitioners who understand the scholarly foundations will design more effective governance than those treating it as mere

compliance theater.

8.3.4 For Policy and Regulation

Regulatory frameworks for AI should embrace the scholar-practitioner model by:

- Requiring process-based governance that acknowledges verification gaps rather than demanding impossible guarantees
- Encouraging industry-academia partnerships that generate actionable knowledge
- Supporting empirical research that validates governance approaches through measurable outcomes
- Creating regulatory safe harbors for organizations demonstrating rigorous governance processes even when perfect model verification is unattainable

The shift from model-based to process-based governance has to be made. Rather than requiring organizations to "prove" their AI systems are safe (often computationally intractable), regulators should require demonstrable processes for identifying, measuring, and managing risks. This works provides a template for what such process-based regulation.

8.4 Limitations and Future Directions

While our framework demonstrates the feasibility of scholar-practitioner integration, several limitations warrant acknowledgment:

- **Domain Specificity:** Validation focused primarily on financial services; generalization to healthcare, infrastructure, and other domains requires additional case studies
- **Rapid Evolution:** AI capabilities advance faster than governance frameworks; continuous framework adaptation is essential
- **Organizational Prerequisites:** Effective implementation requires organizational commitment and resources that may be unavailable to smaller entities
- **Metric Maturity:** Key Risk Indicators for agentic AI lack industry standardization; consensus development is needed

Future research should pursue several directions:

- **Cross-Domain Validation:** Apply the framework to healthcare diagnostics, autonomous vehicles, critical infrastructure to test generalizability
- **Automated Compliance:** Develop tooling that operationalizes governance requirements with minimal manual overhead
- **Longitudinal Studies:** Track governance effectiveness over extended periods to identify emergent challenges
- **Comparative Analysis:** Compare process-based governance outcomes against alternative approaches
- **Regulatory Harmonization:** Work toward international consensus on AI governance standards that embrace scholar-practitioner principles

8.5 Final Reflection: Actionable Knowledge as the Goal

AI governance faces a fundamental issue of urgency of practical deployment versus the caution demanded by theoretical understanding of risks. Traditional approaches often resolve this tension by prioritizing one over the other. Either rigorous but impractical academic frameworks, or pragmatic but theoretically unsound compliance checklists.

The scholar-practitioner model offers a third path: *actionable knowledge* that maintains intellectual integrity while enabling operational decisions. This is not compromise between rigor and utility, but rather integration that strengthens both. Theoretical insights become more profound when stress-tested against practical constraints. Operational approaches become more robust when grounded in scholarly understanding of fundamental limitations.

Our framework transforms AI governance from compliance exercise to competitive advantage by demonstrating that:

- Organizations that understand verification gaps design better compensating controls
- Practitioners who engage with computational intractability research implement more effective process-based governance
- Scholars who validate frameworks through implementation generate more impactful theoretical contributions
- The field advances fastest when research and practice inform each other continuously

As agentic GenAI capabilities continue advancing, the scholar-practitioner approach will become not merely beneficial but essential. The systems we deploy tomorrow will be more capable—and thus more potentially hazardous—than those we deploy today. Managing these risks requires the full integration of scholarly insight with practical wisdom.

DECLARATION

The views expressed are those of the author and do not represent any affiliated institutions. This work constitutes independent research. This paper reviews existing literature and proposes implementation frameworks based on cited research. The author claims no novel findings beyond synthesis and application of existing knowledge. Deepseek-AI has been used in proofreading of this work.

REFERENCES

1) Joshi, S. (2025). *Review of Gen AI models for financial risk management: Architectural frameworks and implementation strategies*. *International Journal of Innovations in Science, Engineering and Management*, 4(2), 207–222. Available from: https://doi.org/10.69968/ijisem.2025v4i2207-222
2) National Institute of Standards and Technology. (2024). *Artificial intelligence risk management framework: Generative artificial intelligence profile* (NIST AI 600-1). U.S. Department of Commerce. Available from: https://doi.org/10.6028/NIST.AI.600-1
3) Benerofe, S. (2025). *AI governance and the verification gap: A framework for law and policy under computational intractability*. *SSRN Electronic Journal*. Available from: https://papers.ssrn.com/abstract=5629290
4) Shneiderman, B. (2020). Bridging the gap between ethics and practice: Guidelines for reliable, safe, and trustworthy human-centered AI systems. *ACM Transactions on Interactive Intelligent Systems*, 10(4), Article 26, 1–31. Available from: https://doi.org/10.1145/3419764
5) Joshi, S. (2025). *Compensating for the risks and weaknesses of AI/ML models in finance*. *Preprints*. Available from: https://www.preprints.org/manuscript/202503.2245
6) Bridging NIST's AI risk management framework with Microsoft Purview data security. (2024). *Severian Blog*. Available from: https://severian.ghost.io/bridging-nists-ai-risk-management-framework-with-microsoft-purview-data-security/
7) Integrating NIST AI RMF into modern risk frameworks. (2024). *GovWhitePapers*. Available from: https://govwhitepapers.com/whitepapers/integrating-nist-ai-rmf-into-modern-risk-frameworks/
8) Themann, S. (2025). *Challenges and strategies used in implementing AI governance: A systematic literature review* (Master's thesis). Stockholm University. https://su.diva-portal.org/smash/get/diva2:1983756/FULLTEXT01.pdf
9) Masood, A. M. (2024). *The AI governance frontier series: Part 6—Open-source tools and initiatives for responsible AI*. *Medium*. Available from: https://medium.com/@adnanmasood/the-ai-governance-frontier-series-part-6-open-source-tools-and-initiatives-for-responsible-and-d41ece940ac1
10) Joshi, S. (2025). *An AI-agentic framework for modernizing pregnancy surveillance systems* (CDC-2025-0750-0404). Available from: https://downloads.regulations.gov/CDC-2025-0750-0404/attachment_1.pdf
11) Joshi, S. (2025). *Regulatory reform for agentic AI: Addressing governance challenges in federal AI adoption* (OSTP-TECH-2025-0067-0401). Available from: https://downloads.regulations.gov/OSTP-TECH-2025-0067-0401/attachment_1.pdf
12) Joshi, S. (2025). *Regulatory frameworks for generative AI-enabled digital mental health devices: Safety, transparency, and post-market oversight* (FDA-2025-N-2338-0070). Available from: https://downloads.regulations.gov/FDA-2025-N-2338-0070/attachment_1.pdf
13) Joshi, S. (2025). *Transformative integration of agentic generative AI in food safety systems: Policy framework, implementation guidelines, and economic impact analysis* (FSIS-2025-0145-0007). Available from: https://downloads.regulations.gov/FSIS-2025-0145-0007/attachment_1.pdf
14) Joshi, S. (2025). *A comprehensive framework for U.S. AI export leadership* (ITA-2025-0070-0042). Available from: https://downloads.regulations.gov/ITA-2025-0070-0042/attachment_1.pdf