



Agentic AI for Risk Assessment Controllers in BFSI: A Technical Framework for Autonomous Risk Mitigation

Gopichand Agnihotram¹
Joydeep Sarkar²
Magesh Kasthuri³✉
Venkatesh Balasubramaniam⁴

^{1,2,3,4}Wipro Limited, Bengaluru, India.
(✉ Corresponding Author)

Abstract

The evolution of agentic AI has fundamentally redefined the Banking, Financial Services, and Insurance (BFSI) sector's approach to risk management. For decades, financial institutions have relied on deterministic risk assessment controllers governed by fixed models, static thresholds, and linear workflows. While effective in stable conditions, these legacy systems struggle to handle the dynamic, interconnected, and data-intensive nature of today's financial ecosystems [1]. Modern BFSI operations generate massive streams of multimodal data, including structured financial metrics, unstructured text, behavioural signals, and real time market data that exceed the processing capacity of traditional risk engines. As a result, many risk controllers remain reactive, discovering threats only after exposure or regulatory breach. The emergence of Agentic AI systems addresses these limitations by introducing autonomy, adaptivity, and explainability into the risk control process. Unlike static models, these systems employ specialized AI agents that collaborate across domains—credit, liquidity, compliance, cybersecurity, and actuarial—using shared context and feedback loops. Each agent continuously perceives, reasons, and acts within its environment to maintain optimal control states. At the core of this evolution lies Reinforcement Learning (RL) and multi-agent orchestration, enabling continuous decision optimization under uncertainty [2]. RL agents learn from environmental feedback, dynamically adjusting thresholds and capital allocations in response to market, operational, or regulatory changes. This paper presents a technical framework detailing how agent based architectures, reinforced by machine reasoning and control theory, can autonomously mitigate risk across BFSI domains. It explores how these systems improve early warning capabilities, enhance model governance, and ensure regulatory compliance all while maintaining explainability and auditability in high stakes environments. In doing so, Agentic AI establishes the foundation for self adaptive risk ecosystems, capable of operating with human oversight yet independent in execution transforming risk management from a reactive function into a predictive and preventive intelligence layer for the modern financial enterprise.

Keywords: Agentic AI, Explainable AI (XAI), Financial services, and insurance), Multi-agent systems (MAS), Reinforcement learning (RL), BFSI (banking, Risk management.

1. Introduction

Risk assessment controllers in the Banking, Financial Services, and Insurance (BFSI) sector form the backbone of modern financial governance. They define how institutions evaluate, monitor, and mitigate exposure to credit losses, market volatility, liquidity shortfalls, operational breakdowns, cyber intrusions, and regulatory non-compliance. These controllers serve as the critical decision fabric that connects business processes, risk models, and regulatory compliance functions into an integrated ecosystem of control.

Historically, BFSI institutions have depended on rule based and statistical models to quantify and manage risk. These systems rely on pre-defined thresholds, deterministic algorithms, and manually maintained control hierarchies. While such frameworks perform well in predictable environments, they often fail to adapt to emerging risk patterns driven by rapidly evolving data sources, interconnected markets, and non-linear dependencies across financial products [3]. The consequences include lagging risk recognition, fragmented decision-making, and operational inefficiency when responding to fast-moving crises or regulatory changes.

Moreover, the modern BFSI environment operates in an era of data explosion, timely regulatory submissions, operational resiliency and cognitive complexity. Every transaction, customer interaction, and market fluctuation generates high-frequency data across multiple channels structured, semi-structured, and unstructured. Traditional

risk systems were not designed to continuously assimilate and reason over such diverse data modalities. As decision cycles shorten from quarterly reviews to near real-time oversight, static risk models lack the contextual intelligence required to monitor the suspicious customers anomalous behaviors and maintain resilience, regulatory compliance dynamically.

This operational gap has accelerated the adoption of Agentic AI systems, a new class of architectures that combine autonomous agents, machine reasoning, and reinforcement learning (RL) to achieve continuous risk awareness and adaptive risk control assessment. These systems emulate human decision-making at scale, embedding intelligence directly into workflows rather than relying solely on centralized supervision. An agent within this framework is an autonomous entity that perceives its environment, reasons about potential actions, and executes control strategies aligned with institutional risk policies.

Through multi agent orchestration, individual risk agents responsible for domains such as credit, liquidity, fraud compliance, cybersecurity, and actuarial analysis collaborate under an orchestrator layer that harmonizes their decisions. This architecture supports dynamic coordination, enabling the institution to react not only to known risks but also to emergent threats detected through behavioural or causal signals. Reinforcement learning techniques allow agents to refine their decision policies based on feedback, optimizing outcomes like loss minimization, capital efficiency, or regulatory compliance.

Furthermore, agentic architectures enable explainability and auditability through continuous context capture. Unlike black-box AI models, these agents operate under defined objectives and maintain traceable reasoning paths, supporting the stringent documentation and validation requirements of regulatory frameworks [4].

In essence, Agentic AI transforms risk assessment controllers from static evaluators into adaptive control systems capable of learning, predicting, and responding in real-time. This transition represents a paradigm shift from reactive monitoring to autonomous prevention and resilience engineering where financial institutions can anticipate disruption before it materializes. As the BFSI industry moves toward intelligent automation and real-time governance, Agentic AI will serve as the foundational architecture that bridges AI driven risk intelligence with human oversight and regulatory assurance, ensuring both performance and compliance in a volatile digital landscape.

2. Agentic AI Architecture for Risk Mitigation

The Agentic AI architecture represents a fundamental shift from rule-based risk management to intelligent, adaptive, and self-regulating systems. In the context of BFSI, this architecture establishes a distributed network of autonomous agents capable of monitoring real-time conditions, reasoning across multidimensional data, and executing corrective or preventive actions. Rather than operating as static risk calculators, these agents function as decision-making entities that continuously learn from feedback and optimize their control strategies.

At its core, the Agentic AI system functions as a multi agent ecosystem, where each agent be it for credit, liquidity, compliance, or cybersecurity perceives its environment, takes actions based on defined objectives, and shares insights through an orchestrator. The architecture ensures that these agents cooperate through shared state awareness, enabling cross-domain reasoning and holistic risk assessment.

2.1. Architectural Overview

The Agentic AI system for BFSI risk control assessment mitigation is designed as a layered architecture comprising five interdependent components:

2.1.1. Data Ingestion and Contextualization Layer

This foundational layer is responsible for aggregating and contextualizing data from diverse financial and operational sources. BFSI systems typically generate terabytes of structured and unstructured data daily ranging from transactional logs, market feeds, and credit exposures to documents, customer communications, and external signals such as economic indicators or news sentiment.

Key functions include:

- Data fusion and normalization: Unifies data from core banking systems, ERPs, trading engines, CRM tools, and regulatory repositories [5].
- Multimodal integration: Ingests both textual and numerical data, as well as behavioural (clickstream, biometrics) and image-based information for insurance and KYC processes.
- Feature streaming: Employs event-driven architectures (e.g., Kafka or Flink) to stream risk indicators and signals in near real-time.

This layer serves as the “sensory interface” for all downstream agents, ensuring that perception is consistent, reliable, and temporally aligned across the enterprise.




Figure 1. Business Process diagram of Agentic AI Architecture for Risk Mitigation.

2.1.2. Risk Orchestrator Agent

The Risk Orchestrator Agent serves as the central intelligence hub coordinating multiple domain specific agents. It maintains a unified view of enterprise-wide risk and ensures consistency in decisions made by individual agents. Technically, the orchestrator is implemented as a message-driven coordination layer, often leveraging APIs, reinforcement signals, or vector-based embeddings for semantic context sharing among agents [6].

Key responsibilities include:

- Cross-agent communication: Mediates data exchange between domain agents using a shared ontology or graph-based knowledge model.
- Regulatory policy harmonization: Aligns agent behaviours with the institution’s overall risk appetite, capital constraints, and compliance mandates.
- Conflict resolution: When multiple agents recommend conflicting actions (e.g., liquidity expansion vs. credit tightening), the orchestrator applies optimization logic or hierarchical weighting to resolve trade-offs.
- Scenario synthesis: Integrates simulation outputs from individual agents to produce consolidated stress-test results and forecasted risk exposure curves.

Technically, the orchestrator is implemented as a message-driven coordination layer, often leveraging APIs, reinforcement signals, or vector-based embeddings for semantic context sharing among agents.

2.1.3. Domain Agents

Domain Agents represent the operational backbone of the Agentic AI system. Each agent autonomously governs a specific risk dimension and is equipped with reasoning and control capabilities aligned to that domain’s unique challenges.

Examples include:

- Credit Risk Agent: Continuously monitors borrower behaviour, market exposure, and portfolio concentrations to predict probability of default (PD) and recommend exposure adjustments.
- Liquidity Agent: Simulates cash flow positions, stress scenarios, and funding sources to maintain capital adequacy and avoid shortfalls.
- Compliance Agent: Uses natural language models and retrieval-augmented generation (RAG) pipelines to interpret new regulatory publications and map them to internal control requirements.
- Cyber Risk Agent: Monitors system telemetry, identity graphs, and anomaly signals to detect and contain security threats autonomously.
- Actuarial Agent: Optimizes pricing and reserve strategies in insurance portfolios using real-time actuarial data and predictive modeling.

Each domain agent possesses:

- Perception (environmental data ingestion and contextualization),
- Reasoning (analytical modeling, causal inference, and simulation), and
- Actioning (execution of control triggers via connected systems or API endpoints).

Together, they enable distributed risk intelligence where each agent contributes to a shared enterprise understanding while maintaining autonomy in its specialized function.

2.1.3.1. Reinforcement Learning (RL) and Optimization Layer

The Reinforcement Learning layer provides adaptive intelligence to the agent network. Traditional supervised models in risk management rely on static datasets and predefined labels, whereas RL enables continuous policy refinement based on reward and penalty outcomes.

In this architecture, agents interact with the environment (e.g., financial markets, transaction systems, operational data streams) to learn optimal actions that balance profitability and compliance risk. The RL framework operates as follows:

- State space: Represents the current environment, including credit exposure, market volatility, or liquidity ratio.

- **Actions:** Correspond to possible interventions—tightening exposure, adjusting hedges, reallocating liquidity, or raising alerts.
- **Rewards:** Quantify desired outcomes such as loss reduction, compliance adherence, or capital efficiency.
- **Feedback:** The system continuously learns from post-action metrics (defaults prevented, breaches avoided) and adjusts its policy network accordingly.

A Multi-Agent RL (MARL) framework allows coordination between agents such as the Credit and Liquidity Agents—enabling collaborative optimization. For instance, one agent might reduce credit exposure while another ensures liquidity sufficiency, balancing the overall risk-reward profile.

2.2. Governance and Explainability Layer

There is no AI system in BFSI can operate without governance, transparency, and auditability. The Governance Layer ensures that agent decisions remain interpretable, compliant, and aligned with human oversight structures.

Its core functions include:

- **Explainability:** Integrates model-agnostic interpreters (e.g., SHAP, LIME) to provide justifications for each agent’s actions.
- **Traceability:** Maintains full decision lineage, linking data inputs to model outputs, reinforcing compliance with regulations like SR 11-7, Basel Model Risk Guidelines [7].
- **Ethical and regulatory oversight:** Embeds business rules that constrain agent autonomy in sensitive decisions, enforcing human-in-the-loop (HITL) review for material actions.
- **Audit interface:** Generates self-documenting reports summarizing risk posture, model decisions, and remediation actions.

This layer transforms Agentic AI from a black-box decision system into a transparent control framework that regulators and auditors can evaluate with confidence.

2.3. Communication and Control Dynamics

Agentic AI systems rely on bidirectional communication between domain agents and the orchestrator. Each agent publishes contextual signals (risk scores, alerts, policy recommendations) while subscribing to global policies and system-wide constraints. The orchestrator interprets and consolidates these signals to update the overall enterprise risk state, feeding it back into the agents for iterative refinement.

Internally, this creates a continuous control feedback loop:

1. Perception (data ingestion) →
2. Reasoning (agent analysis and prediction) →
3. Decision (orchestrator synthesis) →
4. Action (execution of control measures) →
5. Feedback (reward evaluation through RL).

Through this cyclical process, the system evolves toward real-time resilience, capable of detecting deviations, learning from outcomes, and autonomously adjusting its behaviour all within a governance-controlled environment.

Hence the Agentic AI architecture for risk mitigation provides BFSI organizations with a self-learning, explainable, and scalable framework that moves beyond siloed analytics toward continuous, intelligent risk orchestration. By blending autonomous reasoning, multi-agent collaboration, and reinforcement learning, this architecture transforms risk management from a static compliance function into an adaptive control ecosystem that safeguards financial institutions in a constantly evolving digital landscape.




Figure 2. Multi-Agent Orchestration for Risk Control – illustrating orchestration layer, RL feedback loop, and governance interface.

3. How Agents Mitigate BFSI Risks

The Agentic AI framework, each agent operates as a specialized, autonomous component dedicated to monitoring and managing a specific risk domain. These agents function collaboratively through a shared state representation, enabling holistic enterprise-level awareness while preserving domain-level autonomy.

Through continuous perception, reasoning, and reinforcement learning (RL), they dynamically predict, detect, and mitigate risk across the institution's operational and financial landscape.

Agents exchange contextual signals such as credit exposure trends, liquidity stress indicators, compliance alerts, or cyber intrusion graphs through the Risk Orchestrator Layer, ensuring unified situational awareness and coordinated response. This structure transforms risk management into a distributed control network that learns, adapts, and acts in near real-time.

3.1. Credit Risk – Credit Risk Agent

The Credit Risk Agent governs exposure across retail, commercial, and institutional portfolios. Traditional credit assessment models rely on static scoring systems and historical data, limiting their ability to detect emerging borrower risks.

By contrast, the Credit Risk Agent employs reinforcement learning and causal reasoning to adapt to dynamic borrower behaviour and macroeconomic shifts.

Its capabilities include:

- **Dynamic credit scoring:** Continuously recalibrates borrower creditworthiness using live repayment histories, behavioural analytics, and transaction-level data [8].
- **Exposure optimization:** Learns optimal exposure levels by balancing potential profit with default probability under changing market conditions.
- **Feedback learning:** Reinforces decision policies based on actual repayment outcomes, adjusting lending thresholds and interest margins accordingly.

For example, if an early-stage credit delinquency pattern is detected, the agent automatically recommends exposure reduction or restructuring actions—preventing defaults while maintaining profitability.

3.2. Market Risk – Market Risk Agent

The Market Risk Agent manages exposure to price volatility, interest rate changes, currency movements, and derivative sensitivities.

This agent operates in an environment where rapid decision-making is crucial to preserving portfolio value. It leverages multi-scenario simulations and stochastic modeling to evaluate portfolio responses under diverse stress conditions.

Core functions:

- **Sensitivity analysis:** Computes real-time VaR (Value at Risk) and P&L attribution to understand the impact of market shifts.
- **Hedging optimization:** Uses reinforcement learning to identify hedging strategies that minimize loss without overcapitalizing risk buffers.
- **Scenario synthesis:** Continuously simulates hundreds of market paths, using probabilistic forecasting to recommend rebalancing or delta-neutral positions.

The Market Risk Agent interacts with the Credit and Liquidity Agents, ensuring that **capital and exposure** adjustments remain synchronized across the institution's financial ecosystem.

3.3. Liquidity Risk – Liquidity Agent

Liquidity management requires balancing funding adequacy, reserve ratios, and operational flexibility. The Liquidity Agent applies predictive analytics and RL-driven simulation to forecast cash flow gaps and preemptively reallocate capital.

Capabilities include:

- **Funding gap prediction:** Monitors inflows and outflows to predict short-term liquidity pressures based on historical patterns and event-driven triggers.
- **Capital reallocation:** Initiates dynamic rebalancing of cash reserves, adjusting collateral and interbank borrowing levels as required.
- **Stress resilience modeling:** Evaluates liquidity positions under systemic shocks or counterparty failures, optimizing buffers in real time [9].

By operating in tandem with the Credit and Market Agents, the Liquidity Agent ensures enterprise-wide solvency and funding continuity, even under extreme market volatility.

3.4. Operational Risk – Process Control Agent

Operational disruptions often arise from process errors, internal fraud, or automation failures. The Process Control Agent acts as a real-time sentinel for operational integrity. It integrates telemetry data, audit logs, and user behaviour analytics to detect process anomalies and policy breaches.

Functions include:

- **Fraud detection:** Employs graph neural networks and behavioural clustering to identify suspicious transaction sequences.
- **Process monitoring:** Uses machine learning to establish process baselines and detect deviations from standard operating patterns.
- **Automation assurance:** Validates that RPA and AI-based automations execute without unintended exceptions or compliance violations.

The agent can autonomously trigger alerts, rollbacks, or workflow quarantines minimizing cascading failures in complex operational pipelines.

3.5. Cyber Risk – Cyber Defense Agent

In a digital-first BFSI ecosystem, cyber resilience is non-negotiable. The Cyber Defense Agent combines graph reasoning, intrusion detection, and reinforcement response mechanisms to autonomously protect enterprise networks.

Key functions include:

- Threat graph analysis: Builds entity-relationship maps linking user behaviours, access patterns, and device telemetry to detect anomalies.
- Attack surface minimization: Automatically isolates or quarantines compromised nodes upon detecting lateral movement or privilege escalation.
- Adaptive defense: Reinforcement learning enables the agent to update defense strategies based on simulated attack scenarios and observed threat outcomes.

By continuously learning from attempted intrusions, the Cyber Defense Agent evolves into a self-improving digital immune system for BFSI infrastructure [10].

3.6. Compliance Risk – Regulatory Copilot Agent

The Regulatory Copilot Agent addresses one of the most complex challenges in BFSI: aligning operational behaviour with evolving regulatory expectations.

This agent employs large language models (LLMs) integrated with Retrieval-Augmented Generation (RAG) pipelines to parse new regulations, identify obligations, and map them to internal policies [11].

Capabilities include:

- Regulation-to-policy mapping: Uses natural language understanding to correlate regulatory clauses with control processes.
- Gap analysis: Detects control deficiencies and proposes remediation steps.
- Continuous compliance assurance: Monitors regulatory updates in real time and recommends policy or control adjustments proactively.

Through explainable AI mechanisms, it produces audit-ready summaries and compliance dashboards, reducing manual compliance overhead and ensuring ongoing regulatory alignment.

3.7. Model Risk – Model Governance Agent

The Model Governance Agent ensures that AI and quantitative models within the BFSI ecosystem operate under strict control and transparency.

It continuously monitors model drift, bias, and explainability metrics across production systems.

Core functions:

- Drift detection: Flags performance degradation by comparing real-world outputs against training baselines.
- Bias mitigation: Applies fairness diagnostics and causal inference techniques to ensure ethical and regulatory compliance.
- Explainability: Generates interpretable summaries for regulators using SHAP or surrogate model explanations [15].

By enforcing continuous validation and documentation, the Model Governance Agent reduces model risk exposure while supporting compliance with standards like SR 11-7 and the EU AI Act.

3.8. Insurance Risk – Underwriting & Claims Agent

Within insurance operations, the Underwriting & Claims Agent optimizes both policy pricing and claims management. It integrates predictive modeling, behavioural analysis, and fraud detection to enhance profitability and reduce losses.

Key functions:

- Predictive underwriting: Evaluates risk profiles dynamically using customer behaviour, claim history, and external data.
- Claims fraud detection: Leverages anomaly detection and network analysis to identify fraudulent claims patterns.
- Portfolio optimization: Continuously refines pricing models based on market conditions and policyholder feedback.

Through continuous learning, this agent enables insurers to balance risk exposure with competitive pricing, improving both profitability and fairness.

3.8.1. Collaborative Risk Intelligence

While each agent excels within its domain, the real strength of the Agentic AI framework lies in collaboration and shared intelligence. Through the orchestrator layer, agents exchange signals and align decisions to maintain enterprise-wide risk coherence. For example:

- When the Market Risk Agent signals increased volatility, the Credit and Liquidity Agents proactively adjust lending thresholds and reserve ratios.
- The Compliance Agent ensures all responses conform to regulatory frameworks, while the **Governance** Layer provides auditability and explainability.

This multi-agent collaboration enables predictive, adaptive, and explainable risk management, replacing static oversight with autonomous, orchestrated resilience across the BFSI enterprise.

3.8.2. Reinforcement Learning for Continuous Risk Optimization

Reinforcement Learning (RL) introduces a paradigm shift in how financial institutions manage dynamic risk-reward relationships. Unlike static supervised models, RL agents learn optimal policies by continuously interacting

with their environment, observing states, taking actions, and receiving feedback through reward or penalty mechanisms. In the BFSI context, these agents function within a complex decision environment where actions such as adjusting credit exposure, reallocating liquidity, or modifying portfolio hedges directly influence both profitability and institutional stability.

The learning process is guided by reward functions explicitly tied to business and regulatory objectives. Typical reward variables include loss probability reduction, regulatory compliance adherence, capital adequacy ratio improvement, and liquidity buffer optimization. By quantifying these outcomes, RL enables agents to autonomously optimize trade-offs between risk exposure and return efficiency under uncertainty [12].

For example, a Credit Risk Agent may receive positive rewards for lowering default rates while maintaining acceptable yield, whereas a Liquidity Agent gains rewards for ensuring funding resilience without overcapitalizing idle reserves. Similarly, a Compliance Agent is rewarded for reducing regulatory breach probabilities while maintaining operational throughput.

The RL loop comprising state representation, policy selection, action execution, and reward evaluation operates continuously, allowing agents to adapt to new data distributions, market shocks, and policy changes in real time. Over successive episodes, the agent converges toward near-optimal decision policies that balance short-term profitability with long-term resilience.

When extended into a Multi-Agent Reinforcement Learning (MARL) framework, collaboration among agents becomes a critical factor. For instance, the Credit and Liquidity Agents may share environmental signals such as capital availability or macroeconomic indicators to align their decision strategies. The orchestrator agent acts as a central policy coordinator, ensuring that collective actions minimize systemic risk while maximizing enterprise stability.

Through MARL, BFSI systems evolve into self-optimizing control networks, where agents cooperate rather than compete for resource utilization. This results in enhanced enterprise-level equilibrium, reduced volatility, and faster adaptation to both market and regulatory shifts. In essence, reinforcement learning transforms traditional risk management from rule-based supervision into autonomous, feedback-driven optimization, enabling continuous improvement in decision accuracy, compliance assurance, and financial performance.

3.8.3. Technical Use Cases and Live Scenarios

3.8.3.1. Adaptive Credit Risk Control

Reinforcement Learning (RL) agents continuously analyse borrower behaviours, credit utilization, and macroeconomic indicators such as GDP growth, inflation, and interest rate trends. Based on this evolving data, the agent dynamically adjusts credit limits, lending thresholds, and pricing to maintain optimal exposure levels. Feedback from repayment performance and delinquency trends serves as reinforcement signals, enabling the model to refine decision policies over time. This adaptive mechanism ensures proactive risk control while sustaining profitability, reducing non-performing assets, and aligning with Basel credit risk frameworks.

3.8.3.2. Liquidity Simulation Engine

The Liquidity Simulation Engine leverages agent based modeling to forecast short term and intraday funding positions. By ingesting transaction data, settlement flows, and market indicators, the agent performs multi scenario stress simulations to predict liquidity gaps before they occur. Using RL optimization, it dynamically reallocates capital and collateral to maintain reserve ratios and regulatory liquidity coverage. The system operates autonomously during volatility spikes, ensuring continuous solvency and efficient balance sheet utilization. It directly integrates with treasury systems and orchestrator agents for real-time capital coordination.

3.8.3.3. Autonomous AML Compliance

The compliance Agent employs a hybrid of Large Language Models (LLMs) and graph reasoning to detect suspicious activity patterns that traditional rule-based AML systems miss. It merges KYC data, transaction graphs, and sanction watchlists to uncover hidden relationships between entities. Reinforcement learning helps fine tune alert thresholds by minimizing false positives while improving detection accuracy. When anomalies are found, the agent generates transparent, audit-ready justifications for regulators. This continuous learning loop enables proactive compliance aligned with FATF, FINRA, and EU AMLD standards.

3.8.3.4. Insurance Claims Fraud Detection

The Insurance Claims Agent integrates multimodal AI combining computer vision, NLP, and transactional analysis to assess claim authenticity. It scans claim documents, repair invoices, and uploaded images to identify inconsistencies or synthetic fraud patterns. Reinforcement learning feedback helps the agent refine detection logic as confirmed fraud or legitimate cases are reviewed. This results in more precise claim triaging, reduced manual review effort, and faster resolution cycles. The system integrates seamlessly with underwriting and policy management agents, supporting dynamic risk-based pricing in real time.

3.8.3.5. Market Hedging Optimization

Market Risk Agents use multi-agent reinforcement learning to coordinate hedging strategies across trading desks and asset classes. They continuously evaluate real-time portfolio sensitivity to price, rate, and currency movements, running stress tests under alternative volatility regimes. Using live market feeds, the agent recommends optimal hedge adjustments balancing transaction costs, counterparty exposure, and return-on-capital constraints. Over time, policy updates improve hedge efficiency, ensuring consistent market neutrality. This results in enhanced portfolio resilience and compliance with MiFID II and internal risk governance thresholds.

3.9. Human-in-the-Loop Governance

While Agentic AI systems introduce autonomy and self-optimization, human oversight remains an indispensable layer of governance to ensure ethical, regulatory, and strategic alignment. Human-in-the-loop (HITL) mechanisms allow risk officers, compliance analysts, and model validators to supervise, approve, or override agent actions before execution. These interfaces act as control checkpoints, especially for high-impact decisions such as large credit disbursements, fraud alerts, or liquidity reallocations.

Expert feedback is continuously integrated into the learning process through reward shaping, ensuring that reinforcement signals reflect institutional risk appetite, compliance thresholds, and board-level directives. This feedback loop anchors agent behaviours within established frameworks such as Basel III, SR 11-7 (Supervisory Guidance on Model Risk Management), FSB stability implications of AI, and emerging AI governance standards like the EU AI Act and NIST AI Risk Management Framework [13][14].

Moreover, the governance layer provides traceability and explainability, recording every decision rationale and model justification for audit readiness. By blending machine autonomy with human judgment, BFSI organizations achieve controlled intelligence where agents operate adaptively yet remain fully accountable to regulatory and ethical boundaries.

4. Conclusion and Future Outlook

Agentic AI systems are redefining how BFSI institutions assess, manage, and mitigate risks across diverse operational domains. By integrating autonomous agents within a Model Context Protocol (MCP) framework, organizations enable seamless coordination between reasoning models, control systems, and domain specific risk engines. The MCP layer ensures that agents share contextual understanding in real time, allowing risk orchestration to occur dynamically and consistently across credit, financial crime compliance, operational and market environments.

Within this architecture, Reinforcement Learning (RL), particularly using Proximal Policy Optimization (PPO), serves as the optimization backbone for continuous policy refinement [16]. PPO stabilizes learning under uncertainty and supports controlled adaptation of decision policies, ensuring that agents evolve within defined regulatory and ethical limits. This combination of MCP driven context exchange and PPO-based learning creates a self-optimizing ecosystem capable of maintaining enterprise wide risk equilibrium.

Looking ahead, the BFSI sector will move toward federated multi-agent networks that securely share risk intelligence without compromising data privacy. Future innovation will focus on explainable RL governance, causal inference integration, and cross-institutional orchestration, enabling transparent and collaborative responsible risk intelligence. Through such convergence of context-aware orchestration and reinforcement learning, Agentic AI will form the foundation of next-generation autonomous risk management infrastructure, adaptive, auditable, and resilient by design.

References

- Angélod, A., & Ernest, P. (2019). Agent-based modelling for bank's liquidity risk management. *Journal of Economic Interaction and Coordination*, 14(3), 565–591. <https://doi.org/10.1007/s11403-018-0236-8>
- Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012>
- Aziz, S., & Dowling, M. (2019). Machine learning and AI for risk management. In T. Lynn, J. Mooney, P. Rosati, & M. Cummins (Eds.), *Disrupting finance: FinTech and strategy in the 21st century* (pp. 33–50). Palgrave Macmillan. https://doi.org/10.1007/978-3-030-02330-0_3
- Board of Governors of the Federal Reserve System. (2011). *Supervisory guidance on model risk management (SR 11-7)*. Washington, DC:
- Busoniu, L., Babuska, R., & De Schutter, B. (2010). Multi-agent reinforcement learning: An overview. In A. Esposito, A.-M. Esposito, A. Vinciarelli, R. Hoffmann, & V. Müller (Eds.), *Innovations in multi-agent systems and applications-1* (pp. 183–221). Springer. https://doi.org/10.1007/978-3-642-14435-6_8
- European Commission. (2021). *Proposal for a regulation laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) (COM(2021) 206 final)*. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52021PC0206>
- Khaleghi, B., Khamis, A., Karray, F. O., & Razavi, S. N. (2013). Multisensor data fusion: A review of the state of the art. *Information Fusion*, 14(1), 28–44. <https://doi.org/10.1016/j.inffus.2011.08.001>
- Khandani, A. E., Kim, A. J., & Lo, A. W. (2010). Consumer credit-risk models via machine-learning algorithms. *Journal of Banking & Finance*, 34(11), 2767–2787. <https://doi.org/10.1016/j.jbankfin.2010.06.001>
- Kolm, P. N., & Ritter, G. (2019). Modern perspectives on reinforcement learning in finance. *The Journal of Financial Data Science*, 1(1), 26–54. <https://doi.org/10.3905/jfds.2019.1.026>
- Leo, M., Sharma, S., & Maddulety, K. (2019). Machine learning in banking risk management: A literature review. *Risks*, 7(1), 29. <https://doi.org/10.3390/risks7010029>
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-T., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33, 9459–9474. https://proceedings.neurips.cc/paper_files/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html
- Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30. <https://proceedings.neurips.cc/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html>
- National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. U.S. Department of Commerce. <https://doi.org/10.6028/NIST.AI.100-1>
- Nguyen, T. T., & Reddi, V. J. (2021). Deep reinforcement learning for cyber security. *IEEE Transactions on Neural Networks and Learning Systems*, 33(12), 1–15. <https://doi.org/10.1109/TNNLS.2021.3121870>
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). *Proximal policy optimization algorithms*. arXiv preprint arXiv:1707.06347. <https://doi.org/10.48550/arXiv.1707.06347>
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). MIT Press.