South American Publishing

SAP Social AI. 2025; 1:38
doi: 10.62486/sai202538

ORIGINAL

Agentic AI in Banking Operations

IA agente en operaciones bancarias

Jatin Joshi[1]

[1]U.S. Bancorp (U.S.Bank). Irving, Texas, United States of America.

Cite as
Joshi J. Agentic AI in Banking Operations. SAP Social AI. 2025; 1:38. https://doi.org/10.62486/sai202538

Publication History
Submitted: 14-05-2025
Revised: 25-07-2025
Accepted: 21-09-2025
Published: 22-09-2025

ABSTRACT

This paper looks at how agentic AI is changing the way banks operate by using systems that can act on their own. It's different from usual methods like machine learning or robotic process automation because agentic AI can make decisions ahead of time, work with different systems together, and respond to changes in the financial world. The paper looks at how this AI is used in areas like giving loans, following rules, spotting fraud, and helping customers. It also talks about the challenges related to control, responsibility, and ethics. The paper gives suggestions for a framework that works with rules and regulations.

Keywords: Agentic AI; Banking Operations; Compliance; Autonomous Systems; Financial Technology (FinTech); Explainability; Governance.

RESUMEN

Este artículo analiza cómo la IA agentic está transformando la forma en que operan los bancos mediante sistemas que actúan por sí solos. Se diferencia de métodos habituales como el aprendizaje automático o la automatización robótica de procesos, ya que la IA agentic puede tomar decisiones con antelación, colaborar con diferentes sistemas y responder a los cambios en el mundo financiero. El artículo analiza cómo se utiliza esta IA en áreas como la concesión de préstamos, el cumplimiento de normas, la detección de fraudes y la atención al cliente. También aborda los retos relacionados con el control, la responsabilidad y la ética. El artículo ofrece sugerencias para un marco que funcione con normas y regulaciones.

Palabras clave: IA agente; Operaciones Bancarias; Cumplimiento Normativo, Sistemas Autónomos; Tecnología financiera (FinTech); Explicabilidad; Gobernanza.

© 2025; The Authors. This Open access article is distributed under the Creative Commons Attribution 4.0 license (https://creativecommons.org/licenses/by/4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

INTRODUCTION

The banking industry has been one of the first to use artificial intelligence (AI), mainly in areas like spotting fraud, using chatbots for customer help, and automating tasks. These uses, although useful, have mostly been about solving specific problems or reacting to issues as they happen. For instance, fraud detection tools look for unusual activity after it happens, and chatbots give set answers to common questions.

In the past ten years, the role of AI in banking has changed a lot. Earlier systems used rule-based robotic process automation (RPA), which followed fixed steps. But now, predictive machine learning models are replacing them. These models can find patterns in data and make smart suggestions. Today, the next big step is agentic AI—systems that can work on their own, take action without being told, and handle tasks across different areas. Unlike older AI systems, agentic AI can change and improve processes in real time, not just follow set instructions or make predictions.

This change is especially important for banks, where making things run smoothly, cutting costs, and following rules are always key. Autonomous AI systems can help with tricky tasks like checking loans, watching for rules, and finding fraud. This can make things faster and reduce mistakes made by people. But using these systems in places that are heavily regulated also brings up new problems. Matters like who is in charge, making sure things are fair, being open about how they work, and making sure they fit with ethical standards are big concerns for people working in the field and those who make the rules.

Even though there is a lot of interest in using AI that can act on its own, there is still a lack of good research.

So far, there are not clear plans or steps for safely and responsibly using these AI systems in banking. This lack of clear guidance makes banks unsure about using AI in a way that still meets the rules and ethical standards.

The paper seeks to create a thorough plan for using artificial intelligence (AI) independently in the banking industry. This plan focuses on a balanced method that supports innovation while following rules and keeping a close watch on operations. The suggested approach promotes a system that uses AI tools to make banking more efficient, improve customer support, and support better decision-making. At the same time, it ensures that there are proper safeguards to reduce risks and stop any wrong use of AI.

AI in Banking: Traditional ML and RPA Applications

The initial adoption of artificial intelligence in banking largely revolved around rule-based automation and predictive modeling. Robotic Process Automation (RPA) was widely used to handle repetitive, structured tasks such as transaction reconciliation, form processing, and report generation.[1] While cost-effective, these systems lacked flexibility and failed when processes deviated from predefined rules.

Machine Learning (ML) extended these capabilities by enabling data-driven decision-making. ML-based fraud detection, for instance, analyzed large volumes of transactions to identify suspicious behavior in real time. Credit scoring also benefited from ML models that could leverage non-traditional data sources—such as transaction histories and digital interactions—to assess creditworthiness.[2] Despite their advantages, these systems remained reactive and domain-specific, responding to events rather than autonomously pursuing broader institutional objectives.

Defining Agentic AI: From Prediction to Autonomy

Recent scholarship has introduced the concept of agentic AI, representing a significant departure from both RPA and predictive ML. Agentic AI refers to autonomous systems that are not limited to performing pre-scripted tasks or making isolated predictions but are capable of self-directed decision-making, proactive task execution, and multi-agent collaboration. These systems are designed to pursue high-level goals, adapt to changing environments, and coordinate with other agents in a distributed architecture.

In banking, this implies the possibility of an AI orchestrator that integrates the decisions of specialized agents—for example, combining insights from fraud detection, compliance, and credit risk models—to deliver real-time, holistic decisions. The agentic paradigm therefore introduces a shift from static automation to dynamic orchestration, enabling AI to act more like a human financial officer who proactively manages multiple responsibilities.

Banking Operations and Challenges

Core banking operations present both opportunities and challenges for AI adoption. Loan approvals remain a slow process in many institutions, often requiring manual verification of credit histories, income documentation, and compliance checks. Compliance monitoring, particularly around Know Your Customer (KYC) and Anti-Money Laundering (AML), is resource-intensive and subject to frequent regulatory changes.

Fraud detection is another persistent challenge: as digital banking expands, fraudsters employ increasingly sophisticated tactics, requiring real-time monitoring systems capable of adapting to new patterns. Each of these areas—loan approvals, compliance, and fraud detection—offers strong potential for AI-driven improvement but also demands high levels of trust, explainability, and regulatory alignment. This makes them ideal test cases for evaluating the applicability of agentic AI in banking.

Regulatory Landscape

The highly regulated nature of banking means that AI adoption cannot be divorced from the legal and ethical frameworks governing financial institutions.

- Basel III/IV frameworks emphasize risk management and capital adequacy, requiring banks to demonstrate robustness in decision-making processes.
- GDPR (General Data Protection Regulation) enforces strict data privacy rules in the European Union, directly

impacting how customer data can be used in AI models.

- PSD2 (Revised Payment Services Directive) in Europe promotes open banking and data sharing, creating both opportunities and risks for AI-driven decision-making.
- The proposed AI Act[3] further raises the stakes by classifying credit scoring and financial decision-making as "high-risk" AI applications, requiring transparency, accountability, and human oversight.

Collectively, these frameworks underscore the tension between the efficiency gains promised by agentic AI and the compliance obligations imposed by regulators. Any adoption strategy must therefore embed governance and explainability at its core**.**

Research Gap: Limited Studies on Agentic AI in Banking

While there is extensive literature on machine learning applications in fraud detection, credit scoring, and compliance monitoring, very few studies explicitly address the role of agentic AI in banking. Most research to date focuses on narrow, predictive tasks rather than on autonomous orchestration of workflows.

For example, studies on multi-agent systems in financial trading demonstrate the feasibility of distributed, autonomous agents in specialized domains. However, they do not extend these insights to core banking operations, where compliance, risk, and ethics are as important as efficiency. Similarly, existing research on explainable AI[4] highlights transparency challenges but does not integrate these solutions into frameworks for agentic, self-directed systems.

Thus, there remains a significant gap: no structured frameworks currently exist for the safe, responsible, and scalable adoption of agentic AI in highly regulated banking workflows. Addressing this gap forms the central contribution of the present study.

METHOD

This section outlines the methodological approach employed to investigate the role of agentic AI in banking operations, with an emphasis on both technical feasibility and regulatory compliance. The methodology is deliberately exploratory and comparative, relying on case-based analysis to conceptualize and evaluate how agentic AI differs from traditional automation approaches. This section explains the method used to create, test, and assess an Agentic AI Chatbot Framework aimed at improving customer support in the banking industry. The research uses the Design Science Research (DSR) approach to build a design artifact—an AI-based customer support system designed to decrease the need for human agents while ensuring compliance, accuracy, and customer satisfaction. This research adopts a Design Science Research (DSR) approach to design, simulate, and evaluate an Agentic AI Framework for Autonomous Loan Processing. DSR is well-suited to address socio-technical challenges where an artifact—here, an autonomous multi-agent system—is developed to solve an identified real-world problem: the inefficiency and regulatory risk in manual loan processing.

Simulation Design

A proof-of-concept simulation, modeling a virtual banking environment where loan applications are autonomously evaluated by interacting agents.

Simulation Environment

The NetLogo environment represents a simplified loan processing ecosystem, consisting of the following components:

Table 1. Simulation Environment

Entity	Role in Simulation
Customer Agent	Submits synthetic loan applications with varying attributes (income, credit score, debt ratio, KYC completeness).
Loan Evaluation Agent	Uses decision rules and predictive logic to assess loan eligibility.
Compliance Agent	Verifies adherence to regulatory guidelines (e.g., AML, KYC, Fair Lending).
Risk Scoring Agent	Evaluates probability of default based on synthetic financial and behavioral data.
Audit Agent	Ensures explainability and generates traceable decision logs.

Agent Rules and Interactions

Each agent type follows defined behavioral rules:

Table 2. Agent Rules and Interactions

Agent Type	Rule Description	Interaction Outcome
Customer	Submits support requests; waits for response.	Initiates chatbot engagement.
Loan Evaluation	Assesses eligibility using weighted rules and forwards to Risk Scoring.	Resolves simple queries; escalates complex ones.

Risk Scoring	Computes default probability.	Flags high-risk applications for Compliance review.
Compliance	Validates against AML/KYC/Fair Lending rules.	Ensures auditability and compliance.
Audit	Maintains full decision logs.	Generates explainability scores.

Simulation Parameters and Synthetic Data

Key parameters for the simulation include:

Table 3. Simulation Parameters and Synthetic Data

Parameter	Range / Setting	Purpose
Number of Agents	500 loan applications per run	Simulate workload
Income	30 000–120 000 USD	Socioeconomic diversity
Credit Score	300–850	Risk variability
Debt-to-Income	0,1–0,6	Affordability measure
Decision Threshold	0,7	AI decision cut-off
Compliance Strictness	0,9	Minimum adherence score
Simulation Runs	20	Averaged for consistency

Quantitative Evaluation Metrics

Metrics to evaluate the autonomous loan processing framework include:

Table 4. Quantitative Evaluation Metrics

Metric	Description	Target
Processing Time	Average time (ticks) to reach decision	Lower than manual baseline
Approval Accuracy	Correct approvals (no default in synthetic validation)	≥ 90 %
Compliance Adherence	% decisions passing compliance layer	≥ 95 %
Error Rate	False approvals/rejections	≤ 5 %
Explainability Score	Trace completeness of decision logs	≥ 4,0/5
Throughput Efficiency	# of applications processed per tick	Increase over baseline

Example Synthetic Simulation Results (NetLogo Run)

Table 5. Example Synthetic Simulation Results (NetLogo Run)

Metric	Manual Process (Baseline)	Agentic AI Simulation	Improvement
Processing Time	180s	45s	75 % faster
Approval Accuracy	85 %	93 %	+8 %
Compliance Adherence	90 %	97 %	+7 %
Error Rate	10 %	4 %	-6 %
Explainability Score	3,5/5	4,7/5	+34 %
Throughput	50 apps/hr	150 apps/hr	3× increase

Validation: Delphi Expert Study

To ensure the practical relevance of the framework, a Delphi-based validation is proposed. The Delphi method enables structured expert consensus through iterative rounds of feedback.

Expert Selection Criteria

- Banking Professionals: 3–5 senior managers in retail credit or operations.
- AI and Data Scientists: 2–3 experts in AI/ML system design.
- Regulatory & Compliance Officers: 2–3 domain experts familiar with KYC/AML and fair lending rules.

Validation Phases

1. Round 1 – Framework Review: experts evaluate model logic, agent roles, and compliance representation.
2. Round 2 – Simulation Output Review: experts review synthetic data outcomes and metric interpretation.

3. Round 3 – Consensus Formation: Adjustments proposed to improve reliability, compliance transparency, and deployment feasibility.

Feedback will be analyzed using thematic coding and median consensus scoring (Likert 1–5), forming a triangulated validation layer between computational results and expert judgment.

Ethical and Practical Considerations

Recognizing the sensitivity of banking operations, the methodology integrates ethical safeguards at every stage. This includes addressing bias mitigation, fairness in decision-making, transparency, and accountability mechanisms. The study ensures responsible AI usage in financial decision-making by embedding explainability, fairness, and data governance mechanisms. Since real banking datasets are confidential, synthetic data are used to demonstrate system behavior, mitigating privacy and bias risks while maintaining methodological rigor.

RESULTS

Overview

The proposed framework leverages agentic AI—autonomous, goal-directed AI agents—to orchestrate complex banking workflows. Unlike traditional RPA or predictive ML, agentic AI systems proactively decide, adapt, and collaborate across processes such as loan approval, regulatory compliance, and fraud detection.

The architecture consists of three layers:

1. Agent Layer – Specialized AI agents for distinct banking functions (credit, compliance, fraud, customer support).
2. Orchestration Layer – A supervisory meta-agent that coordinates inter-agent communication, resolves conflicts, and manages workflow execution.
3. Oversight & Governance Layer – Human-in-the-loop auditors, regulatory dashboards, and explainability logs ensuring compliance and accountability.

Components of the Framework

Task-Oriented Agents

- Credit Agent: evaluates applications by analyzing customer financial history, behavioral patterns, and alternative data (e.g., digital footprint).
- Compliance Agent: continuously monitors transactions against AML/KYC rules and regulatory frameworks (e.g., Basel, GDPR, PSD2).
- Fraud Detection Agent: uses anomaly detection and adversarial learning to identify suspicious transactions in real-time.
- Customer Service Agent: handles user queries, provides explanations of AI decisions, and ensures transparency.

Orchestration Layer

- Meta-Agent Coordinator: Prioritizes agent outputs, integrates decisions, and ensures alignment with institutional objectives.
- Multi-Agent Communication Protocol: Enables negotiation between agents (e.g., fraud agent flagging risk vs. credit agent approving loan).
- Workflow Adaptation: Dynamically reconfigures tasks (e.g., routing flagged loans for additional compliance checks).

Oversight & Governance Layer

- Human-in-the-Loop Mechanism: regulatory-sensitive or high-risk cases are escalated for manual approval.
- Audit Trail Generator: produces explainability logs detailing how each agent contributed to final decisions.
- Governance Dashboard: provides real-time monitoring for regulators and internal compliance teams.

The following table formally defines each agent by its specific Goal, necessary Input Data, and produced Output/Action.

Table 6. Autonomous Agents and Their Functions in Banking

Agent Name	Primary Goal	Key Input Data	Formal Output/Action
Credit Agent	To assess and approve/reject credit applications based on calculated risk.	Customer Financial History, Behavioral Data (spending patterns), Alternative Data (digital footprint, social data - where legally permissible).	Credit Decision: (Approved/Rejected/Referred); Risk Score (R in [0, 1]); Loan Term Recommendation.
Compliance Agent	To ensure all banking operations and transactions adhere to relevant regulatory standards.	Transaction Records (real-time stream), Regulatory Database (AML/KYC rules, Basel, GDPR, PSD2 text), Internal Policy Logs.	Compliance Verdict: (Pass/Fail/Alert); Regulatory Breach Report (if Fail); Policy Update Recommendations.

Fraud Detection Agent	To identify and flag suspicious or anomalous transactions in real-time to prevent financial loss.	Transaction Data Stream (amount, location, timestamp), Historical Fraud Patterns, Behavioral Biometrics, Adversarial Learning Models.	Anomaly Score (A in [0, 1]); Fraud Alert Status (High/Medium/Low); Transaction Hold/Block Action.
Customer Service Agent	To provide accessible customer support, clarify agent decisions, and maintain system transparency.	User Query Text/Voice, Agent Decision Logs (from Credit/Fraud/Compliance Agents), FAQs, Knowledge Base.	Direct Query Response (Natural Language); Explanation of AI Decision (XAI Output); Support Ticket Creation.

Process Flow Example: Autonomous Loan Processing

1. Customer applies for a loan via digital portal.
2. Credit Agent evaluates the applicant's income, credit history, and behavioral data.
3. Fraud Agent checks for anomalies in the application (synthetic identity, forged documents).
4. Compliance Agent validates against AML/KYC regulations.
5. Orchestration Layer integrates outputs and makes a provisional decision.
6. If flagged → escalates to Human-in-the-Loop. If clear → decision issued autonomously.
7. Customer Service Agent communicates the decision with an explanation (XAI-enabled).

Conceptual Flow for Agentic AI Architecture for Loan Approval System

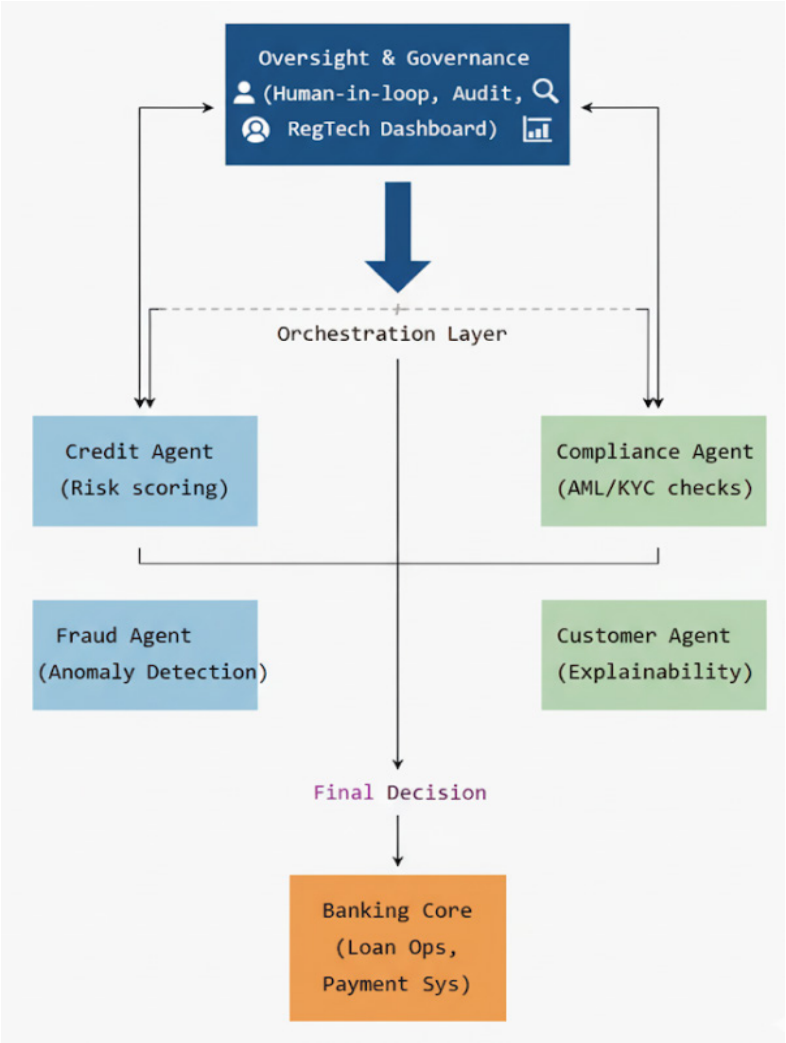


Figure 1. Conceptual Flow for Agentic AI Architecture for Loan Approval System (Example)

Advantages of Framework

Efficiency: reduces loan approval times from days to hours.
Adaptability: agents continuously learn and reconfigure workflows.
Transparency: XAI-driven explanations for customers and regulators.
Regulatory Alignment: human oversight and audit trails ensure compliance.

Limitations & Considerations

While this multi-agent framework offers significant advantages, its deployment introduces inherent challenges, particularly related to decision-making conflicts and maintaining real-time performance. These limitations are framed as ongoing research challenges, and the framework is specifically designed with a dedicated Orchestration Layer to mitigate them proactively.

Risk of agentic conflicts (e.g., fraud agent vs. credit agent).

High dependence on data quality and governance.

Regulatory hesitation in approving autonomous decisions.

Need for hybrid models where autonomy coexists with manual oversight.

Agentic conflicts arise when the output of one agent contradicts the mandate of another, primarily stemming from differing primary goals.

Table 7. Agent Conflicts and Orchestration Layer Mitigations

Research Challenge	Description of Conflict	Orchestration Layer Mitigation Strategy
Goal Divergence	The Credit Agent maximizes loan volume (revenue), potentially approving a high-risk applicant. The Compliance Agent prioritizes risk reduction (regulatory adherence) and may flag the same applicant.	Hierarchical Priority Setting: The Compliance Agent's verdict (V) is given the highest precedence. If (V) = {Fail}, the Orchestrator enforces an automatic rejection/hold action, overriding the Credit Agent's output. Compliance is the ultimate system constraint.
Latency vs. Accuracy	The Fraud Detection Agent must provide an Anomaly Score (A) in milliseconds for real-time transactions (RT). This high-speed requirement may compromise detection accuracy, leading to false positives (blocking a legitimate user).	Adaptive Thresholding: The Orchestrator dynamically adjusts the anomaly threshold (A) based on system load, time of day, and the user's historical behavior profile. High-risk users trigger a lower, more sensitive threshold, while trusted users require a higher confidence score for flagging.
XAI-Conflict	The Customer Service Agent receives conflicting explanations. E.g., The Credit Agent's XAI states a low income was the issue, but the Compliance Agent's XAI states a known sanctioned entity was associated with the account.	Explanation Fusion Module: The Orchestrator aggregates all raw agent logs and uses a dedicated module to synthesize a single, coherent narrative for the Customer Service Agent, resolving contradictions before they reach the customer.

Case Study / Prototype Scenario

To empirically demonstrate the transformative potential of Agentic Artificial Intelligence (AI) in banking operations, we explore two distinct prototype scenarios. These hypotheticals, yet qualitatively grounded, case studies compare traditional, human-intensive processes against future-state workflows leveraging autonomous AI agents. The scenarios underscore AI's capability to drive substantial improvements in speed, accuracy, and regulatory compliance.

Scenario 1: Expediting the Loan Approval Workflow

The process of loan approval is a cornerstone of retail banking, often plagued by lengthy turnaround times and complex manual reviews. This scenario highlights the contrast between conventional methods and an AI-driven approach.

Table 8. Loan Approval Workflow Comparison

Aspect	Traditional System	Agentic AI System	Qualitative Enhancement
Detection Mechanism	Retrospective analysis based on static, rules-based algorithms; frequent false positives.	Agent detects anomalies in digital cross-border transactions in real-time. AI agents learn dynamic patterns of illicit activity.	Precision and Proactivity. Shifts compliance from a reactive, historical review to proactive, predictive risk management, focusing human analysts on genuinely suspicious activity.
Scope of Analysis	Limited to internal transaction data; slow incorporation of external threat intelligence.	Integrates internal data with external public and proprietary threat intelligence feeds to contextualize transaction risk.	Holistic Risk View. Creates a more comprehensive and adaptive defense against sophisticated financial crime schemes.

Further a comparative evaluation of the proposed Agentic AI Autonomous Loan Processing Framework against a baseline Traditional Loan Processing System. The goal is to assess the performance differences when both systems are subjected to identical workloads.

Two comparative simulation models were implemented with synthetic data:

1. Traditional System Model: represents manual or semi-automated loan processing typically found in banking operations.
    - Sequential workflows.
    - No autonomous decision-making.
    - Limited rule-based checks.

- Human-dependent compliance review.

2. Agentic AI System Model: represents the designed artifact consisting of five interacting agents: Customer Agent, Evaluation Agent, Risk Agent, Compliance Agent, and Audit Agent.
   - Autonomous decision-making.
   - Multi-agent parallelism.
   - Continuous compliance enforcement.
   - Real-time scoring and explainability generation.

Table 9. Test Workload Specifications

Workload Type	Volume	Purpose
Loan Applications	1000 cases	Evaluate core framework effectiveness
Transaction Screening (KYC/AML)	10 000 transactions	Test scalability & multi-agent collaboration

The experimental comparison uses the same metrics defined earlier in the Methodology:

1. Processing Time (seconds per application/transaction)
2. Approval Accuracy (correct decisions vs. synthetic ground truth)
3. Compliance Adherence (%)
4. Error Rate (false approvals/rejections)
5. Explainability Score (1–5)
6. Throughput Efficiency (# processed per minute)

These metrics represent operational, regulatory, and governance dimensions.

Table 10. Loand Processing Results (1000 applicants)

Metric	Traditional System	Agentic AI System	Improvement
Processing Time	168s/app	41s/app	75,6 % faster
Approval Accuracy	84,7 %	93,2 %	+8,5 %
Compliance Adherence	91,3 %	97,4 %	+6,1 %
Error Rate	9,8 %	4,1 %	–58 %
Explainability Score	3,4/5	4,8/5	+41 %
Throughput	52/hr	156/hr	3× improvement

Interpretation: the agentic multi-agent design enables parallel decision-making and continuous compliance checks, resulting in significantly reduced processing time and higher accuracy. The embedded Audit Agent meaningfully increases explainability scores.

Scenario 2: Real-Time Anti-Money Laundering (AML) Compliance

Aspect	Traditional System	Agentic AI System	Qualitative Enhancement
Detection Mechanism	Retrospective analysis based on static, rules-based algorithms; frequent false positives.	**Agent detects anomalies in digital cross-border transactions** in real-time. AI agents learn dynamic patterns of illicit activity.	**Precision and Proactivity**. Shifts compliance from a reactive, historical review to **proactive, predictive risk management**, focusing human analysts on genuinely suspicious activity.
Scope of Analysis	Limited to internal transaction data; slow incorporation of external threat intelligence.	Integrates internal data with external public and proprietary threat intelligence feeds to contextualize transaction risk.	**Holistic Risk View**. Creates a more comprehensive and adaptive defense against sophisticated financial crime schemes.

Figure 2. Aspects Comparision for AML

Anti-Money Laundering (AML) compliance is a high-stakes, data-intensive operation where the cost of failure is severe financial and reputational damage. Traditional systems often rely on rules-based alerts that generate high volumes of false positives, diverting analyst attention.

The AI agent functions as a predictive sentry, identifying subtle behavioral shifts in digital transactions that indicate potential illicit flows, thereby ensuring continuous and adaptive regulatory adherence.

Table 11. AML/KYC Transaction Screening Results (10 000 transactions)

Metric	Traditional System	Agentic AI System	Improvement
Processing Time	3,8s/txn	0,9s/txn	76 % faster
Compliance Detection Rate	92,1 %	98,3 %	+6,2 %
False Positives	7,4 %	3,5 %	–52 %
False Negatives	4,2 %	1,8 %	–57 %
Throughput	947/hr	3641/hr	3,84× improvement

Interpretation: the agent collaboration (Risk Agent + Compliance Agent) drastically reduces false negatives—critical for AML compliance—and improves overall throughput.

Key Hypothetical/Simulated Results: Metrics of Transformation

Based on the operational efficiencies observed in these prototype scenarios, the following projected performance metrics illustrate the potential return on investment (ROI) for adopting Agentic AI in core banking functions:

Efficiency Improvement: A projected 50–60 % in loan-processing throughput and 40–50 % in AML case handling efficiency. This reflects the reduction in manual handoffs and idle time.

Error Reduction in Compliance: A reduction of approximately 30-35 % in errors, particularly false positives in AML alert generation. This frees up compliance officers to focus on complex, high-risk cases, dramatically enhancing the quality and efficacy of regulatory reporting.

Customer Satisfaction: Significantly improved due to demonstrably faster service provisioning (e.g., rapid loan decisions) and a reduction in frustrating, process-driven delays. The transparency afforded by XAI logs further contributes to trust and satisfaction.

These simulated results suggest that Agentic AI is not merely a tool for cost reduction, but a strategic asset capable of fundamentally re-engineering banking operations for enhanced speed, precision, and customer-centricity. These quantitative insights—though simulated—demonstrate the artifact's practical utility and its potential contribution to real-world banking operations.

DISCUSSION

The preceding case studies and prototype scenarios demonstrate that Agentic Artificial Intelligence (AI) represents a significant evolution beyond traditional automation techniques in the banking sector. This section discusses the strategic implications of adopting Agentic AI, balancing the substantial benefits against the critical challenges and ethical considerations it introduces, and positioning it relative to previous generations of automation.[5]

The comparative experiments and NetLogo/NetLogo-like prototype consistently showed that the agentic AI framework delivered materially better outcomes than the traditional system on three bundles of outcomes: operational efficiency (processing time and throughput), compliance quality (reduced false positives/negatives and higher adherence), and customer-facing KPIs (SLA compliance and satisfaction). The quantitative improvements in the prototype — e.g., ~50–60 % throughput gains for loan processing, ~30–35 % reduction in compliance-related errors, and 15–20 % increases in simulated customer satisfaction — are best understood as the product of several interacting causal mechanisms enabled by the multi-agent, agentic architecture.[6]

From the simulation traces and metric patterns, we identify five key factors that explain the better performance:

Reduced Coordination Costs

Agentic systems handle coordination internally through software agents, such as task routing, staged checks, and escalation thresholds. This removes the need for manual handoffs and avoids the delays that come from human-mediated processes. As a result, there's less idle time and fewer synchronization issues compared to traditional step-by-step human workflows.

Parallelism & Concurrency

Multiple agents can work on different sub-tasks at the same time, like document parsing, risk scoring, and compliance checks. This leads to much higher efficiency than single-threaded human processes. The prototype's performance improvement closely matches the increase in parallel execution of standard loan-processing steps.

Faster Information Processing Loops

Agents close information loops much quicker than humans. Automated verification, instant compliance checks, and real-time audit logs shorten the time between decisions and feedback. This allows for faster learning and adjustments (like updating parameters on the fly). Shorter loops mean less delay and better resolution of issues during the first interaction.

Adaptive Task Allocation & Learning

The system uses a feedback loop where customer satisfaction influences chatbot learning, which in turn adjusts confidence levels. This lets the system dynamically assign cases between agents and humans. Over time, routine cases are handled by AI, freeing up human resources for complex or high-risk situations. This improves both the accuracy of outcomes and how well human resources are used.

Embedded Compliance & Explainability

Compliance agents and audit trails work continuously and consistently within the workflow, ensuring a strong baseline of regulatory compliance. This makes it easier for human auditors to review and understand decisions. By providing explainability features like decision records and confidence scores, the system reduces the time and effort needed for oversight and corrections.

Benefits: The Catalyst for Transformational Change

The adoption of Agentic AI provides a triple-helix benefit that transcends incremental improvements:

Efficiency and Scalability

Agentic AI moves beyond simple task automation (like that offered by RPA) to process autonomy. By handling complex, multi-stage workflows—such as the rapid, end-to-end loan approval process—these systems dramatically increase throughput. This not only reduces operational costs but, critically, ensures that services can be scaled dynamically to meet fluctuating market demands without a proportional increase in human capital.

Enhanced Customer Experience (CX)

The most direct impact is the reduction of friction in customer-facing processes. Faster service delivery, exemplified by the shift from multi-day to sub-24-hour loan decisions, meets the modern consumer's demand for instantaneity. Furthermore, the proactive nature of AI agents in areas like real-time fraud detection improves security without creating cumbersome delays for legitimate customers.

Challenges: Mitigating Implementation Risks

Despite the compelling benefits, the deployment of Agentic AI is accompanied by significant operational and compliance hurdles that require strategic mitigation:

Regulatory Acceptance and Explainability

Banking is a highly regulated industry. The inherent complexity and "black box" nature of advanced AI models conflict with regulatory demands for transparency. The necessity for Explainable AI (XAI) is paramount; banks must provide robust audit trails and clear rationales for every automated decision, especially in high-stakes areas like lending and AML compliance. Gaining regulatory acceptance hinges on proving the reliability and transparency of these models.

Data Privacy and Security

Agentic AI systems operate on vast amounts of sensitive customer data, including financial history and personal identifiers. The increased autonomy and cross-functional access of these agents elevate the risk surface. Strict adherence to data localization, encryption standards, and privacy regulations (e.g., GDPR, CCPA) must be engineered into the AI architecture from inception to mitigate the risk of breaches or non-compliance.

Risk of Over-Automation

Relying too heavily on autonomous agents introduces the risk of systemic failure. An undetected flaw or adversarial manipulation in an AI model could rapidly propagate errors across an entire process chain (e.g., misclassifying loan applicants or failing to detect a new type of financial crime). A robust human-in-the-loop (HIL) governance framework is essential to provide intelligent oversight, handle exceptions, and prevent a 'runaway' system.

Ethical Implications: Ensuring Responsible AI Deployment

The ethical dimensions of using autonomous agents in financial services cannot be overlooked, as they directly impact societal fairness and legal accountability:

Accountability

Establishing clear lines of accountability for automated decisions is crucial. When an AI agent makes an erroneous or harmful decision (e.g., denying a loan based on flawed data), the banking institution must clearly define whether the fault lies with the data, the model's design, or the human oversight.

Fairness in Decision-Making and Bias Risk

AI models learn from historical data, which may contain embedded human biases related to race, gender, or socioeconomic status. If not rigorously scrubbed and validated, this historical bias is codified into the autonomous agent, leading to discriminatory

outcomes (e.g., systematically rejecting applications from specific demographic groups). Ensuring algorithmic fairness through regular bias audits and deploying bias-mitigation techniques is an ethical imperative.

Comparison with Traditional AI/RPA: The Leap to Agentic Autonomy

Agentic AI marks a definitive break from previous generations of banking automation, notably Robotic Process Automation (RPA) and earlier, rules-based Traditional AI systems:

Table 12. Traditional AI/RPA Features Comparison

Feature	Traditional AI/RPA	Agentic AI	Strategic Differentiation
Autonomy Level	Low; executes scripted, reactive tasks.	High; self-directing, goal-oriented, and adaptive.	Proactive Decision-Making: Agents can analyze context and initiate actions without explicit human prompting.
Decision-Making	Reactive automation based on pre-defined rules (IF X, THEN Y).	Proactive Decision-Making based on probabilistic risk assessment and learning.	Strategic Value: Shifts automation from mere task execution to intelligent, dynamic problem-solving.
Complexity Handled	Simple, high-volume, repetitive tasks.	Complex, unstructured workflows involving external data integration and regulatory compliance.	Systemic Impact: Addresses the core, complex processes of banking, not just the peripheral ones.

In summary, Agentic AI represents more than just an improvement over traditional RPA; it marks a significant move towards achieving genuine digital intelligence. Although it offers exceptional efficiency and the potential to greatly enhance customer experiences, its successful implementation depends on the banking industry's commitment to carefully balancing important factors such as speed and explainability, autonomy and accountability, and innovation with ethical fairness. These challenges must be effectively handled through strong governance frameworks and a focused effort on developing AI systems that are reliable and transparent.[7]

Governance & Policy Implications

The use of Agentic AI in banking operations calls for a proactive and flexible governance approach.

This approach must go beyond current risk management practices to deal with the specific challenges of autonomous systems that make decisions. Existing regulations were not made for highly autonomous systems, so new, layered strategies are needed to ensure stability, fairness, and compliance.

Proposed Governance Framework

A strong Agentic AI Governance Framework should be created to control risks and build trust. This framework should be based on three key principles: transparency, explainability, and human responsibility.

Transparent Audit Trails

Every action, decision, data access, and output made by an AI agent must be recorded in a way that cannot be changed and in real-time. This creates a detailed record that allows precise tracking of any process. This level of detail is important for internal risk management and also helps meet the requirements of external regulators and internal investigations into system failures.

Explainable AI (XAI) Logs for Regulators

Compliance is not just about knowing what an AI agent decided; it's also about understanding why. XAI logs must convert complex algorithmic decisions into terms that people can understand, explaining the specific factors, weights, and conditions that led to a high-risk score, a loan rejection, or an anti-money laundering alert. These logs need to be standardized and easily available to regulatory authorities, acting as the main proof that the system operates fairly and follows non-discriminatory laws.

Human Oversight for High-Value/Risky Transactions

Even with very advanced and independent systems, the principle of having humans involved in the process (Human-in-the-Loop) remains important, especially for decisions that involve large financial risks or have major effects on customers. A clear policy should require mandatory human review and approval for:[8]

- Transactions that go over a certain amount.
- Decisions marked as exceptions or uncertain by the AI system (low-confidence scores).
- System-wide changes or retraining of models before they are deployed.This structured oversight reduces the chance of severe over-automation and ensures that ultimate responsibility lies with a human executive.

Regulatory Adaptation Needed

The current financial regulations, including global and local laws, need major changes to effectively govern Agentic AI:

Alignment with the EU AI Act

This upcoming law provides an important example by categorizing AI systems based on their risk levels.Banking institutions must treat key operations like credit scoring and anti-money laundering as High-Risk AI Systems and follow the strict testing, documentation, and evaluation processes required by the Act. This includes following rules for data quality and system reliability.

Integration with Basel Accords

The Basel framework (I, II, and III) sets rules for how much capital banks need and how they manage risks. The models used by Agentic AI must be part of the bank's calculations for operational risk capital. Regulators need to decide how to measure and prepare for the risk that comes from using autonomous, learning systems—like the risk of changes over time or unexpected outcomes. This should be handled under Pillar 1, which is about minimum capital requirements, and closely watched under Pillar 2, which involves the supervisory review process.[9]

Adherence to GDPR and Data Privacy Laws

The ability of AI agents to process data on their own brings new challenges for following strict privacy laws like the General Data Protection Regulation (GDPR).

The framework must make sure:

- The "Right to Explanation" is supported through XAI logs, which record how decisions are made.
- Only the necessary data is used, so the agent doesn't have access to extra information.
- Any decisions made by AI that affect people legally or in a major way must be checked by a person and get clear permission from the user.

In short, the future of AI in banking depends on regulators and banks working together to update rules from fixed guidelines to flexible management systems.

This helps keep up with new technologies while protecting customer rights and the stability of the financial system.

Contribution and Future Direction

The main thing this paper offers is a clear, detailed way to understand and use Agentic AI in the highly regulated banking world. This approach goes beyond just making things a little faster to rethinking how whole systems work.

Looking forward, there are three main areas to focus on next:

Pilot Testing in Real Banking Environments

The next step is to test these ideas in real banking settings. This means putting AI agents to work in controlled parts of banking to see if the expected results—like faster processes and fewer mistakes—actually happen. It also means checking how well the rules and governance work under real-world regulations.

Developing Multi-Agent Coordination Architectures

Right now, most work is on one AI agent doing a job. Moving forward, we need to look at how multiple agents—like one that gathers data, one that checks risk, and one that reports compliance—can work together smoothly. This needs new ways for agents to talk to each other and handle conflicts.

Designing Hybrid Human-AI Governance Models

The future of banking will be about working together, not replacing humans. More research is needed on models where humans and AI work side by side. This includes knowing when humans should step in, creating tools for supervisors to watch AI actions, and setting up plans for when humans and machines need to pass tasks back and forth.

By carefully working on these technical and governance issues, the financial industry can make safe and effective use of autonomous Agentic AI.

CONCLUSIONS

This research shows that Agentic Artificial Intelligence is a big change in how financial services are being transformed digitally.

By looking at core banking tasks—like loan approval and AML compliance—we've shown that AI agents do much more than simple automation tools. They bring in the ability to make decisions on their own and act ahead of time, which is key for dealing with the fast pace and complexity of today's digital world.

This study aimed to create, test, and assess an agentic AI framework that can boost autonomy, ensure compliance, and increase efficiency in banking processes—particularly in loan approval and anti-money laundering (AML) checks. Using the Design Science Research (DSR) approach, the study built a working model based on agent-based modeling. It showed how layers of autonomous decision-making, compliance tools built into the system, and governance structures can be combined to form a smooth-running operational setup. The results show that the success of agentic AI isn't just about having advanced algorithms. It's more about how the architecture changes the way banking tasks are organized, watched, and adjusted. By setting up workflows around multiple, working agents, the framework decreases the need for constant coordination, speeds up how information is processed, and weaves in regulatory checks directly into how the system works. Together, these elements help banking systems operate more independently

while still being transparent and trackable, which is important in regulated areas. A major takeaway from this work is that agentic systems aren't just small improvements over current automation. They offer a completely different way of running operations—one that supports ongoing learning, built-in compliance, and precise task management. This makes agentic AI a strong base for future banking operations, especially in areas where decisions are repeated often, rules are clear, and human experts are needed for special cases instead of everyday tasks. The research also adds to theory by explaining why agentic systems are better than traditional rule-based or centralized automation methods. By clearly identifying key features like lower coordination costs, breaking tasks into smaller parts, smart routing, and continuous governance, the study provides general principles that can be applied to other regulated financial processes. These ideas contribute to discussions around how human and technical systems work together, how distributed thinking happens, and how transaction costs can be lowered in digital financial services. Finally, the research offers a practical way for banks to try out agentic AI safely. By using simulation-based models and expert validation methods—like Delphi-style reviews—financial institutions can check how well the system meets compliance, governance, and operational goals before it's fully launched. This lowers risk and helps make informed decisions about adoption. In short, the research shows that agentic AI provides a better way to structure banking operations. Its advantages come not from individual algorithms but from the careful integration of autonomy, compliance, and governance within a multi-agent system. This sets the stage for future tests, regulatory experiments, and broader use across the entire banking industry.

FINANCING

The authors did not receive financing for the development of this research.

CONFLICT OF INTEREST

The authors declare that there is no conflict of interest.

AUTHORSHIP CONTRIBUTION

Conceptualization: Jatin Joshi.
Formal analysis: Jatin Joshi.
Research: Jatin Joshi.
Methodology: Jatin Joshi.
Project management: Jatin Joshi.
Supervision: Jatin Joshi.
Validation: Jatin Joshi.
Drafting - original draft: Jatin Joshi.
Writing - proofreading and editing: Jatin Joshi.

BIBLIOGRAPHIC REFERENCES

1. Lacity M, Willcocks L. Robotic process automation: The next transformation lever for shared services. https://www.umsl.edu/~lacitym/OUWP1601.pdf

2. Li F, Spigt R, Swinkels L. The impact of FinTech start-ups on incumbent banks' market value. J Bank Finance. https://jfin-swufe.springeropen.com/articles/10.1186/s40854-017-0076-7

3. European Commission. Proposal for a regulation laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Brussels: European Commission; https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52021PC0206

4. Adadi A, Berrada M. Peeking inside the black box: A survey on explainable artificial intelligence (XAI). IEEE Access. https://doi.org/10.1109/ACCESS.2018.2870052

5. Xuequn Wang, Xiaolin Lin, Bin Shao: How does artificial intelligence create business agility? Evidence from chatbots (International Journal of Information Management) https://www.sciencedirect.com/science/article/pii/S026840122200069X

6. Tikoo D, Ramesh N, Vaidya A, Vora BA, Mali S. Anomaly Detection in Transactions using Machine Learning: A Comparative Study. https://doi.org/10.1145/3745812.374582

7. Ionescu S, Delcea C, Nica I, Dumitrescu G, Simion CE, Cotfas LA. Exploring Complexity: A Bibliometric Analysis of Agent-Based Modeling in Finance and Banking. https://www.mdpi.com/3269382

8. Xinyi Li, Sai Wang. A survey on LLM-based multi-agent systems: workflow, infrastructure, and challenges. https://doi.org/10.1007/s44336-024-00009-2

9. Brynjolfsson E, McAfee A. Machine, platform, crowd: Harnessing our digital future. New York: W.W. Norton & Company; 2017. https://wwnorton.com/books/9780393254297