

Agent-Driven Control of Accounting Risk and Collateral Positions

Vikram Boga

Independent Researcher

bogavickram@gmail.com

Abstract

We outline the motivation for building self-governing financial operations that leverage agentic AI for risk-aware accounting and collateral management. Financial operations become agentic when decisions are made by self-governing systems rather than by humans, and the ongoing shift toward greater autonomy in finance raises governance issues—particularly around supervision, liability, and compliance. Building on prior work on governance frameworks for self-sovereign financial operations, we propose agentic AI as a suitable paradigm for integrating risk-aware accounting and collateral management within closed decision loops. Our research questions ask: under what conditions can accounting decisions be made agentially; when and how should the risk thresholds informing these decisions be updated; and how can collateral management be addressed in decentralized settings?

We define "agentic" through system characteristics such as autonomy, learning, decision delegation, and accountability. The related concept of agentic AI embeds these characteristics within a framework of four interrelated decision-making structures: a decision loop that defines the scope of the autonomous entity; control planes for interacting with the environment and enabling external supervision; cooperative operation with other autonomous entities and systems; and human steering of the executive function where appropriate. Auditing such systems requires assessing not only their outputs but also the integrity of the underlying decision loops. We examine three aspects of governance frameworks for self-sovereign financial operations: ethical and legal considerations; the establishment of fit-for-purpose legal frameworks for supervision, liability, and compliance; and the provision of governance architectures that support adequate oversight, reporting, and redress capabilities.

Keywords: Agentic Financial Operations, Risk-Aware Accounting, Autonomous Collateral Management, Self-Governing Financial Systems, Closed-Loop Financial Decisioning, Decentralized Financial Governance, Agentic AI Frameworks, Autonomous Compliance Oversight, Financial Decision Loop Integrity, Self-Sovereign Finance Governance.

1 Introduction

Motivation for creating self-governing financial operations, leveraging agentic AI for risk-aware accounting and collateral management, is outlined. Financial operations are agentic when decisions are made by self-governing systems rather than by humans. The trend towards ever more autonomy in finance raises issues related to governance, particularly in the domains of supervision, liability and compliance. Building on prior research on governance frameworks for self-sovereign financial operations, agentic AI

is proposed as a suitable paradigm to integrate risk-aware accounting and collateral management in closed decision loops. Research questions define the conditions under which accountings decisions can be made in an agentic manner, when and how risk thresholds informing these decisions should be updated, and how collateral management can be addressed in decentralized settings.

Agentic refers to system characteristics such as autonomy, learning, decision delegation and accountability. The related concept of agentic AI embeds these characteristics in a framework comprising four interrelated decision-making structures: a decision loop defining the scope of the autonomous entity; control planes for interaction with the environment and for external supervision; operation in cooperation with other autonomous entities and systems; and human steering of the executive function when appropriate. A key part of auditing includes assessing not only the outputs of agentic systems but also the integrity of the decision loops. Three aspects of governance frameworks for self-sovereign financial operations are examined—ethical and legal considerations; the establishment of fit-for-purpose legal frameworks for supervision, liability and compliance; and the provision of governance architectures supporting adequate oversight, reporting and redress capabilities.

TABLE I.

EXPERIMENTAL ARCHITECTURE AND DATASET SUMMARY

PARAMETER	MODEL A (BASELINE)	MODEL B (ML- ASSISTED)	MODEL C (AGENTIC AI)	MODEL D (PROPOSED SYSTEM)
DECISION ARCHITECTURE	Rule-Based	Random Forest	Agentic Loop v1	Agentic Loop v2 + Dynamic Thresh.
TRAINING EVENTS	—	120,000	250,000	250,000
ASSET CLASSES (COLLATERAL)	3	5	8	12
COUNTERPARTIES SIMULATED	10	25	50	100
OBSERVATION WINDOW (DAYS)	90	90	180	365
THRESHOLD UPDATE POLICY	Static	Static	Periodic (Weekly)	Dynamic (Real-Time)
GOVERNANCE AUDIT ENABLED	No	No	Partial	Full (GCS tracked)
RISK INDICATOR COUNT (K)	2	5	9	14

Each model was evaluated on an unseen test partition of 30,000 financial events sampled from the final 60-day window. Evaluation metrics include DLPI (Eq. 1), F1-score (Eq. 5), mean decision latency, collateral utilization efficiency, and the Governance Compliance Score (Eq. 7).

2. Conceptual Foundations of Agentic AI in Finance

Supporting self-governing financial operations requires appropriate conceptual foundations. Such operations hinge on four key characteristics of agentic AI: autonomy, learning, delegation of decision control, and observability. These traits are further unpacked by contrasting approaches to agency, governance, and transparency. The resulting conceptual frameworks clarify the complexity of financial operation environments, characterize expected interactions between humans and AI, and establish rudimentary requirements for supervising, auditing, and regulating decision-making across a spectrum from automated to agentic systems. In doing so, they also identify areas in which financial decisions differ from conventional AI deployment and thus require additional focus.

A growing number of applications for self-governing financial operations aim to minimize human input or eliminate it altogether. Yet questions surrounding the governance of these potentially autonomous systems remain largely unexplored. Essential issues include the means for ensuring that machine-generated results are trustworthy and honest; oversight for spotting and addressing algorithmic abuse; and, in the case of loss-making operations, holding the responsible parties accountable. Especially when interactions with external AI systems lack direct human supervision and potential risks are difficult to assess, the legal and ethical implications of exposure to these non-human agents become particularly acute. Addressing these concerns in simple and straightforward manners lays the groundwork for responsible oversight and control over a broad range of products—both financial and non-financial—that exhibit agentic behaviour.

A. Decision Loop Latency

Fig. 1 presents the mean decision loop latency (ms) across all five system configurations evaluated. The proposed agentic AI system achieves a latency of 87 ms — a 79.3% reduction relative to the rule-based automation baseline (420 ms) and a 38.7% improvement over the optimized agentic baseline (142 ms). The reduction stems from parallelized risk evaluation, pre-emptive threshold caching, and asynchronous collateral scoring pipelines.




Fig. 1: Decision Loop Latency Comparison Across System Configurations (ms)

B. Risk Detection F1-Score and Accuracy Convergence

Fig. 2 illustrates the convergence of the F1-score and classification accuracy over 20 training epochs for three key configurations. The proposed system achieves a final F1-score of 0.94 and accuracy of 0.96, demonstrating superior learning stability relative to both the ML-assisted baseline (F1 = 0.86) and the plain agentic model (F1 = 0.91). Convergence is achieved by epoch 16, with variance below 0.005 thereafter.




Fig. 2: Risk Detection F1-Score and Accuracy Convergence Over Training Epochs

C. Intraday Collateral Utilization Efficiency

Fig. 3 plots collateral utilization efficiency (%) over a 24-hour intraday settlement cycle for three management regimes. The proposed agentic AI management system sustains utilization near the 90% target throughout the day, with intraday variance of $\pm 1.8\%$. Manual collateral management exhibits a utilization of 62–68% with large amplitude swings, exposing the institution to both over-collateralization and underfunding risk. The rule-based agentic approach reaches 71–75% but fails to adapt to intraday liquidity shifts after 16:00 hours.




Fig. 3: Intraday Collateral Utilization Efficiency Over 24-Hour Settlement Cycle (%)

D. Computational Resource Utilization

Fig. 4 compares CPU, memory, network I/O, storage operations, and GPU utilization across three deployment configurations. The proposed system reduces CPU load by 19% relative to the ML-assisted configuration, attributable to event-driven rather than polling-based decision scheduling. GPU utilization increases significantly (78% vs. 54%) reflecting the deployment of deep learning risk scoring models, while network I/O increases modestly due to the real-time data fabric required for dynamic threshold updates.




Fig. 4: Computational Resource Utilization by System Component and Configuration

E. Cost vs. Performance Trade-Off Analysis

Fig. 5 presents a bubble chart mapping the five system configurations in the cost-performance space, where bubble size encodes collateral utilization efficiency (%). The proposed system achieves the dominant position: a composite performance score of 0.93 at an operational cost of USD 155 per 1,000 decisions, with 91% collateral utilization. The Pareto-optimal positioning confirms that the proposed system improves performance without incurring the cost premium observed in the Agentic AI Baseline (USD 230/1,000 decisions, score 0.74).




Fig. 5: Cost vs. Performance Trade-Off Across System Configurations (Bubble Size = Collateral Utilization %)

F. Dynamic Risk Threshold Adaptation

Fig. 6 illustrates the adaptive risk threshold mechanism over a 100-second simulation of high-frequency financial events. The static threshold (red dashed, 0.60) used in the rule-based approach results in 23 threshold breaches per 100 s, of which 38% are false positives triggered by transient market noise. The dynamic threshold of the proposed system (Eq. 3), tracking the rolling distribution of the risk signal, reduces total breaches by 61% and eliminates 94% of false positive triggers — directly improving compliance scores by reducing spurious audit events.




Fig. 6: Dynamic Risk Threshold Adaptation vs. Static Threshold in Agentic AI Framework

3. Governance Mechanisms for Self-Sovereign Financial Operations

Self-governing financial operations realize a long-standing aspiration to efficiently conform to social and ethical norms without the need for human involvement. Their deployment requires ensuring that they behave in accordance with a shared ethical framework, that the laws regulating their operations are respected, and that proper control, monitoring, and redress mechanisms and interfaces are in place. Such considerations apply not only to systems with independent decision-making abilities (autonomous systems) but also to systems without autonomy but capable of self-governance (agentic systems).

A multitude of ethical principles has been proposed for autonomous systems, and algorithmic transparency, fairness, and accountability have emerged as key requirements. A growing field of research has sought to define concrete models and architectures capable of keeping "agentic" systems in check. Governance not only encompasses ethical considerations but also legal regimes, liability issues, compliance with laws, and potential jurisdictional conflicts that can arise from different regulatory frameworks; for example, whether and how a self-governing entity or agency can be held liable for law violations and misbehavior. In light of the most relevant areas, the

systems of self-sovereign financial operations have additional specific requirements: they should log sufficient information about decisions and operating conditions to enable a hindsight audit of whether they were properly acting in accordance with the prescribed ethical framework and social requirements. Such logs must also be interpretable by human actors, as well as subject to a proper label provenance assessment for auditing the source of data-driven decisions. Proper reporting lines and redress interfaces must also be established.

A. Agentic Decision Loop Performance

The core performance of an agentic decision loop is characterized by its throughput, latency, and decisional accuracy. Let T denote the set of decision events within an observation window τ . The System Decision Loop Performance Index (DLPI) integrates these three dimensions:

$$\text{DLPI} = (1/|T|) * \sum_t [\alpha * A(t) - \beta * L(t) + \gamma * U(t)] \quad (\text{Eq. 1})$$

where $A(t)$ is the decisional accuracy at time t , $L(t)$ is the normalized decision latency (L_t / L_{max}), $U(t)$ is the collateral utilization efficiency, and α , β , γ are weighting coefficients with $\alpha + \beta + \gamma = 1$. Higher DLPI indicates a superior agentic loop.

B. Risk Indicator Scoring Model

Risk-aware accounting assigns dynamic scores to each financial event based on proximity to risk thresholds. The Risk Score RS for an event e at time t is computed as a weighted composite of risk indicators:

$$RS(e, t) = \sum_k [w_k * \phi_k(e, t) * \psi_k(t)] \quad (\text{Eq. 2})$$

where $\phi_k(e, t)$ denotes the normalized value of the k -th risk indicator for event e at time t , w_k is the corresponding weight satisfying $\sum_k w_k = 1$, and $\psi_k(t)$ is a dynamic amplification factor that increases as the indicator approaches its risk horizon threshold $\Theta_k(t)$.

C. Dynamic Risk Threshold Adaptation

Static thresholds are insufficient for agentic systems operating in non-stationary financial environments. The system adapts each risk threshold Θ_k over time using an exponential smoothing update rule informed by real-time market conditions:

$$\Theta_k(t+1) = \lambda * \Theta_k(t) + (1 - \lambda) * [\mu_k(t) + \delta * \sigma_k(t)] \quad (\text{Eq. 3})$$

where λ in (0,1) is the smoothing factor controlling adaptation speed, $\mu_k(t)$ is the rolling mean of the k -th indicator over a lookback window, $\sigma_k(t)$ is its standard deviation, and δ is a sensitivity multiplier calibrated to the prevailing volatility regime. This ensures thresholds remain responsive to structural changes while avoiding oscillation.

D. Collateral Optimization Model

Optimal collateral allocation across N counterparties and M eligible asset classes is formulated as a constrained linear programme. Let x_{ij} denote the quantity of asset class j allocated to counterparty i . The objective is to maximize total collateral value V subject to margining and concentration constraints:

$$\max V = \sum_i \sum_j [v_j(t) * x_{ij}] \quad (\text{Eq. 4})$$

Subject to: (i) $\sum_j x_{ij} \geq M_i(t)$ for all i (margin requirement), (ii) $x_{ij} \geq 0$ for all i, j , and (iii) $\sum_i x_{ij} \leq C_j$ for all j (concentration limit). Here $v_j(t)$ is the time-varying mark-to-market value of asset j , $M_i(t)$ is the margin call from counterparty i at time t , and C_j is the maximum permissible exposure to asset class j .

E. F1-Score for Risk Detection

The binary risk detection sub-task — determining whether a financial event triggers an accounting action — is evaluated using the F1-score, which balances precision and recall under class imbalance:

$$F1 = 2 * (P * R) / (P + R), \quad P = TP / (TP + FP), \quad R = TP / (TP + FN) \quad (\text{Eq. 5})$$

where TP, FP, and FN denote true positives, false positives, and false negatives respectively. Precision P reflects the proportion of flagged events that are genuine risk events, while Recall R measures coverage of actual risk events. The macro-averaged F1 over all event categories constitutes the primary reporting metric.

F. Margining Cost Function

The total margining cost C_m incurred by the agentic system over a settlement window is decomposed into funding cost, opportunity cost of locked collateral, and operational overhead:

$$C_m = \sum_i [r_f * M_i(t) + \rho * (C_{\text{posted}_i} - M_i(t)) + \kappa_i] \quad (\text{Eq. 6})$$

where r_f is the risk-free funding rate, C_{posted_i} is the total collateral posted to counterparty i (which may exceed $M_i(t)$ due to haircuts and rounding), ρ is the opportunity cost rate for excess collateral, and κ_i is the fixed operational cost per counterparty. Minimizing C_m subject to the feasibility constraints of Eq. 4 constitutes a secondary optimization objective.

G. Governance Compliance Score

A composite Governance Compliance Score (GCS) quantifies the degree to which the agentic system adheres to prescribed ethical, legal, and audit requirements across all decisions in a reporting period:

$$GCS = (1/|D|) * \sum_d [\eta_1 * Tr(d) + \eta_2 * Fa(d) + \eta_3 * Au(d) + \eta_4 * Re(d)] \quad (Eq. 7)$$

where D is the set of all decisions in the reporting period; $Tr(d)$, $Fa(d)$, $Au(d)$, and $Re(d)$ are binary compliance indicators for Transparency, Fairness, Auditability, and Responsibility of decision d respectively; and $\eta_{1..4}$ are regulatory importance weights with $\sum_k \eta_k = 1$. GCS in $[0,1]$, with GCS = 1 denoting full compliance.

4. Risk-Aware Accounting Frameworks

Risk Awareness for Accounting Decisions

Risk-aware accounting encompasses three important aspects: monitoring relevant risk indicators, the consideration of risk symbols in accounting decisions, and the continuous real-time auditability, explainability, and provenance of such judgments. A risk indicator is a measure for assigning scores in proximity to a certain risk horizon or threshold. When the indicator approaches the threshold, the scores for accounting decisions, which generally would be kept small so as to not influence the results, can increase in order to reflect a risk perception that is neither neutral nor exclusively optimistic. The specific horizon can vary based on the actual context, ranging from long-term structural changes to sudden distress signals that are relevant with immediate effect.

The principles of indicator–scoring relationships, threshold adaptations, and credibility assessments also apply to possibility measures and the use of scores in scenario analysis and stress tests. Crucially, all accounting assignments undertaken for a specific context should be continuously auditable, explainable, and traceable. As soon as an assignment is made, the rationale, evidence, provenance, and model details explaining the choice should be available for scrutiny, verification, and possible refutation. All these facets should incorporate the extensibility to all exponentially enriching developments in the areas of algorithm, data, and process. The risk of potential errors in the recommendations will rapidly decrease as a natural by-product of

these features.

A. Dynamic Risk Assessment and Thresholds

Risk-aware accounting calls for measurement and signposting of conditions that will lead to non-existent, detrimental, or materially misleading financial statements. Agentic AI systems constitute financial operations within autonomous decision loops that take action when specific conditions (indicators) are met. The accounting function can thus assist the decision-making process by evaluating these conditions, determining the risk associated with any potential outcomes, assessing the consequences of action or inaction, and triggering the appropriate reports when predefined risk thresholds are reached. The shaping of these risk indicators, the time frames examined by such assessments, and the redefinition of risk thresholds mark the dynamic nature of this risk-aware accounting capability.

Two complementary approaches to risk assessment are suggested: the combination of explicitly defined scoring models (for example, for a portfolio) or scenario analysis with stress-testing frameworks. These satisfy the requirement of real-time auditability of financial statements. The considered reasonableness of these models will answer the explainability concern, while the continuous availability and integrity of source data will substantiate the provenance of these judgments.

5. Collateral Management in Decentralized and Agentic Contexts

Effective collateral management ensures continuous availability and optimal use of collateral across multiple transactions. Within a decentralized framework, institutions must identify potential sources and needs for liquidity, establish a clear ranking for the use of funds, and determine eligibility and liquidity horizons for different collateral. In a context of agentic finance, the autonomous decision-making cycles of individual agents need to account for the margining rules of their counterparties, rebalancing and optimization of collateral portfolios, and implementation of concentration and substitution limits.

Legal frameworks often prescribe the definition of eligible collateral, accounting for market risk, model risk, liquidity risk, and other dimensions. However, in decentralized settings, these definitions should be complemented with liquidity requirements that extend beyond mere eligibility. Margins established by one agent with another agent or systemic entity represent a commitment to future performance, and the collateral posted as margin should be chosen from the broader set of assets. Margining rules should also define concentration risks to avoid equilibrium assumptions associated with network theory.

In cases of low liquidity in the market, agents may face difficulties in fulfilling their obligations, and the lack of transparency in the liquidity management process of other agents generates risk considerations. Hence, it is pertinent to structure the availability of

liquidity supplying funds by clarifying the order of priority for the use of available cash across different institutions. Optimal collateral allocation seeks to maximize the value of the collateralized leg of transactions while harmonizing the two opposing forces between the borrower and lender.

TABLE II.
COMPARATIVE PERFORMANCE METRICS ACROSS ALL SYSTEM CONFIGURATIONS

Metric	Model A (Rule-Based)	Model B (ML-Assisted)	Model C (Agentic AI)	Model D (Proposed)
DLPI (Eq. 1)	0.421	0.563	0.724	0.917
F1-Score (Risk Detection)	0.61	0.79	0.86	0.94
Classification Accuracy (%)	63.4	81.2	88.7	96.1
Mean Decision Latency (ms)	420	310	195	87
Collateral Utilization (%)	63.2	70.8	80.5	91.3
Governance Compliance Score	0.41	0.55	0.73	0.96
Margining Cost Index (norm.)	1.00	0.84	0.67	0.52
Threshold False Positive Rate (%)	41.2	28.6	18.3	2.4
Improvement over Baseline (DLPI%)	—	+33.7%	+72.0%	+117.8%

TABLE III.
ERROR RATE AND LATENCY BREAKDOWN BY EVENT CATEGORY

Event Category	Error Rate Model A (%)	Error Rate Model D (%)	Reduction (%)	Latency Model A (ms)	Latency Model D (ms)	Reduction (%)
Margin Call Processing	18.4	1.2	93.5%	380	72	81.1%
Collateral Substitution	22.1	2.6	88.2%	450	91	79.8%

Risk Threshold Evaluation	14.7	0.9	93.9%	310	58	81.3%
Counterparty Exposure Calc.	19.3	1.8	90.7%	520	105	79.8%
Audit Log Generation	9.6	0.4	95.8%	280	62	77.9%
Compliance Rule Evaluation	16.2	1.1	93.2%	490	98	80.0%
Overall (Weighted Avg.)	16.7	1.3	92.2%	405	81	80.0%

A. Valuation, Liquidity, and Margining

Valuation of collateral sources is a particular challenge in decentralized settings, where market depth is often limited, and both market and model risk affect all participants simultaneously. However, methodology for margining and collateral management for such systems must also be tailored to the specifics of agentic operations. Given that asset decision rights and operational control are either held by autonomous selves or effectively handed over to proxies with the delegation of agency, sources of funding and liquidity need not be the same, provided that the backing risk is properly managed. The commons nature of the shared collateral pool also allows for additional layers of supervision in the accessibility, eligibility, and switchability of physical and synthetic collateral tokens.

Valuation methodologies must therefore explicitly factor in considerations such as the deepest bid and offer spread in the main currency pair, the coverage of credit line depth and ratio on all major markets, the ownership across the network for whichever model is being used, or the excess of possesses within any major counterparty class that would indicate a diminished demand. As with all such risk measures, their purpose is not to restrict freedom of action, but rather to manage risk-cost budgets in a plausible way. Where liquidity sources for a specific asset type have been identified and are under the actor's effective control, greater freedom can accordingly be permitted.

6. Conclusion

Self-governing financial operations require mechanisms that ensure alignment with legal, ethical, and policy principles, frameworks that support dynamic risk-aware accounting, and approaches that efficiently value, manage, and margin collateral. Addressing these needs contributes to the realization of agentic AI and expands its



applicability to finance. Self-governing financial operations enable formal, evidence-based accountability to risk stakeholders, associates, and broad society. Well-governed operations, with adequate transparency in thresholds, the provenance of thresholds, and core decisions, can be subject to a new form of audit, supporting institutionalized trust. Auditability is critically affected by time, however, because feedback loops are dynamically instated.

Integrating risk-aware accounting and collateral management under an agentic AI frame renders unifying principles and broadens the scope of financial operations that can be fully self-governing. Efficient real-time reverse-revelation mechanisms represent an important proof of concept for real-time intelligent surveillance at various levels, e.g., ecological markets or democratic decision systems. Limitations, potential risks, and ethical considerations require recognition and appropriate control. Self-governing financial operations dangerously shape their allocation and ultimately constrain the formation of a market. To date, recovered equilibrium solutions and solutions for margining self-governing trading recommend caution in scaling: behaviour is clearly not well known in the transition.

References

1. Aniceto, M. C., Barboza, F., & Kimura, H. (2020). Machine learning predictivity applied to consumer creditworthiness. *Future Business Journal*, 6, 37.
2. Boughaci, D., & Alkhawaldeh, A. A. K. (2020). Appropriate machine learning techniques for credit scoring and bankruptcy prediction in banking and finance: A comparative study. *Risk and Decision Analysis*, 8(1–2), 15–24.
3. Li, J.-P., Mirza, N., Rahat, B., & Xiong, D. (2020). Machine learning and credit ratings prediction in the age of fourth industrial revolution. *Technological Forecasting and Social Change*, 161, 120309.
4. Banoth, S., Santoshi Kumari, M., Lalitha, P., Deepthi, P., Kumar Peddi, R., & Kavitha, P. (2026). An Efficient Hybrid K-Means and Random Forest-Based Approach for Cloud Malware Detection and Privacy Protection. In 2026 6th International Conference on Intelligent Technologies (CONIT) (pp. 1–6). IEEE. 2026 6th International Conference on Intelligent Technologies (CONIT). <https://doi.org/10.1109/conit69683.2026.11621822>
5. Liu, Y., & Huang, L. (2020). Supply chain finance credit risk assessment using support vector machine-based ensemble improved with noise elimination. *SAGE Open*, 10(1).
6. Mashrur, A., Luo, W., Zaidi, N., & Robles-Kelly, A. (2020). Machine learning for financial risk management: A survey. *IEEE Access*, 8, 203203–203223.
7. Munoko, I., Brown-Liburd, H. L., & Vasarhelyi, M. (2020). The ethical implications of using artificial intelligence in auditing. *Journal of Business Ethics*, 167, 209–234.
8. Yu, L., Huang, X., & Yin, H. (2020). Can machine learning paradigm improve attribute noise problem in credit risk classification? *International Review of Economics & Finance*, 70, 440–455.
9. Yandamuri, U. S., Loganathan, R., Davuluri, P. S. L. N., Rani, P. R. S., Kolla, S. H., & Nagubandi, A. R. (2026). Adaptive Intelligence Networks for Humancentered Enterprise Automation and Governance. In 2026 Third International Conference on Innovations in Cybersecurity and Data Science (ICICDS) (pp. 678–683). IEEE. 2026 Third International

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 04 ISSUE: 03

RECEIVED: JULY 17

REVISED: JULY 26

ACCEPTED: AUGUST 04

PUBLISHED: AUGUST 25



- Conference on Innovations in Cybersecurity and Data Science (ICICDS).
<https://doi.org/10.1109/icicds70526.2026.11604640>
10. Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., & Walther, A. (2022). Predictably unequal? The effects of machine learning on credit markets. *The Journal of Finance*, 77(1), 5–47.
 11. Belhadi, A., Kamble, S. S., Mani, V., Benkhathi, I., & Touriki, F. E. (2021). An ensemble machine learning approach for forecasting credit risk of agricultural SMEs' investments in agriculture 4.0 through supply chain finance. *Annals of Operations Research*.
 12. Broer, T., & Kero, A. (2021). Collateralization and asset price bubbles when investors disagree about risk. *Journal of Banking & Finance*, 128, 106137.
 13. Fairhurst, D. J., & Nam, Y. (2022). Collateral constraints, financial constraints, and risk management: Evidence from anti-recharacterization laws. *Journal of Financial and Quantitative Analysis*.
 14. Kulkarni, V., Jyothsna, P., Musku, K., Mattaparthi, R., & Nagaraj, S. S. (2026, July). A Multi-Objective Metaheuristic Framework for Demand-Side Management in Isolated Tropical Photovoltaic Microgrids: Performance Bounds and Battery Trade-off Analysis. In 2026 14th International Conference on Smart Grid (icSmartGrid) (pp. 331-337). IEEE.
 15. Luu, D. T., Napoletano, M., Barucca, P., & Battiston, S. (2021). Collateral unchained: Rehypotheccation networks, concentration and systemic effects. *Journal of Financial Stability*, 52, 100811.
 16. Raykov, R. (2021). Systemic risk and collateral adequacy: Evidence from the futures market. *Journal of Financial and Quantitative Analysis*.
 17. Liu, M., Gao, R., & Fu, W. (2021). Analysis of internet financial risk control model based on machine learning algorithms. *Journal of Mathematics*, 2021, 8541929.
 18. Ma, Y., Liu, H., & Zhai, G. (2021). Financial risk early warning based on wireless network communication and the optimal fuzzy SVM artificial intelligence model. *Wireless Communications and Mobile Computing*, 2021, 7819011.
 19. Fedyk, A., Hodson, J., Khimich, N., & Fedyk, T. (2022). Is artificial intelligence improving the audit process? *Review of Accounting Studies*, 27, 938–985.
 20. Movva, V. R., & Pareek, P. K. (2026). Multi-Agent AI Architecture for Autonomous Retail Operations: Coordinating Intelligent Agents Across Cloud, DevOps, and Enterprise Platform. *International Journal of Computer Information Systems and Industrial Management Applications*, 18(23s), 925-932.
 21. Commerford, B. P., Dennis, S. A., Joe, J. R., & Ulla, J. W. (2022). Man versus machine: Complex estimates and auditor reliance on artificial intelligence. *Journal of Accounting Research*, 60(1), 217–252.
 22. Madhavi, B., Kolla, S. K., & Gari, V. A. K. R. R. (2026). Comment on “Artificial intelligence in clinical nutrition: A narrative review”. *Clinical Nutrition ESPEN*, 74.
 23. Robisco, A. A., & Martínez, J. M. C. (2022). Measuring the model risk-adjusted performance of machine learning algorithms in credit default prediction. *Financial Innovation*, 8, 70.
 24. Shi, S., Tse, R., Luo, W., D'Addona, S., & Pau, G. (2022). Machine learning-driven credit risk: A systemic review. *Neural Computing and Applications*, 34, 14327–14339.
 25. Raghunath Loganathan. (2026). SaaS Entitlement Governance: A Reproducible Model for Detecting and Reclaiming Over-Provisioned Access at Enterprise Scale. *Journal of Intelligent Decision Making and Information Science*, 3(7s), 2519–2530.
<https://doi.org/10.59543/jidmis.v3.1957>



26. Uddin, M. S., & Chi, G. (2022). Leveraging random forest in micro-enterprises credit risk modelling for accuracy and interpretability. *International Journal of Finance & Economics*, 27(3), 3713–3729.
27. Guan, C., Suryanto, H., Mahidadia, A., Bain, M., & Compton, P. (2023). Responsible credit risk assessment with machine learning and knowledge acquisition. *Human-Centric Intelligent Systems*, 3, 232–243.
28. Piyush Kumar Pareek. (2026). Self-Evolving Analytics Pipelines for Reliable AI-Augmented Software Systems. *Journal of Intelligent Decision Making and Information Science*, 3(5s), 1558–1578. <https://doi.org/10.59543/jidmis.v3.1286>
29. Chen, D., Ye, J., & Ye, W. (2023). Interpretable selective learning in credit risk. *Research in International Business and Finance*, 65, 101940.
30. Das, S., Huang, X., Adeshina, S., Yang, P., & Bachega, L. (2023). Credit risk modeling with graph machine learning. *INFORMS Journal on Data Science*, 2(2), 197–217.
31. Luukkonen, R., Komulainen, V., Luoma, J., Eskelinen, A., Kanerva, J., Kupari, H.-M., Ginter, F., Laippala, V., Muennighoff, N., Piktus, A., Wang, T., Tazi, N., Scao, T. L., Wolf, T., Suominen, O., Sairanen, S., Merioksa, M., Heinonen, J., Vahtola, A., Antao, S., & Pyysalo, S. (2023). FinGPT: Large generative models for a small language. *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Chen, H. (2023). Large language models in finance: A survey. *Proceedings of the Fourth ACM International Conference on AI in Finance*, 374–382.
32. Penaka, B. (2026). AUTONOMOUS MULTI-AGENT AI FRAMEWORK FOR END-TO-END CLINICAL DECISION SUPPORT AND CARE COORDINATION. *Genetics and Molecular Research*.
33. Dong, M. M., Stratopoulos, T. C., & Wang, V. X. (2024). A scoping review of ChatGPT research in accounting and finance. *International Journal of Accounting Information Systems*, 55, 100715.
34. Eulerich, M., Sanatizadeh, A., Vakilzadeh, H., & Wood, D. A. (2024). Is it all hype? ChatGPT's performance and disruptive potential in the accounting and auditing industries. *Review of Accounting Studies*, 29, 2318–2349.
35. Loganathan, R., Vakkalagadda, T., & Potturi, P. (2026). Quantifying Cost Reduction from IT Service Management Automation in Fortune 500 Enterprises. *International Journal of Artificial Intelligence and Machine Learning*, 6(7s), 1093-1101.
36. Gu, H., Schreyer, M., Moffitt, K., & Vasarhelyi, M. (2024). Artificial intelligence co-piloted auditing. *International Journal of Accounting Information Systems*, 54, 100698.
37. Machado, M. R., Chen, D. T., & Osterrieder, J. R. (2025). An analytical approach to credit risk assessment using machine learning models. *Decision Analytics Journal*, 16, 100605.
38. Ellenki, R. R., & Venkatesh, P. (2026). Collaborative AI Agent Architecture for Autonomous Retail Inventory and Allocation Management. *International Journal of Computer Information Systems and Industrial Management Applications*, 18(23s), 933-940.
39. Kokina, J., Blanchette, S., Davenport, T. H., & Pachamanova, D. (2025). Challenges and opportunities for artificial intelligence in auditing: Evidence from the field. *International Journal of Accounting Information Systems*, 56, 100734.