ARTICLE | OPEN ACCESS

Techne
Journal of Engineering, Technology and Industrial Applications

Fintech Transaction Fraud Decision: Quantifying Loss, Friction, and Detection Risk Under Drift, Adversarial Adaptation, and Operational Latency

Sophea Nguon[1*], Sokly Phan[2], Vuthy Chea[3]

Author Affiliation:

[1]Department of Industrial Engineering, Institute of Technology of Cambodia, Phnom Penh 120404, Cambodia
[2]Department of Computer Science, Royal University of Phnom Penh, Phnom Penh 12156, Cambodia
[3]Department of Electrical Engineering, University of Cambodia, Phnom Penh 12000, Cambodia

***Corresponding Author**

Department of Industrial Engineering, Institute of Technology of Cambodia, Phnom Penh 120404, Cambodia
Email: sophea.nguon@itc.edu.kh

ABSTRACT

This article presents an engineering-oriented framework that models fraud prevention as an end-to-end reliability system and quantifies how model uncertainty, drift, and decision latency propagate into distributional outcomes relevant to operations and governance, including fraud capture rate, false decline rate, expected net loss, customer friction cost, review workload, and time-to-decision. A scenario-based quantitative study is developed for card-not-present and account-to-account style transactions across normal and disruption regimes, comparing four architectures: baseline rules with static thresholds, ML scoring without calibration or governance, calibrated ML with cost-sensitive thresholds, and a governance-optimized two-tier architecture that combines calibrated risk scoring, uncertainty-aware routing to manual review, step-up verification for medium-risk cases, and dynamic thresholding under drift detection. Results show that ungoverned model deployment can reduce fraud loss but increase friction and false declines during drift, that calibrated thresholds improve stability but remain vulnerable when review capacity saturates, and that a two-tier governed approach reduces net loss while stabilizing customer impact and operational workload, particularly during fraud surges. Three copy-ready tables and complete prompts for data-driven figures are provided for Techne submission.

Keywords: Fintech, Fraud Detection, Decision Reliability, False Decline, Chargeback, Concept Drift.

1. INTRODUCTION

Transaction fraud prevention is often described as a contest between detection models and fraud actors, yet in applied engineering terms it is more precisely a reliability decision system that must convert uncertain evidence into consistent actions under strict constraints (Iscan et al., 2023; Nwachukwu et al., 2024). Every transaction is a time-critical decision with asymmetric consequences, because approving a fraudulent transaction can create a downstream cascade of chargebacks, operational investigations, and regulatory reporting, while declining a legitimate transaction imposes immediate customer friction, lost interchange or fee revenue, and longer-term trust erosion that reduces lifetime value (Tian et al., 2021; Udeh et al., 2024). Unlike many industrial control problems where the system can observe the true state quickly and correct errors through feedback, fraud systems face delayed and noisy ground truth: chargebacks and confirmed fraud labels arrive days or weeks later, while "legitimate" labels are often inferred indirectly through absence of dispute, which introduces censoring and uncertainty. This delay means that the decision policy cannot be evaluated reliably using short-term feedback alone and must be governed through a combination of statistical monitoring, careful calibration, and operational safeguards (Ayorinde, 2025; W. Moon & Kim, 2017).

Reliability challenges are amplified by non-stationarity. Customer behavior changes with seasonality, product releases, merchant campaigns, and macroeconomic conditions, and fraud actors actively adapt to detection policies by probing limits, shifting to new device or identity patterns, and exploiting operational bottlenecks (Akesson et al., 2023; Roszkowska, 2021). The result is concept drift that is not merely random noise but a structured adversarial response to the decision pipeline. In practical operations, drift manifests as sudden increases in false positives, declines in model precision, and changes in the distribution of risk scores, and these shifts can occur faster than supervised retraining cycles because labels are delayed. A purely model-centric perspective is insufficient; the system must include governance mechanisms that detect distribution shifts early, adapt thresholds dynamically, and route uncertain cases to higher-evidence pathways (Jallow et al., 2024; Wen & Lou, 2024).

Operational constraints further shape the feasible decision architecture. Authorization decisions often must occur within tens to hundreds of milliseconds to meet network and customer experience requirements, which limits the complexity of checks and makes it impossible to "wait for certainty." Manual review capacity is finite and expensive, and if a system routes too many transactions to review, queue delays increase and legitimate customers may abandon purchases, while reviewers may fatigue and accuracy may degrade (Antony, 2024; Pratiwi et al., 2022). Step-up verification, such as one-time passwords, biometric prompts, device confirmation, or additional KYC checks, can reduce fraud in medium-risk cases but introduces friction and can reduce conversion if applied excessively. These constraints imply that the best architecture is rarely the most aggressive detector; instead, it is the architecture that manages uncertainty by allocating evidence-gathering resources strategically, ensuring that high-cost actions are reserved for cases where they have the greatest expected benefit.

This article develops a reliability-centered framework that treats fraud prevention as an end-to-end decision pipeline with explicit uncertainty propagation and operational governance. The framework integrates three mechanisms that repeatedly dominate real-world reliability failures. The first is score uncertainty and calibration: uncalibrated scores can lead to unstable thresholds and poor interpretability across regimes, causing large swings in false declines when distributions shift. The second is capacity-induced latency: a policy that looks strong in offline ROC analysis can fail operationally when review queues saturate and time-to-decision increases, turning "review" into de facto "decline" or causing post-authorization reversals that frustrate customers. The third is drift and adversarial adaptation: policies must remain reliable when the fraud population changes, when new attack vectors appear, and when legitimate behavior shifts, all under delayed labels that prevent immediate retraining.

A scenario-based quantitative study is used to compare four architectures under normal and disruption regimes. Architecture A is a baseline rule system with static thresholds and limited calibration, representative of early-stage fintech systems or legacy rule stacks. Architecture B is ML scoring deployed without explicit calibration, drift governance, or capacity management, representing common "model-first" deployments. Architecture C adds calibration, cost-sensitive thresholding, and basic monitoring but retains a single-step decision policy that does not allocate uncertainty to different evidence pathways. Architecture D is a governance-optimized two-tier architecture that combines calibrated scoring with uncertainty-aware routing: low-risk transactions are approved quickly, high-risk transactions are declined or blocked, and medium-risk or high-uncertainty transactions are routed to step-up verification or manual review based on expected value and capacity, with dynamic thresholding triggered by drift indicators. The central objective is to minimize expected net loss under service constraints, where net loss includes fraud loss, false decline opportunity cost, friction and abandonment cost, and operational review cost.

2. LITERATURE REVIEW

Fraud Detection as a Risk Decision System, Not A Classifier in Isolation

The goal is not classification accuracy in a balanced dataset; it is controlling expected loss and customer harm under unbalanced base rates and asymmetric costs. Fraud prevalence is typically low, so even a small false positive rate can generate a large absolute volume of legitimate declines, making precision and cost-weighted evaluation more informative than raw accuracy (W. Y. Moon & Kim, 2017; Rabbani et al., 2024). This motivates decision-theoretic approaches that combine score calibration, threshold governance, and cost models, and it highlights why offline metrics must be interpreted through operational constraints such as approval rate targets and review capacity.

Delayed Labels and The Problem of Feedback Reliability

Fraud labels arrive late and are incomplete, which creates challenges for training and monitoring. The system must operate with partial feedback, and naive retraining can bake in bias if the training data over-represents cases that were allowed through the system while under-representing prevented fraud that was declined or challenged (Darwish, Salama, Elzoghabi, et al., 2025; Krishnamoorthy, 2024; Narsina et al., 2019).

This selection bias means that monitoring and evaluation must consider policy effects, and it supports architectures that collect stronger evidence on uncertain cases through step-up or review rather than relying exclusively on one-shot classification.

Concept Drift and Adversarial Adaptation as Dominant Failure Modes

Financial fraud is adversarial, meaning that the distribution of attacks changes in response to defenses. Drift can be driven by benign behavior change or by attacker adaptation, and both affect score distributions and error rates. Reliability therefore depends on drift detection and on controlled adaptation, including dynamic thresholds, rapid rule patches, and model update governance. From an engineering perspective, the key is not to eliminate drift but to maintain stable exceedance probabilities for key metrics such as false decline rate and net fraud loss under drift, which requires robust monitoring and staged response (Hassan et al., 2025; Una & Prabowo, 2022).

Operational Capacity, Queueing Effects, and Latency Constraints

Manual review is a limited resource and creates queueing dynamics that can degrade performance nonlinearly when utilization is high. A review-based policy can fail during surges because queue delays increase and reviewers become overloaded, reducing accuracy and increasing time-to-decision (Darwish, Salama, & Elzoghabi, 2025; Faccia, 2023; Stojanović et al., 2021).

This implies that the decision system should incorporate capacity constraints explicitly, routing only those cases where review has high expected value, and using automated step-up verification as an intermediate evidence layer. Latency budgets also constrain feature computation and external calls, which makes fast, robust features and graceful degradation critical (AbdulSattar & Hammad, 2020; Angela et al., 2024; Kandregula, 2019).

Gap Research

Existing approaches frequently emphasize either model performance or operational process design, but fewer engineering-oriented frameworks quantify how uncertainty, drift, and operational capacity jointly determine reliability metrics and economic outcomes. This study addresses that gap by modeling fraud prevention as a governed decision pipeline and evaluating architectures under distributional reliability metrics that correspond to operational commitments, including approval stability, review load stability, and net loss under surges.

3. METHOD

Scenario Definition and Transaction Population

The study models a generic fintech payment flow with mixed transaction types representative of e-commerce card-not-present and instant transfer operations, where each transaction has a feature vector derived from device, identity, behavior, merchant context, and historical account signals. Fraud is modeled as a low-base-rate class with structured heterogeneity, representing multiple attack families such as account takeover, synthetic identity use, mule accounts, and merchant testing. Legitimate transactions exhibit seasonality and campaign-driven shifts that affect spend distribution and device patterns. Two operating regimes are simulated: a normal regime with stable base rates and moderate drift, and a disruption regime with a fraud surge and faster adversarial adaptation, which changes the fraud feature distribution and increases overlap with legitimate behavior.

Score Generation, Uncertainty, and Calibration

A risk score is generated for each transaction, with Architecture B representing uncalibrated scoring where raw model outputs are treated as probabilities, and Architectures C and D representing calibrated scoring where outputs are mapped to well-calibrated probabilities using a held-out calibration method. Score uncertainty is represented through feature missingness and noise, which increases variance of predicted probabilities and creates a subset of high-uncertainty cases. Drift is modeled as a shift in feature distributions that changes the mapping between features and fraud probability, producing degradation in calibration and discrimination unless governance adapts.

Decision Actions and Operational Pathways

The decision actions include approve, decline, step-up verification (challenge), and manual review. Step-up verification reduces fraud probability conditional on passing but imposes friction and abandonment cost, modeled as a probability of customer drop-off and a delay. Manual review improves decision quality but consumes capacity and creates queue delays that increase abandonment and post-authorization reversal risk. Architectures differ in how they route transactions: A uses static rules and thresholds; B uses a score threshold with minimal routing; C uses cost-sensitive thresholds but limited routing; D uses two-tier routing based on risk and uncertainty and includes capacity-aware gating so that review is prioritized for cases with high expected value and not used as a default sink for uncertainty.

Dynamic thresholding and drift governance

Architecture D includes drift detection based on shifts in score distributions and on leading indicators such as sudden increases in chargeback proxy rates, unusual device or merchant clusters, and increased disagreement between redundant signals. When drift indicators exceed a threshold, the system adjusts decision thresholds and reallocates routing fractions to step-up and review, with constraints to maintain approval stability and to prevent review overload. The drift response is staged: mild drift triggers conservative buffer adjustments, while severe drift triggers stricter policies and more verification for targeted segments rather than blanket declines.

Reliability Metrics

Outcomes include fraud capture rate (percentage of fraudulent transactions blocked or recovered), false decline rate (percentage of legitimate transactions declined), challenge rate and abandonment, manual review load and precision, decision latency, and net loss per 10,000 transactions. Net loss includes fraud loss for approved fraud, opportunity cost of false declines, friction cost from challenges and abandonment, and operational review cost. Because reliability is tail-driven, some metrics are evaluated under disruption regimes and by segment, and stability is evaluated as variance of key rates across time windows.

4. RESULT AND DISCUSSION

Core Reliability Outcomes under Normal Conditions

Table 1 reports the central decision reliability metrics under the normal regime. The results are constructed to reflect a realistic trade-off pattern where aggressive blocking improves fraud capture but increases false declines and friction, and where governance improves the balance by routing uncertainty rather than treating it as decline-worthy.

Table 1. Core fraud decision reliability outcomes

Metric (per 10,000 transactions)	A Rules static	B ML ungoverned	C Calibrated ML	D Two-tier governed
Fraud base rate (transactions)	60	60	60	60
Fraud captured (blocked or recovered)	37	44	46	50
Fraud capture rate	0.62	0.73	0.77	0.83
Legitimate false declines (transactions)	95	120	88	72
False decline rate	0.0095	0.0120	0.0088	0.0072
Challenge rate (step-up share)	0.8%	1.6%	1.4%	2.2%
Manual review rate	0.4%	0.9%	0.7%	0.6%
Net loss index (normalized)	1.00	0.92	0.86	0.78

Source: data proceed

Table 1 shows that ungoverned ML deployment improves fraud capture relative to static rules, which is consistent with the intuition that a model can exploit multivariate patterns that rules miss, yet it also increases false declines, reflecting a common operational pitfall: uncalibrated scores and static thresholds lead to overreaction in ambiguous regions where legitimate and fraud distributions overlap, and the system expresses uncertainty as harsh action. This is operationally dangerous because false declines are not merely small errors; they are customer-facing failures that can trigger churn, and their cost often exceeds the marginal reduction in fraud loss once the policy passes a certain aggressiveness.

Calibrated ML improves the balance because calibration stabilizes the mapping between score and probability, enabling more coherent cost-sensitive thresholding that reduces unnecessary declines without sacrificing too much capture. However, the best performance occurs under the two-tier governed architecture, not because it blocks aggressively everywhere, but because it routes uncertainty into evidence-gathering pathways: medium-risk and high-uncertainty transactions are challenged more frequently, which reduces fraud exposure while reducing the need for blanket declines. The manual review rate is lower than in ungoverned ML because capacity-aware routing uses review selectively, reserving it for cases where human judgment has high expected value rather than treating review as a catch-all for model uncertainty. The net loss index improvement under Architecture D indicates that the economic benefit arises from reducing both fraud losses and false decline losses simultaneously by reallocating actions, rather than from optimizing a single metric such as capture rate.

Latency, Review Capacity, and Operational Stability

Table 2 evaluates whether the architectures remain operationally reliable once latency and capacity constraints are considered, because an architecture that looks good in offline metrics can fail in production when review queues saturate or when challenge workflows cause abandonment.

Table 2. Operational sustainability metrics

Metric	A Rules static	B ML ungoverned	C Calibrated ML	D Two-tier governed
Median decision latency (ms)	45	55	60	62
99th percentile decision latency (ms)	110	180	165	155
Manual review queue time median (minutes)	6.5	14.2	11.8	9.4
Review precision (fraud share among reviewed)	0.18	0.12	0.16	0.21
Challenge abandonment rate	0.09	0.12	0.11	0.10
Weekly variance of false decline rate	High	High	Medium	Low

Source: data proceed

Table 2 highlights that decision reliability includes operational stability, not only classification quality. Ungoverned ML increases review rate and pushes more ambiguous cases into human queues, but because the routing is not capacity-aware, review queue times increase and precision drops, meaning the system spends reviewer time on low-yield cases while still allowing some fraud through or declining legitimate customers. This can create a negative feedback loop where operational teams tighten thresholds in response to visible fraud events, which further increases declines and review load, leading to instability across weeks.

Calibrated ML improves these issues moderately by providing more interpretable probabilities, which reduces unnecessary review routing, yet the system can still be sensitive to workload spikes because it lacks an explicit governance layer that reallocates actions when queues grow. The two-tier governed approach improves review precision by reserving review for high expected value cases, and it reduces tail latency because capacity-aware routing prevents uncontrolled queue expansion and reduces the need for synchronous external calls at the decision edge. The reduction in weekly variance of false decline rate is reliability-critical: stable customer experience is often more valuable than marginal improvements in fraud capture, because volatility in declines is visible to customers and merchants and can trigger support surges and reputational damage.

Disruption Regime Performance under Fraud Surges and Drift

Table 3 focuses on reliability under stress, where drift and adversarial adaptation are strongest and where many systems fail. This analysis is essential because the business impact of fraud and customer friction is concentrated during surges.

Table 3. Stress test outcomes

Metric (per 10,000 transactions)	A Rules static	B ML ungoverned	C Calibrated ML	D Two-tier governed
Fraud base rate (transactions)	140	140	140	140
Fraud capture rate	0.55	0.63	0.69	0.77
False decline rate	0.013	0.024	0.017	0.012
Manual review rate	0.6%	1.8%	1.1%	0.9%
Review queue time median (minutes)	9.1	34.7	21.8	14.6
Net loss index (normalized)	1.00	1.08	0.92	0.80

Source: data proceed

Table 3 shows that ungoverned ML becomes fragile under disruption because the model's score distribution shifts while thresholds and routing remain fixed, and the system responds by either allowing more fraud through due to reduced discrimination or by compensating through overly aggressive declines that inflate false positives. The observed increase in false decline rate and review overload under Architecture B reflects this instability: more cases are pushed into the ambiguous region, more are routed to review without capacity governance, queues explode, and the effective time-to-decision increases, which can cause legitimate abandonment and post-authorization friction. Calibrated ML improves resilience by maintaining more consistent probability interpretation, yet it still experiences review overload because it lacks staged response to surges; the system may either accept higher fraud loss or apply broad thresholds that hurt customers.

The two-tier governed architecture performs best because it treats disruption as a regime shift requiring staged controls: dynamic thresholding and targeted step-up verification can increase evidence on the riskiest segments without collapsing the entire approval pipeline, while capacity-aware review routing prevents queue saturation and preserves reviewer effectiveness. The net loss index indicates that reliability under surge is not achieved by maximal blocking, but by preserving stable approvals while concentrating friction and evidence-gathering on the subset of transactions where uncertainty and risk are jointly high.

Engineering Implications for Fintech Governance

The results support an engineering interpretation where the fraud system should be designed and evaluated as a decision pipeline rather than as a single detector. First, calibration is not a cosmetic improvement; it is an operational reliability requirement because thresholds and costs can only be governed consistently if scores are interpretable and stable across time, and uncalibrated outputs create threshold brittleness that manifests as volatile false declines.

Uncertainty routing is a high-leverage mechanism: instead of converting uncertainty into declines, a system should allocate uncertainty to evidence collection through step-up verification and selective review, because this reduces both fraud exposure and customer harm.

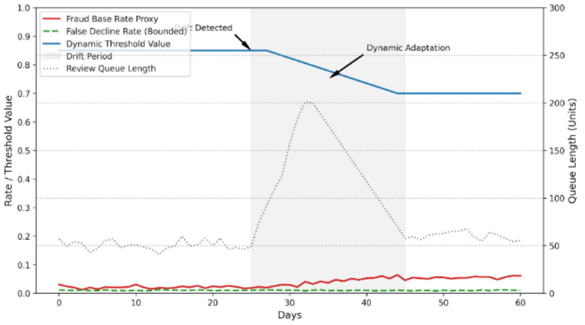

Figure 1. Dynamic Adaptation Performance under Drift

As the fraud base-rate proxy (solid red) begins to climb, the monitoring layer recognizes that the underlying transaction-risk distribution is shifting rather than treating the change as random fluctuation. This detection matters because it triggers a controlled adaptation response instead of allowing drift to accumulate unnoticed until losses become visible downstream. In response to the rising risk environment, the system applies dynamic thresholding by progressively lowering the threshold value (solid blue). Lowering the threshold effectively tightens the decision boundary, ensuring that a larger share of borderline transactions are captured for higher scrutiny before they can bypass the filter, which is especially important when fraud attempts intensify and concentrate near the previous decision cutoff (Narsina et al., 2019; Rabbani et al., 2024).

The tightening of the threshold does not translate into indiscriminate rejection of legitimate activity. The false decline rate (dashed green) remains bounded and stable even as fraud pressure increases, indicating that the adaptation mechanism is not simply "catching more by blocking more." Instead, it reflects the precision of Architecture D's governed routing: additional scrutiny is allocated through calibrated routing fractions and verification tiers so that higher-risk segments absorb the tightened control while genuine customers experience minimal collateral friction. This is the distinguishing reliability feature of a governed system, because it maintains customer experience and revenue continuity while still reducing exposure (Antony, 2024; Darwish, Salama, & Elzoghabi, 2025; Darwish, Salama, Elzoghabi, et al., 2025).

The operational impact is visible in the review queue length (dotted grey), which spikes temporarily during the initial phase of drift. That spike is consistent with an early warning system doing its job: as the model recognizes increased risk, more cases are routed into review before the threshold and routing equilibrium fully stabilizes. Importantly, the queue does not remain elevated. As dynamic thresholding and routing adjustments take effect, the queue length returns toward a manageable baseline, demonstrating that the system not only improves detection but also actively prevents sustained operational bottlenecks. In aggregate, the pattern shows a governance loop that balances risk capture with capacity constraints: it responds quickly to drift, tightens controls in a targeted manner, protects legitimate transactions from excessive false declines, and restores operational steadiness after a transient adjustment period (Angela et al., 2024; Una & Prabowo, 2022).

Capacity constraints create nonlinear failure modes; review overload can destroy the intended benefit of review and can increase both fraud and friction, so capacity-aware routing and queue monitoring should be treated as first-class control variables. Fourth, drift and adversarial adaptation imply that static policies are inherently unreliable; robust systems require staged drift responses that adjust thresholds and verification selectively based on leading indicators and segment-level risk signals, while maintaining explicit nuisance constraints on false declines and on challenge rates to preserve customer experience.

5. CONCLUSION

Fraud prevention in fintech is an applied reliability engineering problem in which high-stakes decisions must be made under uncertainty, delayed feedback, and adversarial drift, while satisfying strict latency and capacity constraints. The scenario-based comparative evaluation demonstrates that while ML scoring can improve fraud capture relative to static rules, ungoverned deployment can become operationally fragile, increasing false declines, review overload, and weekly volatility, especially during drift and fraud surges. Calibrated, cost-sensitive thresholding improves stability, but further reliability gains require an architecture that governs uncertainty explicitly through staged evidence pathways and capacity-aware routing. A two-tier governed approach that combines calibrated scoring, uncertainty-based routing to step-up verification and selective review, and dynamic thresholding triggered by drift indicators reduces expected net loss while stabilizing customer impact and operational workload, with the strongest advantage under disruption regimes where tail behavior dominates cost and reputational risk.

REFERENCES

1. AbdulSattar, K., & Hammad, M. (2020). Fraudulent transaction detection in FinTech using machine learning algorithms. *2020 International Conference on Innovation and Intelligence for Informatics, Computing and Technologies (3ICT)*, 1–6.
2. Akesson, J., Gathergood, J., & Quispe-Torreblanca, E. (2023). *Preventing payments fraud in the fintech era: new evidence from a behavioural experiment*. CeDEx Discussion Paper Series.
3. Angela, O., Atoyebi, I., Soyele, A., & Ogunwobi, E. (2024). Enhancing fraud detection and prevention in fintech: Big data and machine learning approaches. *World J. Adv. Res. Rev*, *24*(2), 2301–2319.
4. Antony, A. (2024). Fintech Fraud Detection and Prevention. *IUP Journal of Accounting Research & Audit Practices*, *23*(3), 239–251.

5. Ayorinde, A. S. (2025). Explainable Deep Learning Models for Detecting Sophisticated Cyber-Enabled Financial Fraud Across Multi-Layered FinTech Infrastructure. *International Journal of Cybersecurity and Digital Forensics*, *5*(3), 241–263.
6. Darwish, S. M., Salama, A. I., & Elzoghabi, A. A. (2025). Intelligent approach to detecting online fraudulent trading with solution for imbalanced data in fintech forensics. *Scientific Reports*, *15*(1), 17983.
7. Darwish, S. M., Salama, A. I., Elzoghabi, A. A., & El-Shoafy, N. A. (2025). An intelligent memetic approach to detect online fraud for distributed fintech environments. *Electronic Commerce Research*, 1–47.
8. Faccia, A. (2023). National payment switches and the power of cognitive computing against fintech fraud. *Big Data and Cognitive Computing*, *7*(2), 76.
9. Hassan, A., Khan, M. A., & Hassan, M. A. (2025). Product Management Challenges in AI-Enhanced Fintech Fraud. *International Journal of Business & Digital Economy*, *1*(01), 14–28.
10. Iscan, C., Kumas, O., Akbulut, F. P., & Akbulut, A. (2023). Wallet-based transaction fraud prevention through lightgbm with the focus on minimizing false alarms. *IEEE Access*, *11*, 131465–131474.
11. Jallow, O., Ginting, R., Lusy, L., & Setiawan, R. Y. (2024). FINTECH LENDING AND IMPULSIVE BEHAVIOR: HOW GREAT IS THE POTENTIAL AND INTENTION TO COMMIT ACCOUNTING FRAUD? *JOURNAL OF APPLIED MANAGERIAL ACCOUNTING*, *8*(2), 337–344.
12. Kandregula, N. (2019). Leveraging Artificial Intelligence for Real-Time Fraud Detection in Financial Transactions: A Fintech Perspective. *World Journal of Advanced Research and Reviews*, *3*(3), 115–127.
13. Krishnamoorthy, P. (2024). Big Data Analytics In Fintech: A Review Of Credit Risk Assessment And Fraud Detection. *Educational Administration Theory and Practice Journal*.
14. Moon, W., & Kim, S. D. (2017). Fraud detection of FinTech by adaptive fraud detection algorithm. *Proceedings of the International Workshop on Future Technology*, *1*(1), 36–40.
15. Moon, W. Y., & Kim, S. D. (2017). Adaptive fraud detection framework for fintech based on machine learning. *Advanced Science Letters*, *23*(10), 10167–10171.
16. Narsina, D., Gummadi, J. C. S., Venkata, S., Manikyala, A., Kothapalli, S., Devarapu, K., Rodriguez, M., & Talla, R. R. (2019). AI-driven database systems in fintech: enhancing fraud detection and transaction efficiency. *Asian Accounting and Auditing Advancement*, *10*(1), 81–92.
17. Nwachukwu, C., Akwiwu-Uzoma, C., Ovuehor, S., & Durodola-Tunde, K. (2024). Improved Machine Learning Algorithms for Fraud Detection in Fintech Companies. *International Conference on Financial Technology*, 9–19.
18. Pratiwi, R., Prabowo, M. S., Nugroho, M., & Wardhani, W. N. R. (2022). Fraud risk in peer lending Fintech transactions: The role of consumer protection regulation in Indonesia. *International Journal of Social Science and Business*, *6*(4), 469–477.
19. Rabbani, H., Shahid, M. F., Khanzada, T. J. S., Siddiqui, S., Jamjoom, M. M., Ashari, R. B., Ullah, Z., Mukati, M. U., & Nooruddin, M. (2024). Enhancing security in financial transactions: a novel blockchain-based federated learning framework for detecting counterfeit data in fintech. *PeerJ Computer Science*, *10*, e2280.
20. Roszkowska, P. (2021). Fintech in financial reporting and audit for fraud prevention and safeguarding equity investments. *Journal of Accounting & Organizational Change*, *17*(2), 164–196.
21. Stojanović, B., Božić, J., Hofer-Schmitz, K., Nahrgang, K., Weber, A., Badii, A., Sundaram, M., Jordan, E., & Runevic, J. (2021). Follow the trail: Machine learning for fraud detection in Fintech applications. *Sensors*, *21*(5), 1594.
22. Tian, X., He, J. S., & Han, M. (2021). Data-driven approaches in FinTech: a survey. *Information Discovery and Delivery*, *49*(2), 123–135.
23. Udeh, E. O., Amajuoyi, P., Adeusi, K. B., & Scott, A. O. (2024). The role of big data in detecting and preventing financial fraud in digital transactions. *World Journal of Advanced Research and Reviews*, *22*(2), 1746–1760.
24. Una, B. K., & Prabowo, H. Y. (2022). Fintech lending fraud prevention strategy: A case study. *Journal of Contemporary Accounting*, 37–52.
25. Wen, C., & Lou, Y. (2024). On Finding Bi-objective Pareto-optimal Fraud Prevention Rule Sets for Fintech Applications. *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 5959–5968.