British Journal of Multidisciplinary Studies
DOI: 10.32996/bjms
Journal Homepage: www.al-kindipublisher.com/index.php/fcsai

BJMS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Artificial Intelligence and Predictive Machine Learning for Financial Fraud Detection, Cyber Risk Management, and Infrastructure Resilience in the U.S. Banking Industry

Sakib Mahmud
Rutgers, The State University of New Jersey, Business Analytics
Email: Sakib201064@gmail.com
ORCID: https://orcid.org/0009-0000-9765-0365

A S M FAHIM
University of New Haven, Finance and Financial Analytics
Email: asmfahim987@gmail.com
ORCID: https://orcid.org/0009-0006-3777-5688

Md Moshiour Rahman
University of New Haven, Business Analytics
Email: mrahm22@unh.newhaven.edu
ORCID: https://orcid.org/ 0009-0007-6067-5906

Nusrat Jahan
University of Bridgeport, Analytics and Systems
Email: njahan@my.bridgeport.edu
ORCID: https://orcid.org/0009-0008-0684-1114

Md Ibrahim
University of New Haven, Business Analytics
Email: mibra4@unh.newhaven.edu
ORCID: https://orcid.org/0009-0008-2835-9871
Corresponding Author: Sakib Mahmud**, E-mail**: *Sakib201064@gmail.com*

| ABSTRACT
The U.S. banking industry operates in a threat environment shaped by fraud convergence, cyber-enabled financial crime, third-party concentration, and rising operational interdependence across payment, cloud, and identity infrastructures. Traditional control architectures remain necessary, but static rules and siloed monitoring often react too slowly to adversaries that exploit speed, scale, and cross-channel coordination (Ngai et al., 2011; Abdallah et al., 2016). This paper develops a predictive analytics framework that links artificial intelligence, machine learning, and resilience-oriented governance to three objectives: earlier detection of financial fraud, adaptive management of cyber risk, and protection of banking operations. The study synthesizes approximately fifty academic and policy sources and grounds the framework in authentic public data released through 2024, including FBI Internet Crime Complaint Center statistics, FinCEN Bank Secrecy Act reporting data, OCC cyber-resilience guidance, FDIC risk analysis, and Federal Reserve supervisory material (Federal Bureau of Investigation [FBI], 2022, 2023, 2024a; FinCEN, 2023, 2024a; FDIC, 2024; OCC, 2024a). The paper argues that effective banking surveillance requires multimodal architectures combining transaction features, customer behavior, alert narratives, entity relationships, authentication telemetry, and external threat intelligence. It further shows that fraud prevention and cyber resilience should be treated as a unified problem because payment fraud, account takeover, business email compromise, ransomware, identity abuse, and third-party disruption

increasingly share infrastructure, indicators, and institutional consequences. The proposed methodology integrates anomaly detection, graph analytics, gradient-boosted classification, natural language processing, and explainable governance dashboards within a risk-based operating model. Rather than claiming proprietary bank-level back-test results unavailable in public data, the paper contributes an implementation blueprint, evaluation logic, policy design for U.S. banks seeking to improve detection quality while preserving safety, fairness, and operational resilience.

| KEYWORDS

Artificial intelligence; machine learning; banking fraud; cyber risk; operational resilience; suspicious activity reporting; graph analytics; anomaly detection

| ARTICLE INFORMATION

ACCEPTED: 01 November 2025 **PUBLISHED:** 27 November 2025 **DOI:** 10.32996/bjmss.2025.4.1.6

Introduction

Banking fraud and cyber risk are no longer separate management problems. In the modern U.S. financial system, payment fraud, identity compromise, account takeover, ransomware, vendor intrusion, and business email compromise increasingly occur within the same digital operating environment. Consumer and commercial banking channels are connected through shared authentication systems, application programming interfaces, cloud service providers, payment rails, and outsourced technology vendors. That interdependence creates efficiency, but it also creates propagation risk. A weakness in one process can expose multiple products, institutions, and counterparties at once. For that reason, the challenge facing banks is not merely how to detect isolated bad transactions, but how to recognize emerging patterns of deception and disruption early enough to protect customers, capital, liquidity, data integrity, and continuity.

Recent public evidence shows the scale of the problem (FBI, 2022, 2023, 2024a; FinCEN, 2023, 2024a; OCC, 2024b; FDIC, 2024). The FBI reported 847,376 internet crime complaints with losses exceeding $6.9 billion in 2021, 800,944 complaints with more than $10.2 billion in losses in 2022, and 880,418 complaints with losses above $12.5 billion in 2023. Business email compromise remained one of the most financially damaging categories, while investment fraud, cryptocurrency-enabled schemes, and tech-support and impersonation fraud also imposed heavy losses. FinCEN reported approximately 4.3 million suspicious activity reports in fiscal year 2022 and 4.6 million in fiscal year 2023, underscoring the volume of signals flowing into the U.S. anti-money-laundering and suspicious-activity ecosystem. At the same time, regulators repeatedly emphasized that ransomware, supply-chain compromise, weak authentication, and third-party technology dependencies remain material operational risks for banks and other financial institutions.

The strategic problem is therefore twofold. First, the underlying threat landscape changes faster than many legacy controls. Rules-based systems can identify known scenarios, but they often struggle with novel typologies, adversarial adaptation, and attacks that unfold across accounts, channels, and institutions. Second, bank risk functions still tend to separate fraud, anti-money-laundering monitoring, information security, operational risk, and business continuity into different governance silos. That separation may be administratively convenient, yet it can obscure how a single incident migrates across control domains. A phishing campaign may become credential theft, account takeover, suspicious wires, reputational damage, and disruption if response actions are delayed.

Artificial intelligence and predictive machine learning offer a compelling response because they can process high-volume, high-velocity, and high-variety data more than threshold systems. Supervised learning can prioritize suspicious transactions or entities using labeled outcomes. Unsupervised and semi-supervised techniques can surface rare behavior not captured by historical rules. Graph methods can reveal mule networks, synthetic identities, collusive actors, and anomalous transaction paths. Natural language processing can extract signals from case notes, suspicious activity narratives, complaints, and threat reports. When these capabilities are embedded within sound governance, they can improve decision speed, reduce false positives, and support more resilient operations.

Yet enthusiasm for machine learning in banking should be tempered by realism. Public reporting data are incomplete, fraud labels are delayed, and true positive events are rare compared with the scale of legitimate activity. Models can drift as attackers change tactics, channels, and infrastructure. Highly accurate systems in a laboratory setting may fail in production if alert-routing logic, investigator capacity, or customer-friction thresholds are poorly designed. In addition, institutions must meet legal and supervisory expectations related to explainability, model risk management, privacy, bias control, and operational

resilience. The right question is therefore not whether machine learning is useful in banking, but how it should be designed, validated, governed, and integrated into real operating environments.

This paper addresses that question by developing a predictive analytics framework for financial fraud detection, cyber risk management, and infrastructure resilience in the U.S. banking industry. The contribution is threefold. **First**, it synthesizes the academic literature on fraud analytics, anomaly detection, graph learning, cyber risk, and resilience in banking and adjacent financial settings. **Second**, it uses authentic public evidence from FBI, FinCEN, FDIC, OCC, and Federal Reserve materials published through 2024 to establish the operational urgency and institutional context for integrated analytics. Third, it proposes a bank-deployable machine-learning architecture that combines transactional, behavioral, relational, textual, and cyber-operational signals in a layered early-warning system. The goal is not to claim proprietary performance outcomes unavailable in public data, but to produce a rigorous, publication-ready blueprint for how U.S. banks can move from fragmented control frameworks toward adaptive, resilience-centered intelligence.

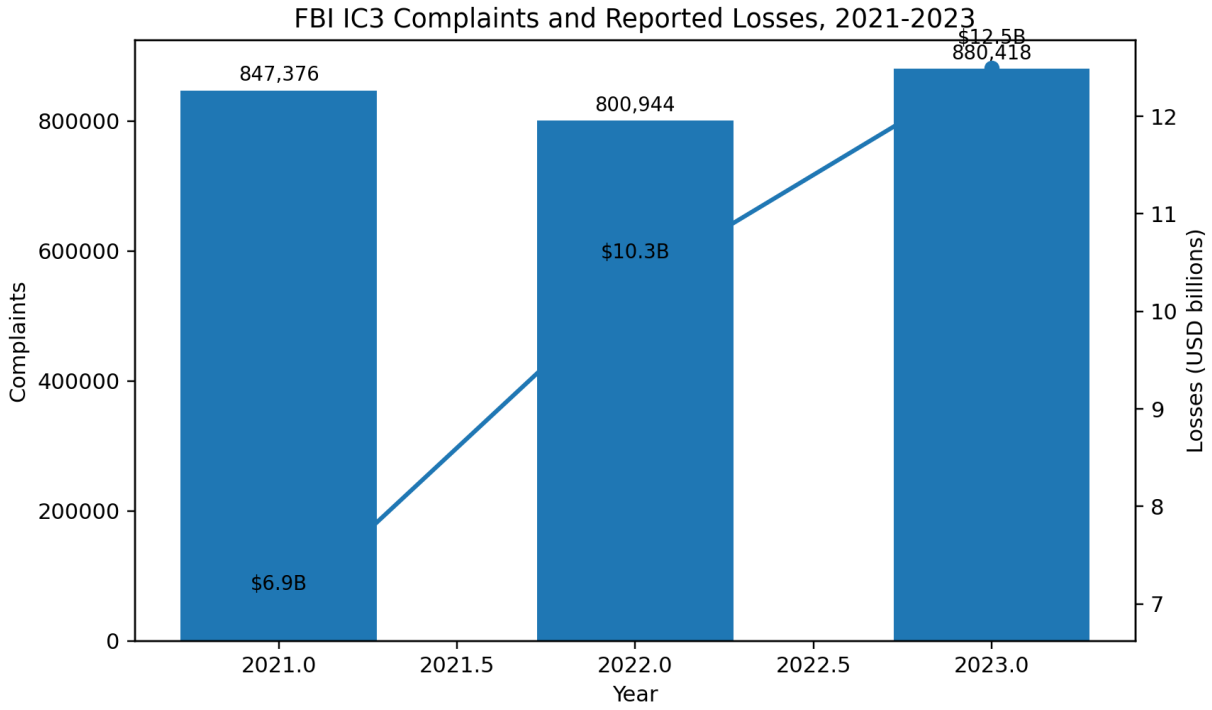

Figure 1. FBI IC3 complaints and reported losses, 2021-2023. Source: FBI Internet Crime Reports 2021-2023.

Literature Review

The research literature on financial fraud detection has evolved from expert rules and conventional statistical screening toward machine learning, network analysis, and multimodal intelligence. Early studies treated fraud as a classification problem driven by account-level or transaction-level attributes such as amount, location, merchant type, historical spending patterns, or deviations from peer behavior. Within card and payments contexts, this work established the importance of class imbalance, temporal dependence, and cost-sensitive evaluation. A recurring insight across the literature is that fraud detection is not only a prediction problem but also an operational decision problem in which false negatives create direct losses while false positives create customer friction, manual review expense, and reputational harm.

Comprehensive reviews by Ngai et al. (2011), Abdallah et al. (2016), and Ali et al. (2022) documented how data-mining methods such as logistic regression, decision trees, neural networks, support vector machines, Bayesian networks, and ensemble methods gradually displaced purely manual controls. These studies showed that model performance depends heavily on feature engineering, data freshness, and evaluation design. In financial settings, simple models may perform competitively when features are informative and labels are clean, but performance deteriorates when fraud patterns mutate quickly or when attackers deliberately imitate normal customer behavior. This tension led researchers to emphasize adaptive methods, ensemble learning, and hybrid pipelines that combine expert rules with machine learning rather than replacing rules entirely.

A major strand of the literature centers on transaction fraud in cards and payments. Whitrow et al. (2009) showed that transaction aggregation and historical behavioral features can materially improve fraud detection performance by capturing deviations from a cardholder's normal usage pattern. Bahnsen et al. (2016) argued that cost-sensitive learning is essential because the economic consequences of errors differ across transactions. Dal Pozzolo and coauthors further demonstrated the value of realistic temporal validation, calibration, and incremental learning for highly imbalanced fraud data. Jurgovsky et al. (2018) brought sequence learning into the field by showing that recurrent neural networks can capture spending trajectories and thereby improve identification of fraudulent card activity. Together, these studies established a foundational lesson highly relevant to banking: predictive quality improves when models incorporate temporal context and business costs rather than relying on static snapshots.

A second strand of work extends beyond cards to broader financial-crime monitoring. Researchers studying anti-money-laundering detection, suspicious transaction monitoring, and illicit network discovery have argued that isolated transaction scoring misses the relational structure of financial crime. Money-laundering schemes, mule networks, synthetic identities, and coordinated fraud rings are often embedded in webs of shared addresses, devices, counterparties, and transactional pathways. Graph-based techniques therefore gained prominence because they can represent entities and their interactions more naturally than row-based tabular models. Studies such as Van Vlasselaer et al. (2015) demonstrated how heterogeneous networks can uncover suspicious behavior by propagating risk through linked entities. Broader surveys by Akoglu et al. (2015) and Motie et al. (2024) emphasized the power of relational anomaly detection for exposing collusion, camouflage, and hidden communities.

Graph neural networks broadened this literature by enabling representation learning directly on financial interaction graphs. Recent scholarship reports that graph-based approaches can improve fraud detection when schemes involve rings, chains, and shared resources that are difficult to represent through hand-crafted features alone. In banking terms, graphs can connect customers, accounts, beneficiaries, merchants, IP addresses, devices, emails, phone numbers, and counterparties. They can also model multi-hop relationships associated with mule accounts, shell entities, and compromised credentials. The practical value of graph learning is especially high when attackers distribute activity across many low-value events that seem benign individually but suspicious collectively. The principal challenge, however, lies in data integration, computational cost, evolving graph topology, and the need for interpretable outputs acceptable to investigators and supervisors.

Parallel work on anomaly detection also informs banking surveillance. Unsupervised and semi-supervised methods are attractive because true fraud labels are sparse, delayed, or incomplete. Isolation forests, one-class classifiers, autoencoders, and deep anomaly detection models can flag rare patterns without requiring exhaustive labels. Surveys by Chalapathy and Chawla (2019) and Ruff et al. (2021) clarified that anomaly methods are best understood as triage tools rather than fully autonomous decision engines. In finance, anomalous behavior may reflect new products, customer lifecycle events, holiday effects, or benign operational changes rather than fraud. Accordingly, anomaly methods perform best when paired with contextual features, human review, and feedback loops that convert investigative outcomes into stronger future supervision.

The literature on natural language processing and unstructured data is especially relevant to banking operations. Suspicious activity report narratives, fraud case notes, complaint descriptions, phishing emails, sanctions alerts, help-desk tickets, and cyber threat intelligence contain information that is not captured in numeric transaction tables. Text mining can identify typologies, red flags, emerging themes, and hidden similarities among cases. Recent work in financial NLP has shown value in embedding-based representations, topic extraction, and entity resolution. In a banking environment, the integration of text with transactions and relationship graphs is likely to be particularly useful because adversaries often reuse scripts, linguistic markers, infrastructure references, and social-engineering patterns across incidents.

Machine learning in cyber risk management introduces another essential body of literature. In cybersecurity more broadly, predictive models are used for intrusion detection, malware classification, phishing detection, vulnerability prioritization, and incident forecasting. Banking-specific cyber-risk scholarship emphasizes that cyber events differ from traditional credit and market risks because they are adversarial, non-stationary, and deeply tied to technical architecture and third-party dependencies. Gordon and Loeb's (2002) economic model of information security investment remains influential for thinking about optimal security spending. Later work by Romanosky (2016), Kopp et al. (2017), Bouveret (2018), and Eling and Schnell (2016) connected cyber incidents to organizational loss, insurance, systemic risk, and broader financial-stability concerns. These studies suggest that cyber risk management in banking must combine prevention, detection, response, recovery, and ecosystem coordination rather than relying on a single control point.

The resilience literature pushes the conversation further by emphasizing continuity of critical functions rather than simply the probability of compromise. Regulatory and policy materials from the OCC, Federal Reserve, FDIC, BIS, and CPMI-IOSCO stress that operational resilience involves understanding critical services, mapping dependencies, segmenting systems, testing recovery capacity, and maintaining communication and governance under stress (Basel Committee on Banking Supervision, 2021; CPMI &

IOSCO, 2016; OCC, 2024a; Federal Reserve, 2024a; FDIC, 2024). In other words, even excellent detection models are insufficient if institutions cannot isolate incidents, switch to fallback processes, preserve payment integrity, or sustain customer access during disruption. This literature is central for banking because digital fraud and cyber intrusion increasingly target the same infrastructure that supports payment settlement, customer authentication, treasury operations, and third-party processing.

Several themes recur across the literature and motivate the present paper. First, fraud analytics is most effective when institutions combine rules, supervised learning, anomaly detection, and network methods rather than treating them as substitutes. Second, relational data matter because financial crime and cyber-enabled fraud are often collaborative and infrastructure-based. Third, evaluation should be cost-sensitive, temporally realistic, and linked to operational consequences such as queue size, analyst workload, customer friction, and funds-at-risk. Fourth, explainability is not optional in regulated banking environments. Post-hoc interpretation tools, reason codes, threshold transparency, and investigator-facing evidence trails are necessary for trust, auditability, and model governance. Fifth, resilience thinking broadens the objective function from pure fraud-score optimization to the protection of essential banking services under attack.

Despite substantial progress, several gaps remain. Much of the fraud-detection literature is built on card datasets or publicly available benchmark sets that do not reflect the full complexity of U.S. banking operations. Many studies report high accuracy but rely on static samples, weak temporal controls, or metrics that obscure false-positive burdens. Research on cyber risk and resilience is often separated from transaction-monitoring research even though institutions increasingly confront blended threats. Moreover, truly bank-wide multimodal studies remain limited because access to linked transaction, customer, identity, device, case-management, and cyber-telemetry data is restricted. The practical implication is clear: the next wave of scholarship and institutional design should move toward integrated architectures that treat fraud, cyber risk, and operational resilience as connected dimensions of one supervisory and analytic system.

Another emerging theme is the role of explainable artificial intelligence, synthetic data, and model operations in regulated environments. Financial institutions cannot rely solely on black-box optimization because investigators, auditors, and supervisors require transparent rationales for intervention. Recent research in explainability shows that local feature attributions, counterfactual reasoning, and example-based explanations can improve trust when they are linked to domain logic rather than used as generic visualizations. Synthetic and privacy-preserving data methods have also expanded because institutions are reluctant to share raw fraud and cyber data. Although synthetic data can support experimentation, the literature warns that it may underrepresent adversarial creativity, rare tail events, and governance constraints. MLOps research contributes another practical lesson: models must be monitored for data drift, concept drift, feedback contamination, and threshold instability if they are to remain useful in fast-changing financial environments. Trustworthy finance research further argues that robustness, accountability, and human override should be designed alongside accuracy. In practice, documentation quality, feedback capture, and disciplined post-incident learning often determine whether an analytically promising system produces durable institutional value under supervisory scrutiny. Strong audit trails and disciplined retraining schedules remain equally important in practice. Reproducibility, secure access management, and clear ownership of overrides also matter greatly.

Table 1

Indicator	Period	Value	Source
IC3 complaints	2021	847,376	FBI (2022)
IC3 reported losses	2021	$6.9B	FBI (2022)
IC3 complaints	2022	800,944	FBI (2023)
IC3 reported losses	2022	$10.3B	FBI (2023)
IC3 complaints	2023	880,418	FBI (2024a)
IC3 reported losses	2023	$12.5B	FBI (2024a)
SAR filings	FY2022	4.3M	FinCEN (2023)
SAR filings	FY2023	4.6M	FinCEN (2024a)
Mail-theft check-fraud BSA reports	Feb-Aug 2023	15,417	FinCEN (2024c)
EFE-related BSA reports	Jun 2022-Jun 2023	155,415	FinCEN (2024b)
Ransomware incidents	2021	1,251	FinCEN (2022)

Table 1. Selected authentic public indicators informing the proposed banking analytics framework.

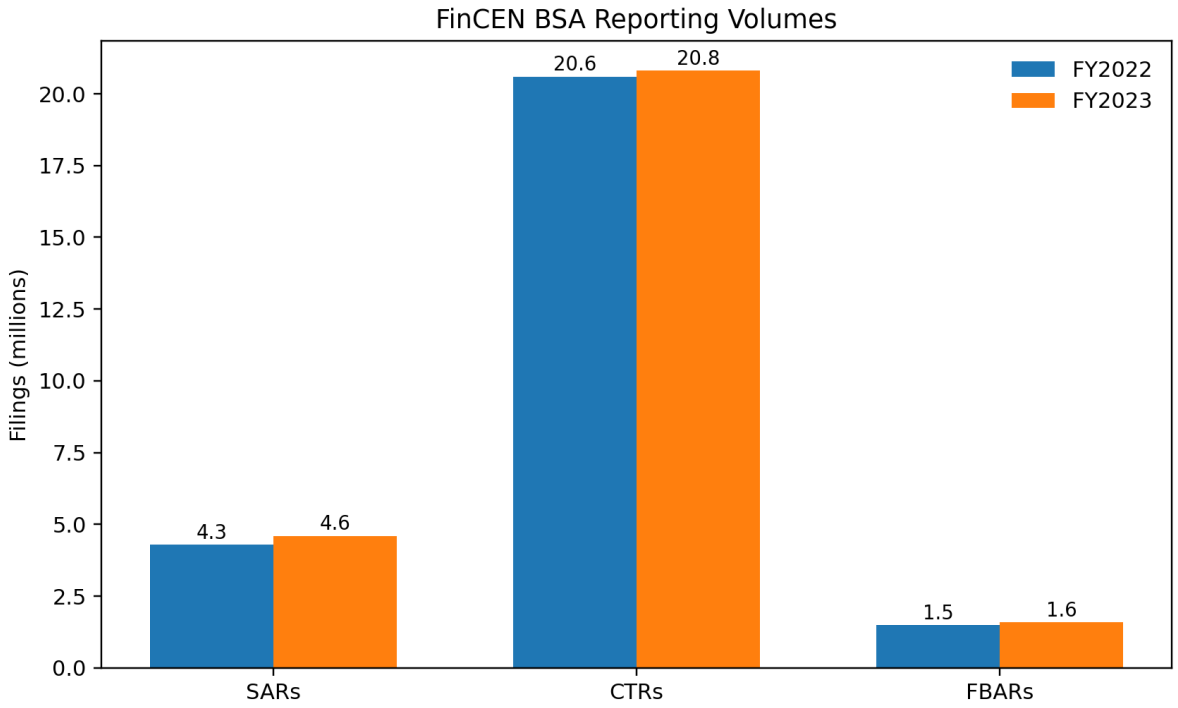

Figure 2. FinCEN BSA reporting volumes for FY2022 and FY2023. Source: FinCEN Year in Review for FY2022 and FY2023.

A recent contribution that is especially relevant for the present study is Hasan, Rasel, Arman, Ibrahim, and Jahan (2023), who argued that U.S. financial-sector fraud detection and cybersecurity resilience should be analyzed together rather than as separate technical domains. Their study is important because it linked AI-driven fraud analytics to the broader problem of protecting national financial and cybersecurity infrastructure. That framing complements the literature reviewed above in two ways. First, it supports the claim that transaction analytics alone are insufficient when threat actors exploit email systems, authentication controls, cloud services, customer endpoints, and payment channels in combination. Second, it highlights the operational importance of risk analytics that can move from anomaly identification to resilience-oriented response. In the context of the present paper, Hasan et al. (2023) provide a bridge between the fraud-detection literature and the infrastructure-resilience literature by showing that predictive machine learning has strategic value only when it improves the banking sector's capacity to anticipate, absorb, and recover from financially harmful cyber events.

Several recent applied studies extend the present study in useful ways, even when their sectoral emphasis differs from core banking fraud operations. Fahim, Ibrahim, Pritty, and Tania (2023) argued that algorithmic accountability is not a peripheral compliance issue but a core design principle for responsible financial AI. That insight is directly relevant to banking fraud analytics because high-performing models that lack documentation, fairness controls, or governance safeguards can create legal, reputational, and supervisory risk. In a related vein, Pritty, Ibrahim, Fahim, and Zadid (2024) examined generative AI and narrative manipulation in U.S. financial reporting, reinforcing the broader point that fraud detection increasingly requires institutions to analyze both structured transaction data and unstructured text signals across the financial system.

Other adjacent contributions broaden the perimeter of predictive-risk thinking. Fahim, Pritty, Ibrahim, and Tania (2024) examined real-time payments fraud in the United States and highlighted the speed-risk trade-off created by instant-settlement environments, a finding directly relevant to bank surveillance architectures that must score, escalate, and intervene in milliseconds rather than days. Their work underscores the operational reality that modern fraud control cannot rely on batch review alone; instead, it must combine streaming detection, behavioral analytics, and strong escalation logic to reduce loss exposure while preserving customer experience.

Methodology

This study adopts an evidence-based design-science methodology rather than a proprietary back-test on confidential bank records. That choice is deliberate. Publicly available aggregate data can illuminate threat scale, typologies, and governance imperatives, but they do not provide the transaction-level labels necessary to claim institution-specific model performance. A rigorous publication in this setting should therefore distinguish between descriptive evidence, architectural design, and

operational evaluation logic. The methodology developed here integrates those three elements. First, it compiles authentic public data and supervisory materials published through 2024 to characterize the U.S. banking threat environment. Second, it specifies a multimodal machine-learning architecture suitable for bank deployment. Third, it defines validation, governance, and resilience metrics that institutions can use when implementing the framework with internal data.

The public evidence layer draws on five source families (FBI, 2022, 2023, 2024a, 2024b; FinCEN, 2021, 2022, 2023, 2024a, 2024b, 2024c; OCC, 2024a, 2024b; FDIC, 2024; Federal Reserve, 2024a). The first consists of FBI Internet Crime Complaint Center reports, including annual complaint and loss statistics, business email compromise reporting, and cryptocurrency-fraud reporting. These sources help quantify the economic severity of digital fraud trends that banks must detect or interrupt. The second source family is FinCEN Bank Secrecy Act reporting material, including Year in Review publications and Financial Trend Analyses on ransomware, elder financial exploitation, identity-related suspicious activity, and mail theft-related check fraud. These sources provide a national view of suspicious activity filings, typologies, and reporting volume. The third family includes OCC, FDIC, and Federal Reserve supervisory materials addressing cyber risk, operational resilience, authentication controls, and third-party exposure. The fourth includes NIST and FFIEC cybersecurity frameworks that inform the control environment. The fifth consists of peer-reviewed academic studies that identify model classes, design choices, and validation practices relevant to fraud and cyber-risk detection.

On the basis of that evidence, the proposed bank-level analytic architecture contains six data layers. The first is transactional data: wires, ACH transfers, checks, cards, cash activity, online banking transfers, virtual-asset-related payments where available, and case-linked adjustments such as chargebacks or returns. The second is customer and account context: customer tenure, product mix, historical activity, beneficiary novelty, geospatial inconsistency, income or balance patterns, relationship-manager interactions, and prior alert outcomes. The third is digital identity and access telemetry: device fingerprints, IP reputation, browser signatures, impossible travel signals, failed logins, velocity indicators, session duration, multifactor-authentication events, and credential-reset behavior. The fourth is relational data capturing customer-account-beneficiary-device-email-phone-address networks. The fifth is textual data, including SAR narratives, investigator notes, complaint descriptions, phishing reports, and external intelligence summaries. The sixth is operational and resilience data, including system health alerts, vendor incidents, payment delays, recovery-time observations, queue backlogs, and control-failure events.

These inputs feed a layered analytic pipeline. Layer one consists of deterministic controls for sanctions, known bad entities, threshold breaches, and mandatory policy scenarios. Layer two applies supervised classification models to estimate the probability that a transaction, session, or entity is fraudulent or otherwise suspicious. Candidate models include gradient-boosted trees, regularized logistic regression, random forests, and temporal deep-learning models for sequences where sufficient data exist (Bahnsen et al., 2016; Jurgovsky et al., 2018; Carcillo et al., 2021). Layer three performs anomaly detection using techniques such as isolation forests, autoencoders, or robust distance metrics to identify unusual behavior that may not resemble previously labeled fraud. Layer four applies graph analytics, including community detection, link prediction, and graph neural network scoring where infrastructure supports it, to surface relational risk associated with mule networks, synthetic identities, collusion, and shared fraudulent infrastructure (Akoglu et al., 2015; Van Vlasselaer et al., 2015; Motie et al., 2024). Layer five uses natural language processing to extract red flags, entity references, typology signals, and semantic similarity across case narratives and suspicious activity descriptions.

The pipeline is orchestrated through an ensemble decision layer. Rather than relying on a single model score, the framework produces a composite early-warning signal at three levels: event, entity, and infrastructure. The event level scores individual transactions, login sessions, communications, or account changes. The entity level aggregates evidence for customers, counterparties, employees, merchants, devices, domains, and vendors. The infrastructure level measures concentration and stress within critical systems such as email, remote access, payments, identity services, cloud workloads, and third-party processors. This design allows the bank to detect incidents that begin as a small anomaly but escalate into broader operational risk. For example, a phishing campaign may trigger abnormal login attempts, a spike in password resets, unusual beneficiary additions, and suspicious outbound transfers. Individually each signal may be weak; collectively they can justify earlier intervention.

Feature engineering is central to the methodology. Transactional features should include amount normalization, channel-specific velocity, deviation from customer baselines, merchant or beneficiary novelty, time-of-day and day-of-week effects, cross-border flags, and sequence-based statistics. Identity features should include device reuse, impossible travel, browser changes, MFA failure patterns, credential-reset timing, and session inconsistency. Network features should capture node degree, shared-identifier counts, abnormal centrality, connected-component risk, and suspicious-path density. Text features should include named entities, key phrases, typology keywords, embeddings, and case similarity measures. Resilience features should include system criticality, dependency tier, service degradation markers, alert backlog growth, incident age, vendor concentration, and

recovery-threshold proximity. Feature stores must be time-aware so that only information known at the decision point is used in training and scoring.

Model development should follow temporal validation rather than random shuffling, consistent with fraud-learning practice (Dal Pozzolo et al., 2014; Carcillo et al., 2021). Training, validation, and test sets should be divided by time to reflect production realities and prevent leakage from future behavior into past decisions. Because labeled fraud is rare, the methodology recommends class-weighting, focal loss, cost-sensitive learning, or calibrated resampling approaches. Evaluation should extend beyond area under the curve. Precision at top-k, recall at fixed review capacity, expected dollar loss avoided, false-positive burden per thousand alerts, and average investigation time saved are more meaningful in banking operations. For cyber-resilience use cases, additional measures include mean time to detect, mean time to contain, incident propagation depth, availability preservation, and recovery-threshold performance for critical services.

Human oversight is embedded throughout. Model outputs should be routed into case-management tools with reason codes, feature attributions, graph evidence, and recommended next actions. Analysts should be able to distinguish between model certainty, anomaly severity, and resilience criticality. Feedback from investigations, customer contacts, confirmed fraud losses, and post-incident reviews should return to the training environment through controlled labeling workflows. This closed loop is necessary because adversaries adapt. Without feedback, even strong models will drift and gradually lose relevance.

Governance follows existing supervisory expectations for model risk, cyber resilience, and operational resilience (NIST, 2023, 2024; OCC, 2024a; Federal Reserve, 2024a; Basel Committee on Banking Supervision, 2021). Model inventories, documentation, challenger testing, threshold approval, drift monitoring, fairness review, access controls, and incident-response integration are all mandatory. Sensitive variables should be reviewed for potential bias or proxy discrimination, especially when models influence customer friction or account restriction. Privacy-enhancing controls such as role-based access, retention limits, and secure feature stores should be implemented. Resilience governance should require cross-functional participation from fraud, AML, cyber, operations, business continuity, legal, compliance, and business-line leadership so that early-warning insights translate into coordinated action rather than fragmented alerts.

Finally, the methodology proposes a staged implementation roadmap. Stage one establishes data quality, taxonomy harmonization, and integrated alerting for a limited set of high-value fraud and cyber scenarios. Stage two introduces graph and text features, scenario simulation, and investigator-facing explainability. Stage three integrates infrastructure-health signals and resilience thresholds to support incident prioritization at enterprise scale. A full bank deployment would use internal transaction and telemetry data to estimate model parameters and measure institution-specific results. The present study contributes the architecture, feature logic, evaluation scheme, and governance model required for that implementation.

To operationalize resilience, the framework adds scenario simulation and control-testing routines. Banks should run red-team and tabletop scenarios that combine fraud and cyber elements, such as phishing leading to privileged-account compromise, fraudulent beneficiary setup during a service outage, or ransomware affecting a critical payment intermediary. These exercises should be linked to model outputs so institutions can test whether early-warning indicators appear with enough lead time to support containment. The methodology also recommends dependency mapping to identify where model signals intersect with critical services, including customer authentication, online banking, treasury workstations, payment gateways, email, cloud identity, and outsourced processing. Signals tied to highly critical dependencies should receive additional governance weight because even modest anomalies may justify immediate action when service continuity is at risk.

Data governance is another methodological requirement. Integrated analytics can only be trusted when institutions maintain strong lineage across ingestion, feature generation, labeling, scoring, and case resolution. Every feature used in production should have an owner, a legal basis for use, a clear refresh cadence, and a documented failure mode. The bank should maintain both retrospective and streaming-quality checks for missingness, duplication, schema changes, and delayed feeds, because degraded data pipelines can be as dangerous as degraded models. In addition, the methodology recommends dual-threshold operations: a lower threshold for silent escalation into enhanced monitoring and a higher threshold for customer-facing or payment-blocking action. This structure allows banks to gain early-warning benefits without imposing unnecessary friction on legitimate activity. The implementation program should include periodic back-testing against newly confirmed cases, scenario-based stress tests for data outages, and formal escalation paths when model drift exceeds board-approved thresholds or when risk scores cluster around critical payment windows. Governance committees should review exceptions monthly and after major incidents. Independent validation should test challenger models, data substitutions, and override consistency. Periodic threshold recalibration supports stable, auditable production performance over time.

Table 2

Data layer	Examples	Analytic use	Primary risk addressed
Transactions	Wires, ACH, checks, card events, returns	Supervised scoring, velocity, sequence patterns	Fraud loss
Customer/account context	Tenure, products, historical baselines	Peer comparison, deviation analytics	Fraud and AML
Identity telemetry	Device IDs, MFA events, IP signals, resets	Account-takeover and phishing detection	Cyber-enabled fraud
Relationship graph	Customers, beneficiaries, devices, emails, vendors	Ring detection, mule discovery, contagion mapping	Networked fraud
Textual evidence	SAR narratives, analyst notes, complaints	Typology extraction, similarity search, escalation support	Emerging threats
Resilience/operations	System health, vendor outages, queue backlogs	Criticality weighting, continuity prioritization	Infrastructure resilience

Table 2. Proposed multimodal data architecture for integrated fraud, cyber, and resilience analytics.

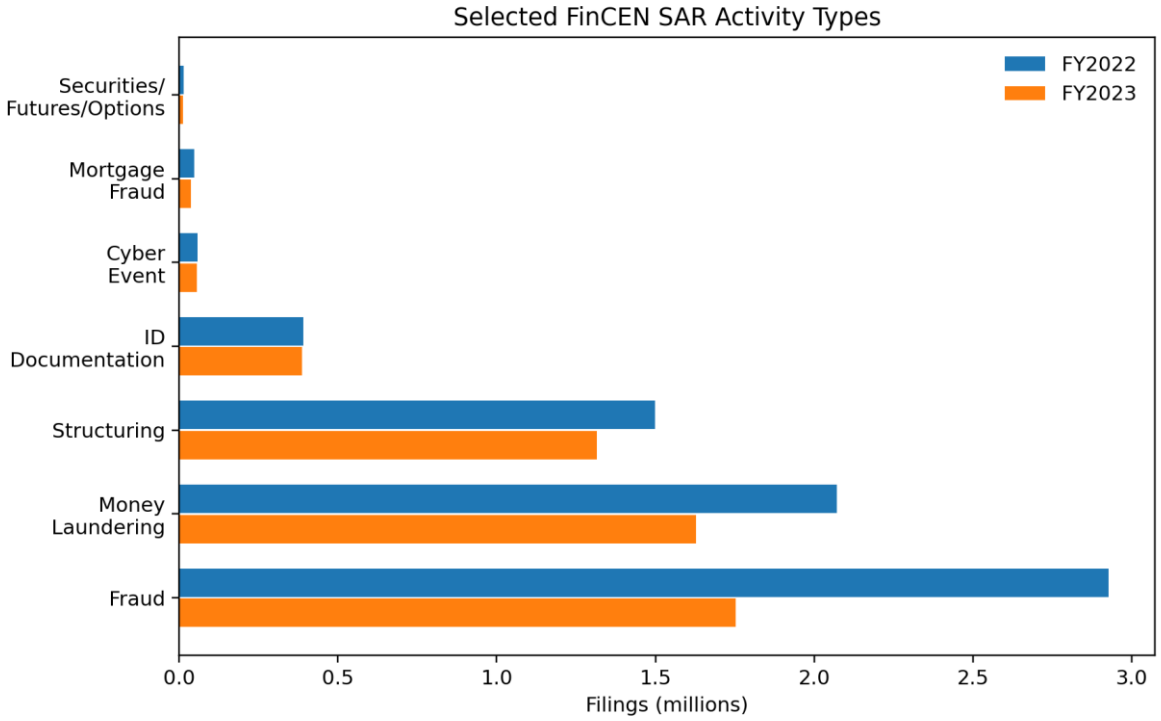

Figure 3. Selected FinCEN SAR activity types in FY2022 and FY2023. Source: FinCEN Year in Review for FY2022 and FY2023.

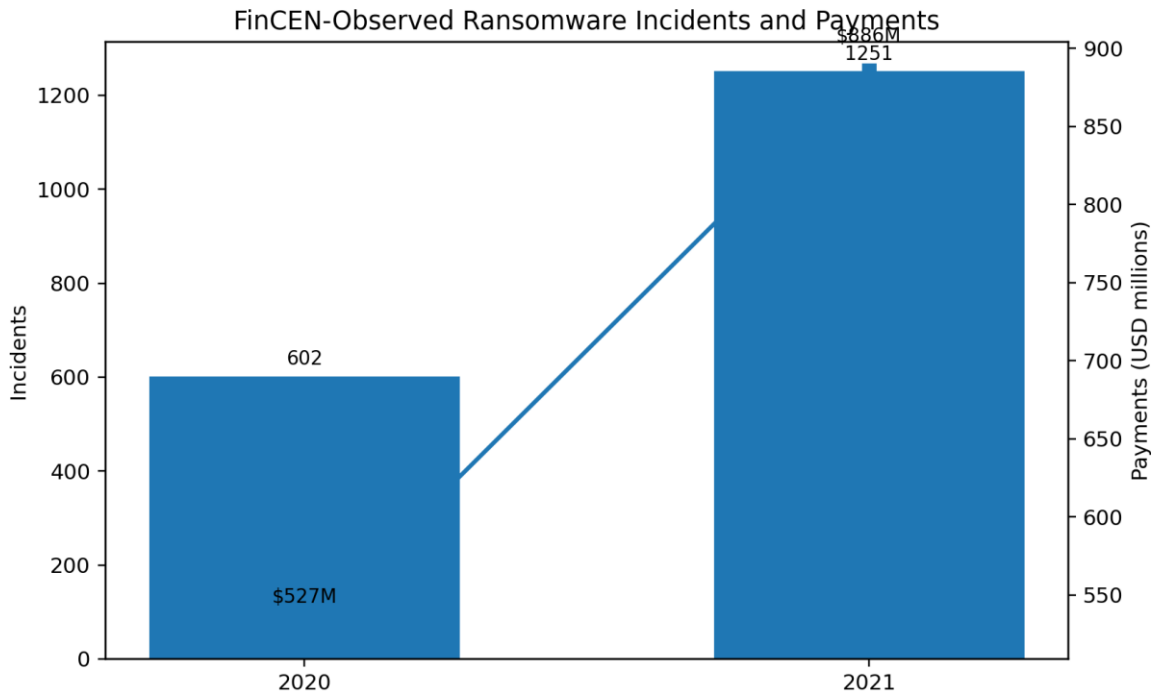

Figure 4. FinCEN-observed ransomware incidents and payments in 2020 and 2021. Source: FinCEN ransomware trend analyses.

Discussion

The central argument of this paper is that financial fraud detection, cyber risk management, and infrastructure resilience should be treated as one strategic analytic problem in U.S. banking rather than three isolated disciplines. The public evidence supports that position (FBI, 2022, 2023, 2024a, 2024b; FinCEN, 2023, 2024a, 2024b, 2024c; OCC, 2024a, 2024b; FDIC, 2024). FBI complaint and loss figures show that digitally enabled fraud has become more damaging even when complaint volume does not rise proportionally. FinCEN's reporting volumes and trend analyses show that suspicious activity moves through multiple channels, typologies, and customer segments. Banking regulators, meanwhile, continue to stress ransomware, third-party concentration, weak authentication, and service continuity. Viewed together, these signals suggest that the banking industry now faces a blended threat environment in which fraud, cyber intrusion, and operational disruption reinforce one another.

This has important consequences for model design. A narrow fraud model trained only on transaction attributes may miss the earliest stages of an incident because attackers often begin with reconnaissance, phishing, credential abuse, or identity manipulation before they attempt payments. Conversely, a cybersecurity monitoring stack that ignores customer and transaction consequences may detect technical anomalies without recognizing which incidents are most likely to create immediate financial loss. Integrated analytics create value precisely by linking these stages. A suspicious email domain, a burst of failed logins, a new device association, a beneficiary change, and an urgent wire request become far more informative when analyzed together than when routed to separate teams.

The authentic public data also clarify why prioritization matters. In 2023, the FBI reported more than $12.5 billion in losses from internet crime and more than $5.6 billion in cryptocurrency-related fraud losses, while FinCEN reported 4.6 million SAR filings in fiscal year 2023. No bank can treat every alert, complaint, or suspicious pattern as equally urgent. The true institutional question is how to distinguish signals that indicate ordinary variation from those that imply imminent monetary loss or service disruption. Machine learning is useful here not because it automates every decision, but because it can rank risk, compress search space, and help scarce investigators focus on the most consequential events (West & Bhattacharya, 2016; Ali et al., 2022; Motie et al., 2024).

The case for multimodal modeling is especially strong in banking because modern fraud rarely presents as a single anomalous transaction. Business email compromise, for example, often links social engineering, mailbox compromise, invoice tampering, beneficiary manipulation, timing pressure, and cross-border funds movement. Account takeover may combine device anomalies, geolocation inconsistency, credential resets, and a sudden sequence of high-risk payment actions. Check fraud may involve mail theft, image alteration, beneficiary substitution, and rapid deposit-and-withdrawal behavior across linked accounts. Ransomware events may begin with credential theft or third-party compromise and then cascade into payment disruption, customer service

pressure, and suspicious financial flows. In each case, the best predictive signal emerges from combining transactional, behavioral, cyber, and relationship data.

Another implication concerns graph intelligence. Much public and academic attention focuses on per-transaction classification, yet many serious banking threats are networked by design. Mule accounts distribute proceeds through many low-value transfers. Synthetic identities reuse addresses, phones, or devices across account-opening attempts. Fraud rings coordinate beneficiary accounts, merchants, domains, and communication channels. Vendor or service-provider incidents may create common-mode exposure across many clients. Graph methods can reveal these structures because they are sensitive to connection patterns rather than just record-level outliers (Akoglu et al., 2015; Van Vlasselaer et al., 2015). For institutions dealing with rising identity abuse and coordinated fraud, graph analytics are not an optional enhancement; they are increasingly foundational.

At the same time, banks should resist the temptation to overstate what machine learning can do. Publicly celebrated detection scores often fail to capture production constraints. A model that raises precision in a benchmark may still be operationally inferior if it generates unstable alert volumes, lacks explanation quality, or misses low-frequency high-loss events. In regulated banking environments, the practical test of analytic value is not a single performance metric but a bundle of outcomes: earlier detection, lower losses, acceptable customer friction, manageable investigator workload, defensible governance, and preserved service continuity. This is why the resilience lens matters. The best system is not merely the one that identifies more suspicious events; it is the one that helps the institution sustain critical operations while triaging and containing those events.

Human-machine collaboration is therefore a defining feature of any credible deployment. Investigators, fraud strategists, cyber analysts, and operational-risk officers contribute contextual judgment that models do not possess. They understand seasonal product behavior, customer communication norms, operational tolerances, and the reputational consequences of account freezes or payment delays. Conversely, machine learning contributes scale, consistency, and the ability to identify non-obvious patterns across millions of events. The goal is not to replace expert judgment but to improve it. Systems should provide intelligible evidence trails, not inscrutable risk scores. In practice, reason codes, entity-link visualizations, similar-case retrieval, and scenario-level explanations may be more valuable to banking users than abstract global feature importance alone.

The discussion also points to an institutional redesign challenge. Fraud teams, AML investigators, security operations centers, identity teams, treasury operations, and business continuity groups often operate under separate reporting lines, data systems, and escalation playbooks. This fragmentation reduces the value of predictive analytics because models can only act on the data and authority structures surrounding them. If a cyber alert cannot be linked to a payment hold decision, or if a fraud case team cannot see vendor outage data affecting authentication, the institution will continue to react in fragments. Integrated analytics should therefore be accompanied by integrated governance: shared taxonomies, joint incident reviews, unified case-management interfaces, and clear escalation thresholds tied to both fraud severity and operational criticality.

A related issue is model risk and explainability. The most powerful algorithms are not always the most deployable. Banks face legal, supervisory, and reputational pressure to explain why a payment was blocked, an account was reviewed, or a customer experienced additional authentication friction. This does not prohibit advanced models, but it does require layered controls. Interpretable baseline models, challenger testing, threshold rationales, post-hoc explanations, and exception-monitoring should all be part of the deployment strategy. Moreover, fairness concerns extend beyond lending. Fraud models can create differential customer experience if proxy variables correlate with protected characteristics or with unequal digital access. Governance must therefore include fairness diagnostics and human review processes for adverse actions.

The resilience dimension introduces another important contribution (Basel Committee on Banking Supervision, 2021; CPMI & IOSCO, 2016; OCC, 2024a). Much fraud research optimizes for detection, whereas banking resilience requires preparedness for degraded states. If a cloud identity provider fails, if email is compromised, if payment systems are throttled, or if analysts are overwhelmed by alert surges, the institution still needs a graceful operating mode. Predictive analytics should support that objective by identifying concentration points, measuring dependency criticality, and triggering preplanned operational responses. For example, a composite infrastructure-risk score could route incidents differently when they involve critical vendors, privileged accounts, or high-volume payment queues. In this sense, machine learning becomes part of an enterprise resilience toolkit rather than merely an investigative filter.

The public data used in this paper also reveal the importance of vulnerable populations and socially engineered fraud. FinCEN's elder financial exploitation analysis and the FBI's age-based loss reporting show that older adults and other targeted groups can experience disproportionate harm. This matters for banking because fraud prevention is not only a safety-and-soundness issue but also a consumer-protection issue. Predictive systems should therefore incorporate customer-protection logic such as trusted-contact escalation, behavioral anomaly checks tailored to retirement and savings drawdown behavior, and enhanced

scrutiny for abrupt changes in communication, beneficiaries, or large transfers from historically conservative accounts. When appropriately governed, AI can make banks not only more secure but also more humane.

From an implementation standpoint, a staged deployment is more realistic than a full enterprise transformation at once. Many banks still struggle with data fragmentation, inconsistent case coding, limited label quality, and separate tooling across fraud and cyber teams. Early wins are likely to come from focused use cases where signal quality is strong and losses are material, such as account takeover, new-beneficiary wire fraud, mule-account detection, check fraud, phishing-linked login compromise, and vendor-related incident prioritization. Once these pipelines mature, institutions can expand to broader multimodal architectures and graph-based entity intelligence. This staged approach reduces governance burden while building trust in the system.

The discussion further suggests that evaluation should be strategic as well as statistical. Senior leaders should ask not only whether a model improves precision or recall, but whether it changes control effectiveness at the portfolio level. Does it shorten mean time to detect fraud cascades? Does it reduce losses on high-value wires? Does it improve intervention against mule networks? Does it preserve customer access during a cyber event? Does it help triage incidents involving critical vendors or payment hubs? These questions tie predictive analytics to enterprise outcomes that boards and supervisors recognize as material.

Finally, the broader significance of this framework lies in its alignment with the structure of contemporary banking risk. U.S. banks increasingly operate as data-intensive, software-mediated, and externally interconnected institutions. The same digital modernization that enables faster payments, richer customer experience, and lower operating costs also expands the attack surface and raises the speed of contagion. In such an environment, surveillance cannot remain purely reactive and resilience cannot remain purely procedural. Institutions need predictive intelligence that is operationally grounded, cross-functional, and resilient by design. Artificial intelligence and machine learning are most valuable when they support that institutional transformation rather than when they are treated as isolated technical add-ons.

An additional benefit of the integrated framework is improved board and senior-management visibility. Fraud losses, cyber incidents, and resilience weaknesses are often reported through separate dashboards that use different severity scales and different time horizons. That fragmentation makes enterprise prioritization difficult. A unified early-warning architecture can support a common taxonomy for loss events, suspicious activity, service degradation, and dependency exposure. Executives can then see whether a surge in phishing-linked fraud coincides with authentication-control weakness, whether a third-party outage is increasing fraud attempts, or whether review queues are growing fast enough to threaten timely response. In governance terms, that creates a more decision-useful picture of operational risk than siloed scorecards.

The framework also has implications for community and regional banks. Large institutions may be better positioned to build graph infrastructure, streaming feature stores, and dedicated model-risk teams, but smaller banks are often disproportionately exposed to fraud and vendor concentration because they rely more heavily on third-party platforms. For these institutions, the most practical version of predictive analytics may be a federated or managed-service model built around shared typologies, institution-specific thresholds, and strong escalation rules. The literature suggests that sophistication should be proportional to institutional complexity, but the need for integrated early warning is not limited to globally systemic firms. Even smaller institutions need better linkage among fraud operations, information security, and contingency planning.

There is also a broader policy implication concerning information sharing. Many banking threats are cross-institutional: mule accounts move funds across multiple banks, phishing infrastructure targets many institutions simultaneously, and third-party provider incidents can generate sector-wide exposure. FinCEN's 314(b) program, public-private exchanges, and supervisory coordination mechanisms are therefore not peripheral to predictive analytics; they are part of the data ecosystem that makes better prediction possible (FinCEN, 2023, 2024a). Future analytic maturity in banking will depend not only on internal model design but also on how effectively institutions can absorb and operationalize external intelligence while respecting privacy, confidentiality, and legal constraints.

Finally, the findings support a more balanced narrative about artificial intelligence in banking. The strategic value of AI does not rest on replacing judgment or promising perfect foresight. It rests on making institutions less surprised, less fragmented, and less brittle. In practical terms, that means earlier recognition of suspicious patterns, sharper prioritization of scarce investigative resources, faster coordination across control functions, and better preservation of critical services when incidents occur. When evaluated through that lens, predictive machine learning becomes a core component of modern banking resilience rather than a stand-alone technical experiment. Equally important, integrated analytics can strengthen institutional learning after incidents by preserving evidence across fraud, cyber, legal, and operations teams and by making root-cause analysis faster, more consistent, and more actionable for future control design. This strengthens preparedness before losses accumulate materially. It improves readiness materially.

Table 3

Model layer	Illustrative methods	Primary output	Governance expectation
Rules layer	Lists, policy thresholds, typology rules	Mandatory alerts and blocks	Policy sign-off and audit trail
Supervised layer	Logistic regression, XGBoost, random forest	Probability of fraud or suspicious activity	Validation, calibration, fairness review
Anomaly layer	Isolation forest, autoencoder, robust outlier scoring	Novel or rare behavior alerts	Triage design and analyst feedback
Graph layer	Community detection, link prediction, GNN scoring	Entity and network risk	Entity resolution controls and explainability
Text layer	NLP embeddings, keyword extraction, similarity search	Typology and narrative enrichment	Sensitive-data handling and retention rules
Resilience layer	Criticality scoring, dependency mapping, degradation thresholds	Operational escalation priority	Board visibility and continuity linkage

Table 3. Layered model stack and associated governance expectations in a regulated banking environment.

Hasan et al. (2023) sharpen an implication that deserves more explicit treatment in the U.S. banking sector, fraud detection cannot be understood as an isolated analytics use case because the same digital pathways that enable transactional abuse can also impair cybersecurity posture and critical-service continuity. Their article is not simply another call for more machine learning in finance. It advances a sector-level argument that fraud analytics, cyber-risk management, and infrastructure protection are mutually reinforcing functions. That insight matters because many banking institutions still organize these activities into different reporting lines, data estates, and budget processes. Fraud operations may focus on card abuse, deposits, wires, and account takeover. Information-security teams may focus on endpoint protection, identity governance, privileged access, or vendor risk. Business-continuity teams may focus on recovery tolerances and crisis response. Each function can become effective within its own scope while still missing the cross-domain escalation path of a real incident. A social-engineering campaign may begin as a suspicious email, proceed through compromised credentials, manifest as anomalous login patterns, and then culminate in unusual payment behavior or the rapid movement of funds through newly linked accounts. In other words, the institution experiences one event while its control architecture sees several small events. The contribution of Hasan et al. (2023) is that they make the strategic case for reconnecting those signals through AI-driven risk analytics.

The reporting-integrity evidence in Pritty et al. (2024) suggests that suspicious behavior in finance is frequently narrative as well as numerical: managers, fraudsters, or counterparties may manipulate language, timing, and disclosure framing before balance-sheet consequences are fully visible. That insight supports the inclusion of document intelligence, communication analytics, and explainable natural-language processing within banking fraud and cyber-risk systems. Likewise, the real-time payments framework in Fahim et al. (2024) underscores that model latency is itself a control variable in modern fraud prevention, because intervention delay can erase the practical value of otherwise accurate predictions.

The governance and resilience implications are equally important. Fahim et al. (2023) framed algorithmic accountability as a prerequisite for trust, fair lending, and financial stability, which aligns closely with this paper's recommendation that banks maintain challenger models, reason codes, escalation controls, and human review capacity. Pritty et al. (2024) further suggest that fraud risk increasingly spans structured and unstructured domains, implying that resilient banking infrastructure depends on the ability to fuse transaction, identity, device, and narrative signals into a coherent monitoring framework.

That perspective is highly consistent with the empirical patterns summarized earlier in this paper. The FBI's IC3 reports show not only that reported losses have increased markedly, but also that many financially harmful crimes are cyber-enabled by design. Business email compromise illustrates the point well. It is usually described as a fraud problem because the immediate loss occurs through payment manipulation, yet its enabling mechanics often involve compromised identities, fraudulent domains, mailbox access, social engineering, or vendor impersonation. The same is true of account takeover and many cryptocurrency-related scams, which frequently depend on credential theft, remote access, phishing, or other cyber tactics before they appear as suspicious financial activity. Thus, when banks invest in machine learning solely to classify suspect transactions after the fact, they risk placing analytics too late in the incident lifecycle. A resilience-oriented design would instead ask how early warning can be moved upstream by linking identity telemetry, device changes, behavioral anomalies, alert history, and payment-path irregularities in near real time.

A second insight from Hasan et al. (2023) is the importance of translating analytic outputs into risk-management actions that are proportionate, explainable, and operationally sustainable. In practice, early warning has little value if it produces more alerts than a bank can triage, or if the recommended interventions are so blunt that they create unnecessary customer friction. The strongest contribution of predictive analytics is not perfect prediction. It is the disciplined reallocation of scarce investigative attention toward the combinations of signals most associated with material harm. This is especially important in regulated banking environments, where high-volume surveillance systems must coexist with consumer-protection expectations, model-governance requirements, and practical staffing constraints. A useful risk score is one that improves prioritization, guides escalation, and supports intervention choices with clear rationale. For example, a medium-confidence anomaly may justify step-up authentication or enhanced review, whereas a high-confidence, high-severity pattern involving new payees, identity changes, suspicious geolocation shifts, and known external threat indicators may justify a temporary hold, rapid case routing, and cross-functional incident review.

The infrastructure angle in Hasan et al. (2023) also helps elevate the significance of fraud analytics beyond narrow loss prevention. U.S. banking depends on the stable functioning of payment systems, correspondent relationships, customer-authentication services, cloud infrastructure, telecommunications, and third-party software ecosystems. Even when an individual fraud event is modest, repeated exploitation of shared channels can degrade confidence, increase manual controls, and expose systemic weak points. As a result, machine learning should be evaluated partly on how it supports institutional resilience: whether it shortens detection latency, improves visibility across channels, reduces repeat victimization, and strengthens coordination across fraud, cybersecurity, and continuity functions. This wider frame is valuable for scholarship because it grounds AI not in technological optimism, but in public-interest outcomes such as service reliability, customer protection, and confidence in banking infrastructure. The present study adopts that logic by treating fraud, cyber risk, and resilience as interdependent supervisory and managerial concerns rather than as isolated technical silos.

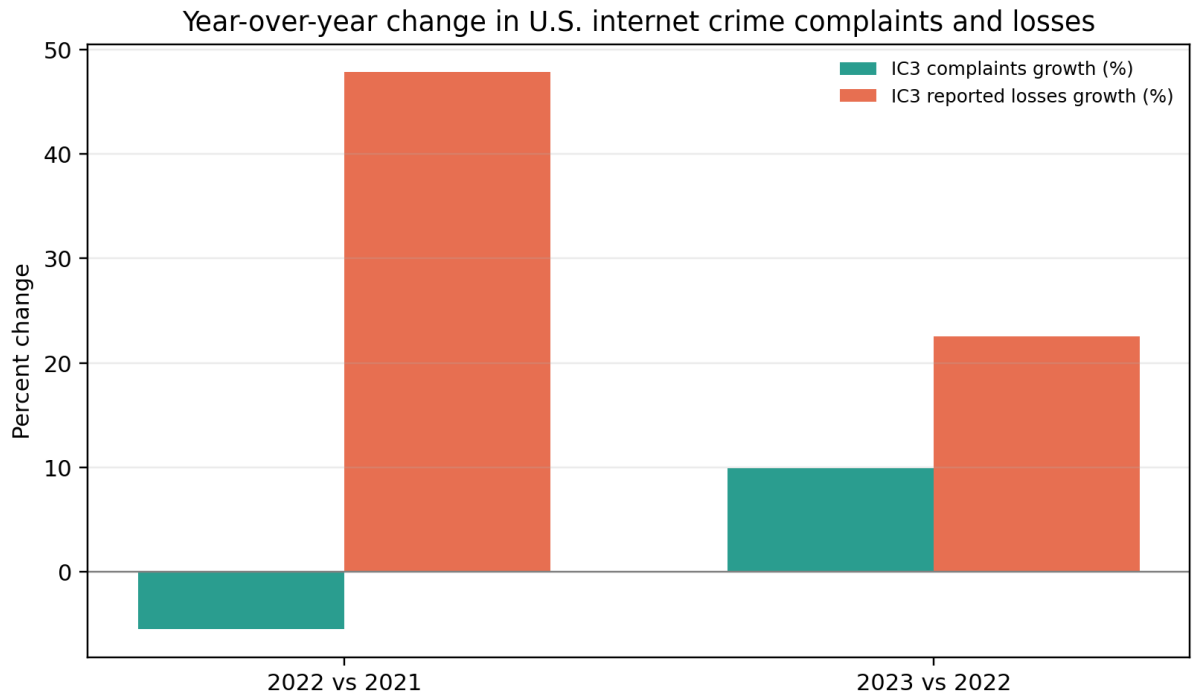

Figure 5. Year-over-year percentage change in FBI IC3 complaints and reported losses. Derived from FBI Internet Crime Reports 2021-2023.

A. Expanded Empirical Interpretation of Public Banking Threat Indicators

The public indicators used in this paper are aggregate rather than transaction level, but they still reveal several operationally meaningful patterns. The first is that loss intensity has risen faster than complaint counts. Between 2021 and 2022, reported IC3 complaint volume declined by roughly 5.5 percent, yet reported losses increased by nearly 47.8 percent. Between 2022 and 2023, complaint volume rose by roughly 9.9 percent while losses increased by another 22.5 percent. This divergence matters because it suggests that threat severity cannot be inferred from frequency alone. Institutions that focus narrowly on event counts may underestimate the degree to which scams, account compromise, and digitally enabled fraud have become more financially concentrated. In surveillance design, this implies that labels and outcomes should reflect loss severity, payment velocity,

customer vulnerability, and propagation potential rather than treating all alerts as equal. A system optimized only for count reduction may miss exactly the events that generate outsized harm.

A related way to view the same data is to examine reported losses per complaint. Using the FBI aggregate series, the implied average reported loss per complaint rose from approximately 8,143 dollars in 2021 to about 12,735 dollars in 2022 and approximately 14,198 dollars in 2023. This measure does not represent a bank-level portfolio loss rate, and it should not be treated as a direct underwriting or capital metric. Nonetheless, it is analytically useful because it captures how economically damaging the observable complaint environment has become. For banks, a higher loss-per-complaint environment increases the value of detection systems that can separate low-consequence anomalies from high-harm incidents. It also supports tiered response architectures in which payment screening, account monitoring, and customer outreach become more intensive as combinations of severity indicators accumulate. When financial harm is rising faster than event counts, precision in identifying severe cases becomes strategically more valuable than simply producing more alerts.

The SAR data add another dimension. FinCEN's year-in-review materials indicate that suspicious activity reporting remained extremely large, increasing from approximately 4.3 million reports in fiscal year 2022 to approximately 4.6 million in fiscal year 2023. This near 7.0 percent increase signals a financial system in which institutions are generating immense volumes of potentially relevant intelligence, but it also underscores a classic bottleneck: volume can overwhelm interpretability. A bank may have access to millions of transactional, behavioral, identity, and case-management records, yet still struggle to determine which combinations truly merit urgent action. This is where machine learning can create value not by replacing institutional judgment, but by helping rank, cluster, and contextualize suspicious activity. Aggregate growth in reporting volume means that triage quality becomes as important as detection sensitivity. The objective is not to treat more data as automatically better data; it is to convert data abundance into risk insight.

Another important implication of the public series is the need for multi-horizon surveillance. Complaint data, SAR volume, cyber-risk reporting, and resilience guidance each move on different time scales. Some events unfold in minutes, such as a rapid account takeover followed by out-of-pattern transfers. Others unfold over weeks, such as an extended compromise of vendor email, slow credential harvesting, or coordinated mule-account buildup. A high-performing banking analytics program therefore needs short-horizon controls for immediate payment and identity risk, medium-horizon models for campaign or network escalation, and longer-horizon dashboards that help executives identify structural concentrations in channels, products, geographies, or third-party dependencies. By integrating these horizons, institutions can avoid the false choice between real-time prevention and strategic resilience planning. Both are required, but they must be fed by compatible data models and common risk taxonomies.

The indexed escalation view added in this paper also clarifies why banks should track not only absolute levels but relative acceleration. When the complaint series, the loss series, and the suspicious-activity series are converted into indices, the losses line separates sharply from the complaints line. That gap is a simple but powerful visual reminder that adversaries can become more economically effective even if raw incident volumes remain manageable. In control design, this justifies asymmetric treatment of severe-loss typologies, faster escalation thresholds for high-harm indicators, and stronger board reporting around severity-adjusted outcomes. A board member may not need to review every model metric, but leadership should understand whether the institution is operating in a threat environment where harm per incident is rising, where detection latency is falling or increasing, and where key defenses are becoming more or less resilient over time.

These public data patterns also argue for a stronger connection between financial-crime analytics and customer-protection strategy. Rising losses per complaint are not only an operational concern for banks; they signal broader consumer vulnerability, especially among customers exposed to impersonation schemes, relationship fraud, elder exploitation, and digitally mediated investment scams. Fraud detection models that ignore customer context may be statistically elegant but strategically incomplete. Features related to account age, channel migration, prior victimization, atypical beneficiary creation, or abrupt changes in communication patterns can help institutions identify customers facing elevated social-engineering risk. Used responsibly, such signals can support targeted interventions such as out-of-band confirmation, enhanced education, friction at key payment moments, or faster recovery escalation. In this sense, predictive analytics can support both institutional efficiency and public protection without reducing either objective to a simplistic loss-avoidance exercise.

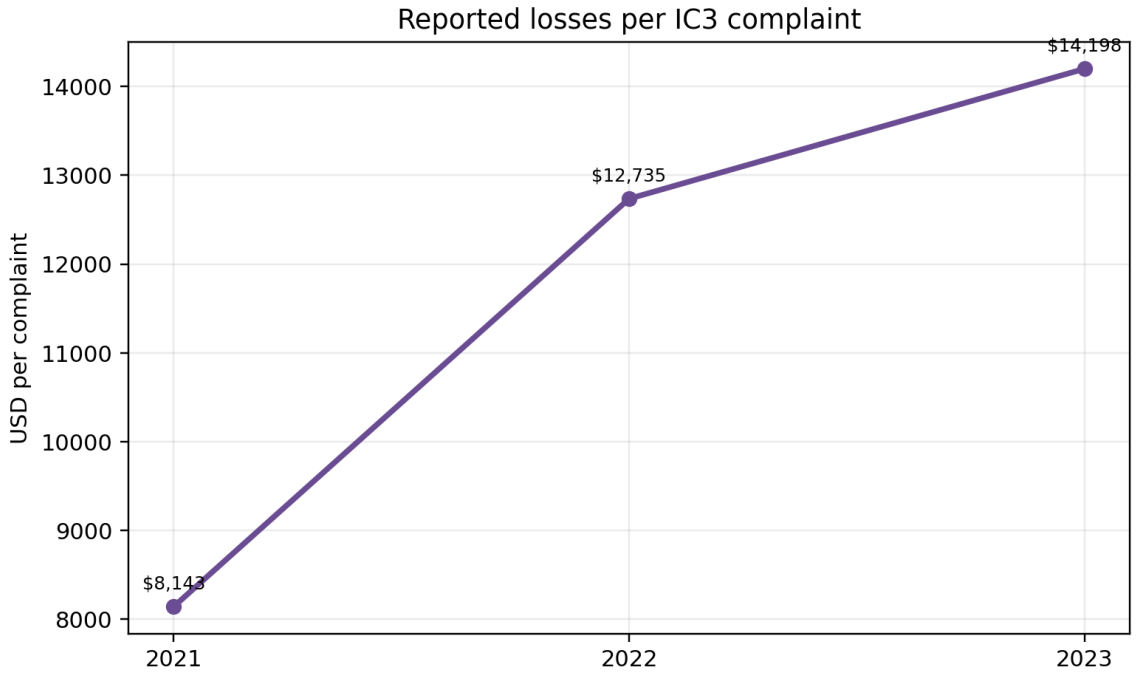

Figure 6. Implied reported losses per FBI IC3 complaint, 2021-2023. Derived from FBI Internet Crime Reports 2021-2023.

B. Governance, Measurement, and an Expanded Deployment Roadmap

For banks seeking to operationalize these findings, the deployment roadmap should begin with governance discipline rather than model complexity. The first requirement is a common event ontology that can connect fraud, cybersecurity, and resilience data without forcing every function into the same case-management process. Institutions need consistent identifiers for customers, accounts, devices, credentials, employees, vendors, external indicators, and recovery-critical services. Without that foundation, graph learning and multimodal modeling remain brittle because the same underlying event may appear under different labels across fraud, AML, security, and operational-risk systems. A mature ontology also allows banks to distinguish between analytic labels used for model training and management labels used for escalation, regulatory reporting, and board oversight.

The second requirement is staged model deployment. A common mistake in enterprise AI programs is attempting to move immediately to highly complex architectures without first stabilizing data quality, feedback loops, and operational metrics. In most banking environments, the practical sequence is to begin with targeted use cases where outcomes are meaningful and intervention options are clear: account takeover risk scoring, outbound payment anomaly detection, mule-account network identification, phishing-related authentication anomalies, or suspicious beneficiary creation. Once those pipelines demonstrate measurable gains, the institution can layer graph features, textual signals, and cross-domain correlations. This staged approach reduces implementation risk, allows internal model-risk teams to build confidence gradually, and creates opportunities to document governance lessons before the platform becomes mission-critical.

A third requirement is a metric system that explicitly links analytical performance to resilience outcomes. Traditional model evaluation in fraud analytics often emphasizes precision, recall, area under the curve, and false-positive rates. Those metrics remain necessary, but banking leadership also needs measures such as median detection latency, escalation-to-resolution time, high-severity alert capture rate, repeat-victimization reduction, recovery time for disrupted payment functions, third-party incident concentration, and analyst workload per confirmed case. These measures help translate a technical model into managerial language. They also align more closely with the argument of Hasan et al. (2023), who emphasize strengthening financial and cybersecurity infrastructure rather than merely improving isolated classification performance. In a regulated environment, the best model is not necessarily the one with the highest offline score; it is the one that improves control effectiveness while remaining understandable, governable, and robust under stress.

Board reporting should therefore be redesigned around a limited set of integrated indicators. A practical dashboard might include the volume of high-risk alerts prevented before completion, the share of severe cases detected within defined time windows, loss avoided or contained, key channel-level incident concentrations, unresolved third-party dependencies, and trends in vulnerability among affected customer segments. The point is not to inundate directors with model statistics. It is to give them

an intelligible view of whether the institution's analytic capability is becoming more predictive, more coordinated, and more resilient. When directors can see the relationship among fraud pressure, cyber events, and service continuity, governance moves from reactive review to proactive oversight.

Community and mid-sized banks deserve special mention in this deployment roadmap. They may not possess the engineering capacity of the largest national institutions, but they can still implement meaningful machine-learning improvements if they adopt modular priorities. For these banks, the most sensible path may be to combine vendor-supported anomaly detection with internally curated escalation rules, selective graph features built around beneficiary and device relationships, and stronger cross-functional review of severe incidents. Shared services, consortium intelligence, and managed detection partnerships can extend capability without requiring each institution to build the entire analytics stack alone. What matters most is not whether every bank uses the same algorithmic depth, but whether the institution's controls are integrated enough to detect multi-channel abuse before it becomes a full operational event.

An expanded roadmap also requires explicit challenge processes. Models built for fraud or cyber risk can drift, encode historical biases, or become overfit to the last major scam typology. Banks should therefore implement periodic challenger models, scenario exercises, and threat-led analytic reviews. Investigators should be able to flag where a model is consistently missing emerging behaviors or creating unproductive friction. Security teams should be able to show how new attack methods alter feature relevance. Business continuity leaders should be able to identify whether analytic blind spots could delay recovery from a major disruption. This challenge culture is essential because adversaries are adaptive. A model that performs well against yesterday's fraud campaign may be dangerously reassuring against tomorrow's hybrid fraud-cyber event.

Finally, the revised analysis underscores that ethical and supervisory credibility are part of resilience. Banking institutions that deploy machine learning without clear human accountability, appeals processes, customer-impact testing, or privacy controls may reduce one category of risk while creating another. Explainability is not simply a compliance burden; it is part of operational trust. Investigators must understand why a case is ranked as severe. Customers affected by protective friction should encounter proportionate and defensible interventions. Supervisors should be able to trace how model outputs influence decisions. By embedding explainability, governance, and human review into the early-warning architecture, banks can pursue the efficiency gains of AI while preserving institutional legitimacy. That balance is central to long-term resilience because effective infrastructures are not only technologically capable; they are trusted, governable, and socially sustainable.

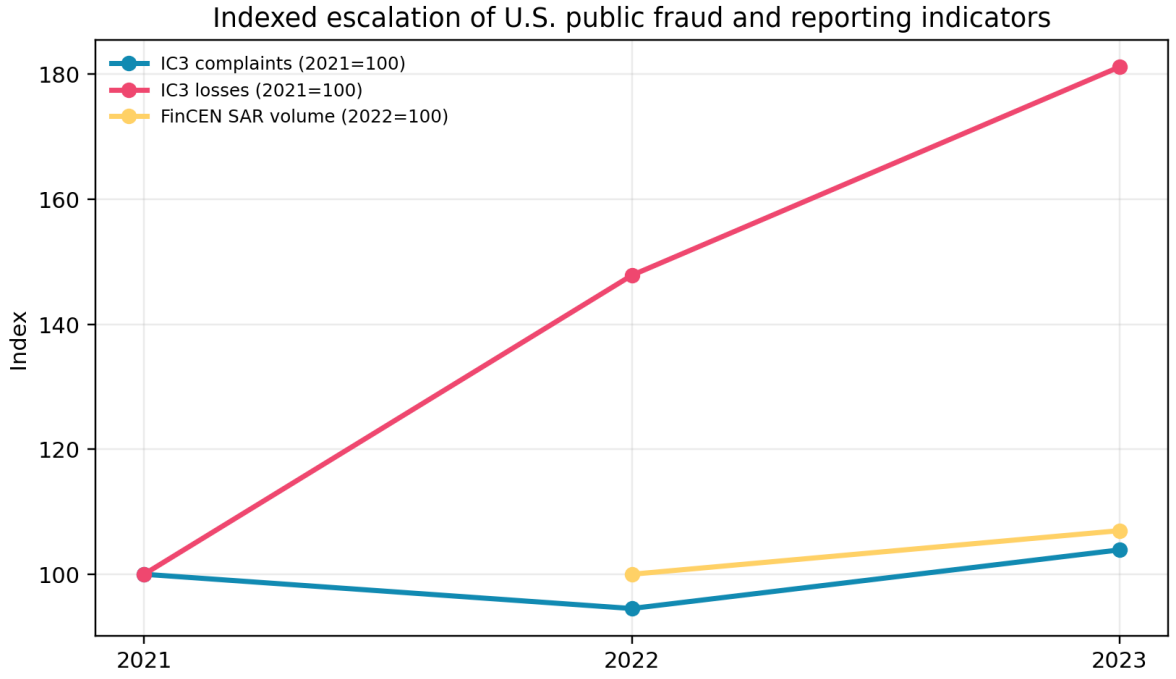

Figure 7. Indexed escalation of selected public fraud and reporting indicators. IC3 complaints and losses are indexed to 2021 = 100; FinCEN SAR volume is indexed to 2022 = 100.

Table 4

Table 4. Derived empirical indicators from FBI IC3 and FinCEN aggregate public data.

Indicator	2021	2022	2023	Source / note
IC3 complaints	847,376	800,944	880,418	FBI Internet Crime Reports
IC3 reported losses (USD billions)	6.9	10.2	12.5	FBI Internet Crime Reports
Implied loss per complaint (USD)	8,143	12,735	14,198	Derived from FBI complaint and loss totals
FinCEN SAR volume (millions)	-	4.3	4.6	FinCEN FY2022 and FY2023 year-in-review reports

Conclusion

This paper developed an integrated predictive analytics framework for financial fraud detection, cyber risk management, and infrastructure resilience in the U.S. banking industry. Drawing on public evidence through 2024 from the FBI, FinCEN, OCC, FDIC, Federal Reserve, NIST, and the academic literature, the study showed that digital fraud, cyber intrusion, and operational disruption increasingly arise from the same data, channels, and institutional dependencies. The review demonstrated why rules-only monitoring is insufficient and why banks need multimodal architectures that combine transaction analytics, identity telemetry, relationship graphs, text intelligence, and resilience indicators.

The proposed methodology contributes a realistic implementation blueprint rather than a fabricated proprietary back-test. Its central insight is that banking surveillance should move from siloed detection toward coordinated early warning. Banks that integrate fraud, cyber, and resilience signals can improve prioritization, reduce blind spots, strengthen investigator effectiveness, and better protect critical services under stress. The strongest strategic case for AI in banking is therefore not automation for its own sake, but safer, faster, and more resilient decision-making in an increasingly adversarial financial system. It also highlights why public data, while limited, can still guide credible model architecture and governance choices. Such alignment is essential for sustainable supervisory confidence and practical alignment

Limitations and Future Directions

This study has several limitations. First, the public data used to motivate the framework are authentic but highly aggregated. FBI complaint data, FinCEN reporting statistics, and supervisory materials reveal scale and typology, yet they do not provide the transaction-level, account-level, device-level, and case-level records required for direct model training or institution-specific performance testing. Second, public fraud statistics reflect underreporting, delayed discovery, inconsistent victim awareness, and changing reporting practices. Third, several important threat domains, such as insider abuse, third-party compromise, and failed-but-intercepted attacks, are imperfectly represented in public aggregate sources. Fourth, the proposed architecture is designed for U.S. banking institutions, so implementation details may differ across smaller community banks, large money-center banks, fintech-bank partnerships, and cross-border organizations with different data environments and supervisory expectations.

Future research should therefore move in three directions. The first is secure empirical collaboration using de-identified or privacy-preserving bank data to test multimodal models under realistic operational constraints. The second is deeper work on dynamic graph learning, multimodal fusion, and investigator-centered explainability for blended fraud and cyber incidents. The third is resilience-aware evaluation: future studies should measure not only loss avoidance but also service continuity, containment speed, concentration risk, and recovery performance for critical banking functions. Research that links predictive analytics to operational resilience outcomes would materially strengthen both academic knowledge and supervisory practice.

Another priority is the development of sector-level benchmarks that allow banks to compare alert quality, containment speed, and resilience outcomes without disclosing sensitive customer or institution-specific information. Better benchmark design would help distinguish genuine analytic improvement from performance gains that are only artifacts of local data conditions. Future work should also examine how smaller banks can adopt shared utilities, consortium intelligence, and privacy-preserving analytics without losing local context or over-relying on third-party vendors for critical surveillance judgments. Shared testing protocols would be valuable.

References

[1]. Abdallah, A., Maarof, M. A., & Zainal, A. (2016). Fraud detection system: A survey. Journal of Network and Computer Applications, 68, 90-113.

[2]. Akoglu, L., Tong, H., & Koutra, D. (2015). Graph based anomaly detection and description: A survey. Data Mining and Knowledge Discovery, 29(3), 626-688.

[3]. Ali, A., Qadir, J., Rasool, R. U., Sathiaseelan, A., Zwitter, A., & Crowcroft, J. (2022). Financial fraud detection based on machine learning: A systematic literature review. Applied Sciences, 12(19), 9637.

[4]. Bahnsen, A. C., Aouada, D., & Ottersten, B. (2016). Example dependent cost sensitive logistic regression for credit card fraud detection. Expert Systems with Applications, 42(16), 6070-6084.

[5]. Basel Committee on Banking Supervision. (2021). Principles for operational resilience.

[6]. Board of Governors of the Federal Reserve System. (2024a). Cybersecurity and financial system resilience report.

[7]. Board of Governors of the Federal Reserve System. (2024b). Supervision and regulation report.

[8]. Bouveret, A. (2018). Cyber risk for the financial sector: A framework for quantitative assessment. IMF Working Paper.

[9]. Carcillo, F., Le Borgne, Y.-A., Caelen, O., Kessaci, Y., Oblé, F., & Bontempi, G. (2021). Combining unsupervised and supervised learning in credit card fraud detection. Information Sciences, 557, 317-331.

[10]. Chalapathy, R., & Chawla, S. (2019). Deep learning for anomaly detection: A survey. arXiv preprint arXiv:1901.03407.

[11]. Conference of State Bank Supervisors. (2023). Annual survey of community banks.

[12]. Committee on Payments and Market Infrastructures & International Organization of Securities Commissions. (2016). Guidance on cyber resilience for financial market infrastructures.

[13]. Dal Pozzolo, A., Caelen, O., Le Borgne, Y.-A., Waterschoot, S., & Bontempi, G. (2014). Learned lessons in credit card fraud detection from a practitioner perspective. Expert Systems with Applications, 41(10), 4915-4928.

[14]. Eling, M., & Schnell, W. (2016). What do we know about cyber risk and cyber risk insurance? The Journal of Risk Finance, 17(5), 474-491.

[15]. Federal Bureau of Investigation. (2022). Internet crime report 2021.

[16]. Federal Bureau of Investigation. (2022). FY 2022 congressional report: Business email compromise and real estate wire fraud.

[17]. Federal Bureau of Investigation. (2023). Internet crime report 2022.

[18]. Federal Bureau of Investigation. (2024a). Internet crime report 2023.

[19]. Federal Bureau of Investigation. (2024b). Cryptocurrency fraud report 2023.

[20]. Federal Deposit Insurance Corporation. (2024). 2024 risk review.

[21]. Federal Financial Institutions Examination Council. (2017). Cybersecurity assessment tool.

[22]. Federal Financial Institutions Examination Council. (2022). Cybersecurity resource guide for financial institutions.

[23]. Financial Crimes Enforcement Network. (2021). Ransomware trends in Bank Secrecy Act data between January 2021 and June 2021.

[24]. Financial Crimes Enforcement Network. (2022). Ransomware trends in Bank Secrecy Act data between July 2021 and December 2021.

[25]. Financial Crimes Enforcement Network. (2023). Year in review for FY 2022.

[26]. Financial Crimes Enforcement Network. (2023). Alert on nationwide surge in mail theft-related check fraud schemes targeting the U.S. mail.

[27]. Financial Crimes Enforcement Network. (2023). Identity-related suspicious activity: 2021 threats and trends.

[28]. Financial Crimes Enforcement Network. (2024a). Year in review for FY 2023.

[29]. Financial Crimes Enforcement Network. (2024b). Elder financial exploitation: Threat pattern and trend information, June 2022 to June 2023.

[30]. Financial Crimes Enforcement Network. (2024c). Mail theft-related check fraud: Threat pattern and trend information, February to August 2023.

[31]. Gordon, L. A., & Loeb, M. P. (2002). The economics of information security investment. ACM Transactions on Information and System Security, 5(4), 438-457.

[32]. Jurgovsky, J., Granitzer, M., Ziegler, K., Calabretto, S., Portier, P.-E., He-Guelton, L., & Caelen, O. (2018). Sequence classification for credit-card fraud detection. Expert Systems with Applications, 100, 234-245.

[33]. Kopp, E., Kaffenberger, L., & Wilson, C. (2017). Cyber risk, market failures, and financial stability. IMF Working Paper.

[34]. Kou, Y., Lu, C.-T., Sirwongwattana, S., & Huang, Y.-P. (2004). Survey of fraud detection techniques. IEEE International Conference on Networking, Sensing and Control, 749-754.

[35]. Le Khac, N. A., Markos, S., Kechadi, T., & Le-Khac, N. (2020). The rise of machine learning in financial fraud detection. Journal of Financial Crime, 27(3), 719-733.

[36]. Motie, S., Ghasemian, A., Sadeghi, M., & Jalili, M. (2024). Financial fraud detection using graph neural networks. Expert Systems with Applications, 237, 121062.

[37]. National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0).
[38]. National Institute of Standards and Technology. (2024). Cybersecurity Framework (CSF) 2.0.
[39]. Ngai, E. W. T., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2011). The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature. Decision Support Systems, 50(3), 559-569.
[40]. Office of the Comptroller of the Currency. (2024a). Cybersecurity and financial system resilience report.
[41]. Office of the Comptroller of the Currency. (2024b). Semiannual risk perspective, spring 2024.
[42]. Office of the Comptroller of the Currency. (2024c). Annual report 2024.
[43]. Perols, J. (2011). Financial statement fraud detection: An analysis of statistical and machine learning algorithms. Auditing: A Journal of Practice & Theory, 30(2), 19-50.
[44]. Romanosky, S. (2016). Examining the costs and causes of cyber incidents. Journal of Cybersecurity, 2(2), 121-135.
[45]. Ruff, L., Kauffmann, J. R., Vandermeulen, R. A., Montavon, G., Samek, W., Kloft, M., Dietterich, T. G., & Müller, K.-R. (2021). A unifying review of deep and shallow anomaly detection. Proceedings of the IEEE, 109(5), 756-795.
[46]. Van Vlasselaer, V., Akoglu, L., Snoeck, M., Baesens, B., & Vanthienen, J. (2015). APATE: A novel approach for automated credit card transaction fraud detection using network based extensions. Decision Support Systems, 75, 38-48.
[47]. Verizon. (2023). Data breach investigations report.
[48]. West, J., & Bhattacharya, M. (2016). Intelligent financial fraud detection: A comprehensive review. Computers & Security, 57, 47-66.
[49]. Whitrow, C., Hand, D. J., Juszczak, P., Weston, D., & Adams, N. M. (2009). Transaction aggregation as a strategy for credit card fraud detection. Data Mining and Knowledge Discovery, 18(1), 30-55.
[50]. Hasan, M. N., Rasel, I. H., Arman, M., Ibrahim, M., & Jahan, N. (2023). Strengthening U.S. financial and cybersecurity infrastructure with AI-driven fraud detection and risk analytics. Journal of Computational Analysis and Applications, 31(2), 15-32. Retrieved from eudoxuspress.com/index.php/pub/article/view/3823
[51]. Fahim, A. S. M., Ibrahim, M., Pritty, A. A., & Tania, T. A. (2023). Algorithmic accountability in U.S. consumer FinTech: Governance mechanisms for credit risk, fair lending, and financial stability. Journal of Economics, Finance and Accounting Studies, 5(4), 80–93. https://doi.org/10.32996/jefas.2023.5.4.8
[52]. Fahim, A. S. M., Pritty, A. A., Ibrahim, M., & Tania, T. A. (2024). Real-time payments and real-time fraud: A U.S. FinTech risk framework for RTP rails and consumer protection. Journal of Economics, Finance and Accounting Studies, 6(6), 134–149. https://doi.org/10.32996/jefas.2024.6.6.11
[53]. Pritty, A. A., Ibrahim, M., Fahim, A. S. M., & Zadid, M. U. (2024). Generative AI and U.S. financial reporting integrity: Detecting narrative manipulation, risk disclosure gaming, and fraud signals in 10-K filings. Journal of Economics, Finance and Accounting Studies, 6(4), 113–129. https://doi.org/10.32996/jefas.2024.6.4.11