**Journal of Economics, Finance and Accounting Studies (JEFAS)**
ISSN: 2709-0809
DOI: 10.32996/jefas
Journal Homepage: www.al-kindipublisher.com/index.php/jefas

JEFAS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

**| RESEARCH ARTICLE**

# Intelligent Pay-by-Bank Architecture: Fraud-Aware Analytics and Risk Modeling for U.S. Account-to-Account Payment Systems

**Md Nurul Huda Razib**
Assistant Professor, Manarat International University
Email: mnhrazib777@manarat.ac.bd
ORCID: https://orcid.org /0009-0000-4010-7636

**Mohammad Mamun Ur Rashid**
Associate Professor, Manarat International University
Email: mamunrashid.dba@manarat.ac.bd
ORCID: https://orcid.org/0000-0002-8067-3621

**Muhammad Helal Uddin**
Assistant Professor, Manarat International University
Email: helal@manarat.ac.bd
ORCID: https://orcid.org / 0009-0009-2216-005X

**Corresponding Author**: Md Nurul Huda Razib, **E-mail**: mnhrazib777@manarat.ac.bd

**| ABSTRACT**

Pay-by-bank is moving from a niche alternative to a strategically important layer in the U.S. payments landscape. Its appeal is clear: direct account-to-account (A2A) transfers can reduce merchant acceptance costs, accelerate funds availability, support recurring and bill-payment use cases, and create a pathway for open-banking-enabled checkout experiences that do not depend on card rails. At the same time, the same structural strengths that make pay-by-bank attractive—irrevocability, speed, richer data exchange, API connectivity, and multi-party orchestration—also create distinctive fraud, operational, and consumer-protection risks. Unlike card networks, many pay-by-bank implementations cannot rely on mature chargeback mechanisms as a primary safety valve, which makes ex ante risk scoring, identity assurance, behavioral analytics, and post-transaction monitoring more important. This paper develops a framework for an intelligent pay-by-bank architecture for the United States, integrating fraud-aware analytics, explainable machine learning, graph-based entity resolution, sanctions and AML screening, and rail-aware control design across ACH, Same Day ACH, RTP, and FedNow-enabled flows. Using official U.S. sources and contemporary scholarly research, the paper synthesizes the structural drivers of A2A growth, maps the major fraud typologies that threaten consumer and merchant adoption, and proposes a layered analytic design that joins customer risk, payment context, device and channel telemetry, open-banking consent events, account validation, and network intelligence into a unified decision stack. The paper also advances a practical governance model centered on explainability, model-risk management, fairness, vendor oversight, and escalation workflows aligned to high-stakes payment decisions. The discussion argues that successful U.S. pay-by-bank adoption depends not only on rail availability but on trusted orchestration: institutions must combine real-time controls, adaptive thresholds, human review for ambiguous cases, and continuous learning from returns, disputes, fraud reports, and customer complaints. The manuscript concludes that intelligent pay-by-bank systems can materially improve the safety and efficiency of U.S. A2A payments when technical architecture, controls, and governance are designed together rather than sequentially.

**| KEYWORDS**

Pay-by-bank, account-to-account payments, ACH, Same Day ACH, FedNow, fraud analytics, explainable AI, open banking, synthetic identity, U.S. payment systems

**| ARTICLE INFORMATION**

**ACCEPTED:** 01 March 2026 **PUBLISHED:** 18 March 2026 **DOI:** 10.32996/jefas.2026.8.4.4

## 1. Introduction

The U.S. payments system is in the middle of a structural transition from batch-oriented, card-dominant retail payments to a more diverse environment in which instant payments, API-mediated account access, and data-rich payment initiation models increasingly coexist.

Within that transition, pay-by-bank has emerged as an important design pattern rather than a single product.

In practice, pay-by-bank refers to checkout, bill-pay, or transfer experiences in which a consumer authorizes a payment directly from a deposit account to a merchant or payee account, usually through ACH, Same Day ACH, RTP, or FedNow-connected processes and often with the help of open-banking-style data access, account verification, or payment initiation tools. The appeal is not merely technological.

For merchants, pay-by-bank promises lower acceptance costs than many card transactions, faster access to funds in some configurations, and reduced exposure to certain card-specific fraud vectors.

For consumers, it can support simplified checkout, direct-from-bank transparency, and potentially lower prices if merchants share part of the cost savings.

For financial institutions and infrastructure providers, it expands the universe of account-based payment use cases and creates a platform for new value-added services around validation, risk scoring, exception handling, and financial-data portability.

The policy and infrastructure backdrop helps explain why this conversation has intensified.

The Federal Reserve's payments data show continued growth in electronic and account-based payment activity, while Nacha reported that the ACH Network handled 31.5 billion payments worth $80.1 trillion in 2023 and 33.6 billion payments worth $86.2 trillion in 2024, with especially rapid growth in Same Day ACH (Board of Governors of the Federal Reserve System, 2023a; Nacha, 2024a, 2025a).

The FedNow Service went live on July 20, 2023, creating another public-sector-supported instant payment rail for U.S. depository institutions (Board of Governors of the Federal Reserve System, 2023b).

Meanwhile, the CFPB's personal financial data rights rulemaking and the broader evolution of consumer-authorized data sharing have reinforced the possibility of more standardized, secure account connectivity in consumer finance (Consumer Financial Protection Bureau [CFPB], 2024).

These developments do not guarantee mass pay-by-bank adoption, but they lower several of the coordination barriers that historically limited direct bank-account payment innovation at the point of sale.

Yet the growth of A2A payments also raises a difficult risk question. Speed, convenience, and directness improve the user experience only when trust is preserved. In the U.S. market, fraud losses associated with bank transfers and payments are already substantial. FTC data for 2023 showed that consumers reported more than $10 billion in total fraud losses, with bank transfers and payments accounting for the largest aggregate reported losses among payment methods at about $1.86 billion (Federal Trade Commission [FTC], 2024a, 2024b).

The risk problem is not confined to one tactic. Account-to-account systems can be exploited through account takeover, synthetic identity onboarding, mule-account networks, business email compromise, invoice redirection, first-party fraud, merchant impersonation, onboarding with falsified business credentials, open-banking consent abuse, and authorized push payment-style deception.

As GenAI capabilities diffuse, impersonation and social engineering risks may intensify because fraudsters can scale persuasive content, document manipulation, and low-cost identity spoofing.

A payment design that reduces card-not-present fraud may still fail if it weakens controls around consent, beneficiary validation, device trust, or anomaly detection.

This paper argues that the central challenge for U.S. pay-by-bank is therefore architectural.

The question is not simply whether A2A payments are cheaper or faster than cards.

The deeper question is whether U.S. institutions can build an intelligent control stack that makes account-based payments safe enough, fair enough, and explainable enough for sustained merchant and consumer adoption.

A viable architecture must be rail-aware because ACH, Same Day ACH, RTP, and FedNow have different timing, revocability,

return, and operational characteristics. It must also be actor-aware because the relevant parties include consumers, merchants, sponsor banks, core processors, payment facilitators, open-banking intermediaries, account-validation vendors, fraud platforms, and investigators.

Finally, it must be lifecycle-aware because risk is not concentrated at one point in time.

It begins at identity proofing and account linking, shifts during consent capture and payment initiation, intensifies during transaction screening and orchestration, and persists after settlement in the form of returns, complaints, investigation outcomes, and model recalibration.

Accordingly, this paper develops a fraud-aware analytics and risk-modeling framework for intelligent pay-by-bank systems in the United States.

This paper makes four contributions.

First, it synthesizes the market, infrastructure, and regulatory conditions that are making pay-by-bank strategically relevant in the U.S. context.

Second, it organizes the main fraud typologies and operational vulnerabilities that arise in A2A payment flows, emphasizing the difference between card-style and bank-account-style control environments.

Third, it proposes a layered analytic architecture that integrates rules, supervised machine learning, graph analytics, explainable AI, and feedback loops from fraud and dispute outcomes.

Fourth, it outlines governance requirements—model risk management, fairness review, human-in-the-loop escalation, third-party oversight, and consumer-protection controls—that are necessary if pay-by-bank is to scale responsibly.

The broader claim is straightforward: intelligent pay-by-bank is not a single fraud model but a coordinated operating model in which payment rail design, data engineering, real-time analytics, and institutional governance reinforce one another.

Figure 1. Growth of U.S. ACH network volume and value

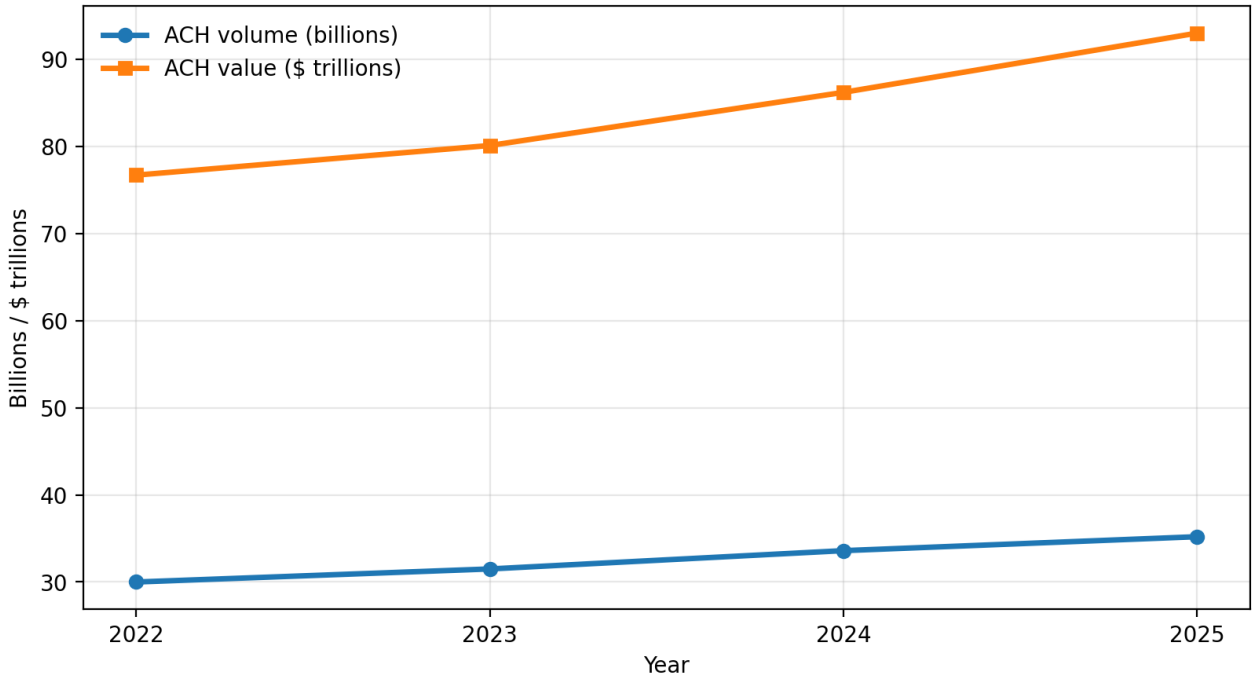


## 2. Literature Review

The literature relevant to intelligent pay-by-bank architecture sits at the intersection of payment-system design, fraud analytics, open banking, explainable AI, and model governance.

No single stream is sufficient on its own. A2A payment risk is not only a classification problem; it is also a systems problem involving incentives, data fragmentation, consumer behavior, settlement timing, and institutional accountability.

A first stream concerns the evolution of payment infrastructures and the economic rationale for direct account-based payments. The Federal Reserve's payments research documents the long-run shift toward noncash electronic payments and provides an essential benchmark for understanding how account-based rails fit into the U.S. payment mix (Board of Governors of the Federal Reserve System, 2023a).

Nacha's volume statistics show the operational depth and resilience of ACH and the especially rapid expansion of Same Day ACH, which is increasingly relevant for faster A2A use cases (Nacha, 2023a, 2024a, 2025a).

Recent Federal Reserve analysis further frames pay-by-bank as an emerging merchant-payments use case whose adoption depends on cost, convenience, perceived security, and merchant integration economics (Hwang, 2025).
This infrastructure literature suggests that pay-by-bank is not merely a theoretical alternative; it is built on rails that are already central to payroll, bill pay, business payments, and treasury operations.
However, the same literature also implies that consumer-facing merchant payments require a different control posture than back-office ACH use cases because checkout transactions are sensitive to latency, abandonment, fraud loss allocation, and user trust.
A second stream addresses financial fraud detection broadly.
Review studies consistently show that fraud detection is characterized by severe class imbalance, adaptive adversaries, limited labels, delayed ground truth, and high error asymmetry.
Ngai et al. (2011) provided one of the foundational reviews of data-mining techniques in financial fraud detection, while West and Bhattacharya (2016) and Abdallah et al. (2016) synthesized the computational and systems challenges faced by fraud detection systems in practice.
More recent reviews continue to emphasize that static rules alone are insufficient when fraud typologies evolve rapidly, especially in online and remote payment settings (Ali et al., 2022; Bockel-Rickermann et al., 2023).
The main implication for pay-by-bank is clear: if the payment experience reduces friction for legitimate users, it may also reduce friction for fraudsters unless institutions deploy continuously updated, data-driven control layers.
A third stream focuses on algorithmic performance in transaction-fraud detection.
Research on credit-card and electronic funds transfer fraud demonstrates the importance of feature engineering, sequence modeling, and ensemble approaches under imbalanced conditions.
Jurgovsky et al. (2018) show that sequence classification can capture user behavior over time rather than evaluating transactions only as isolated events.
Afriyie et al. (2023) compare traditional machine-learning classifiers and find strong performance for random forests on simulated U.S. transaction data, underscoring the continued value of robust tabular models.
Ti et al. (2022) demonstrate that feature design is itself decisive: different categories of recency, frequency, monetary, and anomaly-related features have materially different effects on fraud-detection quality, and statistically grounded, domain-informed features may outperform generic automated feature sets.
Xu et al. (2023) further illustrate that modern boosting architectures can deliver strong fraud-screening performance at scale.
For pay-by-bank systems, this literature indicates that the highest-value features are likely to span not only transaction amounts and velocity but also linked-account history, beneficiary novelty, device continuity, consent timing, account age, retry patterns, cross-merchant graph proximity, and post-transaction return behavior.
A fourth stream concerns synthetic identity, mule-account abuse, and networked fraud.
These threats are especially relevant for pay-by-bank because A2A ecosystems depend on reliable account ownership, beneficiary integrity, and merchant legitimacy.
The Federal Reserve's synthetic identity fraud toolkit argues that no instant check can perfectly confirm that a name, date of birth, and Social Security number correspond to a true individual, and it recommends layered detection that combines machine learning with expert review and operational controls (FedPayments Improvement, 2022).
This observation is crucial for pay-by-bank design.
If onboarding or bank-linking processes are weak, downstream transaction analytics inherit corrupted identities and counterparties.
The graph-oriented logic in Rasel et al. (2023) is relevant here even though the paper addresses mortgage risk rather than payments directly: multimodal and network-based modeling can improve detection where isolated variables fail to capture relational risk.
Similarly, Ibrahim et al. (2024a) emphasize predictive analytics and infrastructure protection in AML settings, highlighting the value of typology-aware monitoring, multi-source data fusion, and risk escalation for suspicious flows.
In A2A payments, those ideas translate into graph analytics for detecting shared devices, common beneficiaries, linked mule accounts, repeated micro-probes, and coordinated cross-channel attacks.
A fifth stream concerns real-time and faster payments.
The U.S. migration toward real-time payment capability changes the economics of fraud control because the time available for review narrows as funds availability accelerates.
The Federal Reserve emphasized from the FedNow planning stage that participating banks remain a primary line of defense against end-user fraud, while the service operator can add infrastructure-level support and fraud-prevention tools over time (Board of Governors of the Federal Reserve System, 2019, 2022, 2023b).
This division of responsibility matters. In faster A2A systems, risk management must shift earlier in the decision process. Rather than relying on ex post dispute remediation, institutions need pre-transaction controls such as confirmation of payee logic where feasible, bank-account validation, device and channel authentication, real-time beneficiary screening, and stepped-up verification for novel or high-risk contexts. Fahim et al. (2024a) frame the problem effectively in their analysis of real-time

payments and real-time fraud: the value proposition of speed can erode quickly if consumer protection and fraud response lag the pace of settlement.

A sixth stream relates to explainability, accountability, and governance. Because payment decisions can deny access, delay funds, or trigger investigations, the use of opaque machine learning creates operational, regulatory, and fairness concerns. Zhou et al. (2023) show that user-centered explainability can connect local and global model explanations to the needs of external stakeholders.

Weber et al. (2024) demonstrate in a broader finance review that explainable AI is especially relevant in regulated domains where decision traceability is necessary for adoption.

These insights align closely with Fahim et al. (2023), who argue that algorithmic accountability in U.S. consumer fintech requires governance structures, documentation, and oversight rather than technical accuracy alone.

For pay-by-bank systems, explainability serves at least four purposes: it supports analyst review of flagged transactions, helps compliance teams defend model use, enables challenge and remediation when consumers dispute outcomes, and improves management understanding of false-positive tradeoffs.

Explainability also matters strategically because merchants and bank partners may resist black-box fraud vendors if approval decisions cannot be audited or tuned against portfolio objectives.

A seventh stream concerns data portability, open banking, and consumer-authorized data sharing.

The CFPB's work on personal financial data rights does not directly create a national pay-by-bank system, but it helps institutionalize the idea that consumers should be able to authorize secure access to covered financial data through standardized, governed mechanisms (CFPB, 2024).

That direction is important because many pay-by-bank experiences rely on third-party connectors, account linking, and verification flows that historically varied in reliability and transparency.

A more standardized data-access environment can strengthen identity and account validation, but it also introduces third-party risk, consent governance challenges, and questions about data minimization.

Thus, open banking can expand the data available for risk modeling while simultaneously enlarging the attack surface if vendors, permissions, or token management are weak.

An eighth stream comes from the user-specified body of recent applied research.

Several of the cited papers are not about payments narrowly, but they contribute design principles that are transferable. Ibrahim et al. (2022) and Ibrahim et al. (2025a) highlight predictive analytics for systemic risk environments and show how risk models gain value when they integrate heterogeneous signals rather than relying on a single domain variable.

Pritty et al. (2024) show how GenAI-related manipulation risks can be detected through narrative anomalies and fraud signals, which is relevant to merchant onboarding, dispute narratives, and synthetic documentation in payments.

Jahan et al. (2024) emphasize early warning analytics and suspicious-signal detection, reinforcing the idea that fraud control must operate as a continuous monitoring process rather than a one-time screening event. Hasan et al. (2023) focus on fraud detection and cybersecurity infrastructure, again underscoring that transaction risk, cyber risk, and system resilience are intertwined rather than separable.

Even Arman and Fahim (2023), though centered on inventory operations, are relevant indirectly because they illustrate how AI can optimize operational decision-making when data, timing, and workflow orchestration are aligned.

Taken together, the literature suggests three gaps that justify this paper.

First, much fraud research remains centered on cards, lending, financial statements, or AML, while merchant-facing U.S. pay-by-bank remains comparatively under-theorized.

Second, payment-infrastructure analysis often under-specifies the analytics and governance stack needed to make A2A payments trustworthy in real consumer and merchant environments.

Third, explainability and control governance are often discussed after model selection rather than being designed into the architecture from the outset. This paper responds to those gaps by treating pay-by-bank as a full-stack risk architecture problem in which rail economics, fraud typologies, machine learning, graph analytics, and governance must be integrated.

**Table 1. Core U.S. rails relevant to pay-by-bank design**

| Rail | Settlement profile | Typical pay-by-bank relevance | Key control implication |
|---|---|---|---|
| ACH | Batch, highly scalable, broad reach | Recurring payments, bill pay, lower-cost merchant flows | Stronger attention to authorization quality, return patterns, and post-payment monitoring |

| | | | |
|---|---|---|---|
| Same Day ACH | Same business day, three settlement windows | Faster checkout, payroll, urgent bill pay, account transfers | Higher value of pre-initiation checks because losses materialize faster than standard ACH |
| RTP | Real-time messaging and settlement | Immediate funds use cases, time-sensitive merchant and commercial flows | Low tolerance for ambiguous risk; requires stronger ex ante identity and beneficiary controls |
| FedNow | Instant payments, 24x7x365 | Broader U.S. instant-payment access through participating banks | Decisioning, authentication, and exception handling must operate in near real time |

Figure 2. Same Day ACH growth in volume and value

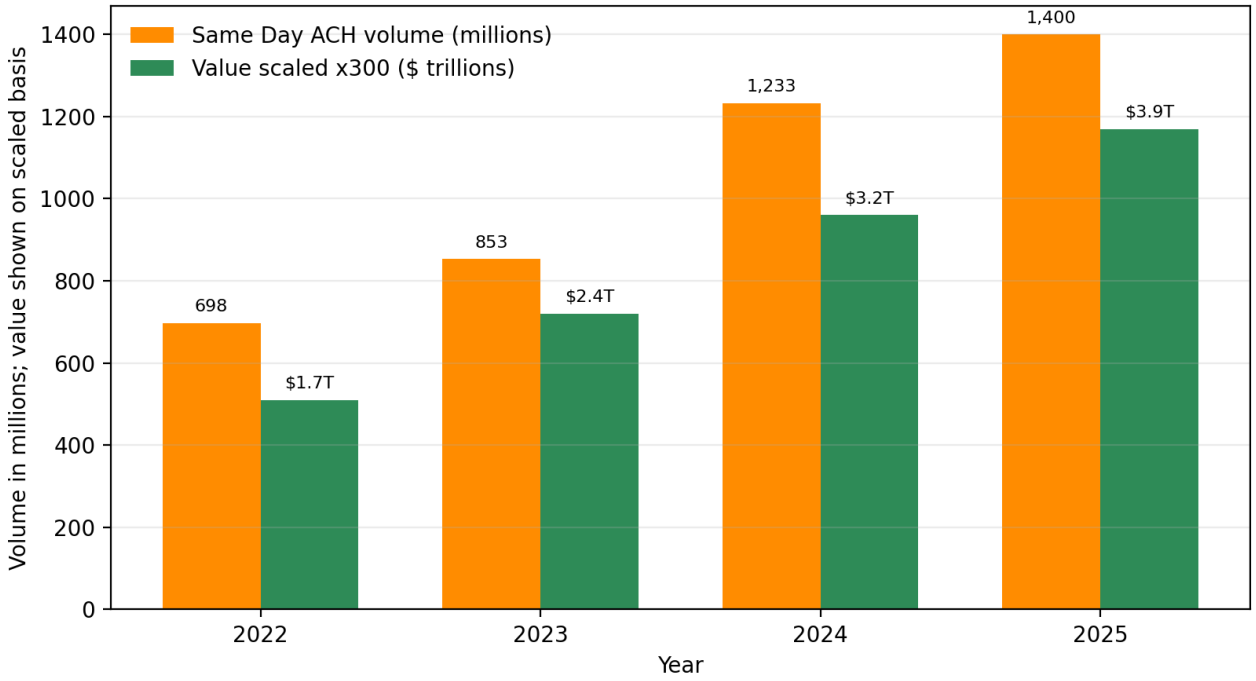


**3. Methodology**

This study uses a design-science-oriented conceptual methodology supported by public payments data, official U.S. policy and infrastructure documents, and contemporary fraud-analytics research.

The goal is not to estimate one proprietary production model on a restricted dataset, but to construct a publication-ready framework for how a U.S. pay-by-bank provider, bank, or merchant-acquiring partner could engineer a fraud-aware architecture suitable for ACH, Same Day ACH, RTP, and FedNow-adjacent A2A use cases.

The methodological stance is therefore synthetic and prescriptive: it combines descriptive evidence on the current U.S. payments environment with an implementation blueprint that can be adapted to institutional context.

The framework is organized around six analytic layers: identity and party resolution, payment-context intelligence, transaction scoring, graph/network analytics, decision orchestration, and post-event learning.

Each layer is tied to specific data elements, control objectives, and explainability outputs. The design principle is defense in depth. No single model is assumed to be sufficient because fraud in A2A environments is heterogeneous.

For example, synthetic-identity onboarding, merchant impersonation, and mule-account routing are structurally different problems even if they all culminate in illicit fund movement.

The architecture therefore uses multiple specialized models and rules that feed a common case-management and policy engine.

Data architecture. The proposed data design includes five major domains. The first is identity and KYC data: legal name, date of birth or business registration details, tax identifiers where permitted, address history, phone and email tenure, documentary verification results, beneficial ownership fields for business merchants, sanctions and politically exposed person screening, and

account-holder-bank verification outputs.
The second is payment-account data: routing and account metadata, account age, prior return history, account validation status, tokenized bank-link metadata, account-link success and failure logs, and ownership-match confidence scores.
The third is transaction-context data: amount, time, merchant category, beneficiary status, initiation channel, purpose code, recurring versus one-off flag, frequency window, payday/bill-day proximity, and whether the payment is refund-like, invoice-like, or a checkout payment.
The fourth is device and session telemetry: browser fingerprint, SDK attributes, IP geolocation consistency, emulator or jailbreak signals, SIM and device change indicators where available, behavioral biometrics, and session path anomalies.
The fifth is network and feedback data: prior fraud labels, disputes, ACH returns, complaint codes, SAR-related escalation outcomes, linked merchants, linked beneficiaries, shared devices, and investigator annotations.
Entity resolution and feature engineering.
The first analytic layer focuses on resolving parties and relationships across fragmented identifiers.
Consumers may appear under multiple devices, email addresses, bank accounts, or tokenized links; merchants may operate through aggregators or affiliates; beneficiaries may recur across otherwise unrelated cases.
A probabilistic entity-resolution module links records using deterministic matches where appropriate and fuzzy or learned matches where necessary. This layer generates features such as identity consistency score, document-attribute conflict count, linked-account multiplicity, beneficiary novelty, merchant-registration age, and historical account-link success ratios. Drawing on the logic of graph-based fraud and synthetic identity detection, the methodology assumes that relational features often outperform purely transactional features when fraudsters coordinate across accounts and channels.
The second layer performs payment-context feature construction.
Here, the framework derives recency-frequency-monetary variables, prior merchant relationship depth, paycheck and bill-cycle alignment, channel-switching patterns, average transaction-amount deviations, rapid retry sequences, and add-payee-to-payment interval measures.
Because pay-by-bank often depends on customer-authorized data access, consent features are also critical: consent creation time, token age, consent scope, refresh frequency, and anomalous re-consent behavior.
A payment initiated immediately after a new bank link, on a novel device, to a new merchant, with unusual consent timing is analytically different from a recurring utility payment from a long-established account.
Model architecture. The framework uses a model ensemble rather than a single classifier. Three model families are recommended.
First, a gradient-boosted tabular model estimates transaction fraud probability using structured features that combine identity, account, context, and device variables. Boosted tree models are practical because they perform strongly on mixed tabular data, handle nonlinearity, and support mature explanation tools such as SHAP.
Second, a sequence model evaluates short-term behavioral trajectories, such as the order and timing of account-linking attempts, beneficiary additions, and payment retries.
This is especially useful for detecting scripted attacks, testing behavior, and rapid changes that may not be visible in static features alone.
Third, a graph model computes relational risk over entities including customers, bank accounts,
merchants, beneficiaries, devices, IP clusters, and phone numbers.
Node- and edge-level features capture shared infrastructure, bridge accounts, centrality of mule-like entities, and proximity to previously confirmed fraud nodes.
For institutions not ready to deploy graph neural networks, graph-derived features can still feed the tabular model effectively.
Rule layer and policy engine.
Machine learning is paired with an explicit rule layer for legal, operational, and commonsense controls.
Examples include sanctions hits, failed account-ownership validation, duplicate payment suppression, velocity caps, embargo screening, return-code blacklists, and merchant-category restrictions.
The rule layer also addresses situations where deterministic business policy is preferred to probabilistic inference.
For example, a newly linked account initiating a high-value first payment to a first-time beneficiary outside a customer's historical geography may require mandatory step-up authentication even if the model score is below the standard decline threshold.
The policy engine merges rule outputs and model scores into one of four decisions: approve, approve with additional verification, hold for review, or decline.
Rail-aware orchestration.
A central contribution of the methodology is rail-aware scoring. Not all A2A payments should be screened identically. ACH debits with established authorization histories and clear return mechanisms may tolerate a different threshold than instant credits over real-time rails.
FedNow and RTP-like contexts require lower tolerance for ambiguous risk because reversal opportunities may be limited and customer expectations for instant confirmation are high.

Same Day ACH can occupy a middle ground in which some timing exists for pre-processing checks, but losses can still crystallize quickly. The architecture therefore includes rail-specific thresholding, feature weighting, and workflow timing.
Transactions are also segmented by use case: consumer checkout, bill pay, payroll-related disbursement, P2P-like transfer, marketplace payout, and commercial invoice payment.
A single universal model is discouraged because fraud incentives differ by use case. Explainability and analyst interface.
Each scored event generates both local and global explanations. Local explanations identify the strongest factors driving the score for a particular transaction, such as unusual device change, first-time beneficiary, failed ownership match, or graph proximity to a known mule cluster.
Global explanations help model owners understand portfolio-level drivers and drift over time.
The analyst interface is designed to show chronology, linked entities, prior history, threshold comparisons, and recommended next actions. This user-centered explainability is important because payment-fraud operations require rapid but defensible decisions. Analysts need to understand not only that a score is high, but why it is high and what evidence should be checked next.
For consumer-facing outcomes, a simplified adverse-action style explanation should be available where appropriate, especially when step-up verification or temporary holds affect legitimate users.
Feedback learning and outcome capture.
The methodology treats outcomes as a delayed and noisy signal rather than a binary truth revealed immediately after payment initiation.
Labels may come from ACH returns, merchant complaints, internal investigations, consumer reports, bank recalls, law-enforcement referrals, or SAR-adjacent escalation.
Because many fraud types are discovered late, the system records time-to-label, label source, and confidence grade.
The retraining pipeline incorporates champion-challenger testing, drift monitoring, calibration checks, and segmented false-positive analysis across rail, merchant, and customer cohorts.
A separate disputes-and-returns dashboard tracks whether fraud controls are shifting losses from outright fraud to customer friction or abandonment.
Evaluation strategy. If implemented empirically, the framework should be evaluated along multiple dimensions rather than a single accuracy score.
Core metrics include precision, recall, F1, PR-AUC, false-positive rate, fraud-dollar capture rate, manual-review rate, investigation cycle time, and customer-friction rate.
In payments, economic measures are essential: expected loss avoided, holdout fraud leakage, incremental revenue saved from reduced abandonment, and analyst-hours consumed per prevented fraud dollar.
Because labels emerge over time, backtesting should use temporal validation rather than random shuffling.
Scenario-based stress tests should also be conducted for spikes in business email compromise, credential stuffing, data-breach replay, and mule-network expansion.
Fairness testing is needed to confirm that proxy discrimination is not being introduced through geography, income-correlated variables, or device-quality proxies.
Governance design.
Model governance follows a three-lines-style structure.
The first line owns operations, thresholds, and day-to-day exception handling.
The second line validates model design, monitors fairness and compliance, and reviews documentation, feature lineage, and third-party dependencies.
The third line or internal audit assesses whether controls operate as documented.
Consistent with NIST's AI Risk Management Framework and established banking model-risk guidance, the framework requires inventorying data sources, documenting intended use, tracking performance limits, recording overrides, and validating that explanations remain faithful and useful (National Institute of Standards and Technology [NIST], 2023; Office of the Comptroller of the Currency, 2021). Third-party risk management is emphasized because many pay-by-bank ecosystems depend on aggregators, identity vendors, risk-scoring providers, and payment orchestrators.
A technically strong model can still fail if vendor service levels, data rights, or escalation protocols are weak.
In sum, the methodology proposes a layered, explainable, rail-aware architecture suitable for production adaptation in U.S. pay-by-bank environments.
It treats fraud management as an operational intelligence system rather than a standalone predictive model.
That orientation is essential because the quality of A2A payment risk management depends on how data, models, rules, investigators, vendors, and payment rails interact in real time.

**Table 2. Fraud typologies relevant to U.S. pay-by-bank systems**

| Fraud typology | Primary stage | Typical signals | Control response |
|---|---|---|---|
| Synthetic identity | Onboarding / linking | Thin-file identity, cross-attribute inconsistencies, shared devices | Layered identity resolution, documentary review, graph linkage checks |
| Account takeover | Bank linking / initiation | Novel device, credential-reset sequence, abnormal location changes | Step-up authentication, session analytics, cooldown logic |
| Authorized push-payment scam | Initiation | First-time beneficiary, urgency pattern, atypical payment narrative | Consumer warnings, beneficiary checks, risk prompts, possible review hold |
| Mule-account routing | Routing / post-settlement | Shared beneficiaries, rapid fan-out, network centrality spikes | Graph analytics, freeze/escalation workflow, AML coordination |
| Return abuse / first-party fraud | Post-settlement | Repeat returns, customer timing anomalies, dispute clustering | Portfolio segmentation, behavioral profiling, merchant-specific policy tuning |

Figure 3. Reported U.S. consumer fraud losses by payment method, 2023

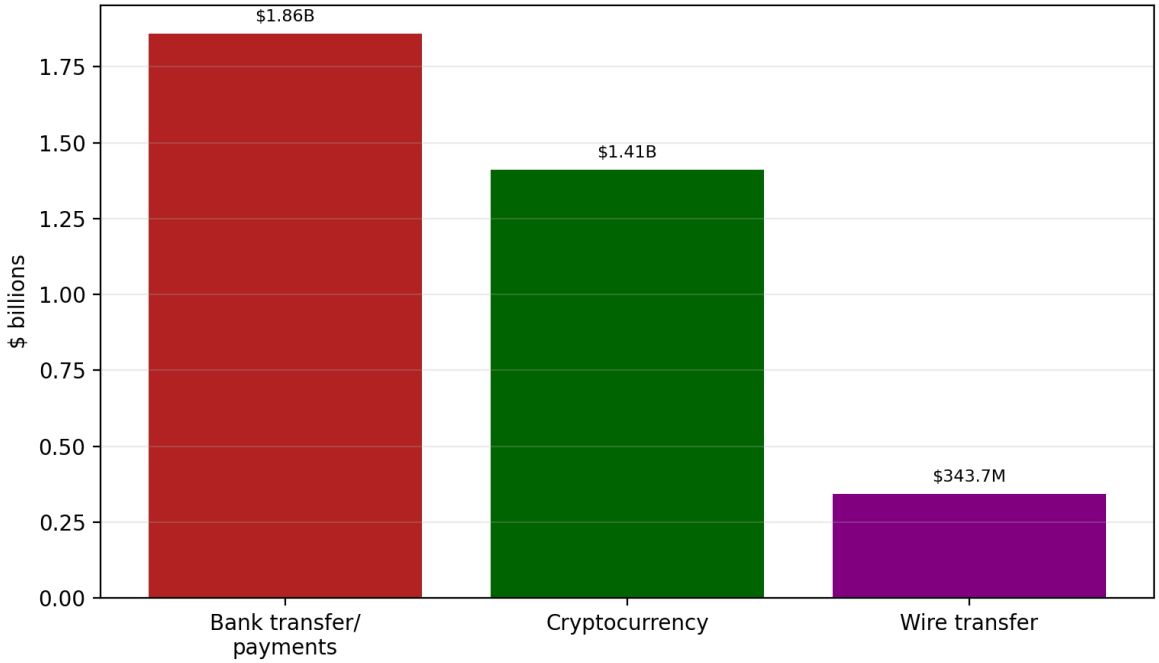

Figure 4. Selected U.S. milestones shaping pay-by-bank readiness

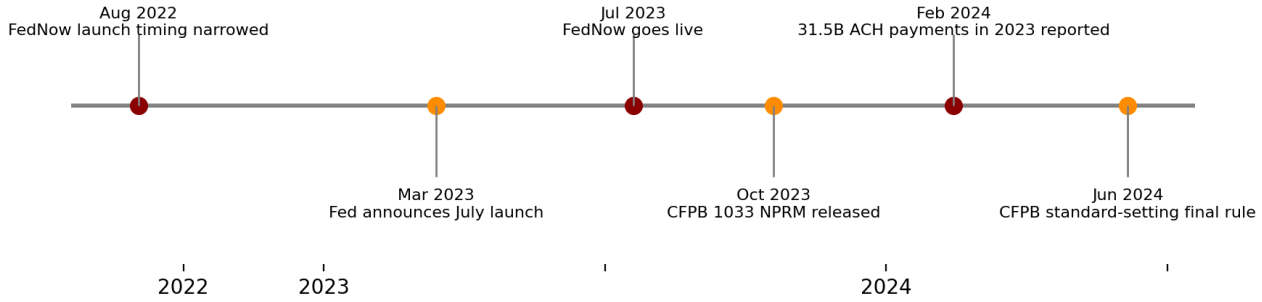


**4. Discussion**

The proposed architecture has practical value because it reframes pay-by-bank risk as a portfolio management problem rather than a binary fraud-filtering problem.

In U.S. A2A systems, the institution is balancing at least five objectives simultaneously: fraud prevention, customer experience, merchant conversion, regulatory compliance, and operational scalability.

Optimizing one objective in isolation usually damages another. Aggressive thresholding may suppress fraud but increase false positives, checkout abandonment, and customer dissatisfaction.

A frictionless flow may boost conversion but invite fraud losses, investigations, and reputational damage.

The argument of this paper is that intelligent pay-by-bank succeeds only when these tradeoffs are measured explicitly and governed continuously.

One major implication is that payment safety must be designed upstream.

Card systems historically trained merchants and consumers to expect some ex post remediation through disputes and chargebacks, even though those processes are costly and imperfect.

Many pay-by-bank flows do not offer an equivalent safety valve, especially on faster rails or on flows where funds move quickly to downstream accounts.

This shifts the center of gravity from dispute-led remediation to ex ante confidence building.

In practice, that means stronger onboarding, beneficiary verification, consent governance, transaction-context analytics, and human review for ambiguous edge cases.

Institutions that try to “bolt on” fraud controls after launching a pay-by-bank experience are likely to discover that the economics deteriorate once loss rates, manual work, and customer complaints accumulate.

A second implication concerns rail heterogeneity. The U.S. market does not have one uniform real-time or A2A rail. ACH, Same Day ACH, RTP, and FedNow differ materially in processing windows, message structures, irrevocability expectations, error-recovery patterns, and operational hours. One smart architecture therefore cannot treat rail choice as a back-end detail. Rail selection is itself a risk variable. Low-risk recurring bill payments from established accounts may be acceptable over lower-cost batch ACH pathways, while urgent but trusted disbursements may justify instant rails.

Higher-risk or ambiguous events may require slower processing paths, stepped-up customer verification, or temporary holds even when the customer requests speed.

The commercial lesson is important: faster is not always better. Optimal orchestration means aligning the rail to the risk, not forcing every payment onto the fastest available route.

Third implication is that open banking expands both capability and exposure. Consumer-authorized data sharing can improve account ownership validation, balance awareness, and user experience, making pay-by-bank more reliable and easier to adopt. However, API-mediated or token-mediated connectivity also creates new control points that must be defended.Consent tokens can be hijacked or misused.

Poorly governed data aggregation relationships can create hidden dependencies. Data minimization failures can increase privacy and cyber risk without improving fraud detection materially.

As a result, institutions should not assume that more data automatically means more safety.

The discussion in the literature on algorithmic accountability and data governance is directly relevant here: only data with a clear control purpose should enter production scoring, and third-party roles must be contractually and operationally clear (Fahim et al., 2023; CFPB, 2024). A disciplined institution will specify which data are used for onboarding, which for transaction scoring, which for post-event investigation, and which are retained only in tokenized or privacy-preserving form.

Fourth implication is that synthetic identity and mule-account risk must be treated as foundational, not peripheral.

Many payment teams focus on transaction anomalies while underinvesting in the integrity of parties.
That is a mistake in pay-by-bank environments. If fraudsters can create or control seemingly legitimate customers, businesses, beneficiaries, or bank links, even sophisticated transaction models will produce unstable results.
The Federal Reserve's synthetic identity guidance is instructive because it highlights the insufficiency of one-shot identity checks and the need for layered evidence across time (FedPayments Improvement, 2022).
For pay-by-bank, this means identity confidence should be dynamic and continuously updated.
A newly registered consumer account, a new merchant relationship, or a newly added beneficiary should not inherit the same trust score as a long-observed relationship with clean historical behavior.
Graph analytics help operationalize this principle by exposing hidden linkages among parties that look normal individually but suspicious collectively.
Fifth implication concerns the role of explainability. In many production settings, explainability is discussed as a regulatory afterthought or a management comfort feature. That understates its value. For payment fraud operations, explainability is part of execution quality. Analysts need interpretable case narratives to decide whether to approve, challenge, or escalate. Business teams need to understand whether a model is suppressing high-value but legitimate traffic. Compliance teams need to confirm that monitoring logic is aligned with policy and that adverse impacts are not emerging in protected or vulnerable groups. Vendors and bank partners often need assurance that the system's behavior is coherent and tunable.
Thus, explainability is not opposed to performance; it is part of a robust performance regime in a high-stakes operational environment.
The literature on user-centered XAI is helpful because it reminds us that explanations must be useful to the actual decision-maker, not merely mathematically available (Zhou et al., 2023; Weber et al., 2024).
In practice, that means explanation interfaces should be role-specific: investigators need evidence trails, managers need portfolio drivers, model validators need calibration and stability diagnostics, and customer-service representatives need simplified reason codes.
A sixth implication is that fraud and AML monitoring should be integrated more tightly than they often are today.
In many institutions, retail fraud teams, AML teams, cyber teams, and payment operations teams remain partially siloed.
That separation is understandable historically, but it is increasingly costly in A2A systems.
A beneficiary account receiving many small payments from newly linked customer accounts may be a fraud concern, a mule-network concern, and a suspicious-activity concern at the same time.
Likewise, device anomalies, geolocation inconsistencies, and linked identities may have both cyber and payment-fraud meaning.
Ibrahim et al. (2024a) argue for predictive analytics frameworks that protect financial infrastructure through typology-aware monitoring and escalation.
Applied to pay-by-bank, that insight supports shared data layers, shared case identifiers, and common graph views across fraud and AML functions. It does not require merging all teams administratively, but it does require shared intelligence so that institutions do not miss patterns visible only across control silos.
A seventh implication is commercial. The future of pay-by-bank in the United States will depend on whether the control model preserves conversion benefits for merchants.
Merchants will not adopt A2A checkout at scale simply because it is technologically possible.
They will adopt it if it reduces costs without creating excessive failed payments, false declines, settlement uncertainty, support burden, or reputational risk.
That means fraud operations must be measured partly as a revenue-protection function.
A model that captures more fraud but damages conversion at key points in the customer journey may be inferior to a slightly less aggressive model that preserves trust and throughput.
Accordingly, decision thresholds should be optimized by merchant segment, ticket size, use case, and risk appetite.
Some merchants may prioritize low friction for loyal customers and accept tighter post-payment monitoring; others may prefer stricter front-end controls for high-ticket orders.
The architecture proposed here supports that flexibility because policy rules and thresholds can be segmented rather than universal.
An eighth implication concerns consumer protection and trust communication. Consumers are more likely to adopt direct-from-bank payment methods if they understand how consent works, how disputes are handled, how to revoke permissions, and what protections apply when something goes wrong.
Technical risk controls alone are not enough. Clear consent screens, transparent revocation options, beneficiary warnings, unusual-payment prompts, and post-payment confirmations all shape the risk environment.
Behavioral fraud prevention can be embedded in design. For example, if a consumer is sending a first-time high-value payment after a recent change in contact details and device, the interface can provide friction that is safety-enhancing rather than merely obstructive. This is especially important in scams and impersonation cases where the customer is socially engineered into authorizing the payment.

In such settings, a well-designed intervention—beneficiary confirmation, cooling-off warning, or explicit scam alert—may outperform purely backend analytics.
A ninth implication is organizational. Institutions often underestimate the workflow burden of advanced fraud systems. Adding more signals, models, and vendors can create analytical richness but also operational noise if review queues are not designed carefully. A successful pay-by-bank program therefore needs not only models but operating discipline: queue prioritization, standardized evidence templates, escalation playbooks, feedback capture from investigators, and clear service-level objectives for merchants and consumers. Human-in-the-loop review should be reserved for cases where it adds distinct value.
If the majority of alerts are obvious false positives, the institution has a thresholding problem.
If the majority are obvious frauds discovered too late, it has a timeliness problem.
If analysts cannot explain why a case was scored highly, it has an explainability problem.
Operational diagnostics should distinguish among these failure modes rather than treating them all as generic "fraud-ops" issues.
A tenth implication is resilience. Payment innovation is often discussed in terms of speed and user experience, but resilience matters equally. Cyber incidents, connectivity failures, vendor outages, and data-feed disruptions can all degrade fraud controls precisely when attackers are most active.
Mature intelligent pay-by-bank architecture therefore needs degraded-mode logic: fallback thresholds, cached identity trust scores, queueing rules during vendor outages, and post-event reconciliation processes.
The Federal Reserve's broader payments and cybersecurity materials emphasize that safety and operational continuity are inseparable in payment systems (Board of Governors of the Federal Reserve System, 2024; NIST, 2023).
For pay-by-bank, resilience means the fraud stack must fail safely, not merely fail visibly.
An institution should know in advance what happens if its account-validation vendor times out, if consent verification is delayed, or if graph features are unavailable during peak processing hours.
Eleventh implication is methodological. The strongest production systems are likely to combine machine learning with institutional knowledge rather than trying to replace human judgment entirely.
The literature consistently warns that fraud is adaptive and labels are delayed.
Models learn from the past; fraudsters exploit the future.
Therefore, the architecture should support rapid policy updates, investigator-driven signal proposals, and champion-challenger experimentation.
This is one reason the proposed framework retains a prominent rule layer instead of assuming that end-to-end machine learning is always superior. Layered intelligence is more robust than monolithic intelligence.
It allows institutions to react quickly to new scams while retraining or recalibrating more complex models in a controlled way.
Finally, the discussion suggests that the broader significance of intelligent pay-by-bank is strategic rather than merely incremental.
If U.S. institutions get the control architecture right, pay-by-bank could become a meaningful complement to cards in merchant checkout, bill payment, subscriptions, and selected commercial flows.
It could also encourage more thoughtful integration of open banking, faster payments, and explainable AI in consumer finance.
But if control design is weak, the opposite could happen: A2A payment innovation could become associated with scam risk, confusing liability, poor recovery experiences, and merchant hesitation.
The difference between those outcomes will depend less on marketing than on execution.
The core message of this paper is therefore practical: pay-by-bank should be engineered as a trust architecture.
Fraud-aware analytics, rail-aware orchestration, clear governance, and consumer-centered design are not supplementary features. They are the conditions under which U.S. account-to-account payments can scale safely and credibly.

A further implication is that evaluation windows must be longer than launch windows.
Many payment programs are judged in their first few months primarily on conversion, partner onboarding, and transaction growth.
Fraud maturity, however, is often revealed later, once attackers discover the new channel, share tactics, and probe institutional response times.
This lag means executives should resist declaring success solely from early loss numbers.
They need cohort-based monitoring that compares early adopters with later cohorts, observes whether fraud migrates from onboarding to transaction stages, and tracks whether preventive controls are merely shifting losses into returns, disputes, or customer attrition.
A pay-by-bank product can appear healthy in quarter one yet become operationally fragile in quarter three if adaptive fraudsters learn where consent, account-linking, or beneficiary controls are weakest.
For that reason, post-launch governance should include formal model and policy reviews at fixed intervals, not only incident-triggered reviews.

There is also a competitive implication for banks and fintech intermediaries. Institutions that can combine lower-cost A2A processing with high-quality fraud orchestration may create a defensible
advantage in merchant acquiring, bill pay, and embedded-finance partnerships. The advantage will not come only from the rail; it will come from the intelligence layer that makes the rail commercially usable.
In that sense, fraud analytics should be viewed not just as a cost center but as a product capability.
Merchants may increasingly choose partners based on who can deliver the highest trustworthy approval rate, the clearest dispute and exception workflow, and the most transparent control framework for regulators and consumers.

Figure 5. Conceptual fraud-risk heatmap across the pay-by-bank lifecycle

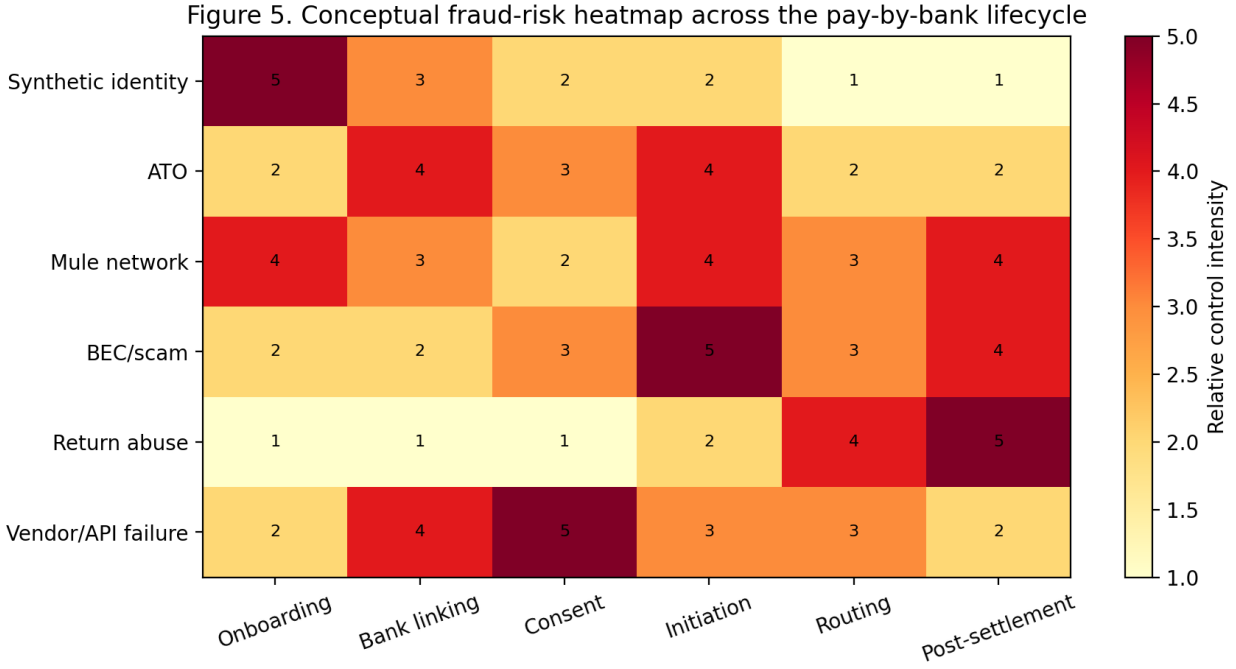


Figure 6. Proposed intelligent pay-by-bank risk architecture

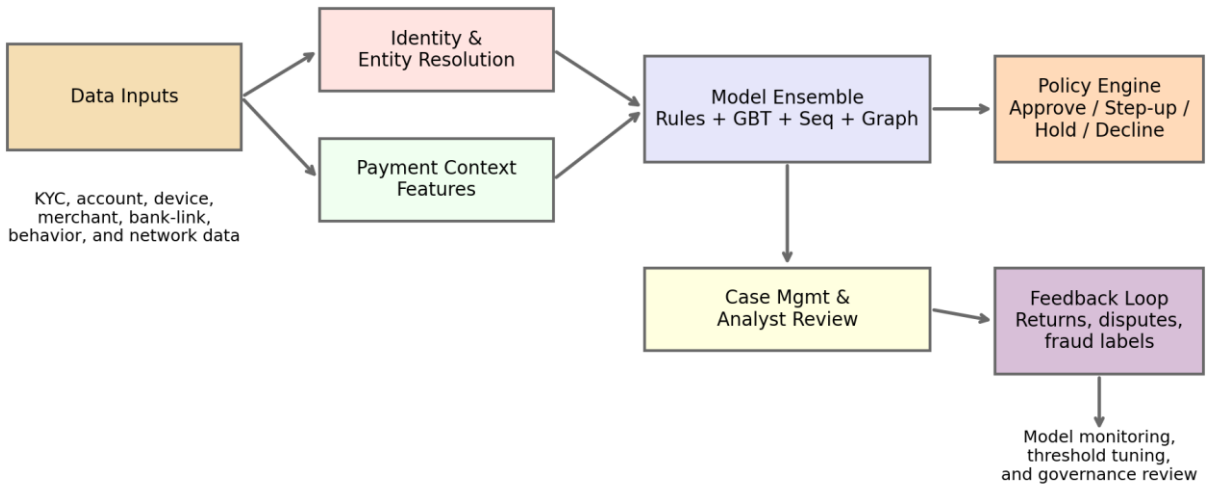

**Table 3. Illustrative data domains and feature families for production scoring**

| Data domain | Examples | Analytic value | Governance concern |
|---|---|---|---|
| Identity/KYC | Name, DOB, EIN, address tenure, BO data | Synthetic-identity and merchant-legitimacy assessment | Data quality, fair use, vendor provenance |
| Bank-link/account data | Ownership match, account age, return history, token age | Account trust scoring and first-payment controls | Consent, token security, third-party access |
| Transaction context | Amount, merchant, beneficiary novelty, timing, use case | Behavioral deviation and scenario-specific risk | Purpose limitation, segmentation bias |
| Device/session | Fingerprint, IP consistency, emulator/jailbreak signals | ATO and scripted attack detection | Privacy, cross-channel data retention |
| Network outcomes | Linked entities, prior alerts, returns, investigator notes | Graph risk, feedback learning, drift detection | Label noise, explainability, access controls |

## 5. Conclusion

Pay-by-bank is best understood as a next-stage operating model for U.S. account-to-account payments rather than as a single checkout feature. Its promise lies in combining lower-cost bank-account routing, faster funds movement, richer data connectivity, and expanding payment-rail choice.

Its vulnerability lies in the fact that direct, often faster, account-based payments can expose merchants, consumers, and banks to fraud patterns that are not well managed by legacy card-era assumptions.

This paper has argued that the path forward is an intelligent architecture that integrates onboarding integrity, payment-context analytics, graph-based relationship detection, explainable machine learning, rail-aware orchestration, and continuous post-event learning. The central message is that trusted pay-by-bank does not emerge automatically from open banking APIs or faster rails. It must be engineered deliberately through layered controls and disciplined governance.

For U.S. institutions, the practical implication is that fraud strategy should be embedded into product design from the first day of a pay-by-bank program.

A2A payments can scale safely when account validation, consent governance, beneficiary risk controls, real-time scoring, and investigation workflows operate as one coordinated system.

Institutions that design for safety, explainability, and operational resilience are more likely to achieve durable merchant adoption, better consumer confidence, and more sustainable payment innovation.

Institutions that prioritize speed without sufficient analytic depth may discover that short-term growth is offset by losses, false positives, and weakened trust.

In that sense, intelligent pay-by-bank architecture is not only a technical framework; it is a strategic blueprint for making U.S. account-to-account payments safer, smarter, and more commercially viable.

## 6. Limitations and Future Directions

This paper has several limitations.

First, it is a conceptual and design-oriented manuscript rather than an empirical production study conducted on a proprietary U.S. pay-by-bank dataset.

As a result, the framework is intended to be operationally realistic, but its specific thresholds, feature weights, and model performance claims remain to be validated institution by institution.

Second, the U.S. pay-by-bank market is still evolving.

Rail capabilities, open-banking standards, fraud typologies, and consumer-protection expectations are changing, which means that some architectural choices discussed here may require recalibration as market structure matures.

Third, public data are stronger for overall payment volumes and reported fraud losses than for merchant-level pay-by-bank conversion, institution-specific false-positive rates, or private vendor detection performance.

That data limitation makes it easier to establish strategic relevance than to benchmark every operational parameter.

Fourth, fraud labels are inherently noisy and delayed. In real deployments, some scam-related authorized payments, first-party fraud events, or mule-account cases may not be identified quickly enough to support clean supervised learning.

Fifth, explainability introduces its own challenges.

Even when SHAP- or rule-based explanations are available, there remains a risk that users over-trust plausible narratives that are descriptive of model behavior but not truly causal.
Sixth, fairness and inclusion tradeoffs are difficult in payments.
Variables that improve fraud capture may correlate indirectly with geography, income, digital access, or thin-file characteristics, so institutions must test carefully for unintended exclusion.
Future research should extend this framework in at least four ways.
One direction is empirical validation using institution-grade A2A transaction data spanning ACH, Same Day ACH, RTP, and FedNow-like flows, with temporal validation and fraud-dollar-based metrics.
Second direction is simulation and agent-based modeling of adaptive fraud migration across rails, merchants, and onboarding channels.
Third is privacy-preserving collaboration, including federated learning or secure multiparty analytics that allow institutions to share fraud intelligence without exposing raw customer data.
A fourth is consumer-centered experimentation on intervention design: warning messages, beneficiary-confirmation interfaces, step-up-authentication timing, and consent controls should be tested not only for fraud prevention but also for fairness, comprehension, and abandonment effects.
Additional work is also needed on governance questions such as model accountability in multi-vendor payment stacks, treatment of GenAI-enabled impersonation, and standardized approaches to explaining high-stakes payment decisions to merchants, analysts, regulators, and end users.
These next steps would move the literature from architectural synthesis toward measurable, institution-ready evidence on how intelligent pay-by-bank can scale safely in the United States.

Appendix A. Implementation and Governance Aids

**Table A. Indicative metric dashboard for intelligent pay-by-bank operations**

| Metric | Why it matters | Owner | Review cadence |
|---|---|---|---|
| Fraud-dollar capture rate | Measures direct loss prevention quality | Fraud strategy | Daily / weekly |
| False-positive rate | Tracks customer and merchant friction | Fraud strategy + product | Daily / weekly |
| Manual-review rate | Signals operational scalability | Fraud operations | Daily |
| Hold-to-approval conversion | Shows whether review adds value | Fraud operations | Weekly |
| Return/dispute rate by merchant cohort | Separates portfolio risk from rail risk | Risk + merchant teams | Weekly / monthly |
| Model drift and calibration | Confirms that scores remain meaningful over time | Model risk / data science | Monthly |

**Funding:** This research received no external funding.
**Conflicts of Interest:** The authors declare no conflict of interest.


**References**

[1]. Abdallah, A., Maarof, M. A., & Zainal, A. (2016). Fraud detection system: A survey. Journal of Network and Computer Applications, 68, 90-113. https://doi.org/10.1016/j.jnca.2016.04.007
[2]. Afriyie, J. K., Tawiah, K., Pels, W. A., Addai-Henne, S., Dwamena, H. A., Owiredu, E. O., Ayeh, S. A., & Eshun, J. (2023). A supervised machine learning algorithm for detecting and predicting fraud in credit card transactions. Decision Analytics Journal, 6, 100163. https://doi.org/10.1016/j.dajour.2023.100163
[3]. Ali, A., Al-Hashedi, K. G., Jamil, M. M. A., Majid, A. H. A., Ariffin, A. F. M., & Sentosa, I. (2022). Financial fraud detection based on machine learning: A systematic literature review. Applied Sciences, 12(19), 9637. https://doi.org/10.3390/app12199637
[4]. Arman, M., & Fahim, A. S. M. (2023). Ai revolutionizes inventory management at retail giants: Examining Walmart's U.S. operations. Journal of Business and Management Studies, 5(6), 145-148. https://doi.org/10.32996/jbms.2023.5.6.15

[5]. Board of Governors of the Federal Reserve System. (2019). Federal Reserve actions to support interbank settlement of faster payments.
[6]. https://www.federalreserve.gov/newsevents/pressreleases/other20190805a.htm
[7]. Board of Governors of the Federal Reserve System. (2020). The future of retail payments in the United States. https://www.federalreserve.gov/newsevents/speech/brainard20200806a.htm
[8]. Board of Governors of the Federal Reserve System. (2023a). The Federal Reserve Payments Study: 2022 triennial initial data release. https://www.federalreserve.gov/paymentsystems/fr-payments-study.htm
[9]. Board of Governors of the Federal Reserve System. (2023b). Federal Reserve announces that its new system for instant payments, the FedNow Service, is now live.

[10]. https://www.federalreserve.gov/newsevents/pressreleases/other20230720a.htm
[11]. Board of Governors of the Federal Reserve System. (2024). Cybersecurity and financial system resilience report. https://www.federalreserve.gov/publications/files/cybersecurity-report-202407.pdf
[12]. Bockel-Rickermann, C., Tesch, J. F., & Gimpel, H. (2023). Fraud analytics: A decade of research: Organizing the domain and identifying future directions. Expert Systems with Applications, 236, 121335. https://doi.org/10.1016/j.eswa.2023.121335
[13]. Consumer Financial Protection Bureau. (2024). Required rulemaking on personal financial data rights. https://www.consumerfinance.gov/personal-financial-data-rights/
[14]. 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
[15]. 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
[16]. Fahim, A. S. M., Pritty, A. A., Ibrahim, M., & Tania, T. A. (2025). Explainable AI for medical debt forecasting: Integrating healthcare and Fintech data for risk prediction. Journal of Management World, 2025(6), 92-103. https://doi.org/10.53935/jomw.v2024i4.1253
[17]. FedPayments Improvement. (2022). Next-level detection through machine learning. https://fedpaymentsimprovement.org/wp-content/uploads/next-level-detection-through-machine-learning.pdf
[18]. Federal Trade Commission. (2024a, February 9). As nationwide fraud losses top $10 billion in 2023, FTC steps up efforts to protect the public. https://www.ftc.gov/news-events/news/press-releases/2024/02/nationwide-fraud-losses-top-10-billion-2023-ftc-steps-efforts-protect-public
[19]. Federal Trade Commission. (2024b). Consumer Sentinel Network data book 2023. https://www.ftc.gov/system/files/ftc_gov/pdf/CSN-Annual-Data-Book-2023.pdf
[20]. 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. https://www.eudoxuspress.com/index.php/pub/article/view/3823
[21]. Hwang, B. H. (2025). Pay-by-bank and the merchant payments use case: Benefits, risks and potential impacts on consumer payment behaviors in the U.S. FEDS Notes. https://doi.org/10.17016/2380-7172.3834
[22]. Ibrahim, M., Fahim, A. S. M., Zadid, M. U., & Pritty, A. A. (2025). FinTech for climate resilience: Measuring insurance gaps, mortgage stress, and household credit risk in the United States. Journal of Economics, Finance and Accounting Studies, 7(4), 190-205. https://doi.org/10.32996/jefas.2025.7.4.15
[23]. Ibrahim, M., Mahmud, S., Zadid, M. U., Jahan, N., Rahman, M. M., & Fahim, A. S. M. (2024). AI-driven predictive analytics framework for anti-money laundering risk management and financial infrastructure protection in U.S. banking systems. Journal of Economics, Finance and Accounting Studies, 6(1), 155-166. https://doi.org/10.32996/jefas.2024.6.6.12
[24]. Ibrahim, M., Razib, M. N. H., Jahan, N., & Rahman, M. M. (2022). Climate risk, financial stability, and global capital allocation: A predictive analytics approach to assessing climate-related financial risk in international investment markets. Journal of Business and Management Studies, 4(4), 264-276. https://doi.org/10.32996/jbms.2022.4.4.34
[25]. Jahan, N., Pritty, A. A., Ibrahim, M., Zadid, M. U., Fahim, A. S. M., & Mahmud, S. (2024). Machine learning-driven early warning analytics for identifying market manipulation, irregular trading activity, and suspicious market signals in U.S. stock markets. Journal of Computer Science and Technology Studies, 6(2), 257-283. https://doi.org/10.32996/jcsts.2024.6.2.26
[26]. 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. https://doi.org/10.1016/j.eswa.2018.01.037
[27]. Mahmud, S., Fahim, A. S. M., Rahman, M. M., Jahan, N., & Ibrahim, M. (2025). Artificial intelligence and predictive machine learning for financial fraud detection, cyber risk management, and infrastructure resilience in the U.S. banking industry. British Journal of Multidisciplinary Studies, 3(1), 58-77. https://doi.org/10.32996/bjmss.2025.4.1.6

[28]. Nacha. (2023a, February 22). ACH network moves 30 billion payments, $77 trillion in 2022 led by growth in Same Day ACH and B2B. https://www.nacha.org/news/ach-network-moves-30-billion-payments-77-trillion-2022-led-growth-same-day-ach-and-b2b
[29]. Nacha. (2024a, February 14). ACH network records strong growth in 2023 as Same Day ACH surpasses 3 billion payments since inception. https://www.nacha.org/news/ach-network-records-strong-growth-2023-same-day-ach-surpasses-3-billion-payments-inception
[30]. Nacha. (2025a, January 30). Same Day ACH passes major milestone in 2024 as the ACH network shows higher growth. https://www.nacha.org/news/same-day-ach-passes-major-milestone-2024-ach-network-shows-higher-growth
[31]. Nacha. (2025b). Same Day ACH. https://www.nacha.org/same-day-ach
[32]. National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework (AI RMF 1.0). https://doi.org/10.6028/NIST.AI.100-1
[33]. 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. https://doi.org/10.1016/j.dss.2010.08.006
[34]. 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
[35]. Rasel, I. H., Ibrahim, M., Pritty, A. A., Fahim, A. S. M., & Jahan, N. (2023). Beyond FICO: Enhancing mortgage default forecasting and inclusive lending via multimodal graph neural networks and urban mobility analytics. Frontiers in Computer Science and Artificial Intelligence, 2(2), 62-81. https://doi.org/10.32996/fcsai.2023.2.2.5
[36]. Razib, M., Ibrahim, M., & Rasel, I. (2025). Predictive analytics and its role in optimizing sustainable supply chain performance. International Journal of Business Management and Economic Review, 8(01), 12-27. http://doi.org/10.35409/IJBMER.2025.3640
[37]. Shah, A., Kabir, S., Razib, M. N. H., & Khan, S. A. (2024). Grid-integrated EV charging systems: Impacts on U.S. power grid stability and resilience. International Journal of Artificial Intelligence Engineering and Transformation, 5(1), 40-45. https://doi.org/10.54660/IJAIET.2024.5.1.40-45
[38]. Ti, Y.-W., Hsu, C.-W., Liao, C.-H., & Lin, S.-D. (2022). Feature generation and contribution comparison for electronic fraud detection. Scientific Reports, 12, 18211. https://doi.org/10.1038/s41598-022-22130-2
[39]. Imran Hossain Rasel, Md Nurul Huda Razib, & Muhaimin Ul Zadid. (2026). Explainable AI for Institutional Fraud Decisions: A Cross-Sector Empirical Study Using Public Healthcare and Financial Transaction Data. Journal of Computer Science and Technology Studies, 8(1), 97-106. https://doi.org/10.32996/jcsts.2025.8.1.7
[40]. Weber, P., Temper, M., & Dudek, C. (2024). Applications of explainable artificial intelligence in finance-a systematic review of finance, information systems, and computer science literature. Management Review Quarterly, 74, 1137-1181. https://doi.org/10.1007/s11301-023-00320-0
[41]. West, J., & Bhattacharya, M. (2016). Intelligent financial fraud detection: A comprehensive review. Computers & Security, 57, 47-66. https://doi.org/10.1016/j.cose.2015.09.005
[42]. Xu, B., Zhan, C., Chen, Y., Zhang, L., & Qi, X. (2023). Efficient fraud detection using deep boosting decision trees. Decision Support Systems, 175, 114038. https://doi.org/10.1016/j.dss.2023.114038
[43]. Zhou, Y., Li, H., Xiao, Z., & Qiu, J. (2023). A user-centered explainable artificial intelligence approach for financial fraud detection. Finance Research Letters, 58, 104309. https://doi.org/10.1016/j.frl.2023.104309