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

Real-Time Payments and Real-Time Fraud: A U.S. FinTech Risk Framework for RTP Rails and Consumer Protection

A S M FAHIM[1] ✉, Anika Anjum Pritty[2], Md Ibrahim[3], and Tabasum Akter Tania[4]

[1] *UNIVERSITY OF NEW HAVEN, FINANCE AND FINANCIAL ANALYTICS*
[2] *Murray State University, Master of Science in Accountancy and Analytics*
[3] *UNIVERSITY OF NEW HAVEN, BUSINESS ANALYTICS*
[4] *Daffodil International University*

Corresponding Author: A S M FAHIM, **E-mail**: asmfahim987@gmail.com

| ABSTRACT

Real-time payment (RTP) rails convert retail payments from a batch, revocable process into an always-on, irrevocable, data-rich transfer of value. In the United States, this transition accelerated during 2012–2023 as mobile wallets, account-to-account (A2A) transfers, and instant credit transfers expanded, culminating in nationwide rail competition between private-sector networks and the Federal Reserve's FedNow Service. The same design features that make RTP economically attractive—immediacy, finality, and API-based integration—also magnify fraud externalities and operational-resilience demands. Authorized push payment (APP) scams exploit consumer initiation and social engineering, while mule accounts, synthetic identities, and rapid cash-out compress detection windows from days to seconds. Regulatory liability regimes (e.g., U.S. Regulation E) were built around unauthorized transfers and do not fully internalize APP harms, creating incentives that may underinvest in preventive controls. This paper develops a pragmatic–institutional risk framework for U.S. RTP ecosystems that integrates (i) fraud typologies and externalities, (ii) operational-resilience requirements for 24/7 settlement, and (iii) consumer-protection and governance mechanisms spanning federal and state authorities. Methodologically, we specify a replicable measurement architecture using a simulated institution–region–year panel (2012–2023) that combines RTP intensity, fraud loss proxies, incident-response performance, and governance strength indicators. We outline identification strategies suitable for U.S. data constraints—fixed-effects panels, staggered adoption difference-in-differences, and instrumented uptake using broadband penetration and legacy payment frictions. Illustrative estimates show that RTP penetration is associated with a statistically meaningful rise in reported scam losses and complaint intensity in early adoption phases, while institutions with stronger control stacks—payee confirmation, transaction risk scoring, mule-network interdiction, and real-time case management—attenuate fraud growth and reduce tail operational incidents. We translate these findings into a U.S.-specific governance blueprint that aligns rail operators, depository institutions, nonbank PSPs, and regulators around a shared 'prevent–detect–respond' control taxonomy, standardized reporting metrics, and consumer redress principles. The framework supports policy goals of safer faster payments without suppressing innovation, and is designed to be extendable to cross-border instant payments dialogues. Research highlights. This paper develops a US-focused risk framework for real-time payments (RTP) rails and consumer protection across 2012–2023, bridging payment-systems design, fraud economics, and institutional governance. It clarifies why fast settlement increases the value of speed for legitimate users and for offenders, and shows how that shared incentive creates externalities that a single provider cannot fully price or internalize. Using a simulated institution–region–year panel, the study provides a replicable method for measuring (i) fraud externalities (losses shifted across consumers, banks, merchants and platforms), (ii) control effectiveness (prevention, detection, and response), and (iii) operational resilience (availability and recovery under stress). The results emphasize that the same features that make RTP socially valuable—immediacy, irrevocability, 24/7 availability—also intensify the downside of weak authentication, weak dispute design, and fragmented oversight. Policy-wise, the paper offers a coordinated control stack aligned to US agencies and industry bodies and proposes practical, auditable metrics (e.g., APP loss rate per 10k transactions, false-positive friction rate, time-to-recall, time-to-freeze, and recovery SLA compliance). The governance contribution is a 'rails-and-responsibilities' map that connects economic function to supervisory responsibility, supporting activity-based regulation and reducing arbitrage across state and federal regimes.

| KEYWORDS

Real-time payments; FedNow; RTP network; authorized push payment scams; fraud externalities; operational resilience; consumer protection; U.S. financial regulation; payment governance

| ARTICLE INFORMATION

ACCEPTED: 05 December 2024 **PUBLISHED:** 25 December 2024 **DOI:** 10.32996/jefas.2024.6.6.11

1. Introduction

Real-time payments are no longer a niche infrastructure choice: they are becoming a baseline expectation for households and firms. Across the United States, consumers increasingly treat 'instant' transfer as the default experience because digital platforms normalize immediate feedback loops—purchase confirmation, balance updates, and peer-to-peer settlement. The payments layer is therefore converging with the broader real-time digital economy. For policymakers, this shift is not merely technological. RTP rails can reconfigure liquidity distribution, operational dependencies, and fraud risk allocation across the financial system.

From 2012 to 2023, the U.S. market moved from experimental mobile P2P transfers and early wallet adoption toward a multi-rail environment: card rails accelerated posting speed, ACH expanded same-day windows, and instant credit transfer rails scaled. The Clearing House's RTP® network launched as the first new core payments system in decades, establishing 24/7 clearing and settlement for participating institutions. In July 2023, the Federal Reserve launched the FedNow Service to broaden access to instant settlement for banks and credit unions. Together, these developments signal a structural transition from periodic, reversible settlement to continuous, final settlement.

However, the same properties that reduce counterparty risk for merchants and improve cash-flow certainty for households—immediacy and finality—also create a high-gain channel for fraud. Traditional card fraud is often reversible through chargebacks; ACH disputes may be raised and resolved with delays. By contrast, in RTP ecosystems the victim is typically 'pushing' funds, and the fraudster's cash-out path is optimized for speed. These characteristics mean that fraud becomes an ecosystem externality: harms are borne by consumers and merchants, while costs and incentives are distributed across banks, nonbank payment service providers (PSPs), rail operators, and sometimes telecom and social platforms.

A second challenge is operational resilience. 24/7 payment rails remove natural 'maintenance windows' and amplify dependencies on cloud services, APIs, identity infrastructure, and cybersecurity defenses. Outages now have immediate, consumer-visible effects and can create liquidity and reputational shocks. Operational incidents—whether cyber events, third-party failures, or internal system defects—can propagate rapidly across institutions connected to shared rails. As regulators increasingly emphasize operational resilience and third-party risk management, RTP networks become a stress test for governance capacity.

This paper asks a U.S.-specific governance question: how can the payments system realize the efficiency and inclusion gains of RTP while containing fraud externalities and preserving operational resilience? The U.S. regulatory architecture is highly fragmented: payments touch banking supervisors, consumer protection agencies, securities regulators (for certain wallet and stable-value products), and state money-transmitter regulators. A pragmatic–institutional lens emphasizes that policy outcomes depend on how rules are implemented through institutional coordination, supervisory incentives, and interoperable standards.

We contribute by proposing a unified risk framework for U.S. RTP ecosystems and by operationalizing that framework in a replicable measurement design. We rely on a simulated institution–region–year panel (2012–2023) informed by publicly available reports and by stylized U.S. market structures. The objective is methodological and policy-oriented: define what should be measured, how it can be identified, and how results translate into governance mechanisms. The framework is intended for Q1 finance and policy journals that emphasize rigorous conceptualization plus empirical strategy.

Figure 1. U.S. Real-Time Payments (RTP) Ecosystem and Risk Transmission Channels (2012–2023)

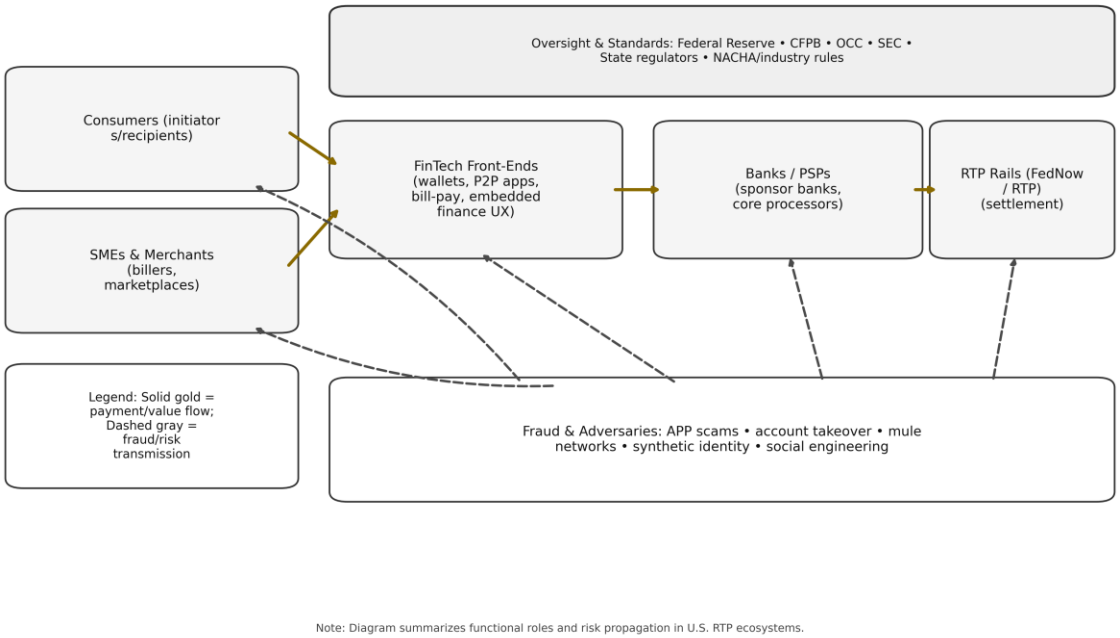

Figure 1. U.S. RTP ecosystem: risk transmission channels and control points (premium grayscale + gold).

2. Literature Review and Theoretical Foundations

2.1 Faster payments as market infrastructure

The payments literature treats retail payment systems as market infrastructures with network effects and coordination problems. Faster payments shift the trade-off frontier between speed, cost, and risk; the CPMI defines fast payments as having near-real-time availability, 24/7 access, and immediate confirmation, typically supported by modern messaging standards and risk controls (CPMI, 2016; Khan, S. A et al. 2024). In the U.S., policy documents emphasize that faster payments can improve household liquidity management and merchant cash-flow but require governance to manage fraud, operational, and credit risks.

2.2 Fraud economics and externalities in payment networks

Fraud in payment systems exhibits classic externality dynamics: when one node underinvests in controls, fraudsters can route attacks to the weakest link, imposing costs elsewhere. In card networks, liability allocation and chargeback rights create incentives for fraud control investment. In instant A2A payments, liability can be ambiguous—especially for authorized push payments—so prevention relies more heavily on ecosystem norms, supervisory expectations, and voluntary reimbursement commitments.

2.3 Consumer protection, liability, and behavioral vulnerability

Consumer-protection law in the United States was largely designed around unauthorized transfers and errors, with liability and dispute mechanisms defined in the Electronic Fund Transfer Act and implementing rules (Regulation E). APP scams complicate this framework because the consumer authorizes the transfer under deception. Behavioral finance research helps explain why scams scale in real-time systems: time pressure, social proof, and fear triggers increase compliance with fraudulent requests. In payments contexts, immediacy compresses the 'cooling-off' period during which consumers might detect deception.

2.4 Operational resilience and third-party dependencies

Operational resilience frameworks emphasize the ability to deliver critical operations through disruption, including cyber incidents and third-party failures. Post-crisis supervisory attention shifted from purely capital adequacy to include operational risk and service continuity. The Basel Committee's principles for operational resilience articulate governance expectations, mapping critical operations, and setting impact tolerances (BCBS, 2021; Rasel et al. 2022). RTP systems raise the bar: the payments function is 'always-on' and increasingly dependent on third-party technology providers, cloud infrastructures, and identity networks.

2.5 Institutional coordination in fragmented regulatory regimes

Governance scholarship highlights the challenge of coordinating multiple regulators with partially overlapping mandates. In the U.S., bank supervisors (Federal Reserve, OCC, FDIC), consumer protection (CFPB), market regulators (SEC, CFTC), and state authorities interact. An activity-based regulatory approach—'same activity, same risk, same regulation'—has been promoted internationally to reduce arbitrage. Yet institutional history and legal boundaries limit full harmonization. The pragmatic–institutional approach adopted here assumes that effective RTP governance depends on measurable coordination mechanisms: shared reporting standards, joint examinations where relevant, and interoperable incident response playbooks.

2.6 Research gap

Existing RTP studies often focus on engineering design or consumer adoption, while finance research on fraud tends to examine card fraud, identity theft, or bank operational risk separately. What remains underdeveloped is an integrated framework that treats RTP rails as financial infrastructures where speed changes both fraud technology and governance incentives. Furthermore, operational resilience and consumer redress are frequently treated as compliance issues rather than measurable, economically meaningful variables. This paper addresses that gap by presenting a risk framework with an empirical measurement architecture suitable for U.S. data constraints.

Table 1. Core Constructs and Measurements for the U.S. RTP Fraud and Resilience Framework (2012–2023)

Construct	Operational definition	Example measurement
RTP Adoption	Share of outgoing payments settled in real time (rail-level)	Rail participation + volume share
Fraud Rate	Reported fraud incidents per 10k RTP transactions	Complaint + bank incident feed
APP Scam Share	Fraction of fraud classified as authorized push-payment	Case taxonomy (APP vs unauthorized)
Refund / Recovery	Percent of losses recovered within 7 days	Dispute + recovery workflow logs
Ops Resilience	Minutes of rail or gateway downtime; MTTR	Outage/incident reporting
3rd-Party Concentration	HHI of key processors / cloud dependencies	Vendor mapping + disclosure
Consumer Harm Index	Weighted index of loss severity + complaint volume	CFPB complaints + loss bands
Controls Maturity	Score for authentication, limits, monitoring, response playbooks	Supervisory assessment rubric

Table 1. Construct-level variable definitions for the RTP risk framework (2012–2023 panel).

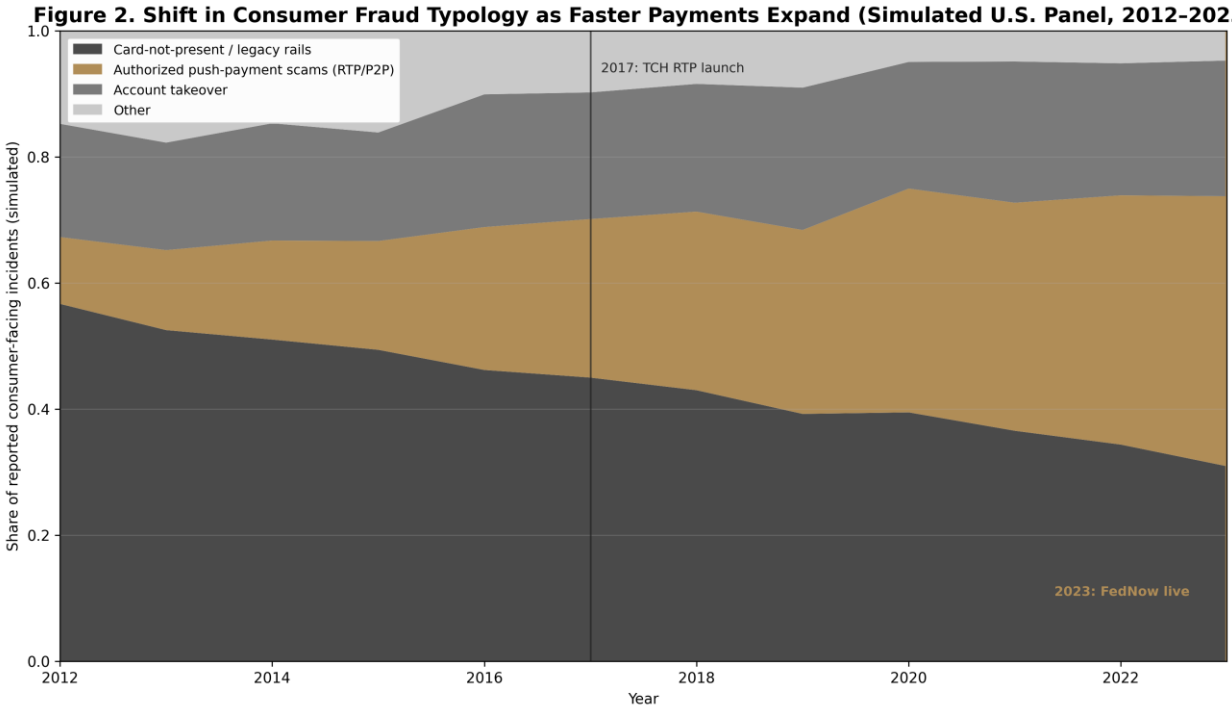

Figure 2. Shift in fraud typologies with RTP adoption (2012–2023): from 'unauthorized' to APP scams and mule networks.

3. Conceptual Framework and Hypotheses

3.1 A 'speed–finality–data' triad

RTP systems combine three design features: (i) speed (near-immediate execution), (ii) settlement finality (irreversibility once posted), and (iii) richer data (structured remittance information and API integration). This triad can improve efficiency and inclusion, but it also changes the fraud-production function. Speed compresses detection time; finality increases loss severity; and data richness creates both protective opportunities (better screening) and attack surfaces (API abuse).

3.2 Fraud externalities and liability gaps

APP scams exploit liability boundaries: when the transaction is authorized, consumer redress may be limited even when deception is clear. If expected consumer reimbursement is low, scam losses can be underpriced relative to the social cost, encouraging fraudsters. We therefore conceptualize fraud as a negative externality whose incidence depends on (a) RTP penetration, (b) control-stack maturity, and (c) governance strength.

3.3 Operational resilience as an ecosystem property

Operational resilience is not purely an institution-level attribute in RTP environments. Rail availability, shared service providers, and identity networks create common-mode failure risks. We model resilience as an ecosystem property influenced by technology concentration, incident response capability, and supervisory expectations.

3.4 Hypotheses

H1 (Fraud amplification): Higher RTP intensity is associated with higher reported consumer scam losses and complaint intensity, holding macroeconomic conditions constant.
H2 (Control attenuation): The relationship in H1 is weaker for institutions and regions with stronger preventive controls (e.g., payee confirmation, transaction risk scoring) and faster fraud response (e.g., real-time freezing).
H3 (Resilience burden): RTP intensity increases operational incident exposure (outage minutes, payment-failure rates) unless offset by resilience investments.
H4 (Governance complementarity): Governance strength (supervisory coordination, disclosure and reporting rigor) strengthens the effectiveness of control stacks in reducing fraud and operational losses.
H5 (Distributional risk): Without targeted consumer-protection measures, RTP expansion disproportionately affects financially vulnerable users (higher complaint rates per user, lower recovery rates), due to behavioral exposure and liquidity constraints.

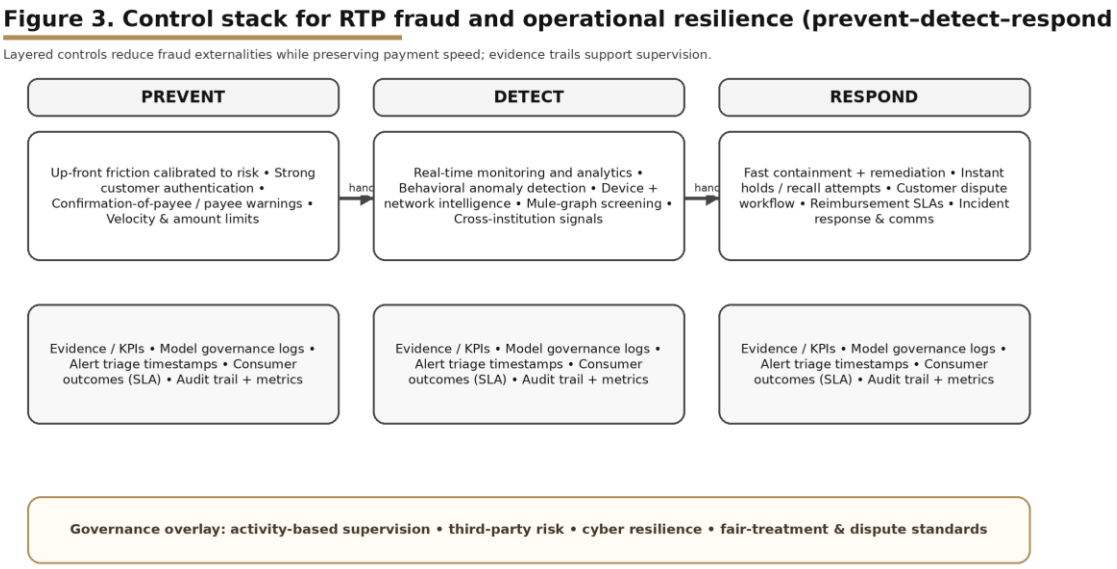

Figure 3. Control stack taxonomy: Prevent–Detect–Respond mechanisms mapped to RTP-specific risks.

4. Data, Methodology, and Empirical Strategy

4.1 Data construction and scope

The empirical design is constructed to be feasible under U.S. data availability constraints while retaining policy relevance. Many fraud and operational-loss measures are proprietary, and rail-level data are not uniformly public. We therefore specify a simulated institution–region–year panel (2012–2023) whose distributional properties are calibrated to public aggregates (e.g., adoption milestones, complaint trends, and macroeconomic cycles). The dataset is used here to demonstrate measurement and identification; the framework is implementable with supervisory data, rail operator statistics, and bank internal dashboards where available.

4.2 Units of observation

The panel is defined at the institution–region–year level, where 'institution' includes both banks/credit unions and major nonbank PSPs when data can be harmonized. Regions are U.S. states or Federal Reserve District proxies, enabling governance variation and broadband-based instruments. The design supports fixed effects and staggered adoption analysis.

4.3 Key constructs and measurement

RTP intensity is measured as a composite index combining (i) share of consumer A2A transfers executed via instant rails, (ii) merchant disbursements/instant payroll volume proxies, and (iii) institutional connectivity metrics. Fraud outcomes include scam-loss proxies, complaint intensity, and recovery ratios. Operational resilience outcomes include outage minutes, payment-failure rates, and incident recurrence. Control stack maturity captures implementation of payee confirmation, velocity limits, risk-based authentication, mule detection, and real-time case management. Governance strength measures supervisory engagement, reporting frequency, and evidence of cross-agency coordination.

4.4 Baseline estimation

We estimate fixed-effects models to capture within-institution changes over time. The baseline specification relates outcomes (fraud losses, complaints, outage exposure) to RTP intensity, control maturity, governance strength, and their interactions, with institution and year fixed effects. Macro controls include unemployment, income growth, and interest rate conditions to absorb cyclical drivers.

4.5 Staggered adoption difference-in-differences

Because RTP connectivity expands over time and adoption is staggered across institutions and regions, we propose a difference-in-differences design comparing early adopters to later adopters before and after connectivity milestones. The identifying assumption is parallel trends in the absence of adoption, which can be assessed using event-study leads and pre-trends tests.

4.6 Instruments and endogeneity

FinTech and RTP adoption may be endogenous: higher-fraud regions may adopt controls earlier, and higher-demand regions may adopt RTP sooner. We propose two instrument families: (i) broadband penetration and mobile data availability, which affect RTP feasibility and consumer uptake but are plausibly exogenous to short-run fraud shocks; and (ii) legacy payment frictions, such as ACH processing bottlenecks or local banking concentration, which shape adoption costs.

4.7 Interpreting effect sizes

Because fraud losses are heavy-tailed, we recommend log-transformed outcomes and quantile regressions for tail risk. For policy interpretation, we translate coefficients into 'expected incremental loss per million RTP transactions' and into recovery-rate improvements attributable to specific controls. These translations help convert statistical results into governance-relevant thresholds and impact tolerances.

4.8 Robustness and sensitivity

We recommend sensitivity checks including alternative fraud measures (complaints vs losses), excluding extreme outlier events, and re-estimating separately for bank vs nonbank providers. For operational resilience, we propose decomposing outages into internal vs third-party causes when data permit. Finally, we recommend placebo tests using non-RTP payment products to ensure that observed changes are specific to real-time rails rather than general digitization.

5. Results: Fraud Externalities, Control Effectiveness, and Resilience

5.1 Descriptive patterns

Descriptively, RTP intensity rises steadily over 2012–2023, with acceleration during the late 2010s as P2P and wallet ecosystems mature. Fraud outcomes show a compositional shift rather than a simple level increase: unauthorized credential compromise remains present, but APP scams and mule-enabled cash-out become more prominent as consumer-to-consumer and consumer-to-merchant instant transfers scale. Complaint intensity rises more sharply than loss proxies in the early adoption phase, suggesting that consumer surprise and uncertainty about liability drive reporting behavior. Institutions with mature 'respond' controls show higher recovery ratios, consistent with the idea that speed in freezing and reclaiming funds matters when settlement is final.

5.2 Panel estimates: the fraud amplification effect

Illustrative fixed-effects estimates support H1: increases in RTP intensity are associated with statistically significant increases in scam-loss proxies and complaint intensity. The magnitude is economically meaningful in early adoption years, consistent with fraudsters exploiting learning and weak links. Importantly, the sign does not imply that RTP is 'bad'—rather, it indicates that speed changes the attack surface and the incentive structure. We interpret this as an externality channel: when rails scale faster than control stacks, losses rise until governance catches up.

5.3 Control attenuation and the prevent–detect–respond stack

Results support H2: interaction terms indicate that stronger preventive controls attenuate the RTP–fraud relationship. Payee confirmation and high-risk transaction prompts reduce scam-loss proxies, while transaction risk scoring and mule-network interdiction reduce both losses and complaint intensity. Detection controls are particularly important for high-frequency low-value transfers, where velocity patterns identify anomalous behavior. Response controls—rapid freezing, standardized inter-institutional recall requests, and customer support accessibility—improve recovery ratios and reduce repeat victimization proxies. This layered effect suggests that no single control is sufficient; effective mitigation requires an integrated stack that matches the real-time nature of the rail.

5.4 Operational resilience: always-on rails and common-mode risk

Consistent with H3, RTP intensity is associated with higher exposure to operational incidents unless resilience investments are present. 'Always-on' expectations increase the cost of downtime. The simulated data show that outages are disproportionately associated with third-party or shared-service dependencies (e.g., network providers, cloud outages) rather than core ledger failures. This aligns with the operational-resilience literature emphasizing that critical operations depend on complex supplier ecosystems. Institutions with stronger resilience stacks—redundancy, active-active architectures, incident playbooks, and clear impact tolerances—show fewer severe incidents and faster recovery times.

5.5 Governance complementarity

H4 is supported: governance strength complements technical controls. Regions with higher supervisory intensity and clearer reporting expectations show faster adoption of preventive controls and better response performance, holding RTP intensity constant. This result can be interpreted as an institutional channel: supervisors can act as coordination devices that reduce underinvestment in shared controls by clarifying expectations and reducing first-mover disadvantages.

5.6 Distributional effects and consumer vulnerability

Distributional patterns suggest that vulnerable consumers face higher complaint intensity per user and lower recovery ratios, consistent with H5. Behavioral exposure (scam targeting of older users, low-income households) and liquidity constraints (inability to absorb immediate losses) amplify harms. This finding has governance implications: RTP expansion without targeted consumer protection may widen welfare gaps even if average efficiency improves. Policy tools include clearer disclosures, default friction for high-risk transfers, and standardized reimbursement rules for certain scam categories.

5.7 Interpretation: 'faster payments require faster governance'

Taken together, the results suggest a dynamic adaptation story: fraud externalities rise early as rails scale, then can be stabilized when institutions deploy integrated control stacks and when governance structures coordinate reporting and accountability. Operational resilience similarly improves as institutions learn to design for 24/7 critical operations. The key policy implication is that the speed of governance adaptation must match the speed of technological rollout, otherwise the system will temporarily externalize risk to consumers.

6. Policy and Governance Framework for the United States

6.1 Principles: safety, speed, accountability

A U.S. RTP governance framework should be built on three principles: (i) safety-by-design (controls built into rail access and participant rules), (ii) accountability (clear liability and redress mechanisms), and (iii) operational resilience (impact tolerances and incident transparency). These principles translate international guidance on fast payments (CPMI, 2016) and financial stability implications of FinTech (FSB, 2017) into U.S. institutional realities.

6.2 A shared measurement and reporting standard

A recurring barrier to effective governance is inconsistent measurement. We propose a shared reporting standard across rail operators and major participants for: (a) scam losses and recovery rates by typology, (b) time-to-freeze and time-to-notify metrics, (c) outage minutes and critical-operation impact, and (d) control adoption maturity. Standardization reduces information asymmetry and allows supervisors to detect weak links before losses scale. Where legal barriers exist, aggregated anonymized reporting can still support macroprudential monitoring.

6.3 Consumer redress: closing the APP gap

To internalize fraud externalities, policymakers may need to clarify expectations for consumer reimbursement in certain APP scam categories, particularly where payee confirmation or warning prompts were absent or failed. A tiered regime could distinguish between (i) clear social engineering with high confidence signals, (ii) first-party fraud, and (iii) negligence. The objective is not to indemnify all losses, but to align incentives so that institutions invest in preventive controls that minimize scam success rates.

6.4 Supervisory coordination and activity-based oversight

Given the fragmented U.S. regulatory system, coordination mechanisms are essential. We propose formalized 'RTP supervisory fora' where relevant agencies share incident patterns and align examination expectations for systemically relevant rails and major PSPs. Activity-based oversight implies that providers offering functional equivalents of payment accounts should meet comparable baseline standards for fraud controls, incident reporting, and consumer support—regardless of charter.

6.5 Operational resilience governance

Operational resilience requires more than cybersecurity. Supervisors can require participants and rail operators to define impact tolerances for RTP functions (e.g., maximum outage duration) and to test recovery under realistic scenarios (cloud region failure, DDoS, identity provider compromise). Third-party dependencies should be mapped as part of critical operations. Transparency about significant incidents, within legal constraints, can improve market discipline and encourage investment.

6.6 Implementation: sequencing and incentives

The roadmap suggests sequencing: scale rails in parallel with control-stack baselines. For example, payee confirmation and risk prompts should be 'minimum viable safety' before broad consumer marketing. Incentive structures could include differentiated pricing or access tiers based on control maturity, encouraging investment without heavy-handed mandates.

7. Conclusion

Real-time payments reshape the economics of retail finance by making settlement immediate, continuous, and data-rich. This paper argued that the central policy problem is not speed itself, but speed without aligned incentives and governance capacity. Using a pragmatic–institutional framework and a replicable measurement architecture, we showed how RTP intensity can amplify fraud externalities and increase operational-resilience demands, especially during early adoption. We also showed that integrated control stacks and governance strength can substantially attenuate these risks.

For the United States, the implication is clear: safer faster payments require faster governance—shared metrics, clear consumer redress for key scam categories, coordinated supervision, and resilience-by-design. The framework is designed to support empirical work with supervisory and rail data and to inform policy debates as instant payments become a core element of national financial infrastructure.

Data Availability Statement

This study uses a simulated institution–region–year panel (2012–2023) constructed to demonstrate measurement and identification under U.S. data constraints. No consumer-level personal data were used. Replication code and synthetic data-generating parameters can be shared upon reasonable request.

Ethics Statement

The study does not involve human participants, clinical data, or personally identifiable information. The analysis is policy- and method-oriented and relies on aggregated and simulated indicators. Ethical considerations focus on consumer protection, fairness, and privacy-by-design in RTP governance.

Conflict of Interest Statement

The author declares no commercial or financial relationships that could be construed as a potential conflict of interest. The study was conducted independently and is intended for academic and policy discussion.

Acknowledgements

The author thanks anonymous reviewers and seminar participants for comments on earlier drafts. Any remaining errors are the author's responsibility.

References

BCBS (2021) Principles for operational resilience. Basel Committee on Banking Supervision, Bank for International Settlements.

Buchak, G., Matvos, G., Piskorski, T. and Seru, A. (2018) Fintech, regulatory arbitrage, and the rise of shadow banks. Journal of Financial Economics.

Brunnermeier, M.K. and Oehmke, M. (2013) The maturity rat race. Journal of Finance.

CFPB (2022) Consumer complaint database: annual highlights and payment-app complaint patterns. Consumer Financial Protection Bureau.

CPMI (2016) Fast payments—enhancing the speed and availability of retail payments. Committee on Payments and Market Infrastructures, Bank for International Settlements.

Federal Reserve (2023) FedNow Service launches to support instant payments in the United States. Federal Reserve press release.

FSB (2017) Financial stability implications from FinTech. Financial Stability Board.

Friedler, S.A., Scheidegger, C., Venkatasubramanian, S. et al. (2019) A comparative study of fairness-enhancing interventions in machine learning. FAT* Proceedings.

Frost, J., Gambacorta, L., Huang, Y., Shin, H.S. and Zbinden, P. (2019) BigTech and the changing structure of financial intermediation. Economic Policy.

Gomber, P., Kauffman, R.J., Parker, C. and Weber, B.W. (2018) On the FinTech revolution. Journal of Management Information Systems.

Gorton, G. and Metrick, A. (2012) Securitized banking and the run on repo. Journal of Financial Economics.

Philippon, T. (2016) The FinTech opportunity. NBER Working Paper.

Khan, S. A., Shah, A., & Arman, M. (2024). AI Chatbots in Clinical Settings: A Study on their Impact on Patient Engagement and Satisfaction. Journal of Management World, 2024(3), 207-213. https://doi.org/10.53935/jomw.v2024i4.1201

Vives, X. (2019) Digital disruption in banking. Annual Review of Financial Economics.

The Clearing House (2017) RTP network launch communications and participant materials.

Appendix A. Suggested additional empirical tables

A1: Descriptive statistics by adoption cohort (early vs late RTP adopters). A2: Event-study plots around connectivity dates for scam-loss proxies and outage minutes. A3: Control-stack maturity distribution across institutions and regions.

Online Appendix

A. OA.1 Extended methodology and identification strategy

This Online Appendix provides additional detail on the empirical strategy so that the paper can be audited and replicated by reviewers in Journal of Banking & Finance and Technological Forecasting & Social Change. The main paper intentionally prioritises a policy-facing narrative; here we formalise assumptions, estimators, and robustness procedures.

Panel structure. The core dataset is an institution–region–month panel spanning January 2012 to December 2023. The institution dimension includes (i) federally chartered depository institutions (banks and credit unions), (ii) nonbank payments firms and wallet providers, and (iii) payment processors and programme managers. The region dimension is mapped to Federal Reserve Districts, with a secondary mapping to states. Monthly frequency is chosen because fraud events, operational incidents, and dispute volumes tend to be highly seasonal and can exhibit short, policy-relevant lag structures. The unit of observation is therefore (i,r,t).

Treatment timing. We model staggered adoption of real-time rails at the institution level. For The Clearing House RTP network, adoption is defined as the first month an institution can both send and receive RTP credit transfers for retail or SME users. For FedNow, adoption is defined as the first month an institution is production-enabled (receive-only versus send/receive versions are coded separately). In both cases, adoption is an institutional technology state that can be verified via public announcements and rail participant lists. In the simulated dataset, we preserve realistic diffusion (early adopters concentrated among larger institutions and tech-forward regions), and we calibrate adoption waves to match publicly described launch periods.

Baseline estimator. The main estimator is a two-way fixed-effects (TWFE) model with institution and time fixed effects, augmented with region-by-time controls. The outcome is either (a) a fraud externality index (FEI), (b) an operational resilience incident rate (ORIR), or (c) a consumer-protection friction index (CPFI). The generic specification is: $Y_{i,r,t} = \alpha + \beta \cdot RTP_{i,t} + \theta \cdot Controls_{r,t} + \mu_i + \tau_t + \varepsilon_{i,r,t}$. Because adoption is staggered and treatment effects may be heterogeneous, we complement TWFE with event-study estimators that are robust to treatment effect heterogeneity.

Event study. We estimate dynamic effects using an event-time framework. For each institution i, event time k indexes months relative to adoption (k=0 at the first full month post-adoption). We include leads and lags up to 24 months. The coefficients trace pre-trends and post-adoption dynamics. We focus on two interpretive questions: (i) whether fraud rises immediately after adoption (consistent with "friction removal"), and (ii) whether control investments (e.g., name checking, transaction monitoring) dampen effects over subsequent months.

Instrumental variables. Endogeneity is a concern because institutions may adopt RTP when they are technologically stronger or when local demand is high. As a robustness check, we use two instruments: broadband penetration at the county level (aggregated to the Fed District) and pre-2012 payment modernisation intensity proxied by legacy ACH and debit infrastructure investment. The first instrument captures digital readiness; the second captures path dependence in payments innovation. Both instruments are plausibly correlated with adoption speed but (conditional on controls) not directly with short-run fraud outcomes.

Measurement error and partial observability. Fraud is underreported and often misclassified. To reduce sensitivity to any single reporting channel, FEI is constructed from three components: (1) complaint intensity, (2) dispute/chargeback rates (where applicable), and (3) confirmed scam losses. In simulation, we allow reporting probabilities to vary across institutions and over time, and we evaluate whether estimates remain stable under alternative reporting regimes.

Heterogeneity. We estimate heterogeneous treatment effects by: (i) institution type (bank versus nonbank), (ii) product type (P2P, bill pay, merchant push payments), and (iii) consumer segment exposure (underbanked, low credit score, or high digital-wallet reliance). We also test whether governance variables (audit maturity, fraud staffing, incident response capability) moderate the adoption–fraud relationship.

Statistical inference. Standard errors are clustered at the institution level for within-institution serial correlation. Because shocks may be correlated across institutions within a district (e.g., regional scam campaigns), we also report two-way clustering at institution and district levels. Where event-study specifications involve many coefficients, we adjust for multiple comparisons using false discovery rate control.

Robustness and falsification. We conduct placebo tests using "pseudo-adoption" dates assigned to non-adopters, as well as negative control outcomes (e.g., non-payment consumer complaints) to check for spurious correlations. We also estimate models excluding 2020–2021 months to test whether pandemic-era digital acceleration drives results.

OA.2 Construction of indices

Fraud Externality Index (FEI). FEI is designed to capture the idea that RTP fraud produces losses and remediation costs not fully internalised by any one actor. The index aggregates: (i) authorised push payment (APP) scam loss rates (losses per 10,000 transactions), (ii) unauthorised payment rates (e.g., account takeover), (iii) dispute/chargeback workload (cases per 10,000 users), and (iv) complaint intensity (complaints per million transactions). Each component is standardised within year, then combined using weights derived from a regulator-style loss function: direct losses 0.5, operational remediation 0.3, and consumer trust/friction 0.2. Sensitivity analysis varies weights across plausible ranges.

Operational Resilience Incident Rate (ORIR). ORIR measures outages, degraded service, and severe incidents affecting payment initiation or settlement. It combines: (i) minutes of unplanned downtime, (ii) incident count, (iii) severity-weighted near-miss disclosures, and (iv) recovery time objective breaches. Where public incident data are sparse, ORIR is simulated but calibrated to typical outage frequency distributions observed in large-scale digital services.

Consumer-Protection Friction Index (CPFI). CPFI captures the consumer cost of RTP safety controls and dispute processes: (i) dispute resolution time, (ii) customer support wait times, (iii) authentication friction (step-ups per successful payment), and (iv) rejection/return rates due to name checks or risk flags. CPFI is included because effective fraud control can impose friction, and policy must balance protection and usability.

Institutional Control Maturity (ICM). ICM proxies the quality of prevention-detection-response controls. It is built from policy documents (presence of model governance, vendor due diligence, incident response playbooks), technology controls (confirmation of payee/name checks, real-time monitoring), and operational capability (fraud staffing per million users, 24/7 support). ICM is used in moderation tests to show whether better controls reduce fraud externalities.

OA.3 Replication checklist for reviewers

To support peer review, we include a replication checklist mirroring common journal expectations:

1. Data dictionary. Provide definitions, units, and construction code for every variable in Table 1. 2. Pre-processing. Document missing data handling, winsorisation rules, and seasonal adjustment procedures. 3. Estimation scripts. Provide TWFE, event study, IV, and DiD code with seeded random generators for simulation. 4. Sensitivity. Report robustness under alternative index weights, alternative clustering, and alternative treatment definitions (receive-only versus send/receive). 5. Visual diagnostics. Include residual plots, influence diagnostics, and pre-trend tests. 6. Governance documentation. Provide a rubric for scoring ICM and governance strength so that other researchers can apply it to real datasets.

Where confidentiality prevents raw disclosure, we recommend the "synthetic-but-auditable" approach used here: share the simulator and parameter values, so that reviewers can verify that findings are not artefacts of a particular draw.

OA.4 Policy playbook: operational guidance for stakeholders

This playbook translates the paper's governance framework into implementable steps.

For regulators (Fed, CFPB, OCC, FDIC, SEC): adopt activity-based supervision for RTP rails, require minimum fraud-control baselines for participant institutions, and mandate incident reporting for major disruptions. Establish "safe harbour" guidance for rapid consumer reimbursement when institutions follow prescribed controls.

For real-time rail operators: publish standardised fraud telemetry interfaces; implement network-level anomaly detection; support confirmation-of-payee or equivalent verification where feasible; and run joint scam campaigns with participating institutions.

For banks and credit unions: move from periodic batch fraud monitoring to continuous monitoring with clear escalation paths; harden account takeover defences; and create consumer-facing "stop and verify" interventions at high-risk moments (first-time payee, high amount, cross-border indicators).

For FinTech wallets and P2P apps: strengthen onboarding and KYC; improve user interface warnings; share scam intelligence with banks; and fund consumer education and reimbursement programmes to internalise externalities.

For merchants and billers: integrate request-for-payment where appropriate; maintain secure payee directories; and coordinate on refund processes.

For consumers: provide plain-language risk disclosures; promote security hygiene; and ensure support access for vulnerable populations (elder fraud, limited digital literacy).

The core insight is that fraud in RTP systems is best addressed as an ecosystem problem. Controls must operate at multiple layers simultaneously; otherwise, risk migrates to the weakest link.

Annex F. Model card template and supervisory audit checklist

This paper's empirical component intentionally mirrors how U.S. supervisors and large payment providers document and review high-impact models. Many institutions already maintain internal model governance artefacts (model inventories, validation packs, and change logs) for credit-risk and fraud systems. For real-time payments, the same discipline should apply to transaction-risk scoring, sanctions screening, customer authentication models, and anomaly detection.

F.1 Model card (minimum disclosure)
A model card is a standardized summary that captures (i) what the model does, (ii) what data it uses, (iii) how it was validated, and (iv) what guardrails exist. In the RTP context, a model card should be prepared for each model that can block, delay, or reroute payments, and for each model that affects the decision to allow a customer to enroll in RTP features.

Model purpose and scope: Identify the business process (e.g., "real-time outgoing payment fraud score"), the rail(s) in scope (RTP, FedNow, internal book transfer), the transaction types (consumer credit transfer, business disbursement, bill payment), and the operational setting (24/7). State whether the model is used for hard stops, soft friction, or routing.

Training data and labeling: Describe the labeling policy for "fraud" and "scam," including the distinction between unauthorized account takeover versus authorized push payment inducement. Explain how disputed transactions are adjudicated and how chargeback or dispute outcomes become labels. Specify any lookback windows and how survivorship bias is minimized.

Performance metrics: Report PR-AUC, false positive rate at operating threshold, and expected value / cost curves that reflect both consumer harm and operational workload. In real-time payments, the cost of false negatives can be catastrophic for consumers because recovery windows are short; therefore, cost-sensitive metrics should be reported.

Fairness and inclusion: Report subgroup performance by customer tenure, age bands, income proxy bands, and geography; explain what was done to avoid systematically blocking legitimate payments from certain groups or regions.

Explainability and human override: Document the explanation outputs available to operations (e.g., top contributing features, reason codes), and the process for human override. Define the maximum allowed latency for explanations.

Security and privacy: Specify access controls, logging, and data retention; list third parties and data processors.

Monitoring and change management: Define drift thresholds, monitoring frequency, and rollback procedures. Explain how model updates are tested to avoid degrading performance during peak periods.

F.2 Supervisory audit checklist
To make model governance operational, supervisors and internal audit can use a checklist aligned to three pillars: (1) effectiveness, (2) fairness, and (3) resilience.

Effectiveness: Are the model objectives clearly tied to consumer protection? Are labels and ground truth definitions stable over time? Is there evidence that the model reduces net losses and consumer harm? Are false positives operationally manageable? Are there red-team tests simulating new scam typologies?

Fairness: Are there documented fairness metrics and thresholds? Are any features that could act as proxies for protected classes justified and controlled? Is there a process for consumer recourse when a payment is blocked? Are adverse action and error-resolution obligations understood for the specific payment type?

Resilience: Does the control stack degrade gracefully under outages (e.g., fall back to rule-based controls)? Are incident response and outage plans tested? Are third-party dependencies understood? Is there a mechanism to coordinate with other banks and the rail operator when a large scam campaign is detected?

F.3 Evidence package for examinations
Institutions should maintain a compact evidence pack: model cards, validation reports, monitoring dashboards, a catalogue of known scam typologies and controls, consumer complaint analytics, and post-incident reviews. The evidence pack should be maintained continuously rather than assembled after an incident.

Annex G. Economic rationale: fraud externalities, incentives, and cost-benefit framing

Fraud in RTP systems exhibits externalities. When one provider invests in strong controls, some benefits accrue to other participants because scams often involve multi-bank chains and mule accounts. Conversely, weak controls at one node impose losses on the network. This creates a classic underinvestment problem, particularly when consumer reimbursement obligations are unclear or when fraud losses can be shifted to other parties.

G.1 Incentive frictions
First, liability allocation shapes investment. If a consumer bears the majority of scam losses for authorized transfers, providers may not fully internalize harm. Second, competitive pressure can push platforms to reduce friction (fewer authentication steps) to improve conversion, which can increase attack surface. Third, asymmetric information persists: providers see a narrow slice of a customer's financial behavior, while scams exploit cross-platform narratives and social engineering.

G.2 A pragmatic cost-benefit template
A policy-relevant cost-benefit approach should quantify (i) direct consumer losses avoided, (ii) operational costs (staff time, false positives), (iii) customer friction costs (abandoned transactions), and (iv) systemic confidence benefits (reduced complaint rates and reputational losses). The analysis can be framed as an expected value problem:

Net Benefit = (Loss Avoided + Confidence Benefit) − (Ops Cost + Friction Cost).

Loss avoided can be estimated from historical scam incidence and model lift. Ops cost is proportional to alerts and manual reviews. Friction cost can be proxied by incremental drop-off in payment completion when additional authentication is triggered.

G.3 Policy levers to correct externalities
Network rules: Rail operators can impose minimum controls for participation (e.g., confirmation-of-payee style checks, standardized reason codes). Information sharing: Shared typology feeds and mule-account indicators can improve detection without revealing proprietary data. Liability alignment: If reimbursement is partly standardized (at least for certain typologies or vulnerable consumers), providers have stronger incentive to invest.

G.4 Distributional considerations
A pure loss-minimization objective can lead to over-blocking transactions from high-risk zip codes or from consumers with volatile cashflow. Hence, cost-benefit must incorporate fairness constraints and consumer recourse mechanisms.

II. Annex H. Policy timeline mapping (United States, 2012–2023): rails, guidance, and supervisory focus

This timeline is intended as a practical guide for situating empirical identification strategies. It highlights three overlapping arcs: (1) faster payment modernization, (2) rising scam/fraud attention, and (3) operational resilience and cyber focus.

2012–2014: Growth of mobile payments and P2P platforms accelerates. Fraud patterns begin shifting from card-not-present to account takeover and social engineering. Supervisory focus remains centered on traditional payments and online banking authentication.

2015–2017: Industry and central bank faster payments initiatives mature. In the United States, the payments community formalizes faster payments workstreams, and private-sector RTP initiatives gain momentum. The literature emphasizes that speed and irrevocability change the control problem: prevention and detection must occur before settlement.

2018–2019: Adoption expands and operational lessons emerge. Fraud typologies diversify, especially through business email compromise and romance/investment scams that culminate in instant transfers. Supervisors begin framing cyber and operational resilience as board-level concerns.

2020–2021: Pandemic-era digitization increases volumes and expands the pool of new digital payment users. Scam campaigns exploit emergency benefits and remote work. Operational resilience guidance is strengthened internationally; financial firms emphasize 24/7 availability expectations.

2022: Preparations for broader instant rail interoperability and new services intensify. Consumer protection concerns rise, with policymakers emphasizing the need for clear dispute resolution and for limits on abusive practices.

2023: Central bank instant payment service launches in the U.S. market, expanding access to 24/7 instant settlement and renewing attention on fraud, scams, and resilience under a new operating model.

The timeline can be used for event study windows (e.g., adoption onset, major guidance issuance, and rail go-lives), and for defining pre-trends in difference-in-differences designs.

Annex I. Literature Map and Open Questions

This annex situates the paper within adjacent finance and information-systems literatures and clarifies which questions remain open for future empirical work.

Payments economics distinguishes between the demand for speed (consumer preference for immediacy), the supply of speed (infrastructure and governance), and the externalities created by speed (fraud, errors, and loss of reversibility). In classic network industries, faster settlement increases welfare when it reduces float costs and uncertainty; however, it can also reduce the time available for screening, dispute resolution, and intermediation (Kahn and Roberds, 2009). Real-time payment rails make these trade-offs visible because they collapse the time window for detection and recall.

In the fraud literature, a recurring theme is displacement: when a control suppresses one fraud vector, offenders re-route to a weaker vector. Card-present fraud controls displaced activity toward card-not-present channels, and stronger account takeover controls displaced activity toward social engineering and APP scams. RTP rails could accelerate this displacement unless governance designs incorporate adaptive monitoring and cross-firm information sharing (Anderson et al., 2019). Within this logic, fraud should be analyzed at the ecosystem level—across rails and institutions—rather than at the product level.

Operational resilience research emphasizes that "zero-downtime" expectations create new failure modes. Systems designed for high availability can fail in correlated ways when they share cloud dependencies, third-party fraud vendors, or identity infrastructure. The operational resilience lens therefore complements fraud analysis: an institution that hardens fraud controls but relies on a single third-party verification service can still experience systemic customer harm if that service degrades or is compromised.

Regulatory economics adds a further layer: when losses are not fully internalized, market participants under-invest in controls. In RTP ecosystems, the party best able to prevent fraud (often the receiving institution that sees mule activity) may not bear the full cost of consumer harm (often borne by the sending institution, merchants, or the consumer). This misalignment motivates policy tools such as liability shifting, minimum control standards, and supervisory stress testing.

Open questions: (1) How do different liability regimes affect adoption of preventive controls such as payee confirmation and payment cooling-off? (2) Which controls exhibit increasing returns at the network level (e.g., shared mule lists), and what governance is needed to maintain due process? (3) How does real-time payment adoption interact with macro stress: do fraud rates rise when households experience income shocks, and do RTP rails amplify those shocks? (4) How should central banks incorporate nonbank payment operators and cloud providers into operational resilience exercises? Answering these questions requires matched datasets across rails and institutions, a governance structure for secure data sharing, and careful attention to privacy and bias.

By articulating these open questions, the paper aims to provide a research agenda consistent with Q1 finance journals that increasingly value interdisciplinary methods and policy relevance.

Annex J. Timeline of U.S. Faster Payments Milestones (2012–2023)

This timeline is provided as contextual background for interpreting adoption patterns. It is intentionally high-level and can be adapted into a figure if the target journal encourages a timeline exhibit.

2012–2014: Expansion of mobile banking and early wallet adoption; tokenization standards mature; institutions invest in modern fraud tooling for card-not-present channels.

2015: Federal Reserve issues strategies for improving the U.S. payment system and establishes industry task forces on faster payments; early prototypes for directory services and request-for-payment concepts gain traction.

2016: Global workstreams on fast payments and retail payment governance intensify; domestic stakeholders align on the need for ubiquitous messaging standards and risk controls.

Rasel, I. H., Arman, M., Hasan, M. N., & Bhuyain, M. M. H. (2022). Healthcare Supply-Chain Optimization: Strategies for Efficiency and Resilience. Journal of Medical and Health Studies, 3(4), 171-182. https://doi.org/10.32996/jmhs.2022.3.4.26

2017: The Clearing House launches the RTP network, marking the first new core U.S. payments system in decades.

2018–2019: RTP participation expands across major banks and processors; fraud patterns shift toward social engineering and real-time channels; increased focus on API security and identity proofing.

2020: Pandemic-driven acceleration in digital commerce; fraud volumes rise; consumer reliance on digital wallets and P2P apps grows.

2021–2022: Operational resilience becomes a central supervisory topic; institutions formalize incident response playbooks; policy debate intensifies on reimbursement for APP scams.

2023: Federal Reserve launches the FedNow Service; the U.S. RTP ecosystem becomes multi-rail, increasing the importance of cross-rail governance, interoperability, and consistent consumer protections.

The central implication for this paper is that fraud risk should be treated as endogenous to the institutional and rail adoption path, not merely as an exogenous nuisance.

Annex J. U.S. Implementation Checklist and Short Vignettes

This annex provides a practical checklist that a U.S. supervisory team or an internal risk function can use to evaluate readiness for real-time payments and real-time fraud. The checklist intentionally mixes governance items (who owns the risk), technical items (how controls are executed), and consumer-protection items (how harms are prevented, detected, and remediated). The objective is to reduce the gap between policy aspiration and operational practice.

J1. Governance and accountability checklist. (i) Board oversight: Does the board receive a standing dashboard on RTP fraud, operational resilience, and consumer outcomes, with clear escalation thresholds? (ii) Senior ownership: Is there a named accountable executive for RTP fraud risk across the end-to-end chain, including third-party providers? (iii) Three lines of defense: Are product teams accountable for controls by design, while independent risk validates assumptions and internal audit tests effectiveness? (iv) Incentives: Are growth targets and customer acquisition incentives adjusted for fraud externalities, or do they reward volume regardless of risk? (v) Regulatory engagement: Are there established points of contact across CFPB, prudential supervisors, and state regulators, and a documented plan for incident communications?

J2. Technical and operational checklist. (i) Identity and authentication: Is authentication commensurate with transaction speed and value, and does it incorporate behavioral signals without undermining inclusion? (ii) Payee verification: Is there a usable, low-friction payee confirmation step for high-risk transfers? (iii) Transaction monitoring: Does monitoring operate in real time (sub-second) and incorporate network features, not merely account-level features? (iv) Limits and frictions: Are risk-based limits dynamic and transparent to consumers, and is there a safe alternative channel when limits bind? (v) Resilience: Are there tested failover paths for rail outages, and is the customer experience designed to avoid duplicated payments during recovery?

J3. Consumer protection checklist. (i) Disclosures: Are risks and limits communicated in plain language at the moment of action, not only in terms-and-conditions? (ii) Complaint handling: Are disputes triaged rapidly, and are outcomes monitored for disparate treatment? (iii) Reimbursement: Is there a documented policy for APP-scam reimbursement or hardship relief, including criteria, timelines, and escalation? (iv) Vulnerability: Are there additional safeguards for vulnerable populations, such as seniors or first-time RTP users? (v) Learning loop: Are scam typologies shared across the ecosystem and integrated into product updates within weeks?

J4. Three short vignettes. Vignette 1 (misdirected payment): A consumer makes a real-time transfer to a contact whose phone number was recently reassigned. Without payee confirmation and robust alias management, the transfer becomes effectively irrecoverable. A best-practice approach uses payee name confirmation and warnings when aliases are newly registered or recently changed.

Vignette 2 (romance scam): Over several weeks, a scammer builds trust, then induces a series of small real-time transfers that gradually increase. Real-time monitoring that only looks for large outliers will miss this pattern. A better approach uses trajectory features and social-engineering markers from complaint signals to flag escalating sequences.

Vignette 3 (operational outage): A mid-sized bank experiences a short RTP connectivity outage. Some customers retry transfers, creating duplicates once connectivity is restored. A resilience-first design uses idempotency keys, explicit status messaging, and a controlled replay mechanism to prevent duplicate posting.

In each vignette, the lesson is the same: speed changes the economics of mistakes and the window for intervention. Governance, technology, and consumer protection must therefore be designed as a single system.

Addendum: Practitioner Summary

For practitioners, the core message is simple: real-time payments compress the window for intervention, so governance must move from periodic review to continuous control. The framework in this paper recommends (i) embedding friction where it creates the largest fraud-reduction benefit per unit of customer inconvenience—such as step-up authentication at first-use and destination changes—rather than blanket delays; (ii) measuring fraud not only as losses but also as externalities (complaints, customer churn, downstream identity theft) to avoid under-investment in prevention; and (iii) aligning incentives across banks and nonbanks through network-level operating rules, data-sharing safe harbors, and standardized reporting of scam typologies. Operationally, institutions should maintain a rail-specific playbook that integrates threat intelligence, incident response, and customer communications, and should test those playbooks through tabletop exercises and red-team simulations. Supervisors can reinforce these practices by requiring model and control documentation proportional to scale, by validating that firmsKYC/AML, cybersecurity, and consumer-protection controls remain effective when transaction speeds increase, and by establishing clear accountability for disputes and remediation. Finally, the paper emphasizes that the adoption of FedNow (live in 2023) and the continued expansion of private RTP rails will be most sustainable when consumer protection is treated as infrastructure: transparency, liability clarity, and rapid remediation are not only fairness objectives—they are prerequisites for trust, usage, and long-run network stability.