Financial Technology (FinTech)

From Cash to Digital: The Transformation of Money Movement — Deep Dive

Advanced Learning Objectives

This deep dive targets the upper tiers of Bloom's Taxonomy:

- **Analyze** interchange fee economics using a game-theoretic framework — why do issuers, acquirers, merchants, and networks reach the equilibrium fees they do? *[Analyze]*
- **Evaluate** the CBDC disintermediation risk to commercial banks under alternative monetary policy regimes *[Evaluate]*
- **Critique** SWIFT's correspondent banking model and assess whether blockchain-based alternatives can achieve systemic scale *[Evaluate]*
- **Compare** UPI, PIX, and FedNow as policy experiments in real-time gross settlement

Assumed Background

You are familiar with: the four-party payment model, authorization/clearing/settlement mechanics, interchange fee basics, and the main CBDC design dimensions. This session interrogates those foundations analytically.

Central Analytical Question

Payment system design is always a political economy problem: every design choice creates winners and losers among issuers, merchants, consumers,

Interchange Fee Economics: A Game-Theoretic Analysis

The Two-Sided Market Structure (Rochet & Tirole, 2003):

- The card network is a *platform* connecting two distinct user groups: cardholders and merchants. Network value requires both sides to participate simultaneously.
- Let b_B = cardholder benefit per transaction (convenience, rewards, credit float); b_S = merchant benefit per transaction (higher sales, reduced cash handling, guaranteed payment).
- Socially optimal usage occurs when total benefit exceeds total cost: $(b_B + b_S) \geq c$, where c is the social cost of processing.
- Interchange a is the transfer from acquirer to issuer that balances platform participation

The Strategic Tension:

- **Issuer preference**: High a maximizes revenue; funds rewards programs that increase cardholder demand.
- **Merchant preference**: Low a minimizes cost of acceptance.
- **Network preference**: Maximise transaction volume; this requires both sides engaged, so the network internalises some of the cross-side externality.

The Waterbed Effect

Capping interchange does not eliminate

Interchange Regulation: Comparative Evidence Across Jurisdictions

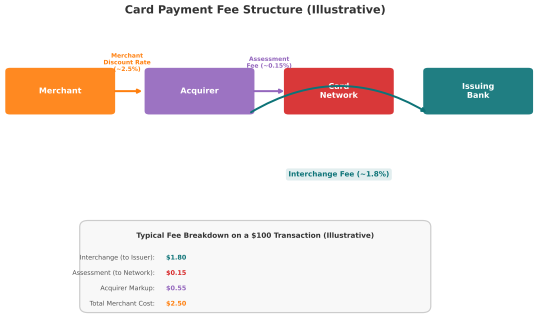

Three Natural Experiments:

(i) EU Interchange Fee Regulation (IFR, 2015):

- Caps: 0.2% debit, 0.3% credit. Calibrated to merchant indifference test.
- Merchant cost reduction: ∼EUR 5B/year (EC estimate).
- Unintended: rewards programs slashed; some issuers introduced annual card fees; consumer price pass-through incomplete over 5-year window.

(ii) Durbin Amendment (USA, 2011):

- Cap: 21c + 0.05% per debit transaction (large issuers only).
- Merchant savings: USD 6–8B/year. But:

CBDC Disintermediation: A Monetary Economics Analysis

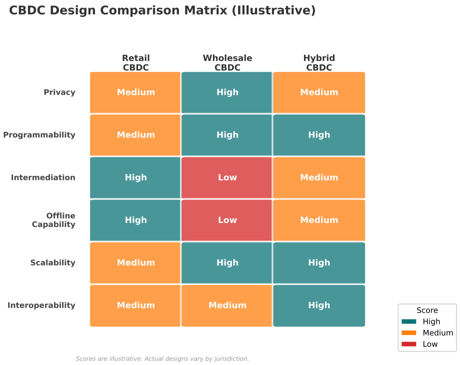

The Disintermediation Mechanism:

- Under a retail CBDC, households can hold direct claims on the central bank rather than commercial bank deposits.
- If the CBDC is interest-bearing at rate $r_{CBDC} \geq r_{deposit}$, rational households shift deposits into CBDC.
- Bank liability reduction $\Rightarrow$ forced asset sales or reduced lending $\Rightarrow$ credit contraction and potential systemic amplification.

Three Design Responses:

CBDC and Monetary Policy Transmission: New Levers and New Risks

Potential Enhancements to Monetary Policy Transmission:

- **Interest rate pass-through**: A programmable CBDC rate provides a direct floor/ceiling for the deposit market, potentially strengthening the transmission of policy rate changes to household borrowing costs.
- **Targeted stimulus**: Programmable money can be designed for specific spending categories or populations (e.g., expiry-date stimulus CBDC to ensure velocity), avoiding leakages via saving.
- **Eliminating the zero lower bound**: An interest-bearing CBDC with a feasible negative rate removes the cash substitution constraint

The Policy Trilemma

Retail CBDC designers face a trilemma: they cannot simultaneously achieve *financial stability* (prevent disintermediation), *monetary policy effectiveness* (meaningful rate setting), and *privacy* (no transaction surveillance). Each design choice sacrifices one vertex.

Empirical Benchmark: Digital Yuan (e-CNY)

China's e-CNY is the world's most advanced large-economy CBDC pilot.

Cross-Border Payment Reform: SWIFT Alternatives and Blockchain Settlement

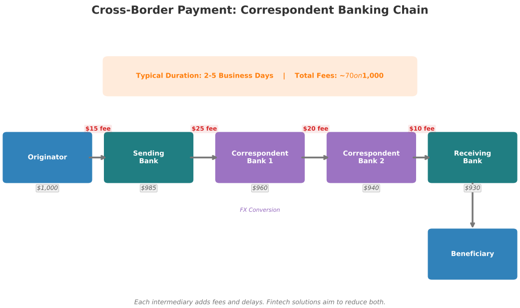

Why SWIFT Is Not the Problem (But Is Part of It):

- SWIFT is a *messaging* network, not a settlement system. The latency and cost of cross-border payments derive primarily from: (a) nostro/vostro pre-funding requirements, (b) AML/KYC compliance at each correspondent hop, and (c) FX conversion spreads.
- Replacing SWIFT messaging alone would not reduce costs materially.

Reform Pathways:

Blockchain Cross-Border Settlement: Scalability and Governance Constraints

The Scalability Trilemma (Buterin, 2014):

- Blockchain-based settlement systems cannot simultaneously achieve all three of: **decentralization**, **security**, and **scalability**.
- *Bitcoin*: decentralized, secure, $\sim$7 TPS (vs. Visa 24,000 TPS peak). Inadequate for payment system scale.
- *Ripple (XRP)*: scalable (>1,500 TPS), secure, but *not* decentralized — Ripple Labs controls validator list. Adoption limited to non-major corridors.
- *Ethereum (post-Merge)*: $\sim$15–30 TPS base layer; Layer 2 rollups achieve $\sim$2,000 TPS but require trust in rollup operator.

The Governance Problem:

- Cross-border settlement requires finality: once settled, funds cannot be recalled. This requires a trusted authority to define finality – precisely what decentralization eliminates.
- A permissioned distributed ledger (central banks as validators) resolves the governance issue but is technically indistinguishable from a shared database with a Byzantine fault-tolerant consensus protocol.

The Honest Assessment

Real-Time Gross Settlement: Comparing UPI, PIX, and FedNow as Policy Experiments

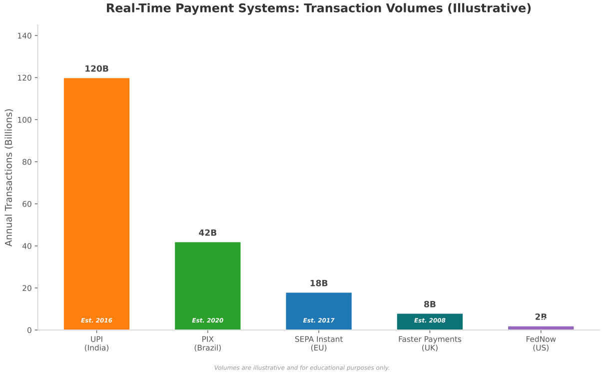

Design Dimensions That Determine Adoption Speed:

(i) Mandate vs. Voluntary Participation:

- UPI/PIX: all licensed banks *must* participate. Instant critical mass.
- FedNow: voluntary. <1,000 banks by end-2024 (of 10,000+). Slow network build-up.

(ii) Pricing Model:

- UPI: government-subsidized zero fee for consumers and merchants. Explicit financial inclusion instrument.
- PIX: free for individuals; small fee for

Payment Fraud and Security Architecture: Layered Defense Design

Fraud Typology in Modern Payment Systems:

- **Card-not-present (CNP) fraud**: Dominant fraud vector in card payments (∼75% of card fraud losses). Card details compromised and used remotely. Mitigated by: 3DS2 authentication, device fingerprinting, behavioral biometrics.
- **Account takeover (ATO)**: Credential stuffing (automated testing of leaked username/password pairs) against payment app logins. Defence: multi-factor authentication, anomaly detection on login behavior.
- **Authorised Push Payment (APP) fraud**: The victim is socially engineered into initiating a legitimate-looking RTGS transfer to the

The Layered Defense Architecture:

1. *Layer 1 — Identity*: Strong customer authentication (SCA), biometrics, device binding. Prevents unauthorized access.
2. *Layer 2 — Transaction monitoring*: Real-time ML scoring of each transaction against behavioral baseline. Flags anomalies within authorization window (100–300ms).
3. *Layer 3 — Network-level*: Cross-bank fraud consortium data sharing (e.g., Cifas in the UK). Pattern recognition across

ML-Based Fraud Detection: Technical Architecture and Evaluation Criteria

Core ML Pipeline for Transaction Fraud:

1. **Feature engineering**: Transaction amount, merchant category, time-of-day, geo-velocity (distance from last transaction divided by time elapsed), device fingerprint hash, session behavioral metrics (typing speed, mouse dynamics).
2. **Model architecture**: Gradient boosting (XGBoost/LightGBM) dominates production systems for tabular transaction data. Neural sequence models (LSTM, Transformer) used for behavioral sequence modeling.
3. **Real-time scoring**: Model must return a fraud probability score within the authorization window ($<$300ms). Requires

The Precision-Recall Tradeoff:

- **False positive**: Legitimate transaction declined. Cost: merchant loses sale, cardholder experiences friction, potential customer churn.
- **False negative**: Fraud approved. Cost: direct financial loss + reputational damage.
- Optimal threshold minimises: $\text{FP cost} \times \text{FP rate} + \text{FN cost} \times \text{FN rate}$, where $\text{FN cost} \gg \text{FP cost}$ in most regimes.

Class Imbalance Challenge

Card fraud rates are typically 0.01–0.10%

Five Open Research Questions in the Field:

1. **RTGS and household financial behavior**: Does instant payment availability increase consumption volatility (by enabling impulsive spending) or reduce it (by enabling faster income access for low-income households)? Causal evidence is scarce.
2. **Interchange regulation and credit supply**: Does reducing interchange cause issuers to tighten credit standards for marginal borrowers (as predicted by the two-sided market model)? US Durbin evidence is mixed; IV strategies remain underdeveloped.
3. **CBDC demand elasticity**: How sensitive is household CBDC adoption to its interest rate,

The Central Unresolved Tension

The efficiency gains of digital payment infrastructure (real-time settlement, near-zero cost, universal access) are achievable only by eliminating the protections embedded in the legacy architecture (chargeback rights, settlement delay as fraud window, batch netting as liquidity buffer).
No jurisdiction has yet found a design that captures all efficiency gains without creating new systemic vulnerabilities. The search for that design is the central engineering and

Payment Economics Glossary

Core Payment Economics Terms:

Acquirer: The merchant's bank. Processes transactions, deposits funds to merchant, manages merchant fraud risk, and pays interchange to the issuer.

Correspondent banking: Arrangement where Bank A holds a nostro account at Bank B (and B holds a vostro account at A) to facilitate cross-border settlement. The backbone of international payments; also its main bottleneck.

Herstatt risk: Settlement risk in FX transactions when one leg settles (currency delivered) but the counterparty fails before delivering the reciprocal currency. Eliminated

Advanced Terms:

Merchant discount rate (MDR): The total fee a merchant pays per transaction, consisting of interchange + network assessment fee + acquirer markup.

Net settlement: At end of day, only net multilateral positions between participants are settled, rather than each transaction gross. Reduces liquidity requirements by 80–90%.

Nostro/vostro accounts: A nostro account is "our money held at your bank"; a vostro account is "your money held at our bank." These bilateral pre-funded accounts are the actual mechanism of correspondent settlement.

RTGS Architecture: Detailed Comparison

Dimension	UPI	PIX	FedNow
Operator	NPCI (private non-profit)	BCB (central bank)	Federal Reserve
Launch	2016	2020	2023
Participation	Mandatory	Mandatory	Voluntary
Consumer price	Zero (gov't subsidy)	Zero (individuals)	Bank-set
Settlement	Gross, real-time	Gross, real-time	Gross, real-time
Alias system	VPA	CPF /	None at

Key Design Lessons:

- **Mandate matters most**: The single largest predictor of RTGS adoption speed is whether participation is mandatory. UPI and PIX achieved mass adoption within 2–3 years; FedNow remains a niche product 18 months post-launch.
- **Price determines use case**: Zero-fee systems (UPI/PIX) displace cash at the low-value, high-frequency end of the transaction distribution. Fee-based systems (FedNow) compete with ACH for bill payment, not with cash at point of sale.

Academic References and Key Papers

Payment Economics and Two-Sided Markets:

- Rochet, J.-C. & Tirole, J. (2003). "Platform Competition in Two-Sided Markets." *JEEA* 1(4), 990–1029
- Baxter, W.F. (1983). "Bank Exchange and Interchange Fees in the Payment Card Industry." *JMCB* 15(4)
- Evans, D.S. & Schmalensee, R. (2005). *Paying with Plastic: The Digital Revolution in Buying and Borrowing.* MIT Press
- Kosse, A., et al. (2017). "The Role of Interchange Fees in Two-Sided Markets." *JMCB* 49(2)

CBDC and Monetary Economics:

Cross-Border Payments and RTGS:

- BIS CPMI (2023). *Fast Payments: Enhancing the Speed and Availability of Retail Payments*
- FSB (2023). *G20 Roadmap for Enhancing Cross-Border Payments: Priority Actions for 2023*
- Agarwal, S., et al. (2023). "The Real Effects of Instant Payments." NBER WP 31129
- BIS (2022). *Project mBridge: Connecting Economies through CBDC*

Fraud, Security, and Regulation:

- PSR (2023). APP Scam Reimbursement Policy Statement PS23/3
- UK Finance (2024). *Annual Fraud Report*

Setup (Rochet & Tirole, 2003):

- Let p_B = price charged to cardholder; p_S = price (MDR) charged to merchant.
- Platform cost per transaction: $c = c_I + c_A + c_N$ (issuer + acquirer + network costs).
- Cardholder transacts iff: $b_B \geq p_B$; merchant accepts iff: $b_S \geq p_S$.
- Platform profit: $\pi = (p_B - c_B + p_S - c_S) \cdot D(p_B, p_S)$

Social Welfare Maximisation:

$$\max_{p_B, p_S} W = (b_B - c_B) + (b_S - c_S)$$

The Interchange as a Rebalancing Instrument:

- Interchange a shifts the price split: issuer net price $= p_B - a$; acquirer net cost $= c_A + a$.
- Optimal: $a^* = c_I - b_S$ (merchant indifference test)
- Interpretation: issuers should be compensated for their costs minus the surplus merchants derive from each transaction.

Why markets don't achieve a^*:

- Issuers compete for cardholders by increasing rewards, which requires

Discussion Questions for Advanced Seminar

Analytical Questions:

1. The Rochet–Tirole model implies that the interchange fee level is not uniquely determined by efficiency considerations — it depends on the relative elasticities of demand on each side. What data would you need to empirically estimate the optimal interchange for a specific payment corridor (e.g., UK contactless debit)?
2. The Durbin Amendment caps debit interchange but not credit interchange. Using the two-sided market framework, predict what happens to: (a) the relative share of debit vs. credit in merchant

Policy and Design Questions:

5. FedNow launched as a voluntary system in a card-dominated market, while UPI and PIX were mandatory. Construct a political economy explanation for why the US Federal Reserve chose a voluntary model. What vested interests, legal constraints, and institutional path dependencies contributed to this choice?
6. Project mBridge uses a distributed ledger operated by multiple central banks for wholesale cross-border settlement. What governance structure would be required to make this system genuinely neutral — i.e., not dominated by any single