Three Myths About Payment Innovation

Wenxin Du
Harvard Business School

Panel Remarks at the Jackson Hole Economic Policy Symposium
August 29, 2026

Thank you to the Federal Reserve Bank of Kansas City for the invitation. It is a pleasure to be here.

The title of this year's symposium is *Financial Innovation: Implications for Payments and Policy*. I would like to use my time today to offer a somewhat different perspective on three widely held propositions that have come to shape the discussion of payment innovation:

1. 24/7 atomic real-time gross settlement is the future of wholesale payments.
2. Retail payment frictions are mainly a legacy-technology problem.
3. Stablecoins are efficient for cross-border payments.

In each case, the claim focuses on the visible payment rail while overlooking the economic architecture around it. Payment innovations are often judged by what is easiest to see, such as whether a system operates around the clock, whether settlement is instantaneous, and whether its ledger is built on a blockchain. But progress in payments involves much more than moving a token quickly. It also requires managing liquidity and risk, complying with regulation, resolving errors and disputes, and earning users' trust.

Myth 1: 24/7 atomic real-time gross settlement is the future of wholesale payments.

Multilateral netting and liquidity savings

There is an understandable appeal to 24/7 real-time atomic settlement: immediate finality, no counterparty exposure during settlement, and no dependence on legacy infrastructure. In wholesale markets, however, the central question is not simply how quickly each transaction can settle. It is also how much liquidity the system needs to settle the full set of transactions safely.

Figure 1. Net vs. Gross Settlements

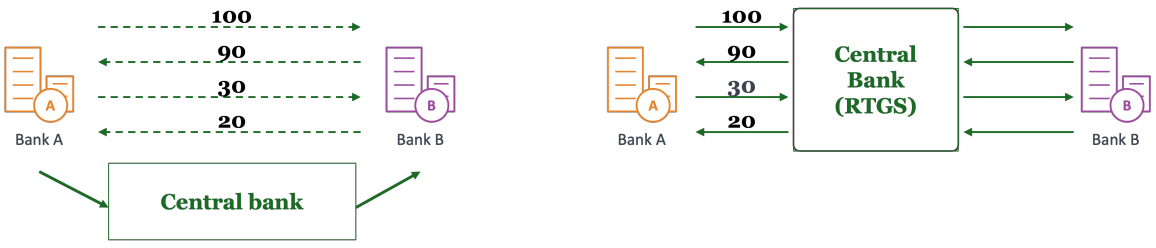

one net transfer: 20 **4 gross transfers: 240**

Figure 1 illustrates the distinction between net and gross settlements. Bank A and Bank B have four payment obligations between them: $100 million in one direction, $90 million in the other, followed by $30 million and $20 million. If each payment settles gross, the system must process and fund payments totaling $240 million. If the obligations are netted, Bank A makes a single net payment

of $20 million to Bank B. The underlying economic obligations are identical, but netting allows them to settle with a fraction of the liquidity.

Most central banks operate real-time gross settlement (RTGS) systems for wholesale payments, but they do not ignore the economic benefits of netting. Most RTGS systems embed liquidity-saving mechanisms (LSMs) that queue payments and continuously search for offsetting flows. Fedwire does not have an embedded LSM, but much of the U.S. financial system nets obligations upstream before the residual payments reach Fedwire. Netting necessarily entails some delay: payments cannot all settle immediately if the system is to identify and offset them. In this setting, delay is not an obsolete feature that technology should eliminate. Instead, it is a core source of efficiency.

Continuous Linked Settlement (CLS), a critical financial market infrastructure at the core of spot foreign exchange (FX) markets, provides a clear illustration. CLS was designed to eliminate FX settlement risk, known as Herstatt risk. By settling both legs of an FX trade simultaneously on its own books, it removes the exposure window in which one party has paid and the other has not. Less appreciated is that this payment-versus-payment (PvP) settlement does not occur at the time of trade. It is deferred to a daily settlement cycle, which allows CLS to net obligations and achieve large funding savings through multilateral netting.

Figure 2. Multilateral Netting in FX: Continuous Linked Settlement (CLS)

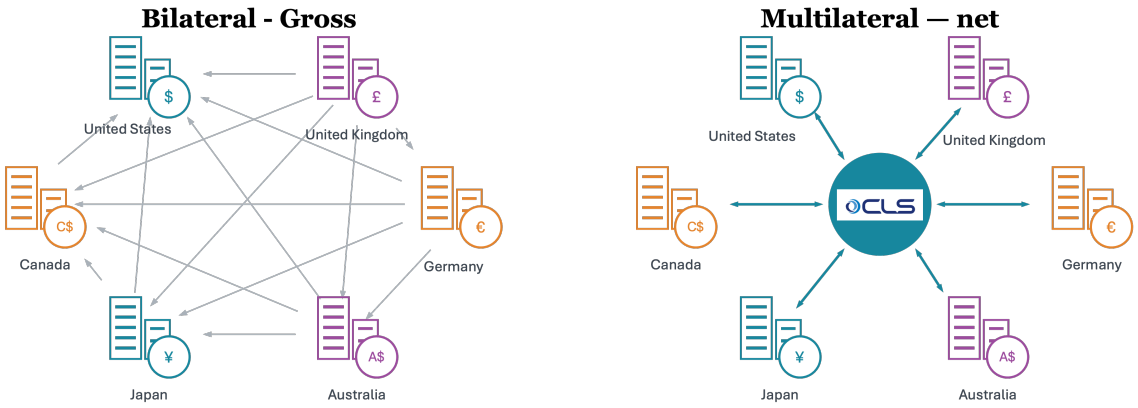

Figure 2 illustrates that six currencies create 15 bilateral currency pairs, and each leg must be funded separately under gross bilateral settlement. Multilateral netting reduces this structure to six net currency positions, with one pay-in per currency. On an average day, CLS settles about $7.9 trillion of FX transactions while requiring approximately $300 billion of actual funding, a netting efficiency of roughly 96 percent.

The same logic applies to several core U.S. payment and capital-market infrastructures. The Clearing House Interbank Payments System (CHIPS), a private-sector high-value payment system used by more than 40 large global banks, reduces roughly $2 trillion of gross daily payments to about $96 billion of net funding flows, for a netting efficiency of 95 percent. The National Securities Clearing Corporation (NSCC), a subsidiary of the Depository Trust & Clearing Corporation (DTCC), reduces around $3 trillion of daily equity and corporate-bond activity to $30 billion to $60 billion of net

settlement obligations, for a netting efficiency of 98 to 99 percent. Fedwire can settle the final residual obligations on a gross basis precisely because so much netting has already occurred upstream.

Implications of real-time gross settlement

Now consider replacing this architecture with trade-by-trade, 24/7 real-time atomic gross settlement. The netting benefit would disappear because every transaction would have to be funded in real time. As Frank La Salla, president and CEO of DTCC, put it "There is not enough money in the galaxy to settle on gross."[1]

There are only a few ways to make such a system work. Central bank balance sheets could expand dramatically, making reserves abundant enough to support gross flows. Alternatively, the private sector could develop a much deeper market for intraday liquidity, but that market remains in its infancy. A 24/7 real-time settlement regime would also raise questions about the central bank's footprint and the design of liquidity backstops. Would central banks provide intraday liquidity around the clock? Would liquidity through central bank swap lines be available overnight and on weekends?

The tension between atomic gross settlement and central clearing also deserves attention. Central clearing has been an organizing principle of post-crisis reform: standard interest-rate swaps and credit-default swaps are centrally cleared, and a clearing mandate for U.S. Treasury cash and repo markets is currently being phased in. A central counterparty (CCP) nets exposures multilaterally and manages risk through margin, default funds, and a defined loss-allocation waterfall. Bilateral, trade-by-trade atomic gross settlement runs in the opposite direction. A CCP has a rulebook, a default-management process, and a legally accountable entity. A public blockchain or decentralized exchange offers no comparable institution to call when settlement goes wrong.

It is also useful to ask how much institutional demand exists for 24/7 settlement. Kinexys, J.P. Morgan's 24/7/365 blockchain unit, processes about $7 billion a day. That is a meaningful business, but it is very small relative to the roughly $12 trillion that moves each day across J.P. Morgan's traditional payment rails.[2]

As discussed in our Kinexys case study, demand for blockchain deposit accounts there remains concentrated among a niche set of users: payment companies, Middle Eastern financial institutions that operate on Sundays, and some manufacturers that value programmability in treasury operations, a function that modern API plug-ins can also provide. J.P. Morgan's deposit token, JPMD, has been live on the public Base blockchain since November 2025 as a proof of concept with a limited set of counterparties, but it has not seen broader take-up.

Myth 2: Retail payment frictions are mainly a legacy-technology problem.

Comparing the speed of cross-border payments

For retail payments, the case for speed is much stronger. Domestic retail payment rails already operate 24/7 in many countries. Remittances are a lifeline for many low-income households, and cross-border transfers often meet urgent liquidity needs. Because payments by individuals and small businesses are typically small in value, they do not create the same system-wide liquidity demands as

[1] Quoted in "Repo tokens won't be cleared. Or will they?" Risk.net, August 17, 2026.
[2] Du, Wenxin, and David Scharfstein. "Kinexys by J.P. Morgan: Building Bank Money on Blockchain." Harvard Business School Case 226-086, March 2026. (Revised June 2026.)

wholesale transactions. Yet retail cross-border payments are also where delays and costs are often highest.

Retail cross-border frictions are often attributed to legacy technology, especially correspondent banking and Society for Worldwide Interbank Financial Telecommunication (Swift). Yet the technology behind correspondent banks and Swift can process payments much faster than the retail experience suggests. According to Swift's 2025 *Spotlight on Speed* report, 75 percent of payments reach the beneficiary institution within 10 minutes. In particular, the cross-border leg of the payment journey is especially fast: less than 20 percent of elapsed time is spent "in flight" between countries, while more than 80 percent is typically spent on the "last mile" within the recipient country. Swift Global Payments Innovation (GPI) provides end-to-end traceability along the way. However, typical cross-border payments for retail customers made through global banks can take days to settle, and retail customers generally cannot themselves track individual payments.

Although large global banks have not widely translated their cross-border payment capabilities into fast, low-cost retail services, fintech money-transfer operators (MTOs) such as Remitly and Wise, both founded in 2011, have transformed how retail customers send money across borders. As discussed in Du, Huang, and Scharfstein (2026),[3] fintech MTOs improved retail cross-border transfers through a technology stack built on prefunded local accounts, internal netting, digitized compliance, and user-friendly mobile-app interfaces. Rather than route each payment through a chain of correspondent banks, an MTO collects funds locally and pays the recipient from an account held with a local banking partner in another country. It then offsets customer flows moving in opposite directions and uses wholesale FX markets or cross-border transfers only to rebalance residual positions. The payment is therefore economically cross-border but operationally resembles two domestic transactions linked by the MTO's internal treasury and netting system. By separating the retail payout from the subsequent rebalancing of liquidity positions, MTOs can deliver funds across a wide range of payment corridors within minutes. This operating model was already well established before stablecoins became a popular topic.

Comparing costs in cross-border payments

Du, Huang, and Scharfstein (2026) show that the cost of sending small cross-border payments differs substantially across global systemically important banks (G-SIBs), fintech MTOs, and community banks. In the U.S.-EU corridor, the average cost across five personal wires through G-SIBs is 364 basis points. For fintech MTOs, the cost is much lower: about 175 basis points on a $200 transfer, 102 basis points on a $500 transfer, and 39 basis points on a $10,000 transfer. A community bank using a business-to-business fintech platform averages about 83 basis points, less than the cost of sending the payment through a G-SIB.

These differences suggest that technology is not the main constraint on payment efficiency. G-SIBs already have the technology to execute cross-border payments more efficiently. What they often lack is a strong economic incentive to deploy that technology for retail customers and small businesses. Fintechs have the opposite strengths and weaknesses: they price aggressively and offer a better user experience, but they may not enjoy the same degree of customer trust and often depend on large banks for wholesale FX and access to RTGS systems. Small and community banks have established customer

[3] Du, Wenxin, Catherine Huang, and David Scharfstein. "Competing Rails for Cross-Border Payments: Banks, Fintechs and Stablecoins." Working Paper, Harvard Business School, February 2026.

relationships and trust. They can rent the fintech stack rather than build it from scratch. An efficient payment model may therefore combine the customer relationship of a bank with the technology and pricing discipline of a fintech.

The lesson from the second myth is that technology matters, but the economics of payments may matter more. Improving retail payments requires better incentives, broader access, and better ways to combine institutional trust with technology, not simply a new ledger.

Myth 3: Stablecoins are efficient for cross-border payments.

The stablecoin sandwich

Cross-border payments are frequently presented as the most important real-economy use case for stablecoins. Are stablecoins indeed efficient for cross-border payments? How do they compare with banks and fintech MTOs?

Using a stablecoin for a fiat-to-fiat cross-border payment requires constructing what we call a *stablecoin sandwich*, illustrated in **Figure 3**. The sender begins with local fiat currency and then converts it to stablecoins. Because roughly 99 percent of stablecoins are denominated in U.S. dollars, this on-ramp can also entail a fiat-to-stablecoin FX conversion. The stablecoin then travels across a public blockchain. At the destination, the stablecoin must be exchanged for the recipient's local currency and transferred through an off-ramp into a bank account or a digital wallet. The blockchain is only the filling. The two slices of bread are the conversions between fiat currencies and stablecoins. These conversions add intermediaries and on- and off-ramp fees, and they are subject to FX liquidity constraints.

Figure 3. Cross-Border Payments with a Stablecoin Sandwich

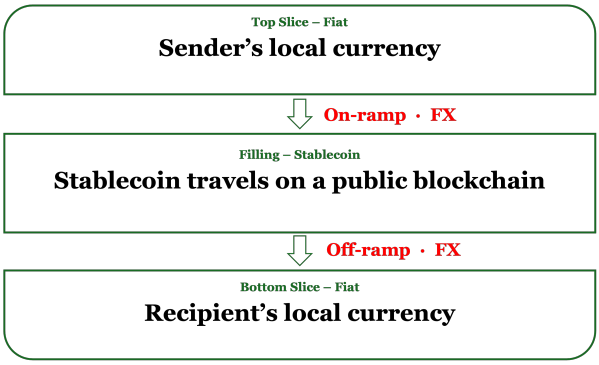

To give a concrete example of the sandwich, consider a $100 transfer from a U.S. bank account to a euro-denominated bank account using USDC. The sender can fund a Coinbase account via ACH at no explicit cost, although settlement may take one business day. Near-instant funding requires bank wires, which would incur a $10 incoming wire fee and a $15–$30 outgoing wire fee. Once the funds reach Coinbase, the sender converts the dollars into USDC and transfers the USDC to Kraken, incurring an average Ethereum gas fee of approximately $0.19. After the USDC reaches Kraken, it is converted into euros at a 1 percent trading fee, and the resulting euros are transferred to the recipient's bank account through SEPA for an additional €1. Therefore, even before considering whether the USDC-euro FX conversion rate is favorable, these explicit charges from the stablecoin sandwich total

at least 219 basis points on a $100 transfer, excluding the U.S. wire fees. That total already exceeds the all-in costs charged by MTOs in the U.S.-EU corridor.

Even if explicit on- and off-ramp fees decline, a stablecoin-based cross-border payment system would require liquid markets between stablecoins and a wide range of fiat currencies. At present, these markets are extremely small relative to conventional FX markets (Du, Huang, and Scharfstein, 2026). The average daily turnover involving the euro in conventional FX markets is approximately $2.7 trillion. Average daily euro turnover against USDC and USDT on centralized crypto exchanges is approximately $489 million, equivalent to only 0.018 percent of conventional euro FX turnover. Across all currencies, fiat-to-stablecoin FX activity therefore remains a rounding error relative to the markets it would need to supplement or replace at scale. Thin liquidity means wider effective spreads, more volatile prices, and greater price impact on larger transactions.

Stablecoin-implied exchange rates can appear substantially more favorable than conventional benchmarks for dollar-outbound conversions into currencies of economies with capital controls, such as the Argentine peso and Nigerian naira. This apparent advantage should not be interpreted as evidence that stablecoin technology has created a dramatically more efficient FX market. Rather, it reflects a stablecoin access premium: in economies with volatile exchange rates, high inflation, and capital controls, residents may be willing to pay a substantial premium for dollar-denominated stablecoins, both as a store of value and as a means of circumventing capital controls.

The stablecoin sandwich alone is not a full payment product

Fintech companies can build on the stablecoin sandwich to deliver cross-border payment services, but the payment rail alone is not a full payment product. Beyond the stablecoin sandwich, a full payment product includes verifying customers, screening transactions, preventing fraud, correcting errors, resolving disputes, and providing a good user experience and customer support. These other services are costly, but they are also central to the take-up of the product.

For a mature fintech MTO, the mechanics of moving money across the rail are among the least expensive components of the service. Customer support and compliance, by contrast, account for the largest fixed-cost categories. Wise's cost breakdown for a pound-to-euro transfer illustrates the point. The estimated fixed cost per transaction is £0.11 for customer support, £0.06 for compliance (including anti-money-laundering, fraud prevention, know-your-customer procedures, and sanctions screening), and £0.05 for other servicing and operations. The actual costs of moving money, such as bank fees, treasury, and FX, together account for only £0.02 of the £0.24 total.[4]

Stablecoin rails can provide fintech payment companies with a foundation for developing alternative operating models, and concrete applications have begun to emerge, as illustrated by our Félix Pago case study.[5] However, the stablecoin sandwich alone is unlikely to make a product competitive with mature fintech MTOs in small- and medium-value transfers or with large global banks in institutional and wholesale payments. Achieving that competitiveness will require innovations in other elements of payment products outside the blockchain rail.

[4] Unit cost breakdown directly provided by Wise in 2026.

[5] Du, Wenxin, and David S. Scharfstein. "Félix Pago: Reinventing Remittances using Stablecoins." Harvard Business School Case 226-087, June 2026.

Conclusion

Separating myths from realities

When separating myths from realities, we reach the following three conclusions.

First, 24/7 atomic real-time gross settlement is unlikely to be the future of wholesale payments. Without netting and robust intraday liquidity management, a fully tokenized, trade-by-trade architecture would introduce major risks and inefficiencies and require a much larger central bank footprint.

Second, retail payment frictions are not simply a technology problem. The economics of payments can matter more than the underlying payment rail. Consumer-facing fintechs, such as Remitly and Wise, have substantially improved the speed and reduced the cost of cross-border payments, and they did so before the arrival of stablecoin and blockchain rails.

Third, stablecoins are not a magic solution for cross-border payments. On- and off-ramp fees are significant, fiat-to-stablecoin FX liquidity is thin, and apparent cost advantages may reflect capital-control frictions. A full payment product also requires compliance, dispute resolution, and customer support, all of which lie outside the blockchain rail.

Implications for central banks

Central banks should not build a new ledger merely to be compatible with the crypto ecosystem. They should instead focus on improving core payment systems and financial market infrastructures.

Central banks can consider extending RTGS operating hours while strengthening liquidity-saving mechanisms; preserving multilateral netting and central clearing; opening central bank payment rails to fintechs that meet compliance standards; and interlinking domestic fast-payment systems so that the benefits of fast domestic payments can extend across borders.

For central banks, the right question is not just whether a payment can move faster on a new ledger. Some of the most valuable innovation may be less visible than a new token or a new rail: better netting, better access, better interoperability, better resilience, and better competition should be the priorities.