Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

The Fragility of Perfectly Safe Digital Money

Elizabeth C. Klee, Arazi Lubis, Chase P. Ross, Sharon Y. Ross, and Alexandros P. Vardoulakis

2026-037

Please cite this paper as:
Klee, Elizabeth C., Arazi Lubis, Chase P. Ross, Sharon Y. Ross, and Alexandros P. Vardoulakis (2026). "The Fragility of Perfectly Safe Digital Money ," Finance and Economics Discussion Series 2026-037. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2026.037.

The Fragility of Perfectly Safe Digital Money*

Elizabeth C. Klee[†1], Arazi Lubis [‡1],
Chase P. Ross[§1], Sharon Y. Ross[¶1], and Alexandros P. Vardoulakis[‖1]

[1]Federal Reserve Board

June 2, 2026

Abstract

Digital money differs from previous forms of money in an important way: it unbundles trust. Instead of relying on a trustworthy institution to settle payments, it relies on decentralized verification, whose cost is priced separately through congestion-sensitive gas fees. This arrangement creates a novel fragility from the interaction of two opposing forces: network externalities, which make digital money more valuable as adoption rises, and congestion fees, which make it more costly to use. We show that these forces generate strategic complementarities in redemption decisions and can create runs even when digital money is backed by perfectly safe reserves.

JEL Codes: C72, E42, E44, G01, G23

Keywords: Stablecoins, Digital Assets, Payments, Runs, Financial Stability, Global Games

*For comments and suggestions, we thank Marcus Opp (the editor), Toni Ahnert (discussant), Dirk Niepelt, James Chapman, and participants at Economics of Payments XI conference and The Future of Payments conference. The views expressed in this paper are those of the authors and do not necessarily represent those of the Federal Reserve Board, or anyone in the Federal Reserve System.

[†]elizabeth.c.klee@frb.gov
[‡]arazi.lubis@frb.gov
[§]chase.p.ross@frb.gov
[¶]sharon.y.ross@frb.gov
[‖]alexandros.vardoulakis@frb.gov

1 Introduction

Digital money differs from previous forms of money in an important way: it unbundles trust. Money traditionally relies on a trustworthy institution, such as a sovereign or bank, for payment settlement. Trust in the money is bundled with trust in the institution. Digital money strips trust out and prices it separately, one transaction at a time. On a permissionless blockchain, the cost of trust is the cost of decentralized verification, paid in gas fees, which incentivize validators to process each transaction. That cost depends on demand for blockspace across the entire network. Higher network congestion makes it more expensive to use digital money for payments. The arrangement stands in stark contrast to the traditional payment system, where settlement prices are set by policy and do not fluctuate with real-time demand. We show that the unbundling of trust introduces a novel fragility—one that remains even when reserves are safe and liquid.

The fragility of digital money arises from the interaction of two opposing forces inherent to permissionless blockchains: network externalities and blockchain congestion. Digital money is more useful for payments when many agents use it, a classic network effect. But the adoption that strengthens network externalities also brings heavier blockchain use, which drives up congestion. High congestion and low network externalities combine to produce strategic complementarities in redemption decisions: whether an individual user should redeem depends on whether other users also redeem. Each redemption erodes the network externalities for those who remain, prompting further redemptions. The interaction can generate a discrete jump in redemptions rather than a smooth, continuous response to higher congestion—and this regime-switching prediction is what we take to the data.

We formalize the mechanism in a global games model and empirically test it using Ethereum-based stablecoins. The model delivers a unique redemption threshold as a function of congestion, which falls as network externalities weaken. Below the threshold, redemptions remain limited; above it, redemptions jump. The mechanism generates two empirical predictions. First, congestion should drive redemptions only when network externalities are low. Second, the congestion-redemption relationship should be concentrated in the high-congestion regime, not spread smoothly across the congestion distribution. We test both predictions.

Our contribution is the identification—both theoretically and empirically—of a source

of fragility in digital money that is distinct from the credit and liquidity risk that drives instability in the existing literature. A large body of work studies stablecoin runs that arise when reserve assets are risky or illiquid, following the banking tradition of Diamond and Dybvig (1983); see, for example, Ahmed, Aldasoro, and Duley (2023), Bertsch (2023), Gorton, Klee, Ross, Ross, and Vardoulakis (2025), and Ma, Zeng, and Zhang (2025). Bolt, Frost, Shin, and Wierts (2024) study stablecoin instability in a global games model with network externalities, but the source of fragility in their paper is still asset-side risk. Sockin and Xiong (2023) show that network externalities in cryptocurrency platforms create strategic complementarities between user participation and speculator sentiment. We show that instability can arise even when reserve assets are perfectly safe and liquid, driven entirely by the interaction of network externalities and the congestion costs inherent to decentralized settlement.

Our empirical approach is motivated by Goldstein, Jiang, and Ng (2017), who establish strategic complementarities in the redemptions of bond mutual funds. In their setting, investors redeem more aggressively from funds with low liquidity after poor performance because they anticipate other investors will redeem first, imposing liquidation costs on remaining holders. In our setting, congestion plays the role of bad performance, and network externalities play the role of liquidity. When network externalities are high, the stablecoin is like a liquid fund: individual redemptions do not erode the payment network enough to trigger further exits. When network externalities are low, the stablecoin is like an illiquid fund: each redemption erodes the payment network that makes the stablecoin valuable, amplifying the incentive for others to redeem.

We show that stablecoins with low network externalities experience disproportionately higher redemptions when congestion rises. This pattern is consistent with strategic complementarities. The effect of congestion alone is not statistically significant: higher gas fees do not generate redemptions for stablecoins with normal or high network externalities. Redemptions arise only from the interaction of the two forces. We can distinguish our mechanism from the alternative of smooth, pre-emptive redemptions. When we decompose the interaction across the congestion distribution, the entire effect is concentrated in the high-congestion regime. No smooth relationship exists between congestion and redemptions at lower levels. This threshold-like pattern is the hallmark of strategic complementarities.

A natural concern is that the price of congestion, the gas fee, is endogenous. We address

this by exploiting plausibly exogenous variation in Ethereum blockspace supply: empty slots caused by random validator failures reduce the supply of blockspace for reasons unrelated to stablecoin demand. We also show that supply-driven congestion shocks shift stablecoin circulation away from Ethereum toward competing blockchains for stablecoins with low network externalities, and that the effect is again concentrated in the upper tail of the shock distribution. We complement this analysis with direct transaction-level evidence, showing that higher Ethereum gas fees increase net transfers of USDT from Ethereum to Tron. Stablecoin holders respond to congestion not by abandoning the stablecoin but by migrating to a blockchain with lower congestion costs, consistent with fragility stemming from the rail rather than the liability.

To gauge the economic magnitude of decentralized trust, we compare the cost of settlement on Ethereum with the counterfactual cost of routing the same transactions through the traditional payment system, the Federal Reserve's Fedwire system. Gas fees can exceed the cost of equivalent Fedwire settlement by an order of magnitude, and they are far more volatile. This comparison illustrates why the congestion channel we identify has material consequences for the stability of digital money.

Our results point to a distinction that discussions have largely overlooked: the difference between the liability and the rail. Focus has been almost exclusively on the liability—the quality of issuers' reserve assets, whether issuers can honor redemptions at par, and how stablecoins compare to deposits and money market funds. We show that even when the liability is perfectly safe, the rail on which the money circulates can be a source of fragility. This distinction matters beyond stablecoins. Our mechanism applies to any digital money that circulates on a permissionless blockchain with congestion-sensitive fees: stablecoins, tokenized deposits, tokenized Treasuries, or a central bank digital currency. A stablecoin fully backed by Treasuries would still be exposed to the congestion-driven fragility if it settles on a permissionless blockchain with volatile gas fees, as would a tokenized insured deposit or even a central bank digital currency issued on Ethereum.[1] Designs that decouple transaction costs from network-wide congestion—such as *trusted* Layer 2 solutions or dedicated application-specific chains—could mitigate the fragility, but they may come with their own impediments

[1] According to the Financial Times, the European Central Bank is exploring running a digital euro on a public blockchain like Ethereum rather than a private one. `https://www.ft.com/content/8ad60169-d1e5-4d2c-b928-d53d668f0ec6`.

which are beyond the scope of our paper.

These questions matter now more than ever. Stablecoins have existed for over a decade and appear to be a lasting feature of the digital financial landscape. A wave of regulatory changes around the world suggests that the boundary between traditional finance and digital finance is narrowing. Such developments point toward a new regime, one in which digital money and its fragilities play a more central role.

Related Literature. Our paper relates to several strands of the literature. First and foremost, we contribute to the literature studying run risk on stablecoins. Because stablecoins promise redemption at par and on demand, a central concern has been their ability to maintain the peg to fiat currency. This concern has led to a focus on the reserve assets of stablecoin issuers and redemption frictions, following the banking tradition starting with Diamond and Dybvig (1983). Fragility arises when reserve assets are risky or illiquid, weakening confidence in the ability of the stablecoin to maintain its peg in stress periods. See, for example, Ahmed et al. (2023), Bertsch (2023), Gorton et al. (2025), and Ma et al. (2025).

All these papers build on global games techniques (Carlsson and van Damme 1993; Morris and Shin 2003) to pin down a unique equilibrium, which is important for tying self-fulfilling beliefs to fundamentals and making the models useful for comparative statics and policy analysis. However, we depart from standard bank-run global-game models where instability is driven by credit and liquidity risk on the asset that financial institutions with runnable liabilities hold (Goldstein and Pauzner 2005; Kashyap, Tsomocos, and Vardoulakis 2024).[2] In our model, perfectly safe and liquid assets back the stablecoins, and credit and liquidity risk do not drive instability. Instead, instability can arise from the uncertainty about the level of on-chain transactions, which is linked to uncertainty regarding blockchain network

[2]See also Ahnert, Anand, Chapman, and Gai (2019), Infante and Vardoulakis (2020), Carletti, Leonello, and Marquez (2023), Matta and Perotti (2023), and Schilling (2023) for recent global game bank-run models that study fragility accruing from the riskiness and illiquidity of a financial institution's assets. Anadu, Azar, Cipriani, Huang, Landoni, Spada, Macchiavelli, Malfroy-Camine, and Wang (2024) compare runs on stablecoins with illiquid assets to runs on prime money market funds. Ahnert, Hoffman, Leonello, and Porcellachia (2024) and Carapella, Chang, Infante, Leistra, Lubis, and Vardoulakis (2025) for global game models to study the financial stability risks from central bank digital currencies. Fernández-Villaverde, Sanches, Schilling, and Uhlig (2021) show that a CBDC can transform the central bank into a deposit monopolist that crowds out private financial intermediation, while Schilling, Fernández-Villaverde, and Uhlig (2024) identify a trilemma whereby central banks offering CBDCs cannot simultaneously achieve allocative efficiency, financial stability, and price stability.

congestion, and crucially interacts with network externalities in stablecoin usage. Moreover, contrary to the aforementioned papers on stablecoin runs, we empirically test our mechanism. Our strategy follows Goldstein et al. (2017), who identify strategic complementarities in bond mutual fund redemptions; we adapt their approach to the stablecoin setting, replacing fund performance with congestion and fund liquidity with network externalities.

Bolt et al. (2024) also study stablecoin instability in a generic global games model where the use of stablecoins for payment generates network externalities. In particular, each agent's payoff from holding bank money increases with the aggregate adoption of bank money, creating strategic complementarities. This feature is also present in our work. However, the source of instability in their paper is risk from the asset holding of stablecoins, similar to the aforementioned literature. In particular, what triggers the run is a sufficiently negative realization of fundamentals that amplifies asset illiquidity and an inability to credibly issue money with negative equity. We differ in this aspect as in our model stablecoins hold perfectly safe and liquid assets, and are always solvent. To the best of our knowledge, ours is the first paper to show stablecoin instability can arise from the interaction between network externalities and network congestion abstracting from the risks of stablecoins' reserves.

Second, we contribute to the literature studying network externalities in the use of money. In the contract denomination and international macro literature (Doepke and Schneider 2017; Gopinath and Stein 2020; Chahrour and Valchev 2021), these externalities reflect incentives to align assets, liabilities, and transactions in a common currency. Coppola, Krishnamurthy, and Xu (2026) apply this logic to financial markets, showing that greater use deepens asset market liquidity, improves matching in money markets, and lowers settlement costs through reduced search frictions, generating increasing returns and persistent dominance. They use a search-based model with endogenous market depth and increasing returns to scale to generate network effects giving rise to a dominant currency for debt denomination. By contrast, Arseneau, Rappoport, and Vardoulakis (2020) show, in a similar framework, that such network effects do not obtain under a constant returns to scale matching technology. Sockin and Xiong (2023) show that network externalities in cryptocurrency platforms create strategic complementarities between user participation and speculator sentiment, which can trigger platform breakdown when user optimism falls sufficiently low. Boissay, Cornelli, Doerr, and Frost (2022) discuss how the limits to blockchain scalability result in fragmentation, preventing network effects from taking place. More related to stablecoins, Gorton, Ross, and Ross (2026)

study how stablecoins develop convenience yields and identify technology, reputation, and dollar demand as key forces. Our paper abstracts from the exact microfoundations of network externalities in the use of money as a medium of exchange or the source of convenience yields for stablecoins.[3]

Third, we relate to the literature on financial market instability with endogenous liquidity and congestion, where trading frictions and intermediary balance sheet constraints make market depth state-dependent and can give rise to runs and fire-sale dynamics. Bernardo and Welch (2004) show how the anticipation of worsening liquidity can generate pre-emptive selling: investors choose to liquidate early to avoid selling into a congested market with lower prices. In their framework, however, prices adjust continuously with aggregate sales, so that incentives to sell are smoothed and the equilibrium features gradual liquidation and is unique. Our model shares the congestion channel, arising in our case from network congestion and the costs of decentralized trust (Budish 2025), but augments it with network externalities in adoption. This introduces strategic complementarities in agents' selling decisions, giving rise to multiple equilibria, which we address using global game techniques. Eisenbach and Phelan (2026) show that sufficiently strong balance sheet costs of intermediaries can also generate strategic complementarities and equilibrium multiplicity. Our mechanism differs in that multiplicity arises from the interaction between congestion and network externalities, rather than from specific assumptions on intermediary balance sheet costs.

2 Institutional Background

This section describes the institutional features of stablecoins that are central to our analysis: how stablecoin transactions are settled, why settlement costs differ from traditional payment systems, and how network externalities and congestion interact on permissionless blockchains.

[3]The literature has also studied the presence of strategic complementarities in investment decisions and technology adoption (see, for example, Cooper and John (1988) and Angeletos and Pavan (2004)), i.e., individual agents are more likely to invest in a technology if others do. The concepts are related, though our work focuses on externalities from the use of private monies in payments and their interaction with network congestion.

2.1 Stablecoins and Digital Money

Stablecoins have emerged as the dominant store of value and medium of exchange in the digital asset ecosystem. Their defining feature is their peg to fiat currency, typically the U.S. dollar, which insulates users from the volatility of unbacked cryptocurrencies such as Bitcoin. The stablecoin market has grown to over $300 billion in market capitalization as of April 2026, with Tether (USDT) and USDC accounting for the vast majority.

Both bank deposits and stablecoins can be used to make payments, but the two settle differently. When a depositor pays someone at another bank, the banks exchange central bank reserves. Settlement is in central bank money and is final, and the cost of settlement is stable and predictable. Stablecoin transactions, by contrast, are settled directly on a blockchain, without any central bank money changing hands.[4] Decentralized third parties verify each transaction, and the cost of doing so depends on network-wide demand for blockspace. This contrast between *institutional* and *decentralized* settlement is the source of the congestion dynamics at the center of our analysis.

2.2 Institutional Trust and Decentralized Trust

In the traditional payment system, the central bank's trustworthiness and reputation make settlement final. The cost of settling a transaction with central bank money is related to the fixed and marginal costs of operating the payment system, and these costs can be significant. Estimates of the cost of payment systems can climb to nearly one percent of GDP (Hayashi and Keeton 2012). Yet, depositors today give little thought to the exact mechanism for clearing and settlement, enjoying low-friction transfers at stable, predictable prices that are largely decoupled from the volume of transactions on the network. This *institutional trust* in the institutions that intermediate the payments system is absent in permissionless blockchain systems.[5]

Stablecoins rely instead on decentralized trust embedded in blockchain protocols. In a permissionless blockchain, decentralized third parties (validators or miners) verify transactions and require compensation to do so. Validators maintain the integrity of the blockchain by registering payments in a distributed ledger using cryptographic techniques, a system we

[4]Stablecoins are bearer instruments: the token itself is sufficient for final settlement once transferred on-chain.
[5]There are also permissioned blockchains controlled by the developers or a central party, such as Onyx by JPM Chase, which are not the focus of our paper.

refer to as *decentralized trust*.[6]

The cost of verifying a blockchain transaction depends on the computational complexity of the transaction, the volume of transactions competing for block space, and the capacity of validators to process them (Budish 2025). The *gas fee* paid for a single transaction is the product of the *gas used*, a measure of the resources required to execute the transaction, and the *gas price*, the per-unit cost of that work denominated in ETH. The gas price moves with aggregate demand for block space and is determined by network congestion and protocol rules, while the gas used is fixed by the nature of the transaction itself: a standard USDC transfer consumes roughly 65,000 units of gas, whereas a decentralized exchange swap or a new contract deployment can consume several times that amount.[7] Unlike the traditional payment system, the cost of using digital money for a payment is not set by an institution but determined endogenously by congestion on the network.

Gas fees are volatile both over time and across transactions at a given point in time. Figure 1 illustrates both dimensions. The top-left panel plots the daily average gas fee in U.S. dollars; fees spiked sharply during the 2021–2022 crypto boom and have trended down since, though periodic spikes remain. Even within a single day, however, different users can face very different fees because gas used varies by transaction type. The daily average in the top-left panel therefore reflects both time-series variation in the gas price and cross-sectional variation in the complexity of transactions settling that day.

The remaining panels document a second dimension of heterogeneity: gas fees are regressive with respect to transaction size. Because gas fees do not scale individual transaction value, the fee-to-value ratio mechanically falls as the size of the transfer rises. Splitting USDC transfers at the daily median transaction size makes this clear. For below-median transfers (bottom-left), the 75th percentile of the fee-to-value ratio frequently exceeds 100 percent—the fee paid to settle the transaction is larger than the amount being transferred. For above-median transfers (bottom-right), the same ratio almost never exceeds 5 percent. The settlement infrastructure thus accommodates large-value transfers at modest cost but imposes prohibitive

[6]Decentralized trust may be an abuse of terminology, but it contrasts with institution-based trust. Strictly speaking, Bitcoin's creator, Satoshi Nakamoto, originally described the need for "an electronic payment system based on cryptographic proof instead of trust." We refer to the trust accruing from cryptographic proof as decentralized trust and the trust in the institutions intermediating payments as institutional trust. Decentralized trust, as opposed to trust in an intermediary, can endogenously arise under certain conditions.

[7]Although we refer to these verification fees as gas fees throughout the paper, in practice this terminology narrowly refers to certain blockchains, including Ethereum.

relative costs on the small, high-frequency payments that a medium of exchange needs to support.

The volatility of gas fees stems from network congestion driven by all transactions submitted for settlement on the blockchain, not just stablecoin transactions. Figure 2 plots gas fees against the number of transactions on Ethereum, showing an increasing and convex relationship. Because stablecoins settle on permissionless blockchains, they share the underlying infrastructure with all other digital-asset transactions—cryptocurrency speculation, NFT minting, decentralized finance protocols. As network activity rises, so do congestion and fees for stablecoin use. A surge in demand for block space from cryptocurrency trading can raise the cost of a routine stablecoin payment, even though the stablecoin's own fundamentals have not changed.

Table 1 illustrates the quantitative importance of this distinction. We compute the counterfactual cost of processing Fedwire transactions on Ethereum. Fedwire fees are set by Federal Reserve policy and do not fluctuate with network demand: a bank making a payment on a busy day faces the same fee as on a quiet day.[8]

On Ethereum, the cost of the same transaction can vary by orders of magnitude depending on network-wide activity.[9] In 2021, processing Fedwire's annual volume on Ethereum would have cost between $1.0 billion and $7.5 billion, against $34 million to $172 million on Fedwire itself—the Ethereum upper bound forty times larger. Gas fees have since trended down. By 2025, Ethereum's lower bound of $23 million falls below Fedwire's $42 million, but Ethereum fees remain far more volatile: the upper bound of $223 million still exceeds Fedwire's $210 million.

[8]Current Fedwire fees can be found on the Federal Reserve Services website: `https://www.frbservices.org/resources/fees/wires-2025`. Fees follow a tiered, volume-based pricing schedule. As of 2025, the lowest fee for transfers over $90,000 is $0.195 and the highest fee is $0.97.

[9]For each year, we compute lower and upper bounds using the annual average of the daily 25th and 75th percentile Ethereum gas fees, and the lowest and highest fees in Fedwire's volume-based schedule. Each bound is multiplied by the year's total Fedwire transfers. We use percentiles for Ethereum because gas fees spike with congestion and minima or maxima would reflect outliers; Fedwire fees follow a fixed tiered schedule, so the lowest and highest tiers are the full range a bank can face. On days of peak congestion, Ethereum fees exceed the 75th percentile by an order of magnitude; the table therefore understates the dispersion between the two systems.

3 Model

Two institutional features from Section 2 interact to generate a distinctive source of fragility. Digital monies inherit network externalities from their use as a medium of exchange (Coppola et al., 2026), but the congestion costs of settling on a permissionless blockchain are shared with all other users of the blockchain, most of whom are not holding digital money. We now formalize this interaction in a simple model and show that it generates strategic complementarities in redemption decisions, giving rise to runs even when the digital money is backed by perfectly safe reserves.

3.1 Economy

Consider an economy with two types of money: old money and new money. Both types of money can be used for the same type of payments but differ in the following two aspects. First, old money uses old payment rails which are less efficient but cheaper. Second, new money uses new payment rails which are more efficient but can be more expensive to use, as specified below. The two types of money can be exchanged one-to-one always but there are costs associated with switching from one money to another, also specified below. To the extent that these monies are liabilities of financial institutions, assume the institutions invest in perfectly safe and liquid assets, such as central bank reserves. This choice is intentional as we want to abstract from the typical bank-run instability of private liabilities and focus on a novel aspect of money fragility.

There is a continuum of agents $i \in [0, m]$, i.e., of mass m, each of whom holds one unit of new money and wants to make payments. Each agent has two options. They can either hold the new money and make the payments, or exchange the new money for old money and use the old money to make the payments. Neither type of money depreciates nor does it pay any interest. Also there is no time discount.

Agents receive an extra benefit from using the new money for a payment. Importantly, the extra benefit depends on how many agents actually use the new money, which proxies for the importance of network externalities. If more than $\hat{\lambda} \in (0, m)$ agents end up using the new money, the extra benefit per payment is $v_h > 0$; otherwise it is $v_l > 0$ with $v_h > v_l$. The latter assumption captures the increasing returns to scale from the use of the new money: the more people buy into this technology, the more valuable it will be.

We normalize the cost of using old money for transactions to zero. By contrast, agents need to pay a transaction fee when they transact through the payment network on which the new money circulates. This fee accrues both when exchanging new money to old money and when using new money for payments, it is levied on each individual transaction, and it does not depend on the face value of the transaction. The fee is a function $\alpha \cdot c(x)$ of the aggregate number of transactions in the network that new money circulates, denoted by x, with $\alpha > 0$, $c'(x) > 0$, $c''(x) \geq 0$, and $c(0) = 0$.[10] These cost properties are consistent with the empirical relationship between Ethereum gas fees and transaction volume shown in Figure 2. Section IA.B formally tests for convexity by estimating a second-order polynomial of the daily gas fee on the number of transactions; the coefficient on the quadratic term is positive and highly significant, supporting the $c''(x) \geq 0$ assumption.

We assume that each agent needs to transact $n > 1$ times, each requiring $1/n$ units of money. Hence, the incremental cost from using new versus old money is equal to $(n-1)\alpha c(x)$; recall that exchanging new to old money also entails a transaction fee but just once.

3.2 Agent strategies

We consider switching strategies under which an individual agent either decides to retain their new money or exchange all of it for old money. Denote by $\lambda \in [0, m]$ the mass of agents that decide to hold on and use their new money. Moreover, denote by γ the aggregate number of transactions on the same network where new money circulates, which nevertheless do not require new money to settle. These transactions correspond to other activities such as trading crypto-currencies or other digital assets with each other without new money serving as the medium of exchange. We assume that $\gamma \sim U[0, \Gamma]$. Nature chooses the true realization of γ, but agents only receive a noisy signal $s_i = \gamma + \epsilon_i$, with $\epsilon_i \sim U[-\epsilon, \epsilon]$ and i.i.d.

The aggregate network congestion is $x = m - \lambda + n\lambda + \gamma = m + (n-1)\lambda + \gamma$, i.e., equal to the number of three types of transactions: (i) those exchanging new to old money, $m - \lambda$; (ii) those using new money directly for payments, $n\lambda$; and (iii) those that are unrelated to the use of new money, γ.

Denote by $\pi(\gamma, \lambda)$ the payoff differential for an individual agent between using new and old money for given γ and λ. If $\lambda \geq \hat{\lambda}$, then $\pi(\gamma, \lambda) = nv_h - (n-1)\alpha c(m + (n-1)\lambda + \gamma)$.

[10] $c'' = 0$ corresponds to a special case of linear cost, which will be useful later to derive a closed form solution. But our analysis is general for strictly convex cost functions.

If $\lambda < \hat{\lambda}$, then $\pi(\gamma, \lambda) = nv_l - (n-1)\alpha c(m + (n-1)\lambda + \gamma)$.[11,12] Given private signal s_i, an individual agent will update their posterior about γ, which will be uniform in $\gamma \in [s_i - \epsilon, s_i + \epsilon]$ and compute their expected payoff differential given by

$$\Delta(s_i) = \int_{s_i-\epsilon}^{s_i+\epsilon} \pi(\gamma, \lambda) \frac{d\gamma}{2\epsilon}. \tag{1}$$

An individual agent holds on to the new money if $\Delta(s_i) > 0$. Because $\Delta(s_i)$ depends on the actions of other agents via λ, an individual agent may need to form beliefs about λ. We first derive two extreme regions of the fundamental value of γ where individual actions are independent of beliefs about the actions of others; these regions are known as dominance regions. Then, we show that there is an intermediate region where beliefs about λ matter and we use global game techniques to pin down such beliefs and the unique actions of individual agents.

We derive two levels of γ, $\underline{\gamma}$ and $\overline{\gamma}$, such that $\pi(\gamma, \lambda) > 0$ for $\gamma \leq \underline{\gamma}$ and $\pi(\gamma, \lambda) < 0$ for $\gamma \geq \overline{\gamma}$ for any λ. To derive $\underline{\gamma}$, focus first on $\lambda < \hat{\lambda}$ such that the low benefit v_l accrues. Given that the transaction cost is increasing in λ and γ, it suffices that $\pi(\underline{\gamma}, \hat{\lambda}) = 0$ yielding

$$\underline{\gamma} = c^{-1}\left(\frac{1}{\alpha}\frac{n}{n-1}v_l\right) - m - (n-1)\hat{\lambda}. \tag{2}$$

To guarantee that $\pi(\gamma, \lambda) > 0$ also for $\lambda \geq \hat{\lambda}$ when $\gamma \leq \underline{\gamma}$, we consider the higher possible

[11] Recall that we have assumed that investors pay the transaction fee when converting new money to old money and, hence, that the aggregate quantity of new money converted, $m - \lambda$, enters the aggregate network congestion. These assumptions are meant to capture aspects of reality, because an investor would need to transfer their new money (stablecoins or other tokenized monies) from their own wallet to that of an authorized participant in order to redeem to fiat. However, we can easily relax this assumption such that network congestion is $x = n\lambda + \gamma$ and investors' payoffs are $nv_k - n\alpha c(n\lambda + \gamma)$, with $k \in \{l, h\}$. As will be clear in the subsequent analysis, the behavior at the boundaries of the strategy space and the monotonicity properties of the payoff functions are preserved under this specification. Thus, our results in the following sections continue to hold under this modified framework.

[12] We have considered a situation where investors already hold new money and decide whether to hold on to it. Arguably this analysis relates to a more mature state of new money adoption. However, our analysis can easily be applied to an alternative situation where investors initially hold old money and, then, choose to adopt new money. In this case investors' payoffs could be written as $nv_k - n\alpha c(n\lambda + \gamma) - \bar{q}$, with $k \in \{l, h\}$, where $\bar{q}$ is the cost of conversion to new money and λ the mass of investors choosing to adopt it. As will be clear in the subsequent analysis, the behavior at the boundaries of the strategy space and the monotonicity properties of the payoff functions are preserved under this specification. Thus, our results in the following sections continue to hold under this modified framework.

cost accruing for $\lambda = m$ and $\gamma = \underline{\gamma}$, and require that

$$v_h > \frac{n-1}{n}\alpha c(nm + \underline{\gamma}). \tag{3}$$

To derive $\overline{\gamma}$, we focus on v_h such that the high benefit accrues. Setting $\lambda = \hat{\lambda}$, such that the transaction cost is at its lowest level, $\pi(\overline{\gamma}, \hat{\lambda}) = 0$ yields

$$\overline{\gamma} = c^{-1}\left(\frac{1}{\alpha}\frac{n}{n-1}v_h\right) - m - (n-1)\hat{\lambda}. \tag{4}$$

To guarantee that $\pi(\gamma, \lambda) < 0$ also for $\lambda < \hat{\lambda}$ when $\gamma \geq \overline{\gamma}$ we consider the lowest possible cost accruing for $\lambda = 0$ and $\gamma = \overline{\gamma}$, and require that

$$v_l < \frac{n-1}{n}\alpha c(m + \overline{\gamma}). \tag{5}$$

Note that all these conditions only depend on exogenous variables. We will now derive the joint set of these exogenous variables under which dominance regions exist and there is an intermediate region where equilibrium multiplicity pertains. Observe that $\overline{\gamma} > \underline{\gamma}$, since c^{-1} is strictly increasing and $v_h > v_l$. Hence, to obtain the dominance regions, we just need to characterize the set of parameters such that $\underline{\gamma} > 0$ given that we can choose any *exogenously set* value Γ to satisfy $\overline{\gamma} < \Gamma$. Using (2), $\underline{\gamma} > 0$ requires $v_l > \frac{n-1}{n}\alpha c(m + (n-1)\hat{\lambda})$. Hence, in conjunction with (5), we need $v_l \in \left(\frac{n-1}{n}\alpha c(m + (n-1)\hat{\lambda}), \frac{n-1}{n}\alpha c(m + \overline{\gamma})\right)$, which is only possible if $\overline{\gamma} > (n-1)\hat{\lambda}$. From (4), the latter condition requires that

$$v_h > \frac{n-1}{n}\alpha c(m + 2(n-1)\hat{\lambda}). \tag{6}$$

Hence, combining (3) and (6) we require that

$$v_h > \max\left(\frac{n-1}{n}\alpha c(nm + \underline{\gamma}), \frac{n-1}{n}\alpha c(m + 2(n-1)\hat{\lambda})\right). \tag{7}$$

Note that if $\hat{\lambda} < m/2$, then (6) is trivially satisfied from (3), which implies that network externalities obtain even if fewer than 50 percent of agents decide to use the new money.

The following Proposition characterizes the equilibrium decision in the dominance regions.

Proposition 1. *For v_l and v_h satisfying* (5) *and* (7)*, and Γ sufficiently high, there exist $0 < \underline{\gamma} < \overline{\gamma} < \Gamma$ given by* (2) *and* (4) *such that: (i) any individual agent always holds on to their new money for $\gamma \leq \underline{\gamma}$; and (ii) any individual agent always redeems their new money for old money for $\gamma \geq \overline{\gamma}$.*

Proposition 1 requires conditions (5) and (7) on (v_l, v_h), which involve the thresholds $\underline{\gamma}$ and $\overline{\gamma}$ that are themselves functions of (v_l, v_h). Appendix IA.A establishes that for any exogenous structural parameters $(m, \hat{\lambda}, n, \alpha)$ and cost function $c(\cdot)$, there exists an open set of (v_l, v_h, Γ) satisfying these conditions, ensuring well-defined dominance regions.[13]

For intermediate realizations of fundamentals $\gamma \in (\underline{\gamma}, \overline{\gamma})$, the individual actions depend on the beliefs about λ. We focus our analysis on threshold strategies according to which an individual agent redeems only if their signal is above a common threshold γ^*. This common strategy enables agents to form well-defined beliefs about the actions of others, i.e., λ. Using global-game techniques, we derive γ^* and show that it is unique. Finally, we show that a threshold strategy around γ^* is indeed an equilibrium.[14]

3.3 Unique equilibrium threshold

Following standard steps from the global games literature we can derive the following condition that determines γ^* (see Appendix IA.A for detailed derivations):

$$\begin{aligned}\Delta^* = &\int_0^{\hat{\lambda}} [nv_l - (n-1)\alpha c(m + (n-1)\lambda + \gamma^*)]d\lambda \\ &+ \int_{\hat{\lambda}}^{m} [nv_h - (n-1)\alpha c(m + (n-1)\lambda + \gamma^*)]d\lambda = 0. \end{aligned} \tag{8}$$

From the dominance regions, we know that the expression in (8) is positive for $\gamma \leq \underline{\gamma}$ and negative for $\gamma \geq \overline{\gamma}$. Moreover, it is strictly decreasing in γ, hence there exists a unique γ^* satisfying (8). So, we just need to show that an agent with $s_i > \gamma^*$ will choose to redeem

[13] For example, with $n = 2$, $m = 0.6$, $\hat{\lambda} = 0.2$, $\alpha = 1$, and $c(x) = x$, setting $v_l = 0.50$ and $v_h = 0.85$ yields $\underline{\gamma} = 0.20$ and $\overline{\gamma} = 0.90$. Conditions (5) and (7) are satisfied, with $0 < \underline{\gamma} < \overline{\gamma} < \Gamma$ for any $\Gamma > 0.90$.

[14] Although there exists only one equilibrium threshold strategy, there may also exist non-threshold strategies, which we have abstracted from. Arguably additional assumptions may be required to exclude such strategies. See, among others, Morris and Shin (2003) and Matta and Perotti (2023) who also touch upon this issue within the scope of the Goldstein and Pauzner (2005) model and suggest focusing on threshold strategies for studying more elaborate models.

and an agent with $s_i < \gamma^*$ will choose to hold on and use their new money. We focus on $s_i > \gamma^*$; the proof is symmetric for the other case. Similarly to Goldstein and Pauzner (2005), the absence of global strategic complementarities, i.e., the fact that $\pi(\gamma, \lambda)$ is not strictly monotonic in λ complicates the proof. To add to the complexity, the single-crossing property they employ to establish the threshold-strategy equilibrium does not need to hold in our framework as $\pi(\gamma, \lambda)$ is positive for γ close to $\underline{\gamma}$ and λ close to zero, and turns negative at some point before $\hat{\lambda}$, given the aforementioned assumptions. Nevertheless, as it will become clear in a bit, we can still implement the logic of the proof in Goldstein and Pauzner (2005) under the additional assumption that $\pi(\gamma, 0) < \pi(\gamma, m)$, which requires that

$$v_h > v_l + \frac{n-1}{n}\alpha \left[c(nm + \gamma) - c(m + \gamma)\right]. \tag{9}$$

In other words, v_h needs to be high enough.[15] The within-regime monotonicity of $\pi(\gamma, \lambda)$ in λ means the largest low-v-regime payoff is at $\lambda = 0$ and the smallest high-v-regime payoff is at $\lambda = m$, so (9) says that payoffs in the low-v-regime are always lower than in the high-v-regime.

Under condition (9), $\pi(\gamma, \lambda)$ would be either like the left-chart, the middle-chart, or the right-chart in Figure 3. In all cases, any value of π for $\lambda < \hat{\lambda}$ is strictly smaller than any value of π for $\lambda \geq \hat{\lambda}$. Then, the logic of the Goldstein and Pauzner (2005) goes as follows: An agent that observes a signal equal to γ^* is indifferent between redeeming or not, i.e., the integral of $\pi(\gamma, \lambda)$ in Figure 3 is zero. Observing a slightly higher signal about γ shifts the beliefs about λ lower, as the agent would expect that more people got higher signals and, hence, decide to redeem. In Figure 3, these beliefs shift the probability distribution from $\lambda's$ above $\hat{\lambda}$ to $\lambda's$ below. Given the latter are associated unambiguously with lower π from (9), a higher signal shifts probability from higher π to lower π, resulting in a negative expected payoff differential and a decision to redeem. The following proposition summarizes the existence and uniqueness of the redemption threshold γ^*, proven above.

Proposition 2. *For v_l and v_h that satisfy* (5), (7), *and* (9), *and Γ sufficiently high, there exists a unique threshold $\gamma^* \in (\underline{\gamma}, \overline{\gamma})$ such that every agent redeems their new money if and only if their private signal s_i is higher than γ^*.*

[15] Since the right-hand side of (9) is increasing in γ, the condition is most binding at $\gamma = \Gamma$. At this point, the right-hand side is independent of v_h, so choosing v_h sufficiently large satisfies the inequality. For instance, with $n = 2$, $m = 0.6$, $\hat{\lambda} = 0.2$, $\alpha = 1$, $c(x) = x$, $v_l = 0.50$, $v_h = 0.85$, and $\Gamma = 1$, the condition holds and Proposition 1 is satisfied (as shown above).

3.4 Comparative statics

For limiting noise $\epsilon \to 0$, the threshold γ^* also converges to the level of fundamentals below which every agent holds on to their new money. Hence, the probability that the new money is stable and used for payments is equal to γ^*/Γ.

Totally differentiating (8), we can compute how γ^* changes in response to changes in other parameters p from

$$\frac{d\gamma^*}{dp} = -\frac{d\Delta^*/dp}{d\Delta^*/d\gamma^*}. \tag{10}$$

The following Proposition summarizes the results.

Proposition 3. *New money is more stable, i.e., γ^* is higher if: (i) network externalities accrue for smaller money usage (lower $\hat{\lambda}$); (ii) the incremental benefit of using new money in either regime is higher (higher v_l or v_h); (iii) agents need to make fewer payments (lower n); (iv) transaction costs are smaller (lower α).*

These comparative statics are intuitive and straightforward to establish (see section IA.A). To establish intuition, we consider here the special case of a linear cost to obtain a closed form solution for γ^*.

Closed-form solution. Suppose $c(x) = x$. Then, solving (8) for γ^* yields the following closed-form solution:

$$\gamma^* = \overbrace{\frac{1}{\alpha}\left[\frac{\hat{\lambda}}{m}\frac{n}{n-1}v_l + \left(1 - \frac{\hat{\lambda}}{m}\right)\frac{n}{n-1}v_h\right]}^{\substack{\text{Positive weighted effect} \\ \text{from network externalities}}} \overbrace{-\frac{m(n+1)}{2}}^{\substack{\text{Negative contribution to congestion} \\ \text{from new money transactions}}} \tag{11}$$

A higher incremental benefit from using new money either in the high or low regime enhances its stability by pushing up the level of congestion γ^* that the new money can sustain without collapsing—the first term in (11). Similarly, lower $\hat{\lambda}$, α, and n make new money more stable because the network externalities obtain earlier, the fees are lower, and fewer new-money transactions incur a fee, respectively. Interestingly, a bigger new money, i.e., higher m, has an ambiguous effect on stability. On the one hand, a higher m makes the region with high

network effects $m - \hat{\lambda}$ bigger, pushing γ^* up—the first term in (11). On the other hand, a higher m contributes to network congestion itself, pushing down γ^*—the second term in (11). The latter force dominates for $m > \sqrt{\frac{v_h - v_l}{\alpha} \frac{2n}{n^2-1} \hat{\lambda}}$, resulting in higher fragility as m continues to increase above this level. Moreover, if $\hat{\lambda}$ grows proportionally with m, i.e., network externalities require larger adoption when new money is bigger, then the first force is muted and the new money is unambiguously more fragile.

Apart from these direct effects on stablecoin fragility, we can also compute the cross derivatives. For example, $d^2\gamma^*/(dv_h\, d\alpha)$ captures how the positive effect of a higher benefit from using new money changes when fees are higher. Using (11), we get $d^2\gamma^*/(dv_h\, d\alpha) = -\frac{n}{n-1}\alpha^{-2}(1 - \hat{\lambda}/m) < 0$. When congestion costs are high, network externalities are less effective at stabilizing the stablecoin.

3.5 Fundamental tradeoff and connection to empirics

Before turning to our empirical analysis, there is an important point to make when taking the theoretical prescription to the data. The model derives the probability of a run which is not observed in the data. Instead, we will utilize the model predictions to infer from the data how the different parameters matter for new money fragility. In doing so, we will adapt in our framework the approach of Goldstein et al. (2017), who test for strategic complementarities in bond mutual funds' redemptions. Like us, they motivate their empirical work with a global games model. Because the economics driving instability in our model are different from theirs, we investigate the implications of the interaction of network externalities from the use of money and network congestion under decentralized trust for what we might expect to see in the data vis-à-vis what we would expect under redemptions accruing solely from congestion.

Congestion induces agents to move to the old network, but in the absence of network externalities, it cannot generate the type of strategic complementarities and abrupt regime change that we describe. Instead, agents would smoothly, yet increasingly, redeem their new money for old money as their posterior about γ moves from a low to a high level.

Without network externalities from wider adoption, network congestion by itself generates incentives to redeem of the type analyzed in Bernardo and Welch (2004): an individual agent redeems their new money if $E[\pi(\gamma, m)|s_i] < 0$. In equilibrium, an agent stops redeeming for an aggregate level of redemptions equal to $m - \lambda^*(\gamma)$, which is an increasing function of γ. Thus, in their setting the decision to preemptively redeem new money for old is more akin

to strategic substitutabilities: the more agents redeem, the lower the marginal incentive to redeem as congestion costs fall. As γ increases, the aggregate redemptions increase, acting as an equilibrating force to the higher gas fees. Yet, $m - \lambda^*(\gamma)$ is a continuous function of γ and can take any value in $[0, m]$, which is in sharp contrast with the level of redemptions in our framework that are either zero for $\gamma \leq \gamma^*$ or m for $\gamma > \gamma^*$.

Figure 4 presents stylized examples of aggregate redemptions as a function of γ for the two cases. This theoretical fact will also be useful to distinguish this redemption motive from strategic complementarities in the data. In particular, the next section examines whether redemption behavior is concentrated, discontinuously, at higher levels of congestion.

4 Data

We use blockchain transaction data from the Amazon Web Services (AWS) Public Blockchain Open Data Set, which contains historical Ethereum data on blocks, transactions, logs, token transfers, and contracts. From the transaction and token transfer data, we compute gas fees, stablecoin transfer volumes, and stablecoin circulation on Ethereum. We complement these with prices and market capitalizations from CoingGecko, daily minimum prices from CryptoCompare (used to identify depegging episodes), and blockchain-level circulation from DefiLlama. Reserve asset composition comes from each issuer's public disclosures.

We focus on five stablecoins on Ethereum: Tether (USDT), USDC, DAI, PYUSD, and BUSD.[16] These five coins represent the large majority of stablecoin circulation on Ethereum. Our main weekly panel spans from November 2017 to December 2025, where available, and daily data are used for the empty-slots and transfer analyses in Sections 5.4 and 5.5.

4.1 Key Variables

Congestion. We proxy for network congestion using Ethereum gas fees. We use the average Ethereum gas fee in week t as our baseline measure of congestion. The measure is the same across all stablecoins in a given week because congestion reflects network-wide activity, not

[16] For stablecoin transaction data, we focus on the following contract addresses: USDC (0xa0b86991c6218b36c1d19d4a2e9eb0ce3606eb48), USDT (0xdac17f958d2ee523a2206206994597c13d831ec7), DAI (0x6b175474e89094c44da98b954eedeac495271d0f), PYUSD (0x6c3ea9036406852006290770bedfcaba0e23a0e8), and BUSD (0x4fabb145d64652a948d72533023f6e7a623c7c53).

stablecoin-specific demand. Gas fees are volatile, with spikes driven by surges in network activity, and have been trending downward in recent years.

For robustness, we also construct a stablecoin-specific congestion measure using the average gas fee paid on each stablecoin's own transfer transactions. This measure varies across stablecoins and over time, allowing us to include date fixed effects in specifications where congestion is no longer a collinear time-series variable.

Network Externalities. We measure the benefit of using a stablecoin for payments with its *velocity*, defined as the ratio of total daily transaction value to total circulating amount for each stablecoin. Higher velocity indicates greater transactional use of the stablecoin, which we interpret as reflecting stronger network externalities. We remove transactions associated with the treasury accounts of each stablecoin's issuer so that the measure captures transactional use rather than issuer operations.[17]

We construct the indicator $\mathbb{I}(\text{Low Network Externality})_{i,t-1}$, which equals one when stablecoin i's velocity in week $t-1$ falls more than one standard deviation below its 8-week moving average, where the standard deviation is calculated across all stablecoins. The indicator captures periods when a stablecoin's transactional use is unusually low relative to its recent history—precisely the states in which our model predicts that users are closest to the tipping point. The indicator equals one roughly 7.5 percent of the time.

Redemptions. We calculate stablecoin i's redemptions in week t as the negative change in the log of its circulation on Ethereum:

$$\text{Redemptions}_{it} = -\Delta \log(\text{Circulation on Ethereum}_{it}) \times 100.$$

We remove any outliers where the absolute value of redemptions exceeds 100. A positive value indicates a decline in circulation through tokens leaving Ethereum either due to redemptions to fiat currency or due to migration to other blockchains.

[17] Treasury addresses are identified using public documentation from each stablecoin issuer. See the Internet Appendix for details.

4.2 Controls

We include several control variables to account for confounding factors. To control for general conditions in crypto markets, we include Bitcoin's weekly return (r_t^{BTC}) and its 30-day realized volatility ($BTC\ Vol_t$). We control for fluctuations in stablecoin demand stemming from dollar demand using the stablecoin basis from Gorton et al. (2026), defined as the deviation of the stablecoin's exchange-rate-implied dollar price from the conventional FX rate.

To account for the possibility that changes in stablecoin supply reflect concerns about backing rather than the stablecoin's value as a means of payment, we include two stablecoin-specific controls. First, the *capital ratio*, defined as the ratio of (reserves−liabilities) to total liabilities, measured using the most recent end-of-month disclosure and lagged one period. Second, the *liquid reserve share*, calculated as the share of reserves held in cash, repurchase agreements, and U.S. Treasuries relative to total reserves.[18] We also include $Depeg_{i,t-1}$, an indicator equal to one when the stablecoin's price falls below \$1, to control for episodes where reserve quality concerns may drive redemptions independent of our mechanism.

4.3 Summary Statistics

Table 2 reports summary statistics for the main variables in our weekly panel. The data are an unbalanced panel, with 312 USDT observations and 312, 305, 146, and 109 observations for USDC, DAI, BUSD, and PYUSD, respectively.[19]

The average stablecoin-week exhibits modest growth in Ethereum circulation, with mean redemptions of −1.8 percent, although the distribution is wide: the interquartile range spans roughly −3.1 percent to 0.9 percent, and the most extreme episode in the sample reaches 44 percent in a single week.

The average gas fee over the sample is \$7.56, but the distribution is heavily right-skewed,

[18] For DAI, which is a decentralized stablecoin backed by crypto collateral, the liquid reserve share is set to zero throughout the sample. For the remaining stablecoins, we use each issuer's most recent end-of-month attestation report as of the observation date and carry forward the most recent disclosure until a new one becomes available. For observations before the first attestation, we backward-fill the first observation for the capital ratio and liquid reserve share variables.

[19] We include BUSD observations through February 13, 2023, when the New York Department of Financial Services (NYDFS) ordered Paxos to stop minting new BUSD tokens, and results are similar if we end the BUSD sample when Binance ends support for BUSD on December 15, 2023.

with a median of $3.39 and a maximum of $55.22. This skewness is important because our identification relies on episodes of unusually high congestion. The low-network-externality indicator equals one in 7.5 percent of observations, confirming that the condition is an infrequent but recurring state rather than a permanent characteristic of any individual stablecoin. The stablecoin basis, which helps proxy for dollar demand for stablecoins by comparing stablecoin-fiat pricing against foreign exchange rates, begins in 2019. A positive average stablecoin basis of 1.1 percent indicates that local fiat currency buys fewer stablecoins than in traditional FX markets, i.e., users pay a premium for dollar-like stability. The average liquid reserve share of 0.63 reflects heterogeneity in issuer transparency: the median is 0.79, indicating that more than half of observations correspond to issuers holding largely liquid reserves, while the lower tail captures issuers with opaque or illiquid backing.

Our cross-chain analysis uses a daily panel with a different set of variables. Table IA.2 reports summary statistics for this sample used in the cross-chain analysis. The average empty slots rate is 0.7 percent, meaning that validators fail to propose a block in roughly one out of every 135 slots. The distribution is right-skewed, with a median of 0.6 percent but a maximum of 9.6 percent, reflecting rare episodes of severe validator disruption. The AR(1) innovation has mean zero by construction and a standard deviation of 0.3 percentage points; its range from -3.0 to 7.8 percentage points indicates that the largest positive shocks—the events that generate our identifying variation—are substantially larger in magnitude than the largest negative ones.

Ethereum dominates the cross-chain distribution, with an average circulation share of 58.7 percent compared with 8.5 percent for Solana. The difference variable, which serves as our dependent variable, has a mean of 0.50 and ranges from -0.32 to 0.84. The negative minimum indicates that some stablecoins in our sample have more of their circulation on Solana than on Ethereum by the end of the sample period. Solana's share is itself highly skewed: the median is just 2.5 percent but the 75th percentile reaches 12.4 percent and the maximum exceeds 65 percent, consistent with rapid Solana adoption concentrated among a subset of stablecoins and in the later portion of the sample.

5 Empirical Results

We now test whether the data bear out the fragility described in our model. First, comparative statics indicate that stablecoins are more stable and face fewer redemptions when there are higher network effects and lower congestion costs. Second, we show that these two forces—network effects and congestion—interact to make digital monies circulating on permissionless blockchains fragile. We show that the combination of these two forces can lead to an abrupt halt in the use of digital money and redemptions to traditional money or other blockchains, resembling a regime switch. Third, we show that exogenous supply-driven variation in congestion causes sharp shifts in stablecoin circulation across blockchains. Last, we discuss the policy implications of our results.

5.1 Comparative Statics

We first examine whether the univariate relationships in the data are directionally consistent with the model's comparative statics in Table 3. Proposition 3 predicts that the stablecoin is more stable when network externalities are stronger (higher v_l or v_h), when the critical mass for high network externalities is easier to reach (lower $\hat{\lambda}$), and when congestion costs are lower (downward shift in c). We examine each prediction, mapping model parameters to observable proxies and reporting simple panel correlations with stablecoin fixed effects. These regressions are descriptive, not causal: many of the proxies are themselves equilibrium outcomes. We interpret the correlations as preliminary evidence of directional consistency with the model.

Network externality threshold ($\hat{\lambda}$). In the model, a lower $\hat{\lambda}$ means that network externalities accrue for smaller money usage, making the stablecoin harder to destabilize. We proxy for this with three measures. First, we use a stablecoin's market share on Ethereum relative to its total circulation. A stablecoin with a larger market share has, by revealed preference, already passed the critical mass threshold and is further from the tipping point. Next, we use the number of unique addresses that use the stablecoin scaled by Ethereum circulation and by total circulation. These ratios measure the density of the user base per dollar of stablecoin outstanding.

Columns 1–3 report the results. All three measures are negatively correlated with

redemptions. We interpret these correlations as descriptive rather than causal. Market share and the number of active addresses are themselves equilibrium outcomes shaped by first-mover advantages, platform listing decisions, and survivorship. The signs are consistent with the model: stablecoins with broader adoption face fewer redemptions. The coefficients on the unique addresses are particularly large because stablecoins have much larger circulation than unique addresses using the stablecoin.

Network externalities (v). Proposition 3 predicts that a higher incremental benefit from using the stablecoin for payments stabilizes the stablecoin. We proxy for v with velocity, defined as total transaction value divided by circulating supply, which captures the intensity of transactional use. Column 4 uses the level of velocity; column 5 uses the deviation of velocity from its 8-week moving average, which isolates within-stablecoin variation. Higher velocity is associated with fewer redemptions, consistent with the model: stablecoins that are more useful for payments are more stable.

Congestion costs (c). The model predicts that higher congestion costs make the stablecoin less stable. We proxy for c with two measures: the network-wide average gas fee (column 6) and the stablecoin-specific average gas fee, calculated from each stablecoin's transfer transactions (column 7). Both measures are positively correlated with redemptions but they are not statistically significant.

The unconditional effect of congestion on redemptions is weak because congestion only destabilizes the stablecoin when network externalities are low. In weeks when network externalities are high—the vast majority of the sample—users tolerate higher gas fees, attenuating the unconditional correlation. In Section 5.2, we test the interaction of congestion with network externalities.

Payment frequency (n). Proposition 3 also predicts that the stablecoin is more stable when agents need to make fewer payments (lower n). We do not test this prediction separately because payment frequency is not independently identifiable from the other two forces in our setting. In the model, n enters on both the benefit and cost sides simultaneously. On the benefit side, agents who transact more frequently receive the per-payment network externality benefit v more often, so higher payment frequency is associated with greater transactional

use—exactly the channel captured by velocity. On the cost side, n scales both the individual's incremental cost of using the stablecoin, $(n-1)\alpha c(x)$, and aggregate congestion itself, since total network activity is $x = m + (n-1)\lambda + \gamma$. The data confirm the dominance of the cost channel: in unreported regressions, the number of stablecoin transactions per unit of circulation is *positively* correlated with redemptions, the opposite sign from what one would expect if higher payment frequency primarily reflected the benefit of transactional use. The positive correlation is consistent with the model's prediction that more transactions raise congestion and destabilize the stablecoin, but it cannot be interpreted as a clean test of Proposition 3 because the same variation also increases $c(x)$ and is mechanically correlated with the redemption transactions themselves. Our measures of network externalities and congestion costs already absorb the empirically relevant channels through which n operates.

Taken together, the correlations in Table 3 provide preliminary evidence that the forces highlighted by the model—network externalities, adoption breadth, and congestion costs—are associated with stablecoin stability in the predicted directions. The individual forces, however, do not work in isolation. The model's central prediction is that fragility arises from the *interaction* of network externalities and congestion: neither force alone is sufficient to generate the strategic complementarities and regime-switch dynamics that distinguish our mechanism from the pre-emptive redemptions of Bernardo and Welch (2004). We test this prediction next.

5.2 Redemption Dynamics and Strategic Complementarities

Our model predicts that the interaction of network externalities and network congestion generates strategic complementarities in the use of digital money. When network externalities are high, users value the stablecoin for payments and tolerate congestion costs. When network externalities are low, users are closer to the tipping point at which redeeming becomes optimal and an increase in congestion can push them over the edge. The empirical implication is that redemptions should be disproportionately higher when congestion is high and network externalities are low.

Our empirical approach is motivated by Goldstein et al. (2017), who establish strategic complementarities in the redemptions from bond mutual funds. Goldstein et al. (2017) show that, following bad performance, investors redeem more aggressively at funds with low liquidity because they are concerned that other investors will redeem first and impose

liquidation costs on remaining investors (see, also, Chen, Goldstein, and Jiang 2010). In our setting, congestion plays the role of bad performance and is the shock that makes holding the stablecoin less attractive. Network externalities play the role of illiquidity, creating interdependence across users that makes one investor's redemption costly for others. When network externalities are high, the stablecoin is like a liquid fund: individual redemptions do not materially affect the remaining holders. When network externalities are low, the stablecoin is like an illiquid fund: each redemption erodes the payment network that makes the stablecoin valuable.

We test this prediction by estimating a weekly panel regression:

$$\begin{aligned}\text{Redemptions}_{it} = {} & \beta_1 \text{Congestion}_{t-1} \\ & + \beta_2 \text{Congestion}_{t-1} \times \mathbb{I}(\text{Low Network Externality})_{i,t-1} \\ & + \text{Controls}_{i,t-1} + \text{Controls}_{i,t} + a_i + \varepsilon_{i,t}\end{aligned} \tag{12}$$

where Congestion_{t-1} is the average Ethereum gas fee in week $t-1$, common across all stablecoins. $\mathbb{I}(\text{Low Network Externality})_{i,t-1}$ equals 1 when stablecoin i's velocity falls more than one standard deviation below its 8-week moving average, where the standard deviation is calculated across all stablecoins, and zero otherwise. We include stablecoin fixed effects a_i to account for time-invariant heterogeneity across stablecoins. We do not include time fixed effects in this specification because the congestion measure is a time-series variable common to all stablecoins; in the Internet Appendix, we show the results are robust to using stablecoin-specific congestion fees and time fixed effects to control for common cross-sectional shocks.

Our coefficient of interest is β_2, which captures the marginal effect of congestion on redemptions when network externalities are low. A positive β_2 indicates that low-network-externality stablecoins are more sensitive to congestion than high-network-externality stablecoins—consistent with strategic complementarities. The regression also directly tests for strategic substitutabilities: under the pre-emptive redemption mechanism of Bernardo and Welch (2004), higher congestion should generate redemptions for all stablecoins regardless of their network externalities, implying a positive and significant β_1.

Table 4 reports the results. Column 1 includes no controls beyond stablecoin fixed effects. Columns 2–4 progressively add controls for general crypto-market conditions (Bitcoin return

and volatility in week t, the stablecoin basis from Gorton et al. 2026), the lagged capital ratio, the liquid reserve share, and a depeg indicator interacted with low network externalities. Column 5 adds lagged redemptions.

The interaction coefficient β_2 is positive across all five specifications and statistically significant at the 5 percent level in columns 2–5. With controls, the coefficient is around 0.083 in columns 2–4, and the estimate is stable across all specifications, ranging from 0.069 to 0.084. A one standard deviation increase in the average gas fee ($10.83) corresponds to a roughly 0.9 percentage point increase in weekly redemptions when network externalities are low, a state that occurs about 7 percent of the time. Gas fees can spike well beyond one standard deviation, leading to proportionally larger redemptions in the low-network-externality state.

The coefficient on Congestion_{t-1} alone (β_1) is small and statistically insignificant in every specification. Higher congestion does not generate redemptions for stablecoins with normal or high network externalities. Similarly, $\mathbb{I}(\text{Low Network Externality})_{i,t-1}$ on its own is insignificant: low network externalities without congestion do not trigger redemptions. Redemptions arise only from the interaction of the two forces, consistent with the strategic complementarities predicted by our model and inconsistent with the strategic substitutabilities of Bernardo and Welch (2004), under which β_1 alone would be positive.

Our theoretical analysis assumed that stablecoin reserves are perfectly safe and liquid. In practice, stablecoins have been exposed to non-trivial run risk associated with the quality of their reserve assets, particularly during episodes like the TerraUSD collapse and USDC's depegging after Silicon Valley Bank's failure. We account for this confounding factor—that changes in stablecoin supply may reflect concerns about backing rather than the incremental value of the stablecoin as a means of payment—in two ways. First, we control for the time-varying stablecoin-specific liquid reserve share, calculated as cash, repo, and Treasuries as a share of total reserves (column 3). Second, we include $Depeg_{i,t-1}$, an indicator for when the stablecoin's price falls below $1, interacted with $\mathbb{I}(\text{Low Network Externality})_{i,t-1}$ to allow depeg-driven redemptions to differ for low network externality stablecoins (column 4). Our results remain significant after controlling for both concerns. Column 5 additionally controls for lagged redemptions to account for persistence; β_2 falls modestly to 0.069 but remains significant at the 5 percent level.

In the Internet Appendix, we show several robustness checks. The results hold when we use 12-week instead of 8-week rolling windows to calculate summary statistics (Table IA.4),

when we replace the network-wide gas fee with stablecoin-specific congestion costs computed from each stablecoin's own transfer fees (Table IA.5), and when we use daily rather than weekly redemptions (Tables IA.6 and IA.7). The stablecoin-specific congestion measure allows us to include date fixed effects, since congestion is no longer a collinear time-series variable, and the results are robust to this more demanding specification.

5.3 Regime Shift

Table 4 establishes that congestion interacted with low network externalities generates redemptions. We now ask where in the congestion distribution this interaction operates. The answer distinguishes between two mechanisms that our model highlights.

The first mechanism is strategic complementarities. In our model, redemptions are small for moderate congestion and then jump discontinuously once congestion crosses a threshold. The discontinuity arises because each agent's incentive to redeem depends on the redemption decisions of others: once enough users begin to redeem, the resulting decline in network externalities makes the stablecoin less valuable for the remaining holders, inducing further redemptions. The second mechanism is pre-emptive redemptions akin to Bernardo and Welch (2004): as congestion increases, stablecoin investors individually and incrementally redeem their tokens until the lower resulting supply eliminates congestion-driven redemption incentives. Under this mechanism, the relationship between congestion and redemptions is smooth and continuous. Higher congestion always generates proportionally more redemptions across the entire congestion distribution.

As mentioned in the model, Figure 4 illustrates the distinction. Under strategic complementarities (left panel), redemptions are zero below a congestion threshold γ^* and then jump discretely. The figure is stylized such that redemptions are flat levels of zero below the threshold and positive above the threshold, but if both horizontal lines were upward sloping, or if the line above the threshold is upward sloping, that too would be consistent with strategic complementarities if there remains a distinct regime switch at the threshold. Under strategic substitutabilities (right panel), redemptions increase smoothly and continuously with congestion. The key empirical difference is that strategic complementarities predict no smooth relationship between congestion and redemptions, the effect is concentrated entirely at the threshold, while pre-emptive redemptions predict a positive and significant relationship throughout the distribution.

We test between these two patterns by decomposing the interaction effect from Table 4. Rather than estimating a single interaction of congestion with low network externalities, we allow the sensitivity of low-network-externality stablecoins to congestion to differ across three congestion regimes:

$$\begin{aligned}\text{Redemptions}_{it} = {} & \gamma_1 \, \text{Congestion}_{t-1} \\ & + \gamma_2 \, \text{Congestion}_{t-1} \times \mathbb{I}(\text{Congestion}^M)_{t-1} \\ & + \gamma_3 \, \text{Congestion}_{t-1} \times \mathbb{I}(\text{Congestion}^H)_{t-1} \\ & + \beta^M \, \text{Congestion}_{t-1} \times \mathbb{I}(\text{Congestion}^M)_{t-1} \times \mathbb{I}(\text{Low Network Externality})_{i,t-1} \\ & + \beta^H \, \text{Congestion}_{t-1} \times \mathbb{I}(\text{Congestion}^H)_{t-1} \times \mathbb{I}(\text{Low Network Externality})_{i,t-1} \\ & + \text{Controls}_{i,t-1} + \text{Controls}_{i,t} + a_i + \varepsilon_{i,t}\end{aligned}$$

where $\mathbb{I}(\text{Congestion}^H)_{t-1}$ equals 1 if Ethereum's average gas fee exceeds its 8-week moving average by more than one standard deviation, and $\mathbb{I}(\text{Congestion}^M)_{t-1}$ equals 1 if the gas fee is above its moving average but below the high-congestion threshold. The omitted category is weeks in which the gas fee is at or below its moving average. The two triple interactions, β^M and β^H, decompose the single interaction from Table 4 into separate contributions from moderate and high congestion. Controls, stablecoin fixed effects, and standard error clustering are the same as in Section 5.2.

We also include specifications with level interactions $\mathbb{I}(\text{Low Network Externality})_{i,t-1} \times \mathbb{I}(\text{Congestion}^M)_{t-1}$ and $\mathbb{I}(\text{Low Network Externality})_{i,t-1} \times \mathbb{I}(\text{Congestion}^H)_{t-1}$ to absorb mean differences in redemption levels for low-network-externality stablecoins across congestion regimes. Within each group, successive columns add depeg controls and then the liquid reserve share and lagged redemptions.

Table 5 reports the results. The table provides two findings that together support the strategic complementarities interpretation. First, there is no smooth, continuous relationship between congestion and redemptions. The coefficient on Congestion_{t-1} (γ_1) is statistically indistinguishable from zero in every specification. The slope does not steepen at moderate congestion: $\text{Congestion}_{t-1} \times \mathbb{I}(\text{Congestion}^M)_{t-1}$ is also insignificant. The same is true at high congestion: $\text{Congestion}_{t-1} \times \mathbb{I}(\text{Congestion}^H)_{t-1}$ is insignificant. Across the entire congestion distribution, higher gas fees do not generate proportionally more redemptions for the average

stablecoin. This pattern is inconsistent with the pre-emptive redemption mechanism of Bernardo and Welch (2004), which predicts a positive and significant γ_1.

Second, the interaction effect documented in Table 4 is entirely concentrated in the high-congestion regime. β^H, the sensitivity of low-network-externality stablecoins to congestion when gas fees are in the upper tail, is positive and statistically significant across all eight specifications. In columns 1 and 5, β^H is 0.094 and 0.090, respectively. The estimate is stable: it ranges from 0.089 to 0.094, and remains significant even in the most demanding specifications, which include level interactions, depeg controls, the liquid reserve share, and lagged redemptions. By contrast, β^M, the corresponding coefficient for moderate congestion, is not significant in any specification. In columns 3 and 4, β^M is economically small. In the Internet Appendix, we re-estimate the regime-shift specification replacing Ethereum redemptions with redemptions on Solana and Tron (Table IA.8). The high-congestion triple interaction is insignificant on Solana and flips sign on Tron, indicating that our main results reflect Ethereum-specific congestion rather than common shocks affecting all blockchains.

The combination of these two findings—no smooth slope and an interaction effect concentrated in the upper tail—is consistent with the threshold response predicted by strategic complementarities. Redemptions do not increase gradually with congestion; instead, the effect appears only when the combination of high congestion and low network externalities crosses a critical level, producing a discrete shift in redemption behavior.

5.4 Supply Shocks to Congestion

One inherent challenge of using gas fees as a measure of congestion is that gas fees are endogenous. Gas fees are the equilibrium outcome of the supply of and demand for block space, so they could partly reflect stablecoin redemption demand itself. For example, if many stablecoin holders simultaneously try to redeem their tokens on Ethereum, the resulting surge in transaction demand would push gas fees higher—creating a mechanical correlation between redemptions and gas fees that has nothing to do with our mechanism. To address this concern, we exploit plausibly exogenous day-to-day variation in Ethereum blockspace supply to test whether supply-driven congestion shifts stablecoin circulation away from Ethereum.

The Ethereum Merge in September 2022 transitioned the network from proof-of-work (mining) to proof-of-stake (staking). Under proof-of-stake, time is divided into 12-second *slots*, and each slot is one opportunity to add a new *block*—a bundle of transactions written to the

ledger together. For each slot, the protocol randomly picks a validator (who has staked ETH in exchange for the right to help run the network) to propose the block. Proposing a block means assembling a set of pending transactions, ordering them, and broadcasting the result to the network within the 12-second window. The validator earns a reward if it proposes a valid block on time and nothing if it does not. If the selected validator is offline, encounters a software error, or fails to propose the block in time, the slot goes empty. Empty slots reduce the effective supply of block space available for transactions. They are not driven by the demand for stablecoin transactions, but rather by idiosyncratic validator issues: random proposer downtime, relay problems, or infrastructure failures.

With 12-second slots, Ethereum can produce a maximum of 7,200 blocks per day. In practice, it is uncommon for all 7,200 to be created. We calculate the daily empty slots rate as:

$$\text{Empty Slots Rate}_t = \frac{7,200 - \text{Blocks}_t}{7,200}.$$

Figure 5 plots the daily empty slots rate since September 2022. On average, 0.7 percent of slots go empty, and the rate is volatile, with occasional spikes driven by validator coordination failures. A simple back-of-the-envelope estimate shows this can translate to large dollar values. If we consider the daily trading volume on Ethereum, which is roughly $5.4 billion at the end of 2025, and evenly divide it across the maximum number of blocks, each block would account for an average of $750,000. Since the median number of empty slots per day is 48, this could correspond to $36 million in transactions.

We isolate the unexpected component of the empty slots rate by estimating a simple AR(1) model:

$$\text{Empty Slots Rate}_t = \alpha + \beta \text{ Empty Slots Rate}_{t-1} + u_t.$$

The residual u_t, which we refer to as the supply *Shock*, captures the day-to-day innovation in blockspace supply that is not predicted by the previous day's empty slots rate. A positive shock value indicates that more slots went empty than expected, which is a contraction in effective blockspace supply. Summary statistics for the daily sample appear in Table IA.2.

Two properties of the supply shock support our identification strategy. First, higher supply shocks correspond to higher gas fees. A day with more empty slots than expected

reduces blockspace supply, which raises the price of transacting. Table 6 confirms that both the raw empty slots rate and the AR(1) innovation are positively and significantly associated with Ethereum gas fees, with and without controls.

On average, a one standard deviation increase in the empty slots rate of 0.004 corresponds to a $0.77 increase in gas fees. However, the empty slots rate is right skewed, so a large change could lead to dramatically higher gas fees. In the extreme scenario of the blockchain moving from no empty slots to all empty slots, gas fees would increase by roughly $200, above the maximum gas fee that we observe.

Second, the supply shock is orthogonal to stablecoin demand conditions. Table 7 shows that the innovation is not significantly correlated with Bitcoin returns, Bitcoin volatility, or the stablecoin basis for the payment stablecoin USDC. The supply shock thus captures variation in congestion driven by the supply side of the blockspace market rather than by speculative or redemption activity.

If congestion costs spike due to a supply-driven reduction in block space, stablecoin holders can shift away from Ethereum to other blockchains. They can move tokens by using a cross-chain bridge or by redeeming on Ethereum and purchasing the stablecoin on another blockchain. Either channel would reduce the stablecoin's share of its total circulation on Ethereum relative to competing chains. We measure a stablecoin's Ethereum circulation share using data from DefiLlama on the stablecoin's circulation on each blockchain:

$$\text{Ethereum Share}_{it} = \frac{\text{Circulation on Ethereum}_{it}}{\text{Total Circulation}_{it}}.$$

We similarly calculate Solana Share_{it} and limit the sample to stablecoins with at least $100 million in average circulation on Solana over the sample period, ensuring that circulation shares are precisely measured and that day-to-day variation reflects genuine reallocation rather than noise from individual transactions. Three stablecoins meet this criterion—USDT, USDC, and PYUSD—with average Solana circulation of $3.8 billion, $1.3 billion, and $354 million, respectively. The remaining stablecoins have Solana circulation below $2 million, too small for circulation shares to be meaningfully estimated. The supply shock is common across all three stablecoins, and identification comes from its interaction with $\mathbb{I}(\text{Low Network Externality})_{i,t-1}$, which varies both across stablecoins and over time as each coin's velocity fluctuates relative to its own rolling average.

Our dependent variable is the difference $\textit{Ethereum Share}_{it} - \textit{Solana Share}_{it}$. A decline in this variable indicates a relative shift in circulation from Ethereum to Solana.

We estimate the following daily panel regression:

$$\begin{aligned}\text{Ethereum Share}_{it} - \text{Solana Share}_{it} = {} & \beta_1 \, \text{Shock}_{t-1} \\ & + \beta_2 \, \text{Shock}_{t-1} \times \mathbb{I}(\text{Low Network Externality})_{i,t-1} \\ & + \text{Controls}_{i,t-1} + \text{Controls}_{i,t} + a_i + \varepsilon_{i,t}\end{aligned}$$

where Shock_{t-1} is the lagged AR(1) innovation to the empty slots rate, $\mathbb{I}(\text{Low Network Externality})_{i,t-1}$ is defined as in the previous sections using the stablecoin's velocity relative to its 90-day moving average, and a_i are stablecoin fixed effects.

We estimate this regression at the daily frequency because the supply shock is inherently short-lived: each empty slot removes 12 seconds of block space, and the resulting gas fee spike typically dissipates within hours as subsequent validators successfully propose blocks. Weekly aggregation would blend these transient supply disruptions with days of normal block production and with the endogenous demand-side responses that follow—precisely the confounding variation the supply shock is designed to avoid. The dependent variable can also respond quickly because cross-chain transfers settle within minutes, so both sides of the regression operate at a frequency well below one week.

Table 8 reports the results. Column 1 presents the baseline specification. The coefficient on the interaction $\text{Shock}_{t-1} \times \mathbb{I}(\text{Low Network Externality})_{i,t-1}$ is -8.13: when a stablecoin's network externalities are low, a supply shock to congestion causes circulation to shift away from Ethereum toward Solana. This is consistent with our model's prediction that the combination of high congestion and low network externalities triggers reallocation.

The positive coefficient on Shock_{t-1} indicates that an unexpected reduction in blockspace supply shifts stablecoin circulation toward Ethereum on average; however, this is more than offset by the effect when the stablecoin also has low network externalities on Ethereum.

Column 2 tests whether this effect is concentrated in periods when the supply shock is particularly large, paralleling the regime-shift analysis in Section 5.3. We add an indicator $\mathbb{I}(\text{Shock}^H)$ for when the shock exceeds its 90-day moving average plus one standard deviation. The triple interaction $\text{Shock}_{t-1} \times \mathbb{I}(\text{Shock}^H) \times \mathbb{I}(\text{Low Network Externality})_{i,t-1}$ is -14.94. When the supply shock is smaller, the interaction with low network externalities is no

longer significant, indicating that the reallocation is concentrated in the upper tail of the supply-shock distribution. This threshold-like pattern is consistent with the regime-switching mechanism in our model rather than a smooth, continuous response.

Column 3 provides an additional specification that interacts the supply shock with indicators for whether realized gas fees are in the medium ($\mathbb{I}(\text{Gas}^M)$) or high ($\mathbb{I}(\text{Gas}^H)$) range of their 90-day distribution. The triple interaction with high gas fees and low network externalities is -40.68, confirming that the effect is strongest precisely when the supply shock translates into the highest realized congestion costs. The medium-gas triple interaction is not significant, further supporting the nonlinear, threshold-like character of the effect.

To interpret the economic magnitudes, it's helpful to scale the coefficients by a typical shock. The standard deviation of Shock is 0.3 percentage points of the daily empty slots rate. Consider a one-standard-deviation shock during a day of high realized gas fees and low network externalities. The column 3 triple interaction implies a shift of roughly 12 percentage points in the Ethereum share relative to Solana. Evaluated at more extreme shocks, the magnitude grows commensurately.

5.5 Transfers to Tron

The previous section used plausibly exogenous variation in Ethereum blockspace supply to show that congestion shifts stablecoin circulation from Ethereum to competing blockchains. We now complement that analysis with direct transaction-level evidence. When Ethereum gas fees rise, stablecoin holders who still value the stablecoin as a means of payment, but find Ethereum too costly, can migrate to a cheaper blockchain. We test this by tracing individual USDT transfers from Ethereum to Tron, a blockchain with substantially lower transaction costs. These transfers rule out a potential confounding mechanism: that redemptions reflect doubts about USDT's backing rather than the cost of transacting on Ethereum. A user who moves USDT from Ethereum to Tron keeps her exposure to the token; the shift is across blockchains, not out of USDT.

We construct a matched-transfer dataset by linking USDT transfer events on Ethereum to USDT transfer events on Tron that have the same dollar amount and occur shortly after one another. The intuition is that a cross-chain transfer leaves a footprint on both blockchains: a user (or intermediary) moves USDT off Ethereum and, shortly after, an equal amount of USDT appears on Tron.

For example, on November 19, 2022, 14:35 UTC, $51 million of USDT was redeemed on Ethereum, while the same amount was minted on Tron, both transactions occurring within 41 seconds of one another. The matched amount and tight timing identify the pair as a single cross-chain transfer: the holder's USDT exposure is unchanged; only the blockchain hosting the tokens has changed.

We classify the direction of each matched pair based on which chain's transfer occurred first: if the Ethereum transfer precedes the Tron transfer, we label it an Ethereum-to-Tron flow, and vice versa. We identify 1,230 instances of matched USDT transfers between May 2020 and December 2025, which span 616 transfer days. Roughly 70 percent of days with transfers have only 1 transfer, and around 55 percent are transfers from Ethereum to Tron. Transfers occur within 11 minutes of each other on average. The average transfer is $176 million, slightly larger from Ethereum to Tron ($183 million) than the other way ($167 million).

We aggregate the matched transfers to a daily frequency and compute the net transfer amount from Ethereum to Tron on each date. We focus on USDT because it is the largest stablecoin and circulates widely on both Ethereum and Tron. To isolate the relationship between congestion and cross-chain transfers from secular trends in gas fees, we demean Ethereum's average gas fee using its trailing 365-day moving average. The demeaned gas fee captures short-run deviations in congestion costs from their recent norm.

Figure 6 plots the daily net transfer amount from Ethereum to Tron. The series is volatile and lumpy, with large positive spikes indicating episodes when substantial USDT balances moved from Ethereum to Tron. Negative values indicate net flows in the opposite direction. Figure 7 plots a binned scatter of the daily net transfer amount against the lagged demeaned gas fee. The positive relationship is evident: days following above-average congestion on Ethereum tend to coincide with larger net transfers of USDT from Ethereum to Tron.

We formalize this relationship with a daily time-series regression:

$$y_t = \alpha + \beta \, \text{Gas Fee}_{t-1} + \gamma' X_t + \varepsilon_t$$

where y_t is either $\ln(\text{Transfer Amount}_t)$, the log of the net daily transfer from Ethereum to Tron, or $\mathbb{I}(\text{Transfer Ethereum to Tron}_t)$, an indicator equal to one if the net transfer from Ethereum to Tron is positive on day t. Gas Fee_{t-1} is the lagged demeaned gas fee, and controls

X_t include the contemporaneous Bitcoin return and Bitcoin volatility to account for general conditions in cryptocurrency markets that may affect stablecoin demand independently of Ethereum congestion. We lag the gas fee by one day so that congestion conditions are predetermined relative to the transfer decision.

Table 9 reports the results. The top panel examines the intensive margin. Our baseline specification in columns 1–2 uses matched transfers occurring within 60 minutes of each other. The coefficient on Gas Fee_{t-1} is 0.036, indicating that a one-dollar increase in the demeaned gas fee corresponds to a roughly 4 percent increase in the daily transfer amount.

One concern is that the one-hour window identifies spurious matches when two unrelated transfers of the same rounded dollar amount happen to coincide across chains. To address this, we shorten the window to 15 minutes. We also exclude transfers at the two most common round-number amounts, \$100 million and \$500 million, which each appear more than 100 times among the 1,230 matches. These common-number transfers could reflect routine treasury management rather than gas-driven reallocation and could produce disproportionately many false positives.

The estimates from columns 1 and 2 are stable across specifications with these restrictions: columns 3–4 tighten the matching window to 15 minutes, and columns 5–6 further exclude transfers at common round-number amounts. Across all six columns, the coefficient on the lagged gas fee remains positive and statistically significant at the 5 percent level, ranging from 0.030 to 0.039. Bitcoin volatility enters with a negative and significant coefficient, consistent with stablecoin holders being less likely to execute cross-chain transfers during periods of broad market stress.

The bottom panel examines the extensive margin. The dependent variable is an indicator for whether the net daily transfer from Ethereum to Tron is positive. The coefficient on the lagged gas fee is positive across all specifications, though statistically weaker than the intensive-margin results. Higher gas fees relative to recent history modestly increase the probability that a given day features net migration from Ethereum to Tron, but the primary effect operates through larger transfer volumes on days when migration does occur.

Taken together, the results in Table 9 provide direct, transaction-level evidence that Ethereum congestion costs push stablecoin balances toward competing blockchains. Stablecoin holders respond to high gas fees not by abandoning the stablecoin entirely, but by shifting the blockchain on which they hold it. The test supports the mechanism underlying our

model: congestion raises the cost of using digital money on a given blockchain, and users who value the stablecoin's payment function migrate to cheaper alternatives. The cross-chain migration documented here is a concrete channel through which the aggregate circulation shifts identified in the empty-slots analysis (Section 5.4) occur.

5.6 Implications

Our results point to a distinction that many discussions largely overlook: the difference between the liability and the rail. Proposed regulatory frameworks have focused almost exclusively on the quality of the reserve assets backing stablecoins, whether issuers can honor redemptions at par, and how stablecoins compare to bank deposits and money market funds. These are important concerns. But our paper shows that even when the liability is perfectly safe, the rail on which the money circulates can be a source of fragility.

The fragility we document stems from the reliance on decentralized trust for settlement. In the traditional payment system, settlement finality comes from institutional trust: the central bank operates the payment infrastructure, and the cost of settling a transaction is stable, predictable, and largely decoupled from the volume of transactions on the network. The cost of operating Fedwire, for example, is borne by the Federal Reserve and recovered through fixed per-transaction fees that do not vary with network congestion (Table 1). By contrast, settlement on a permissionless blockchain requires decentralized trust, priced through gas fees that fluctuate with network-wide demand for block space. This introduces a fundamentally different source of instability in which the cost of using money depends on the behavior of all other users of the blockchain, most of whom are not stablecoin holders.

Stablecoin Regulation. Recent U.S. and European stablecoin proposals have largely focused on the issuer side of the stablecoin arrangement—especially reserve portfolio limits, disclosures and attestations, and holders' redemption rights—but that framing may be narrow and does not fully address broader prudential and operational issues.

A stablecoin fully backed by Treasuries and subject to bank-like regulation would still be exposed to the congestion-driven fragility we document if it settles on a permissionless blockchain with volatile gas fees. Consideration should be given not only to the quality of reserves but also to the properties of the settlement infrastructure, including the volatility and level of transaction costs, the degree to which costs are shared with unrelated blockchain

activity, and the capacity of the network to absorb surges in demand.

Central Bank Digital Currencies on Public Blockchains. Several central banks, including the European Central Bank, have explored issuing central bank digital currencies (CBDCs) on public blockchains rather than on proprietary infrastructure. Our results suggest that this design choice carries infrastructure-related risk. A CBDC circulating on a permissionless blockchain would inherit the congestion externality we document: its usability would depend on network-wide activity, and its holders would face gas fees driven by demand for block space from cryptocurrency speculation, NFT minting, and other activity unrelated to the CBDC itself. A CBDC, a liability of a central bank so ostensibly the safest possible money, could become fragile not because of any concern about its backing but because the rail on which it circulates is congested and expensive.

Tokenized Deposits and Treasuries. The same logic applies to tokenized deposits and tokenized Treasury securities issued on permissionless blockchains. Several large financial institutions have recently announced plans to tokenize traditional financial assets on Ethereum and other public blockchains. Our mechanism implies that these instruments would be exposed to congestion-driven fragility regardless of the creditworthiness of the issuer. A tokenized insured deposit is as safe as the guarantee, but if the cost of transferring it spikes because of unrelated blockchain congestion, holders may migrate to cheaper rails, triggering the same network-externality dynamics we document for stablecoins.

Addressing the Rails. Our analysis also points toward design features that could mitigate the fragility. The key condition for instability is that the cost of settlement is volatile, congestion-sensitive, and shared across all users of the blockchain. Designs that break any of these links would reduce fragility. A central-bank-operated blockchain, for example, would replace decentralized trust with institutional trust, eliminating the congestion externality. Dedicated blockchains or application-specific chains that reserve block space for financial transactions would prevent congestion from unrelated activity from spilling over into the cost of using digital money. Layer 2 solutions that batch transactions off-chain and settle periodically on the base layer could decouple the user's transaction cost from real-time network congestion. Each of these approaches involves tradeoffs with the decentralization,

interoperability, and permissionlessness that make public blockchains attractive in the first place. For example, to be trusted, Layer 2 chains would likely require some degree of centralization, operated by a financial entity resembling CCPs in traditional finance, under established contractual and regulatory arrangements to address moral hazard and other risks. Our results suggest that these tradeoffs deserve more explicit attention in the design of digital payment infrastructure.

6 Conclusion

Digital money differs from earlier forms of money in how it provides trust. Rather than relying on a trusted intermediary, digital money circulating on permissionless blockchains relies on decentralized verification, priced one transaction at a time through gas fees. We show that this unbundling of trust introduces a novel source of fragility, one that persists even when the reserves backing the money are perfectly safe and liquid.

The fragility arises from the interaction of two forces inherent to permissionless blockchains: network externalities in the use of digital money and congestion in the supply of decentralized trust. Network externalities make the money more valuable as more users adopt it, but heavier use of the blockchain congests the network and drives up transaction costs. We use a global games model to show that this interaction makes individual redemption decisions strategic complements, and we document the resulting fragility empirically using stablecoins on Ethereum. Redemptions are disproportionately higher when congestion is high and network externalities are low, and the effect is concentrated in the upper tail of the congestion distribution, consistent with a threshold-driven regime shift rather than a smooth, continuous response. Using plausibly exogenous variation in Ethereum block-space supply from missed validator slots, we show that supply-driven congestion shifts stablecoin circulation from Ethereum to competing blockchains, and direct transaction-level evidence confirms that higher gas fees push USDT balances from Ethereum to Tron.

Our results separate two components of digital money: the liability and the rail. Even when the liability is unquestionably safe—fully backed by central bank reserves or high-quality sovereign bonds like Treasuries—the rail itself can be a source of fragility if it relies on costly and volatile decentralized trust for settlement. This distinction matters for policy. Stablecoins, tokenized deposits, tokenized Treasuries, and central bank digital currencies

issued on permissionless blockchains would all be exposed to the mechanism we document, regardless of the safety of their backing. Alternatives that reduce reliance on decentralized trust could mitigate these fragilities: for example, with a central-bank-operated blockchain, a fully trustworthy Layer 2, or designs that decouple settlement costs from network congestion.

Most observers have thus far largely focused on what stablecoin issuers themselves can do to keep their tokens trading at par. Our analysis shows that digital monies depend not only on the quality of their backing but also on the capacity and design of the rails on which they circulate.

References

Ahmed, R., I. Aldasoro, and C. Duley (2023). Par for the course: Public information and stable coin runs. *Working Paper*.

Ahnert, T., K. Anand, J. Chapman, and P. Gai (2019). Asset encumbrance, bank funding, and fragility. *Review of Financial Studies 32*(6), 2422–55.

Ahnert, T., P. Hoffman, A. Leonello, and D. Porcellachia (2024). Central bank digital currency and financial stability. *working paper*.

Anadu, K., P. Azar, M. Cipriani, C. Huang, M. Landoni, G. L. Spada, M. Macchiavelli, A. Malfroy-Camine, and J. C. Wang (2024). Runs and flights to safety: Are stablecoins the new money market funds?

Angeletos, G.-M. and A. Pavan (2004). Transparency of information and coordination in economies with investment complementarities. *American Economic Review 94*(2), 91–98.

Arseneau, D. M., D. E. Rappoport, and A. P. Vardoulakis (2020). Private and public liquidity provision in over-the-counter markets. *Theoretical Economics 15*(4), 1669–1712.

Bernardo, A. E. and I. Welch (2004). Liquidity and financial market runs. *The Quarterly Journal of Economics 119*(1), 135–158.

Bertsch, C. (2023). Stablecoins: Adoption and fragility. *Sveriges Riksbank Working Paper Series*.

Boissay, F., G. Cornelli, S. Doerr, and J. Frost (2022). Blockchain scalability and the fragmentation of crypto. *BIS Bulletin No 56*.

Bolt, W., J. Frost, H. S. Shin, and P. Wierts (2024). The bank of amsterdam and the limits of fiat money. *Journal of Political Economy 132*(12), 3919–3941.

Budish, E. (2025). Trust at scale: The economic limits of cryptocurrencies and blockchains. *The Quarterly Journal of Economics 140*, 1–62.

Carapella, F., J.-W. Chang, S. Infante, M. Leistra, A. Lubis, and A. P. Vardoulakis (2025). Financial stability implications of cbdc. *International Journal of Central Banking*.

Carletti, E., A. Leonello, and R. Marquez (2023). Loan guarantees, bank underwriting policies and financial stability. *Journal of Financial Economics 149*(2), 260–295.

Carlsson, H. and E. van Damme (1993). Global games and equilibrium selection. *Econometrica 61*(5), 989–1018.

Chahrour, R. and R. Valchev (2021, 10). Trade finance and the durability of the dollar. *The Review of Economic Studies 89*(4), 1873–1910.

Chen, Q., I. Goldstein, and W. Jiang (2010). Payoff complementarities and financial fragility: Evidence from mutual fund outflows. *Journal of Financial Economics 97*(2), 239–262.

CoinGecko (2026). *Cryptocurrency Price and Market Data*. Retrived from `https://www.coingecko.com`.

Cooper, R. and A. John (1988). Coordinating coordination failures in keynesian models. *The Quarterly Journal of Economics 103*(3), 441–463.

Coppola, A., A. Krishnamurthy, and C. Xu (2026). Liquidity, debt denomination, and currency dominance. *Journal of Finance*.

CryptoCompare. Enterprise PRO API and Custom Measures Data. `https://www.cryptocompare.com/`.

Diamond, D. W. and P. H. Dybvig (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy 91*(3), 401–419.

Doepke, M. and M. Schneider (2017). Money as a unit of account. *Econometrica 85*(5), 1537–1574.

Eisenbach, T. M. and G. Phelan (2026). Fragility of safe asset markets. *Review of Financial Studies*.

Fernández-Villaverde, J., D. Sanches, L. Schilling, and H. Uhlig (2021). Central bank digital currency: Central banking for all? *Review of Economic Dynamics 41*, 225–242.

Goldstein, I., H. Jiang, and D. T. Ng (2017). Investor flows and fragility in corporate bond funds. *Journal of Financial Economics 126*(3), 592–613.

Goldstein, I. and A. Pauzner (2005). Demand-deposit contracts and the probability of bank runs. *Journal of Finance 60*(3), 1293–1327.

Gopinath, G. and J. C. Stein (2020, 10). Banking, trade, and the making of a dominant currency*. *The Quarterly Journal of Economics 136*(2), 783–830.

Gorton, G. B., E. C. Klee, C. P. Ross, S. Y. Ross, and A. P. Vardoulakis (2025). Leverage and stablecoin pegs. *Journal of Financial and Quantitative Analysis*, 1–65.

Gorton, G. B., C. P. Ross, and S. Y. Ross (2026). Making money. *Journal of Finance*.

Hayashi, F. and W. R. Keeton (2012). Measuring the costs of retail payment methods. *Economic Review, Federal Reserve Bank of Kansas City 97*.

Infante, S. and A. P. Vardoulakis (2020, 12). Collateral Runs. *The Review of Financial Studies 34*(6), 2949–2992.

Kaplan, S. N. and A. Schoar (2005). Private equity performance: Returns, persistence, and capital flows. *Journal of Finance 60*(4), 1791–1823.

Kashyap, A. K., D. P. Tsomocos, and A. P. Vardoulakis (2024). Optimal bank regulation in the presence of credit and run risk. *Journal of Political Economy 132*(3), 772–823.

Llama Corp. DeFiLlama. `https://defillama.com/`.

Ma, Y., Y. Zeng, and A. L. Zhang (2025). Stablecoin runs and the centralization of arbitrage. *Review of Financial Studies*.

Matta, R. and E. Perotti (2023, 11). Pay, Stay, or Delay? How to Settle a Run. *The Review of Financial Studies 37*(4), 1368–1407.

Morris, S. and H. S. Shin (2003). Global Games: Theory and Applications. *In Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress, Vol. 1, ed. Mathias Dewatripont, Lars P. Hansen, and Stephen J. Turnovsky. Cambridge: Cambridge University Press*.

Schilling, L. (2023). Optimal forbearance of bank resolution. *The Journal of Finance*.

Schilling, L., J. Fernández-Villaverde, and H. Uhlig (2024). Central bank digital currency: When price and bank stability collide. *Journal of Monetary Economics 145*, 103554.

Sirri, E. R. and P. Tufano (1998). Costly search and mutual fund flows. *Journal of Finance 53*(5), 1589–1622.

Sockin, M. and W. Xiong (2023). A model of cryptocurrencies. *Management Science 69*(11), 6684–6707.

7 Figures

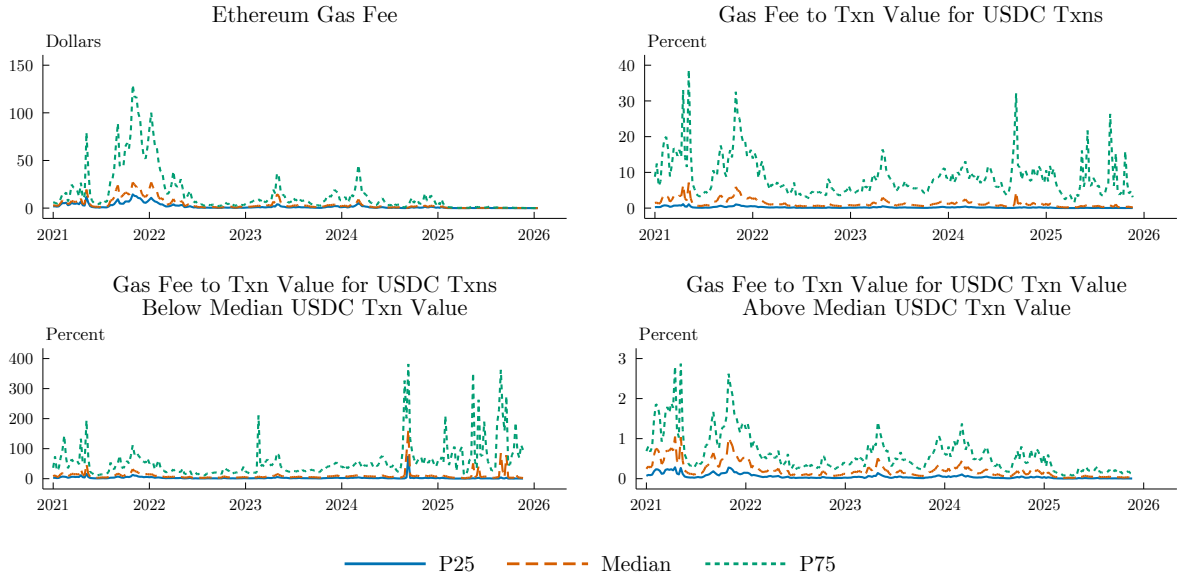

Figure 1: Gas Fees and Ratio to Transaction Value. Figure plots weekly Ethereum gas fees and the ratio of the gas fee to transaction value for USDC transfers from 2021 through 2025. The gas fee paid for a single transaction equals the *gas price* (the per-unit cost of blockspace, denominated in ETH) multiplied by the *gas used* (a unitless measure of the computational work the transaction requires). The resulting fee is converted to U.S. dollars at the prevailing ETH/USD exchange rate. Each series shows the 25th, 50th, and 75th percentiles across transactions within a day, averaged to the weekly frequency. The top-left panel plots the Ethereum gas fee in U.S. dollars. The top-right panel plots the gas-fee-to-transaction-value ratio for all USDC transfers on Ethereum. The bottom-left panel restricts to USDC transfers with values below the daily median transaction size. The bottom-right panel restricts to USDC transfers with values above the daily median transaction size.

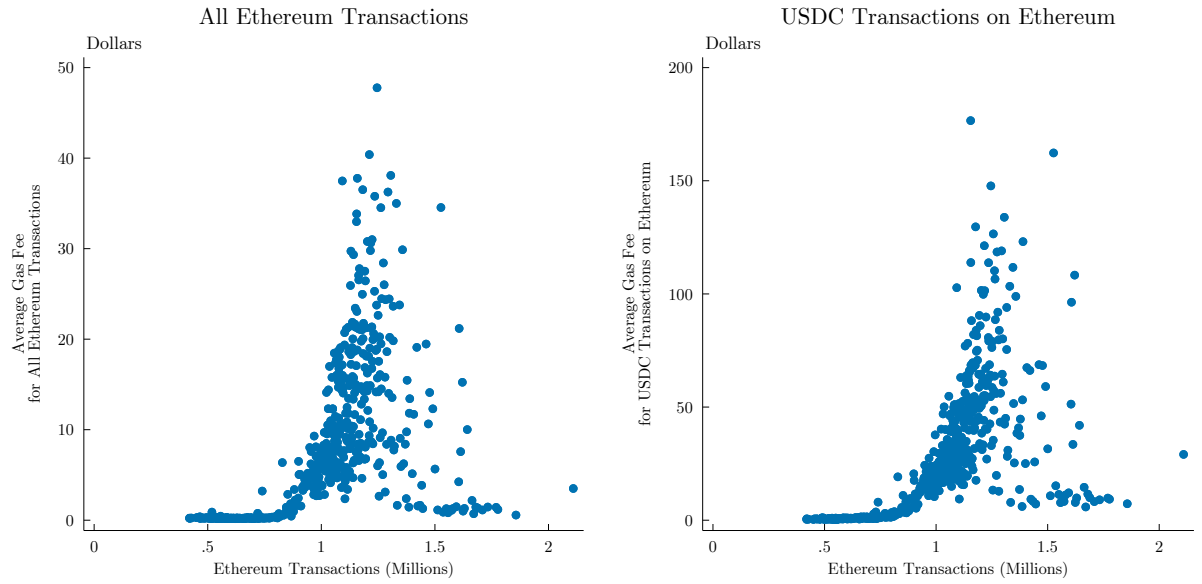

Figure 2: Gas Fees vs. Number of Transactions. Figure plots a binned scatter of Ethereum gas fees against the number of daily transactions. The left panel shows the gas fees for all Ethereum transactions; the right panel shows gas fees for USDC transactions.

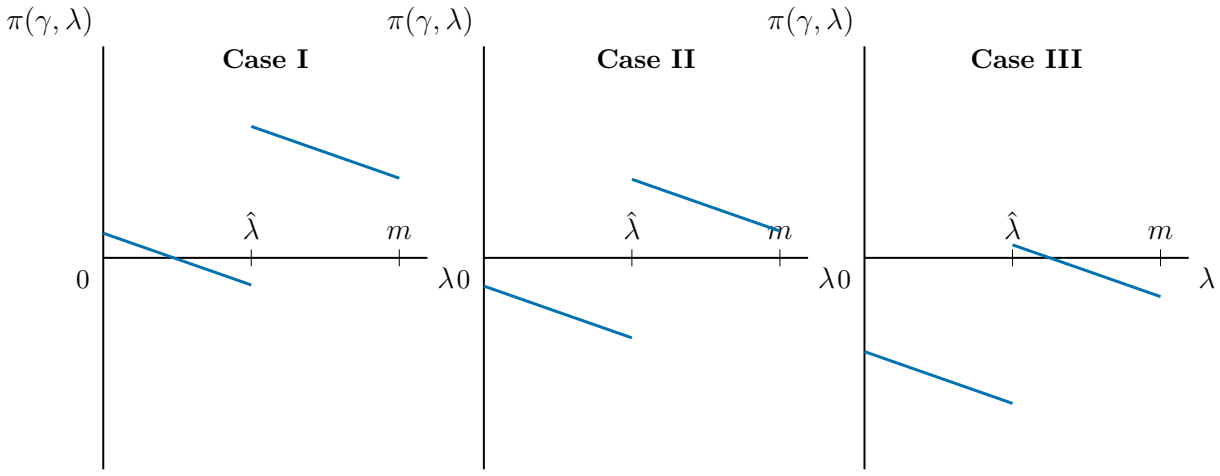

Figure 3: Possible Payoff differential $\pi(\gamma, \lambda)$ as a function of λ for $\gamma \in (\underline{\gamma}, \overline{\gamma})$.

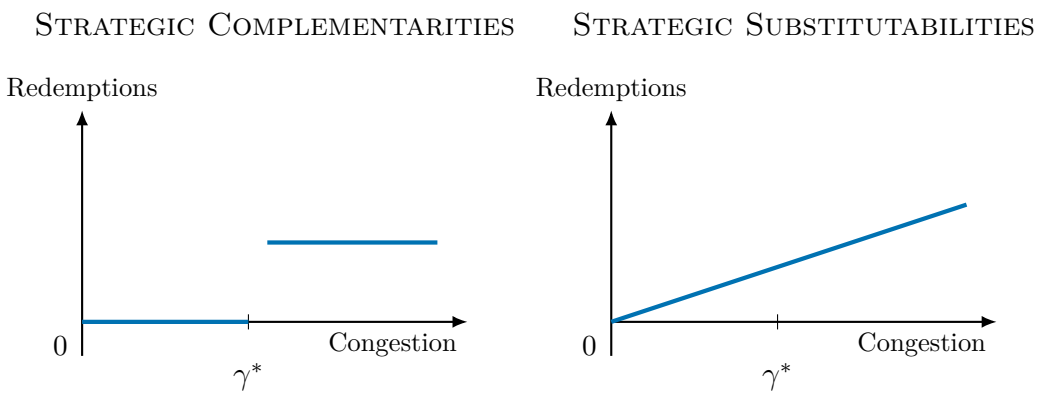

Figure 4: Redemptions under strategic complementarities in our model and pre-emptive redemptions in Bernardo and Welch (2004).

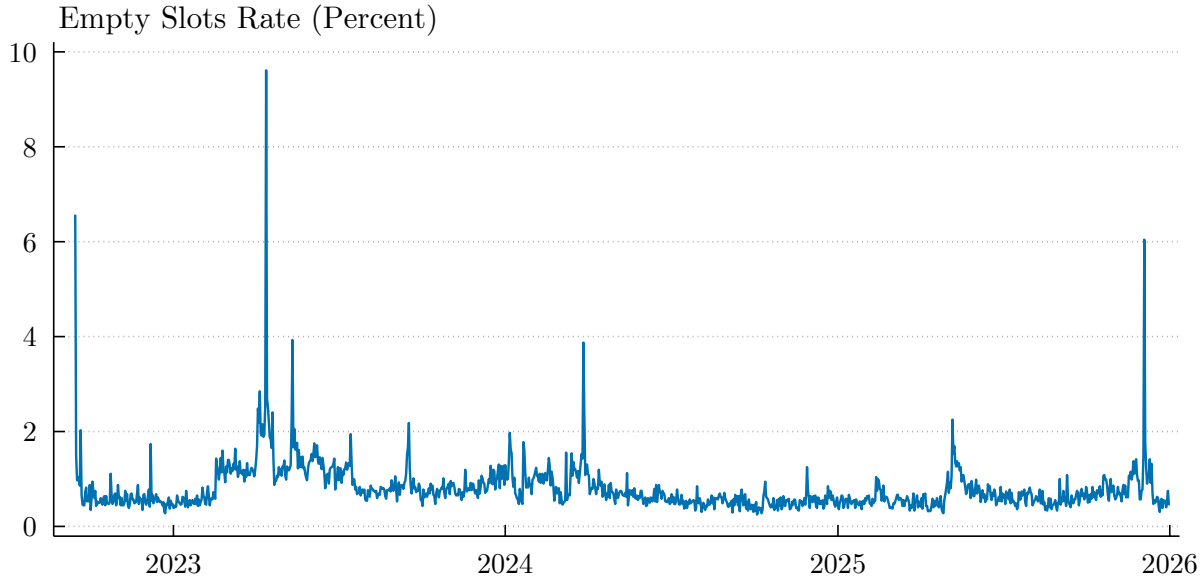

Figure 5: Empty Slots Rate on Ethereum. Figure plots the daily share of empty slots on the Ethereum blockchain since the Merge in September 2022.

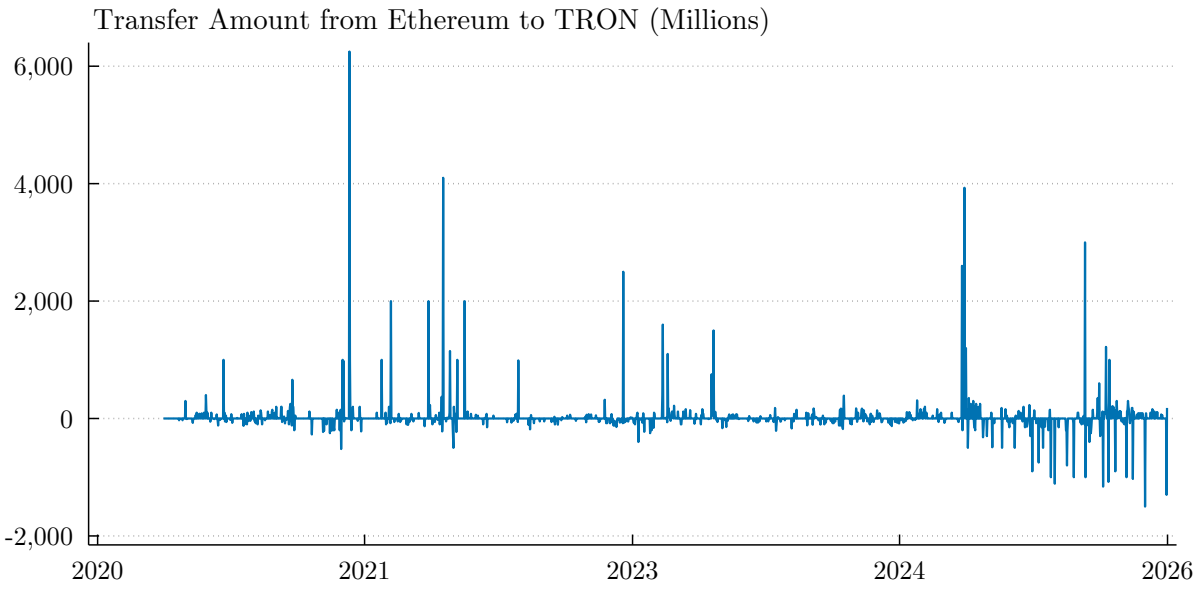

Figure 6: USDT Transfers from Ethereum to Tron. Figure plots the daily net transfer amount of USDT from Ethereum to Tron, in millions of U.S. dollars. Transfers are identified by matching USDT transfer events on Ethereum to USDT transfer events of the same dollar amount on Tron occurring within 60 minutes of each other. Positive values indicate net flows from Ethereum to Tron; negative values indicate net flows in the opposite direction.

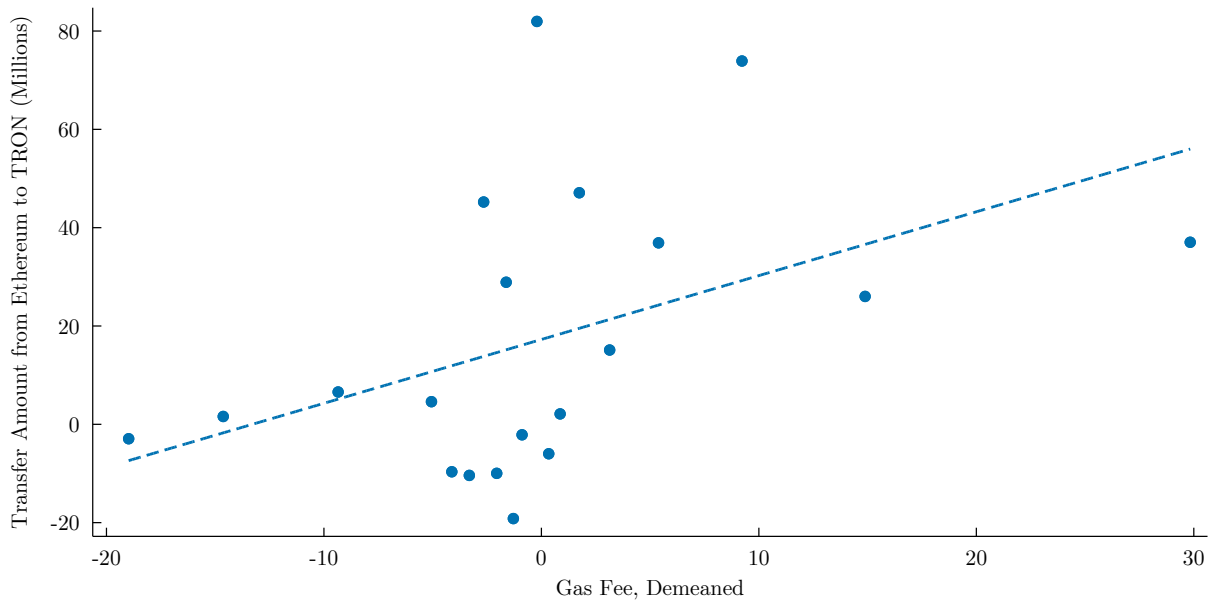

Figure 7: USDT Transfers from Ethereum to Tron vs. Gas Fees. Figure plots a binned scatter of the daily net USDT transfer amount from Ethereum to Tron (in millions) against the demeaned Ethereum gas fee, where the gas fee is demeaned by its 365-day trailing moving average. The positive relationship indicates that days following above-average congestion on Ethereum tend to coincide with larger net transfers of USDT from Ethereum to Tron.

8 Tables

$ millions	Fedwire Transaction Costs		Ethereum Transaction Costs	
Year	Lower Bound	Upper Bound	25th Percentile	75th Percentile
2021	33.74	171.77	1,013.43	7,512.20
2022	33.33	172.53	390.58	3,568.22
2023	34.80	177.85	214.28	1,924.37
2024	39.88	197.32	209.81	1,890.20
2025	42.37	210.78	22.82	222.52

Table 1: Comparison of Fedwire and Ethereum Transaction Costs. All values are in millions of U.S. dollars. Each column reports the cost of processing the year's Fedwire transfer volume under the relevant fee schedule. Fedwire lower and upper bounds multiply annual Fedwire transfers by the lowest and highest fees in Fedwire's tiered, volume-based schedule. Ethereum lower and upper bounds multiply annual Fedwire transfers by the annual average of the daily 25th and 75th percentile Ethereum gas fees. Fedwire volumes are from `https://www.frbservices.org/resources/financial-services/wires/volume-value-stats/monthly-stats.html`.

	N	Mean	SD	Min	p25	Median	p75	Max
Redemptions (%)	1,184	−1.823	7.146	−78.096	−3.089	−0.447	0.937	44.195
ln(Circulation on Ethereum)	1,184	22.547	1.806	15.518	21.707	22.756	24.065	25.117
Congestion ($)	1,184	7.564	10.831	0.078	1.410	3.393	7.606	55.224
Velocity	1,184	0.537	1.148	0.002	0.110	0.245	0.518	25.125
$\mathbb{I}$(Low Network Externality)	1,184	0.075	0.264	0.000	0.000	0.000	0.000	1.000
Controls								
Bitcoin Return (%)	1,184	1.253	8.592	−36.469	−3.884	0.361	4.838	36.006
Bitcoin Volatility (%)	1,184	2.990	1.192	1.093	2.183	2.739	3.504	8.344
Stablecoin Basis (%)	1,184	1.100	3.356	−3.853	0.025	0.396	1.069	55.373
Capital Ratio	1,184	0.008	0.014	0.000	0.000	0.002	0.004	0.057
Reserves Liquid Share	1,184	0.626	0.427	0.000	0.000	0.787	1.000	1.000

Table 2: Summary Statistics. Table reports summary statistics for the main estimation sample used in Table 4. The unit of observation is a stablecoin-week. Redemptions is the negative weekly percentage change in the stablecoin's circulation on Ethereum. Congestion is the average Ethereum gas fee in dollars over the week. Velocity is the ratio of the stablecoin's weekly on-chain transfer volume to its average circulation. $\mathbb{I}$(Low Network Externality) equals one when a stablecoin's velocity falls more than one standard deviation below its 8-week trailing moving average. Controls include the weekly Bitcoin return, 30-day Bitcoin volatility, the stablecoin basis, the stablecoin issuer's capital ratio, and the issuer's liquid reserve share.

	λ			v		c	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Market Share	-7.120^{***}						
	(-3.07)						
$\frac{\text{Unique Addresses}}{\text{Ethereum Circulation}} \times 100$		$-2{,}484.659^{***}$					
		(-4.12)					
$\frac{\text{Unique Addresses}}{\text{Circulation}} \times 100$			$-2{,}799.701^{***}$				
			(-4.03)				
Velocity				-0.994^{**}			
				(-2.38)			
Velocity $-$ Velocity^{MA}					-0.411^{*}		
					(-1.69)		
Gas Fee						0.014	
						(0.85)	
Gas Fee in Stablecoin Transfers							0.000
							(0.05)
N	1,372	1,368	1,368	1,368	1,367	1,372	1,368
R^2	0.02	0.04	0.04	0.01	0.00	0.00	0.00
Stablecoin FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 3: Comparative Statics. Table presents univariate panel regressions of weekly stablecoin redemptions on lagged proxies for model parameters. The dependent variable is $\text{Redemptions}_{it} = -\Delta \log(\text{Circulation on Ethereum}_{it}) \times 100$. Columns are organized by the model parameter they proxy for. Columns 1–3 proxy for the network externality threshold $\hat{\lambda}$: $\text{Market Share}_{i,t}$ is the stablecoin's Ethereum circulation as a share of total stablecoin circulation; $\frac{\text{Unique Addresses}}{\text{Ethereum Circulation}} \times 100$ and $\frac{\text{Unique Addresses}}{\text{Circulation}} \times 100$ measure the number of distinct addresses transacting in the stablecoin scaled by its Ethereum or total circulation, respectively. Columns 4–5 proxy for the network externality benefit v: $\text{Velocity}_{i,t}$ is total transaction value divided by circulating supply, excluding issuer treasury transactions; $\text{Velocity} - \text{Velocity}^{MA}$ is the deviation of velocity from its 8-week moving average. Columns 6–7 proxy for congestion costs c: Gas Fee_t is the average Ethereum gas fee in dollars; $\text{Gas Fee in Stablecoin Transfers}_{i,t}$ is the average gas fee on stablecoin i's own transfer transactions. All explanatory variables are lagged one week. The sample covers five stablecoins (USDT, USDC, DAI, PYUSD, and BUSD) at a weekly frequency. All specifications include stablecoin fixed effects. t-statistics are reported in parentheses using robust standard errors clustered by week, where $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Redemptions				
	(1)	(2)	(3)	(4)	(5)
Congestion_{t-1}	0.006	0.000	0.000	0.000	−0.002
	(0.34)	(0.01)	(0.01)	(0.02)	(−0.17)
$\mathbb{I}(\text{Low Network Externality})_{i,t-1}$	−1.275	−0.665	−0.665	−0.690	−0.313
	(−0.94)	(−0.77)	(−0.78)	(−0.78)	(−0.37)
$\text{Congestion}_{t-1} \times \mathbb{I}(\text{Low Network Externality})_{i,t-1}$	0.084*	0.083**	0.082**	0.083**	0.069**
	(1.79)	(2.56)	(2.55)	(2.53)	(2.00)
$\text{Reserves Liquid Share}_{i,t-1}$			3.340	3.367	2.345
			(1.43)	(1.41)	(0.96)
$\text{Depeg}_{i,t-1}$				0.053	−0.081
				(0.12)	(−0.19)
$\text{Depeg}_{i,t-1} \times \mathbb{I}(\text{Low Network Externality})_{i,t-1}$				0.104	0.805
				(0.07)	(0.59)
$\text{Redemptions}_{i,t-1}$					0.268***
					(4.46)
N	1,368	1,184	1,184	1,184	1,184
R^2	0.00	0.02	0.02	0.02	0.09
Stablecoin FE	Yes	Yes	Yes	Yes	Yes
Controls	No	Yes	Yes	Yes	Yes

Table 4: Stablecoin Redemptions are Higher When Congestion is High and Network Externalities are Low. Table presents weekly panel regressions of stablecoin redemptions on Ethereum gas fees interacted with an indicator for low network externalities. The dependent variable is Redemptions_{it}, the negative change in the log circulation of stablecoin i on Ethereum from week $t-1$ to t, multiplied by 100. Congestion_{t-1} is the lagged weekly average Ethereum gas fee. $\mathbb{I}(\text{Low Network Externality})_{i,t-1}$ equals one when stablecoin i's velocity is more than one standard deviation below its 8-week moving average, where the standard deviation is calculated across all stablecoins. Column 1 includes stablecoin fixed effects and no controls. Column 2 adds the stablecoin basis, Bitcoin return, Bitcoin volatility, and lagged capital ratio. Column 3 adds the lagged liquid reserve share. Column 4 adds $\text{Depeg}_{i,t-1}$ and its interaction with $\mathbb{I}(\text{Low Network Externality})_{i,t-1}$. Column 5 adds lagged redemptions. t-statistics are reported in parentheses using robust standard errors clustered by week, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Redemptions on Ethereum							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Congestion_{t-1}	0.012	0.012	0.007	0.008	0.010	0.010	0.002	0.002
	(0.43)	(0.42)	(0.27)	(0.32)	(0.52)	(0.52)	(0.13)	(0.12)
$\text{Congestion}_{t-1} \times \mathbb{I}(\text{Gas}^M)_{t-1}$	0.003	0.003	−0.000	0.005				
	(0.07)	(0.08)	(−0.01)	(0.15)				
$\text{Congestion}_{t-1} \times \mathbb{I}(\text{Gas}^H)_{t-1}$	−0.010	−0.009	−0.007	−0.008	−0.008	−0.008	−0.002	−0.003
	(−0.28)	(−0.27)	(−0.20)	(−0.24)	(−0.28)	(−0.28)	(−0.09)	(−0.10)
$\mathbb{I}(\text{Low Network Externality})_{i,t-1} \times \text{Congestion}_{t-1} \times \mathbb{I}(\text{Gas}^M)_{t-1}$	0.087	0.087	0.042	−0.059				
	(1.14)	(1.15)	(0.56)	(−0.74)				
$\mathbb{I}(\text{Low Network Externality})_{i,t-1} \times \text{Congestion}_{t-1} \times \mathbb{I}(\text{Gas}^H)_{t-1}$	0.094**	0.094**	0.091**	0.094**	0.090**	0.089**	0.089**	0.094**
	(2.24)	(2.22)	(2.10)	(1.99)	(2.16)	(2.14)	(2.05)	(1.99)
$\mathbb{I}(\text{Gas}^M)_{t-1}$	−0.589	−0.598	−0.557	−0.741				
	(−0.76)	(−0.78)	(−0.86)	(−1.11)				
$\mathbb{I}(\text{Gas}^H)_{t-1}$	−0.369	−0.374	−0.293	−0.310	−0.243	−0.241	−0.171	−0.154
	(−0.56)	(−0.57)	(−0.47)	(−0.46)	(−0.39)	(−0.39)	(−0.29)	(−0.24)
$\mathbb{I}(\text{Low Network Externality})_{i,t-1}$	−0.432	−0.396	−0.059	−0.600	−0.255	−0.229	0.035	0.074
	(−0.54)	(−0.50)	(−0.08)	(−0.64)	(−0.33)	(−0.29)	(0.05)	(0.09)
$\mathbb{I}(\text{Low Network Externality})_{i,t-1} \times \mathbb{I}(\text{Gas}^M)_{t-1}$				2.793*				
				(1.85)				
$\mathbb{I}(\text{Low Network Externality})_{i,t-1} \times \mathbb{I}(\text{Gas}^H)_{t-1}$				0.329				−0.209
				(0.24)				(−0.16)
$\text{Depeg}_{i,t-1}$		−0.025	−0.123	−0.129		0.026	−0.075	−0.076
		(−0.06)	(−0.29)	(−0.31)		(0.06)	(−0.18)	(−0.18)
$\text{Depeg}_{i,t-1} \times \mathbb{I}(\text{Low Network Externality})_{i,t-1}$		−0.133	0.621	0.997		−0.074	0.667	0.670
		(−0.09)	(0.45)	(0.73)		(−0.05)	(0.49)	(0.49)
$\text{Reserves Liquid Share}_{i,t-1}$			2.091	2.074			2.233	2.243
			(0.85)	(0.84)			(0.91)	(0.91)
$\text{Redemptions}_{i,t-1}$			0.269***	0.267***			0.269***	0.269***
			(4.44)	(4.39)			(4.46)	(4.45)
N	1,184	1,184	1,184	1,184	1,184	1,184	1,184	1,184
R^2	0.02	0.02	0.10	0.10	0.02	0.02	0.09	0.09
Stablecoin FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 5: The Interaction of Congestion and Low Network Externalities is Concentrated in the Upper Tail of the Congestion Distribution. Table decomposes the interaction effect from Table 4 by allowing the sensitivity of low-network-externality stablecoins to congestion to differ across congestion regimes. The dependent variable is Redemptions_{it}, the negative change in the log circulation of stablecoin i on Ethereum from week $t-1$ to t, multiplied by 100. Congestion_{t-1} is the lagged average Ethereum gas fee. $\mathbb{I}(\text{Gas}^M)_{t-1}$ equals one when the gas fee exceeds its 8-week moving average but is less than one standard deviation above it. $\mathbb{I}(\text{Gas}^H)_{t-1}$ equals one when the gas fee exceeds its 8-week moving average by more than one standard deviation. $\mathbb{I}(\text{Low Network Externality})_{i,t-1}$ equals one when the stablecoin's velocity is more than one standard deviation below its 8-week moving average. All specifications include stablecoin fixed effects and controls for the stablecoin basis, Bitcoin return, and Bitcoin volatility. t-statistics are reported in parentheses using robust standard errors clustered by date, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Gas Fee					
	(1)	(2)	(3)	(4)	(5)	(6)
Empty Slots Rate$_t$	191.837***	221.481***				
	(3.43)	(3.30)				
Shock$_t$			162.617***	202.372**		
			(2.58)	(2.51)		
Shock$_{t-1}$					151.853***	190.215**
					(2.63)	(2.51)
N	1,204	1,126	1,203	1,125	1,202	1,124
R^2	0.06	0.08	0.02	0.04	0.02	0.04
Controls	No	Yes	No	Yes	No	Yes

Table 6: Empty Slots Increase Ethereum Gas Fees. Table presents regressions of Ethereum gas fees on the empty slots rate and the AR(1) innovation to the empty slots rate. The empty slots rate is the share of slots without a proposed block relative to the daily maximum of 7,200. The innovation is the residual from an AR(1) regression of the empty slots rate. Controls include Bitcoin return, Bitcoin volatility, and the stablecoin basis. The sample begins in September 2022 with the Ethereum Merge. t-statistics are reported in parentheses using robust standard errors clustered by date, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Shock$_t$			
	(1)	(2)	(3)	(4)
Stablecoin Basis$_t$	0.000			
	(1.00)			
$\mathbb{I}$(Depeg)$_t$		−0.000		
		(−0.01)		
Bitcoin Return$_t$			0.000	
			(0.35)	
Bitcoin Volatility$_t$				0.000
				(1.26)
N	1,125	1,203	1,203	1,203
R^2	0.00	0.00	0.00	0.00

Table 7: The Empty Slots Supply Shock is Orthogonal to Demand Conditions. Table presents regressions of the AR(1) innovation to the empty slots rate on variables that proxy for stablecoin demand conditions: Bitcoin return, Bitcoin 30-day realized volatility, and the USDC stablecoin basis. Each column reports a separate univariate regression. t-statistics are reported in parentheses using robust standard errors, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Ethereum Share − Solana Share		
	(1)	(2)	(3)
Shock_{t-1}	3.370***	7.815**	7.054***
	(2.61)	(2.51)	(4.35)
$\mathbb{I}(\text{Low Network Ext.})_{i,t-1}$	0.134***	0.143***	0.131***
	(19.60)	(20.29)	(19.05)
$\text{Shock}_{t-1} \times \mathbb{I}(\text{Low Network Ext.})_{i,t-1}$	−8.127**	−1.064	
	(−2.16)	(−0.36)	
$\mathbb{I}(\text{Shock}^H)_{t-1}$		−0.015**	
		(−2.55)	
$\text{Shock}_{t-1} \times \mathbb{I}(\text{Shock}^H)_{t-1}$		−4.822	
		(−1.38)	
$\text{Shock}_{t-1} \times \mathbb{I}(\text{Shock}^H)_{t-1} \times \mathbb{I}(\text{Low Network Ext.})_{i,t-1}$		−14.938**	
		(−1.97)	
$\text{Shock}_{t-1} \times \mathbb{I}(\text{Gas}^M)_{t-1}$			−5.670***
			(−3.16)
$\text{Shock}_{t-1} \times \mathbb{I}(\text{Gas}^H)_{t-1}$			−2.743
			(−1.00)
$\mathbb{I}(\text{Gas}^M)_{t-1}$			0.009*
			(1.82)
$\mathbb{I}(\text{Gas}^H)_{t-1}$			0.015***
			(2.99)
$\text{Shock}_{t-1} \times \mathbb{I}(\text{Gas}^M)_{t-1} \times \mathbb{I}(\text{Low Network Ext.})_{i,t-1}$			−3.599
			(−1.50)
$\text{Shock}_{t-1} \times \mathbb{I}(\text{Gas}^H)_{t-1} \times \mathbb{I}(\text{Low Network Ext.})_{i,t-1}$			−40.682***
			(−5.24)
N	2,723	2,723	2,723
R^2 (adjusted)	0.62	0.62	0.62
Stablecoin FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes

Table 8: Supply Shocks to Ethereum Congestion Shift Stablecoin Circulation to Solana. Table presents daily panel regressions of the difference between a stablecoin's Ethereum and Solana circulation shares on supply shocks to Ethereum congestion. The dependent variable is the difference between a stablecoin's Ethereum circulation share and its Solana circulation share. Shock_{t-1} is the lagged residual from an AR(1) regression of the daily empty slots rate on Ethereum. $\mathbb{I}(\text{Low Network Ext.})_{i,t-1}$ equals one when stablecoin i's velocity is more than one standard deviation below its 90-day moving average. $\mathbb{I}(\text{Shock}^H)_{t-1}$ equals one when the previous day's shock exceeded its 90-day moving average plus one standard deviation. $\mathbb{I}(\text{Gas}^M)_{t-1}$ and $\mathbb{I}(\text{Gas}^H)_{t-1}$ equal one when the previous day's gas fee was above its 90-day moving average (medium) or above the moving average plus one standard deviation (high). Controls include the stablecoin basis, Bitcoin return, Bitcoin volatility, lagged capital ratio, and lagged liquid reserve share. The sample begins in September 2022 with the Ethereum Merge. t-statistics are reported in parentheses using robust standard errors clustered by date, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Intensive Margin	ln(Transfer Amount)					
	Baseline		Within 15 Min		Within 15 Min, Excluding Routine Amounts	
	(1)	(2)	(3)	(4)	(5)	(6)
Gas Fee_{t-1}	0.033**	0.036**	0.036**	0.039**	0.030**	0.032**
	(2.15)	(2.33)	(2.30)	(2.47)	(2.03)	(2.17)
Bitcoin Return_t		0.049		0.044		0.037
		(1.49)		(1.37)		(1.19)
Bitcoin Volatility_t		−0.491***		−0.484***		−0.384***
		(−7.68)		(−7.74)		(−6.58)
N	2,662	2,552	2,671	2,561	2,698	2,588
R^2	0.00	0.02	0.00	0.02	0.00	0.01

Panel B: Extensive Margin	$\mathbb{I}$(Transfer Ethereum to Tron)					
	Baseline		Within 15 Min		Within 15 Min, Excluding Routine Amounts	
	(1)	(2)	(3)	(4)	(5)	(6)
Gas Fee_{t-1}	0.0012*	0.0013*	0.0013*	0.0014**	0.0011*	0.0012*
	(1.75)	(1.89)	(1.89)	(2.02)	(1.69)	(1.81)
Bitcoin Return_t		0.0021		0.0019		0.0016
		(1.31)		(1.17)		(1.04)
Bitcoin Volatility_t		−0.0224***		−0.0225***		−0.0183***
		(−6.94)		(−7.09)		(−6.13)
N	2,950	2,840	2,950	2,840	2,950	2,840
R^2	0.00	0.01	0.00	0.01	0.00	0.01

Table 9: Higher Ethereum Congestion Increases USDT Transfers to Tron. Table presents daily time-series regressions of USDT cross-chain transfers from Ethereum to Tron on lagged Ethereum gas fees. The dependent variable in the top panel is $\ln(\text{Transfer Amount}_t)$, the log of the net daily matched transfer amount from Ethereum to Tron. The dependent variable in the bottom panel $\mathbb{I}(\text{Transfer Ethereum to Tron}_t)$ is an indicator equal to one if the net matched transfer from Ethereum to Tron is positive on day t. Gas Fee_{t-1} is the lagged Ethereum gas fee demeaned by its 365-day trailing moving average. t-statistics shown using robust standard errors clustered by date where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Internet Appendix to

The Fragility of Perfectly Safe Digital Money

Elizabeth C. Klee, Arazi Lubis,
Chase P. Ross, Sharon Y. Ross, Alexandros P. Vardoulakis

The Internet Appendix consists of five sections. Section IA.A provides model derivations. Section IA.B formally tests gas fee convexity. Section IA.C presents robustness checks for the baseline results. Section IA.D presents robustness checks for the regime-shift results. Section IA.E provides additional exhibits.

IA.A Derivations

IA.A.1 Proof of Proposition 1

We show that there is an open set of (v_l, v_h, Γ) with $0 < v_l < v_h$ and $\Gamma > 0$ such that $0 < \underline{\gamma} < \overline{\gamma} < \Gamma$. The last inequality is trivially satisfied by setting Γ high enough, which does not affect any other conditions. To establish the rest we just need to show that, for any cost function $c(x)$ satisfying $c(0) = 0$, $c'(x) > 0$, and $c''(x) \geq 0$, and for any choice of parameters $m > 0$, $\hat{\lambda} \in (0, m)$, $n > 1$, and $\alpha > 0$, there exist values $0 < v_l < v_h$ such that (5) and (7) hold, ensuring well-defined dominance regions.

First, pick any $v_l > 0$. This determines $\underline{\gamma}$ uniquely from equation (2). Given v_l, the lower bound for v_h is given by the right-hand side of (7):

$$v_h^{lb}(v_l) \equiv \max\left(\frac{n-1}{n}\alpha c(nm + \underline{\gamma}), \frac{n-1}{n}\alpha c(m + 2(n-1)\hat{\lambda})\right), \tag{13}$$

which is positive and depends on v_l but not on v_h. Hence, we can always choose $v_h > v_h^{lb}(v_l)$ such that (7) holds. The last step is to show that (5) holds for sufficiently large v_h, i.e., $v_l < \frac{n-1}{n}\alpha c(m + \overline{\gamma})$. Note that as v_h increases, the right-hand side of this expression increases, because $d\overline{\gamma}/dv_h > 0$ from (4); recall that $c^{-1'} > 0$. As $v_h \to \infty$, $\lim_{v_h \to \infty} \overline{\gamma} = \infty$ and $\lim_{v_h \to \infty} c(m + \overline{\gamma}(v_h)) = \infty$. Hence, for any fixed $v_l > 0$, we can choose high enough v_h such that both $v_h > v_h^{lb}(v_l)$ and (5) are satisfied.

IA.A.2 Detailed Steps for Derivation of Unique γ^*

We now show the detailed steps to go from equation (1) to (8). Given the private signal, an individual agent updates their posterior about γ, which is uniform in $[s_i - \varepsilon, s_i + \varepsilon]$ and computes the expected payoff differential (equation (1)). If $s_i \geq \overline{\gamma} + \varepsilon$, the individual agent can conclude that $\gamma \geq \overline{\gamma}$ and will redeem the new money for old money, independent of their belief about λ ($\Delta(s_i) < 0$). Similarly, if $s_i < \underline{\gamma} - \varepsilon$, the individual agent can conclude that $\gamma < \underline{\gamma}$ and will not redeem, independent of their belief about λ ($\Delta(s_i) > 0$). These are the regions for γ, where the individual action is independent of the beliefs about the actions of others. For intermediate $s_i \in [\underline{\gamma} - \varepsilon, \overline{\gamma} + \varepsilon)$, the sign of $\Delta(s_i)$ depends on the beliefs about λ.

To pin down these beliefs, we focus on a threshold strategy that all agents follow. Consider a signal threshold γ^*, such that every agent redeems if their private signal $s_i > \gamma^*$ and does not redeem if $s_i < \gamma^*$. Given this threshold, an individual agent can form well-defined beliefs about the total mass of investors not redeeming, denoted by $\lambda^b(\gamma, \gamma^*)$.

The mass of investors who do not redeem equals the mass receiving signals below γ^*, i.e., $\lambda^b(\gamma, \gamma^*) = m \cdot \Pr(s_i < \gamma^* \mid \gamma)$. Since $\varepsilon_i \sim U[-\varepsilon, \varepsilon]$, we have $s_i \sim U[\gamma - \varepsilon, \gamma + \varepsilon]$. Under the threshold strategy, $\lambda^b(\gamma, \gamma^*)$ is given by:

$$\lambda^b(\gamma, \gamma^*) = \begin{cases} m & \text{if } \gamma < \gamma^* - \varepsilon \\ m \cdot \frac{\gamma^* - \gamma + \varepsilon}{2\varepsilon} & \text{if } \gamma^* - \varepsilon \leq \gamma \leq \gamma^* + \varepsilon \\ 0 & \text{if } \gamma > \gamma^* + \varepsilon \end{cases}. \tag{14}$$

If $\gamma < \gamma^* - \varepsilon$, all agents get signals $s_i < \gamma^*$, none redeem, and $\lambda^b(\gamma, \gamma^*) = m$. If $\gamma > \gamma^* + \varepsilon$, all agents get signals $s_i > \gamma^*$, all redeem, and $\lambda^b(\gamma, \gamma^*) = 0$. If $\gamma^* - \varepsilon \leq \gamma \leq \gamma^* + \varepsilon$, some agents get signals $s_i > \gamma^*$ while others get signals $s_i < \gamma^*$; thus, under the threshold strategy, $\lambda^b(\gamma, \gamma^*) = m(\gamma^* - \gamma + \varepsilon)/(2\varepsilon)$.

Using the beliefs in the equation above, an agent can compute the expected payoff differential using their posterior about γ, given the signal s_i and an assumed value for γ^*:

$$\Delta(s_i, \gamma^*) = \int_{s_i - \varepsilon}^{s_i + \varepsilon} \frac{\pi(\gamma, \lambda^b(\gamma, \gamma^*))}{2\varepsilon} d\gamma. \tag{15}$$

Unlike in equation (1), beliefs are now uniquely determined and pin down the payoff differential. Under a threshold strategy, an agent does not redeem ($\Delta(s_i, \gamma^*) > 0$) if $s_i < \gamma^*$ and redeems ($\Delta(s_i, \gamma^*) < 0$) if $s_i > \gamma^*$. By continuity, the agent that receives the threshold signal γ^* is indifferent between not redeeming and redeeming, i.e.,

$$\Delta(\gamma^*, \gamma^*) = \int_{\gamma^* - \varepsilon}^{\gamma^* + \varepsilon} \frac{\pi(\gamma, \lambda^b(\gamma, \gamma^*))}{2\varepsilon} d\gamma = 0. \tag{16}$$

As the noise vanishes ($\varepsilon \to 0$), we change variables from γ to λ to express the integral in terms of redemption levels. From $\lambda^b(\gamma, \gamma^*) = m(\gamma^* - \gamma + \varepsilon)/(2\varepsilon)$, as γ increases from $\gamma^* - \varepsilon$ to $\gamma^* + \varepsilon$, λ^b uniformly decreases from m to 0. We obtain $d\gamma = -2\varepsilon/m d\lambda$.

Substituting into the indifference condition:

$$\Delta^* = \int_0^m \frac{\pi(\gamma^*, \lambda)}{m} d\lambda = 0, \tag{17}$$

where, as $\varepsilon \to 0$, $\gamma \to \gamma^*$ (the threshold value).

The payoff π depends on whether λ crosses the network externality threshold $\hat{\lambda}$. For $\lambda < \hat{\lambda}$, the low benefit v_l accrues, while for $\lambda \geq \hat{\lambda}$, the high benefit v_h accrues. Therefore:

$$\begin{aligned} \Delta^* = \frac{1}{m} \Bigg[& \int_0^{\hat{\lambda}} \left[n v_l - (n-1)\alpha c(m + (n-1)\lambda + \gamma^*) \right] d\lambda \\ & + \int_{\hat{\lambda}}^m \left[n v_h - (n-1)\alpha c(m + (n-1)\lambda + \gamma^*) \right] d\lambda \Bigg] = 0, \end{aligned} \tag{18}$$

or, equivalently,

$$\Delta^* = \int_0^{\hat{\lambda}} \left[nv_l - (n-1)\alpha c(m + (n-1)\lambda + \gamma^*) \right] d\lambda + \int_{\hat{\lambda}}^{m} \left[nv_h - (n-1)\alpha c(m + (n-1)\lambda + \gamma^*) \right] d\lambda = 0, \tag{19}$$

which is equation (8).

IA.A.3 Proof of Proposition 3

Because $\frac{d\Delta^*}{d\gamma^*} = -(n-1)\alpha \int_0^m c'(m + (n-1)\lambda + \gamma^*)d\lambda < 0$, the sign of $\frac{d\gamma^*}{dp}$ equals the sign of $\frac{d\Delta^*}{dp}$, from equation (10). For $\hat{\lambda}$, we have $\frac{d\Delta^*}{d\hat{\lambda}} = n(v_l - v_h) < 0$, so $\frac{d\gamma^*}{d\hat{\lambda}} < 0$. For network externalities, $\frac{d\Delta^*}{dv_l} = n\hat{\lambda} > 0$ and $\frac{d\Delta^*}{dv_h} = n(m - \hat{\lambda}) > 0$. For congestion costs, $\frac{d\Delta^*}{d\alpha} = -(n-1)\int_0^m c(m + (n-1)\lambda + \gamma^*)d\lambda < 0$. For n, the sign of the derivative is ambiguous for any γ, but focusing on the equilibrium point we can utilize (8) to get $\frac{\partial\Delta^*}{\partial n} = -\frac{1}{n-1}[\hat{\lambda}v_l + (m - \hat{\lambda})v_h] - (n-1)\alpha \int_0^m c'(m + (n-1)\lambda + \gamma^*)\lambda \, d\lambda < 0$.

IA.B Gas Fee Convexity

To formally test for convexity in the gas fee–transaction volume relationship, we use a nested polynomial regression framework similar to Kaplan and Schoar (2005), who test for a concave relationship between venture capital fund size and performance by sequentially estimating linear and quadratic specifications, and Sirri and Tufano (1998), who document a convex flow-performance relationship for mutual funds. We regress:

$$\text{Gas Fee}_t = \beta_0 + \beta_1(\text{Transaction Count}_t) + \beta_2(\text{Transaction Count}_t)^2 + \varepsilon_t. \tag{20}$$

We estimate this regression for the full sample period and separately for the period before the Pectra hard fork (May 7, 2025), which substantially altered Ethereum's capacity to process transactions. Over the course of 2025, the network underwent a series of capacity expansions: validators increased the gas limit from 30 million to 36 million in February, the Pectra hard fork in May doubled blob throughput and raised calldata costs, and a further gas limit increase to 45 million followed in July.[20] These expansions approximately doubled Ethereum's effective transaction capacity within a few months. The capacity expansions were themselves partly a response to high and volatile congestion costs—precisely the convexity we document. Ethereum's developers explicitly justified the gas limit increases as necessary to reduce congestion and stabilize fees during peak usage.[21] The fact that the platform's governance responded to convex congestion costs by expanding capacity reinforces the economic significance of the pre-expansion fee structure for stablecoin fragility: convexity

[20] The Pectra upgrade was Ethereum's most comprehensive hard fork since the Merge. The most economically significant changes were EIP-7691, which doubled blob throughput from 3 to 6 target blobs per block; EIP-7623, which raised calldata costs to incentivize the use of blobs and reduced worst-case block sizes; and EIP-7702, which introduced account abstraction allowing gas payments in tokens other than ETH.

[21] See, e.g., `https://vitalik.eth.limo/general/2025/02/14/l1scaling.html`.

was severe enough to prompt a protocol-level intervention, and it will reemerge as demand grows toward the new capacity constraint.

Table IA.1 reports the regression estimates. In the pre-Pectra sample before May 2025, the quadratic term is positive and statistically significant, indicating a convex relationship between gas fees and transaction volume (column 2). An F-test rejects the linear specification in favor of the quadratic. Estimating the same specification on the full sample without controlling for the structural break yields a negative quadratic term (column 3), reflecting the level shift in the fee–volume relationship rather than a change in the underlying congestion technology.

Columns (4) and (5) account for the break. Including a post-Pectra indicator restores a positive and significant quadratic term (column 4). Column (5) tests whether the entire polynomial shifts after Pectra by interacting both the linear and quadratic terms with the post-Pectra indicator. The pre-Pectra coefficients in column (5) are identical to those in column (2), confirming that the convex relationship is unchanged within the pre-Pectra regime. The interaction on the quadratic term is negative and significant, indicating that the post-Pectra fee–volume relationship is significantly less convex. This is consistent with the capacity expansion: when the network operates well below its constraint, congestion costs no longer accelerate with transaction volume, and the fragility channel identified in Section 3 is less likely to bind. Figure IA.1 illustrates the shift: the left panel shows the convex fee–volume relationship before Pectra, while the right panel shows the flattened relationship after the capacity expansion.

IA.C Robustness of Baseline Results

This section presents robustness checks for the baseline regression in Table 4, which establishes that stablecoin redemptions are disproportionately higher when congestion is high and network externalities are low. We show that the result is robust to alternative rolling window lengths, alternative congestion measures, and alternative data frequencies.

Alternative Rolling Window. Our baseline specification uses 8-week rolling windows to calculate the moving averages and standard deviations underlying $\mathbb{I}(\text{Low Network Externality})_{i,t-1}$. Table IA.3 replicates the baseline using 52-week rolling windows. The interaction coefficient β_2 is positive and significant in all specifications with controls, with magnitudes that are considerably larger than the baseline—ranging from 0.331 to 0.501—reflecting the fact that a longer lookback window identifies more persistent departures from normal velocity. Table IA.4 replicates the baseline using 12-week rolling windows instead. The interaction coefficient β_2 remains positive and significant in nearly all specifications with controls, with magnitudes comparable to the baseline. The result is not sensitive to the choice of lookback window.

Stablecoin-Specific Congestion. Our baseline uses the network-wide average Ethereum gas fee as the congestion measure, which is common to all stablecoins in a given week. One concern is that this measure may not capture the congestion costs faced by individual stablecoin users, since gas fees can vary across transaction types. Table IA.5 replaces the network-wide gas fee with a stablecoin-specific congestion measure computed from the average

gas fee paid on each stablecoin's own transfer transactions. Because the congestion measure now varies both across stablecoins and over time, it is no longer collinear with time, and we include both stablecoin and time fixed effects. The time fixed effects absorb any common cross-sectional shocks that might confound the baseline specification. The interaction coefficient remains positive and significant, confirming that the result is not driven by the choice of congestion measure or by the absence of time fixed effects in the baseline.

Daily Frequency. Our baseline is estimated at the weekly frequency. Tables IA.6 and IA.7 replicate the baseline at the daily frequency using 365-day and 90-day rolling windows, respectively, for the moving averages and standard deviations underlying $\mathbb{I}(\text{Low Network Externality})_{i,t-1}$. The daily specifications provide roughly seven times as many observations and test whether the result holds at higher frequency. The interaction coefficient β_2 is positive and significant across all specifications in both tables. The 90-day window is the daily-frequency analog of the 8-week window used in the weekly regressions; the 365-day window provides a longer lookback that is less sensitive to short-term fluctuations in velocity. The consistency of results across both windows and both frequencies supports the robustness of the baseline finding.

IA.D Robustness of Regime-Shift Results

Placebo: Redemptions on Solana and Tron. If the regime-shift results in Table 5 reflect Ethereum-specific congestion rather than a common shock affecting all blockchains, then Ethereum gas fees should not drive redemptions on other chains. Table IA.8 estimates the regime-shift specification replacing Ethereum redemptions with redemptions on Solana (columns 1–3) and Tron (columns 4–6), where redemptions on each chain are defined as the negative change in the log circulation of the stablecoin on that chain. The triple interaction β^H—the coefficient on congestion interacted with high congestion and low network externalities—is not significant for Solana redemptions in any specification. For Tron, the coefficient is negative and significant, indicating that high Ethereum congestion combined with low network externalities is associated with increases in circulation on Tron rather than decreases—consistent with stablecoin holders migrating from Ethereum to Tron when Ethereum congestion is high, as documented in the main text.

Daily Frequency. Table IA.9 replicates the regime-shift specification at the daily frequency using 90-day rolling windows. The daily results are stronger than the weekly results: both β^M and β^H are positive and significant for low-network-externality stablecoins, with t-statistics above 2.4 across all specifications. The stronger results at the daily frequency likely reflect greater statistical power from the larger sample, as well as the fact that daily data better capture the short-lived congestion episodes that drive the regime-shift dynamics.

IA.E Figures and Tables

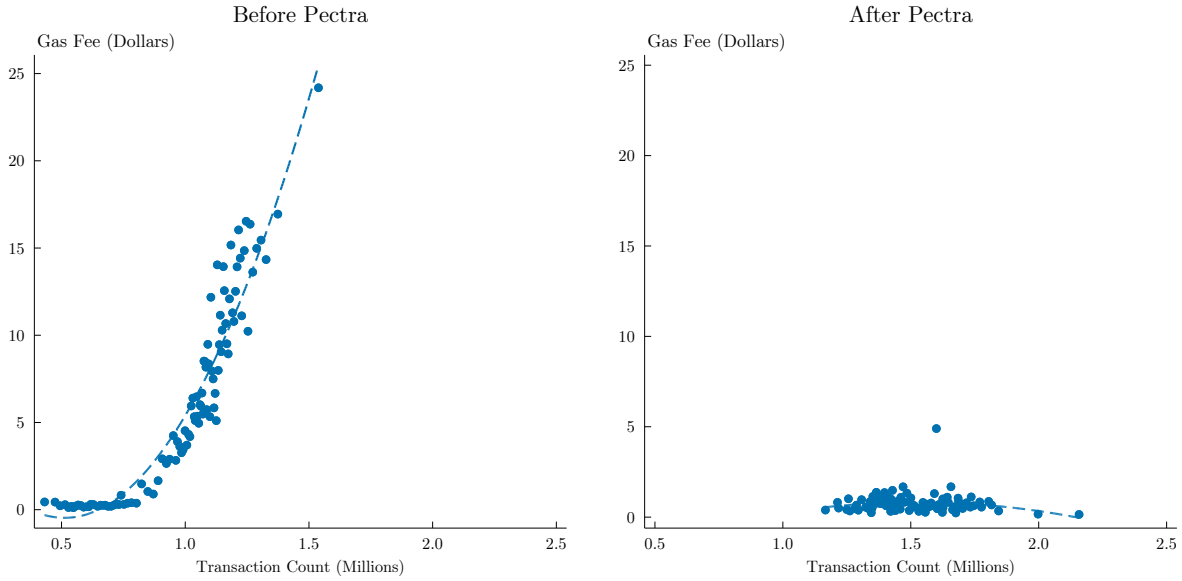

Figure IA.1: Gas Fees vs. Number of Transactions: Before and After Pectra. Figure plots binned scatters of Ethereum gas fees against daily transaction counts with quadratic fitted lines. The left panel uses data before the Pectra hard fork (May 7, 2025); the right panel uses data after the Pectra hard fork. Both panels use the same axis scales. The convex relationship visible in the left panel flattens after the capacity expansion.

	Before Pectra		Full Sample		
	Gas Fee				
	(1)	(2)	(3)	(4)	(5)
Transaction Count	19.808***	−25.100**	45.762***	9.717**	−25.100**
	(25.79)	(−2.30)	(19.22)	(2.08)	(−2.30)
Transaction Count^2		24.544***	−17.105***	4.939*	24.544***
		(3.92)	(−14.16)	(1.84)	(3.91)
$\mathbb{I}$(Post Pectra)				−17.406***	−10.052**
				(−14.85)	(−2.15)
$\mathbb{I}$(Post Pectra) × Transaction Count					31.569***
					(2.84)
$\mathbb{I}$(Post Pectra) × Transaction Count^2					−26.669***
					(−4.23)
N	2,715	2,715	2,953	2,953	2,953
R^2	0.22	0.24	0.13	0.23	0.26

Table IA.1: Gas Fee Convexity. The dependent variable is the average Ethereum gas fee in USD, Transaction Count is the total number of Ethereum transactions, measured in millions. Columns (1)–(2) restrict the sample to the period before the Pectra upgrade (May 7, 2025), which restructured Ethereum's fee mechanism and expanded blockspace capacity. Columns (3)–(5) use the full sample. $\mathbb{I}$(Post Pectra) equals one for dates after May 7, 2025. t-statistics are reported in parentheses using robust standard errors, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	N	Mean	SD	Min	p25	Median	p75	Max
Empty Slots Rate	2,723	0.007	0.004	0.002	0.005	0.006	0.008	0.096
Innovation in Empty Slots Rate	2,723	−0.000	0.003	−0.030	−0.001	−0.001	0.001	0.078
Ethereum Share − Solana Share	2,723	0.502	0.188	−0.316	0.393	0.448	0.614	0.843
Ethereum Share	2,723	0.587	0.166	0.342	0.413	0.612	0.720	0.862
Solana Share	2,723	0.085	0.118	0.006	0.011	0.025	0.124	0.658

Table IA.2: Summary Statistics for Daily Cross-Chain Analysis. Table reports summary statistics for the estimation sample used in Table 8. The unit of observation is a stablecoin-day. Empty Slots Rate is the fraction of Ethereum validator slots with no proposed block. Innovation in Empty Slots Rate is the residual from a daily AR(1) model of the empty slots rate. Ethereum Share and Solana Share are the fractions of the stablecoin's total circulation on their respective blockchains. The sample spans 2022 to 2025 and covers stablecoins with at least $100 million in circulation on Solana.

	Redemptions				
	(1)	(2)	(3)	(4)	(5)
Congestion_{t-1}	0.004	0.000	0.000	0.001	−0.001
	(0.23)	(0.03)	(0.02)	(0.09)	(−0.07)
$\mathbb{I}(\text{Low Network Externality})_{i,t-1}$	−3.379	−2.662	−2.650	−1.272	−1.180
	(−1.64)	(−1.48)	(−1.47)	(−1.10)	(−1.08)
$\text{Congestion}_{t-1} \times \mathbb{I}(\text{Low Network Externality})_{i,t-1}$	0.501**	0.430**	0.431**	0.422**	0.331**
	(2.50)	(2.20)	(2.21)	(2.33)	(2.10)
$\text{Reserves Liquid Share}_{i,t-1}$			3.335	3.415	2.395
			(1.43)	(1.43)	(0.98)
$\text{Depeg}_{i,t-1}$				0.280	0.091
				(0.67)	(0.23)
$\text{Depeg}_{i,t-1} \times \mathbb{I}(\text{Low Network Externality})_{i,t-1}$				−2.948	−1.790
				(−1.25)	(−0.88)
$\text{Redemptions}_{i,t-1}$					0.261***
					(4.54)
N	1,368	1,184	1,184	1,184	1,184
R^2	0.01	0.02	0.02	0.03	0.10
Stablecoin FE	Yes	Yes	Yes	Yes	Yes
Controls	No	Yes	Yes	Yes	Yes

Table IA.3: Strategic Complementarities Robustness: 52-Week Rolling Window. Table replicates the baseline regression from Table 4 using 52-week instead of 8-week rolling windows to calculate the moving averages and standard deviations underlying $\mathbb{I}(\text{Low Network Externality})_{i,t-1}$ and the congestion indicators. All specifications include stablecoin fixed effects and controls are the stablecoin basis, Bitcoin return, and Bitcoin volatility. t-statistics are reported in parentheses using robust standard errors clustered by date, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Redemptions				
	(1)	(2)	(3)	(4)	(5)
Congestion_{t-1}	0.008	0.002	0.002	0.002	0.001
	(0.47)	(0.16)	(0.15)	(0.16)	(0.04)
$\mathbb{I}(\text{Low Network Externality})_{i,t-1}$	−0.866	−0.192	−0.198	−0.315	0.213
	(−0.62)	(−0.24)	(−0.25)	(−0.36)	(0.24)
$\text{Congestion}_{t-1} \times \mathbb{I}(\text{Low Network Externality})_{i,t-1}$	0.075	0.074**	0.074**	0.075**	0.051
	(1.43)	(2.05)	(2.05)	(2.02)	(1.39)
$\text{Reserves Liquid Share}_{i,t-1}$			3.334	3.365	2.329
			(1.43)	(1.41)	(0.96)
$\text{Depeg}_{i,t-1}$				0.045	−0.072
				(0.10)	(−0.17)
$\text{Depeg}_{i,t-1} \times \mathbb{I}(\text{Low Network Externality})_{i,t-1}$				0.445	1.058
				(0.36)	(0.86)
$\text{Redemptions}_{i,t-1}$					0.269***
					(4.48)
N	1,368	1,184	1,184	1,184	1,184
R^2	0.00	0.02	0.02	0.02	0.09
Stablecoin FE	Yes	Yes	Yes	Yes	Yes
Controls	No	Yes	Yes	Yes	Yes

Table IA.4: Strategic Complementarities Robustness: 12-Week Rolling Window. Table replicates the baseline regression from Table 4 using 12-week instead of 8-week rolling windows to calculate the moving averages and standard deviations underlying $\mathbb{I}(\text{Low Network Externality})_{i,t-1}$ and the congestion indicators. All specifications include stablecoin fixed effects and controls are the stablecoin basis, Bitcoin return, and Bitcoin volatility. t-statistics are reported in parentheses using robust standard errors clustered by date, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Redemptions				
	(1)	(2)	(3)	(4)	(5)
$\text{Congestion}_{i,t-1}$	−0.003	0.017	0.018	0.018	0.018
	(−0.55)	(0.80)	(0.85)	(0.88)	(0.89)
$\mathbb{I}(\text{Low Network Externality})_{i,t-1}$	−1.431	−1.909**	−1.848*	−1.465	−1.296
	(−1.12)	(−1.98)	(−1.90)	(−1.52)	(−1.36)
$\text{Congestion}_{i,t-1} \times \mathbb{I}(\text{Low Network Externality})_{i,t-1}$	0.024**	0.021**	0.021*	0.019*	0.018*
	(2.09)	(2.03)	(1.95)	(1.87)	(1.75)
$\text{Reserves Liquid Share}_{i,t-1}$			−5.976**	−5.801*	−4.675
			(−2.09)	(−1.95)	(−1.61)
$\text{Depeg}_{i,t-1}$				0.394	0.291
				(0.78)	(0.59)
$\text{Depeg}_{i,t-1} \times \mathbb{I}(\text{Low Network Externality})_{i,t-1}$				−1.025	−0.259
				(−0.81)	(−0.20)
$\text{Redemptions}_{i,t-1}$					0.185***
					(2.97)
N	1,368	1,184	1,184	1,184	1,184
R^2	0.00	0.01	0.01	0.02	0.05
Stablecoin FE	Yes	Yes	Yes	Yes	Yes
Time FE	No	Yes	Yes	Yes	Yes
Controls	No	Yes	Yes	Yes	Yes

Table IA.5: Robustness: Stablecoin-Specific Gas Fees. Table replicates the baseline regression from Table 4, replacing the network-wide average gas fee with a stablecoin-specific congestion measure computed from each stablecoin's own transfer fees. All specifications include stablecoin fixed effects and controls are the stablecoin basis, Bitcoin return, and Bitcoin volatility. $\mathbb{I}(\text{Low Network Externality})_{i,t-1}$ is defined using 8-week summary statistics. t-statistics are reported in parentheses using robust standard errors clustered by date, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Redemptions				
	(1)	(2)	(3)	(4)	(5)
Congestion_{t-1}	0.003	0.001	0.001	0.001	0.001
	(1.37)	(0.29)	(0.31)	(0.32)	(0.32)
$\mathbb{I}(\text{Low Network Externality})_{i,t-1}$	-1.457^{***}	-0.668^{***}	-0.663^{***}	-0.423	-0.395
	(-3.79)	(-2.65)	(-2.64)	(-1.46)	(-1.37)
$\text{Congestion}_{t-1} \times \mathbb{I}(\text{Low Network Externality})_{i,t-1}$	0.113^{***}	0.053^{***}	0.053^{***}	0.052^{***}	0.046^{***}
	(3.82)	(3.48)	(3.47)	(3.53)	(3.57)
$\text{Reserves Liquid Share}_{i,t-1}$			0.623^{**}		0.582^{*}
			(2.01)		(1.84)
$\text{Depeg}_{i,t-1}$				0.017	0.017
				(0.31)	(0.33)
$\text{Depeg}_{i,t-1} \times \mathbb{I}(\text{Low Network Externality})_{i,t-1}$				-0.526	-0.420
				(-1.47)	(-1.22)
$\text{Redemptions}_{i,t-1}$					0.074^{**}
					(2.23)
N	9,689	8,314	8,314	8,314	8,314
R^2	0.01	0.01	0.01	0.01	0.01
Stablecoin FE	Yes	Yes	Yes	Yes	Yes
Controls	No	Yes	Yes	Yes	Yes

Table IA.6: Robustness: Daily Frequency, 365-Day Rolling Window. Table replicates the baseline regression from Table 4 at the daily rather than weekly frequency. The dependent variable is daily redemptions, defined as the negative change in the log circulation of stablecoin i on Ethereum from day $t-1$ to day t, expressed as a percentage. $\mathbb{I}(\text{Low Network Externality})_{i,t-1}$ is defined using a 365-day rolling window for the moving average and standard deviation of velocity. All specifications include stablecoin fixed effects and controls are the stablecoin basis, Bitcoin return, and Bitcoin volatility. t-statistics are reported in parentheses using robust standard errors clustered by date, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Redemptions				
	(1)	(2)	(3)	(4)	(5)
Congestion_{t-1}	0.004*	0.001	0.001	0.001	0.001
	(1.69)	(0.38)	(0.40)	(0.42)	(0.41)
$\mathbb{I}(\text{Low Network Externality})_{i,t-1}$	−0.898***	−0.427**	−0.428**	−0.266	−0.246
	(−3.06)	(−1.99)	(−1.99)	(−1.12)	(−1.03)
$\text{Congestion}_{t-1} \times \mathbb{I}(\text{Low Network Externality})_{i,t-1}$	0.028***	0.016**	0.016**	0.015**	0.013**
	(2.62)	(2.21)	(2.19)	(2.11)	(1.98)
$\text{Reserves Liquid Share}_{i,t-1}$			0.650**		0.598*
			(2.08)		(1.88)
$\text{Depeg}_{i,t-1}$				0.008	0.011
				(0.15)	(0.21)
$\text{Depeg}_{i,t-1} \times \mathbb{I}(\text{Low Network Externality})_{i,t-1}$				−0.415	−0.340
				(−1.34)	(−1.16)
$\text{Redemptions}_{i,t-1}$					0.076**
					(2.28)
N	9,689	8,314	8,314	8,314	8,314
R^2	0.00	0.01	0.01	0.01	0.01
Stablecoin FE	Yes	Yes	Yes	Yes	Yes
Controls	No	Yes	Yes	Yes	Yes

Table IA.7: Robustness: Daily Frequency, 90-Day Rolling Window. Table replicates the baseline regression from Table 4 at the daily frequency using a 90-day rolling window for the moving average and standard deviation of velocity underlying $\mathbb{I}(\text{Low Network Externality})_{i,t-1}$. The 90-day window is the daily-frequency analog of the 8-week window used in the weekly specifications. All specifications include stablecoin fixed effects and controls are the stablecoin basis, Bitcoin return, and Bitcoin volatility. t-statistics are reported in parentheses using robust standard errors clustered by date, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Redemptions on Solana			Redemptions on Tron		
	(1)	(2)	(3)	(4)	(5)	(6)
Congestion_{t-1}	−0.485*	−0.469	−0.268	0.048	0.028	0.055
	(−1.66)	(−1.59)	(−1.18)	(0.44)	(0.25)	(0.69)
$\text{Congestion}_{t-1} \times \mathbb{I}(\text{Gas}^M)_{t-1}$	0.560*	0.595*		0.036	0.050	
	(1.75)	(1.86)		(0.37)	(0.56)	
$\text{Congestion}_{t-1} \times \mathbb{I}(\text{Gas}^H)_{t-1}$	0.111	0.094	−0.103	0.023	0.036	0.010
	(0.38)	(0.32)	(−0.43)	(0.21)	(0.31)	(0.11)
$\mathbb{I}(\text{Low Network Externality})_{i,t-1} \times \text{Congestion}_{t-1} \times \mathbb{I}(\text{Gas}^M)_{t-1}$	0.300	0.245		−1.195***	−1.179***	
	(0.83)	(0.66)		(−4.62)	(−4.44)	
$\mathbb{I}(\text{Low Network Externality})_{i,t-1} \times \text{Congestion}_{t-1} \times \mathbb{I}(\text{Gas}^H)_{t-1}$	−0.475	−0.479	−0.528	−1.642**	−1.749**	−0.458
	(−0.72)	(−0.72)	(−0.79)	(−2.06)	(−2.13)	(−0.61)
$\mathbb{I}(\text{Gas}^M)_{t-1}$	−1.144	−1.344		0.303	0.417	
	(−0.79)	(−0.94)		(0.26)	(0.37)	
$\mathbb{I}(\text{Gas}^H)_{t-1}$	4.704*	4.696*	5.035*	−0.710	−0.302	−0.374
	(1.78)	(1.77)	(1.89)	(−0.45)	(−0.19)	(−0.25)
$\mathbb{I}(\text{Low Network Externality})_{i,t-1}$	0.290	0.400	0.831	13.495***	13.269***	7.681**
	(0.17)	(0.23)	(0.53)	(4.15)	(4.14)	(2.59)
$\text{Reserves Liquid Share}_{i,t-1}$		19.501	21.021		11.893***	11.696***
		(0.77)	(0.84)		(2.80)	(2.76)
$\text{Redemptions}_{i,t-1}$		0.111	0.105		0.092	0.093
		(0.71)	(0.67)		(0.45)	(0.45)
N	641	641	641	490	490	490
R^2	0.05	0.05	0.05	0.03	0.04	0.04
Stablecoin FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes

Table IA.8: Placebo: Redemptions on Solana and Tron and Ethereum Congestion. Table estimates the regime-shift specification from Table 5 replacing Ethereum redemptions with redemptions on Solana (columns 1–3) and Tron (columns 4–6). Congestion is measured using stablecoin-specific gas fees from Ethereum transfer transactions. $\mathbb{I}(\text{Low Network Externality})_{i,t-1}$ is defined using 8-week summary statistics. All specifications include stablecoin fixed effects and controls. t-statistics are reported in parentheses using robust standard errors clustered by date, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Relative to 90 Days				Relative to 365 Days			
	Redemptions on Ethereum							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Congestion_{t-1}	0.006	0.005	0.005	0.003	0.021***	0.020***	0.019***	0.006*
	(1.48)	(1.46)	(1.39)	(1.23)	(3.25)	(3.25)	(3.03)	(1.80)
$\text{Congestion}_{t-1} \times \mathbb{I}(Gas^M)_{t-1}$	−0.003	−0.003	−0.004		−0.014*	−0.014*	−0.013*	
	(−0.64)	(−0.62)	(−0.69)		(−1.86)	(−1.78)	(−1.69)	
$\text{Congestion}_{t-1} \times \mathbb{I}(\text{Gas}^H)_{t-1}$	−0.004	−0.004	−0.004	−0.002	−0.012*	−0.012*	−0.011	0.002
	(−0.88)	(−0.84)	(−0.80)	(−0.44)	(−1.76)	(−1.71)	(−1.61)	(0.39)
$\mathbb{I}(\text{Low Network Externality})_{i,t-1} \times Congestion_{t-1} \times \mathbb{I}(\text{Gas}^M)_{t-1}$	0.022**	0.019**	0.018**		0.075***	0.070**	0.065**	
	(2.53)	(2.19)	(2.09)		(2.69)	(2.44)	(2.42)	
$\mathbb{I}(\text{Low Network Externality})_{i,t-1} \times Congestion_{t-1} \times \mathbb{I}(\text{Gas}^H)_{t-1}$	0.016*	0.018**	0.015*	0.014*	0.051***	0.057***	0.052***	0.048***
	(1.91)	(2.05)	(1.87)	(1.79)	(3.72)	(4.04)	(3.96)	(3.92)
$\mathbb{I}(Gas^M)_{t-1}$	−0.061	−0.060	−0.041		−0.093	−0.101	−0.090	
	(−0.72)	(−0.72)	(−0.50)		(−1.03)	(−1.12)	(−1.00)	
$\mathbb{I}(\text{Gas}^H)_{t-1}$	−0.111	−0.113	−0.094	−0.087	−0.434***	−0.441***	−0.405***	−0.409***
	(−1.28)	(−1.29)	(−1.08)	(−1.03)	(−3.91)	(−4.01)	(−3.59)	(−3.69)
$\mathbb{I}(\text{Low Network Externality})_{i,t-1}$	−0.403**	−0.247	−0.230	−0.182	−0.603**	−0.350	−0.342	−0.254
	(−2.00)	(−1.07)	(−1.00)	(−0.85)	(−2.51)	(−1.18)	(−1.15)	(−0.92)
$\text{Depeg}_{i,t-1}$		0.005	0.008	0.011		0.009	0.004	0.006
		(0.09)	(0.15)	(0.20)		(0.17)	(0.07)	(0.11)
$\text{Depeg}_{i,t-1} \times \mathbb{I}(\text{Low Network Externality})_{i,t-1}$		−0.421	−0.344	−0.366		−0.552	−0.443	−0.461
		(−1.33)	(−1.16)	(−1.24)		(−1.52)	(−1.28)	(−1.33)
$\text{Reserves Liquid Share}_{i,t-1}$			0.526	0.555*			0.071	0.186
			(1.64)	(1.74)			(0.21)	(0.56)
$\text{Redemptions}_{i,t-1}$			0.076**	0.076**			0.070**	0.071**
			(2.27)	(2.27)			(2.10)	(2.13)
N	8,314	8,314	8,314	8,314	8,314	8,314	8,314	8,314
R^2	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.02
Stablecoin FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table IA.9: Regime Shift: Daily Frequency. Table replicates the regime-shift specification from Table 5 at the daily rather than weekly frequency. The dependent variable is daily redemptions on Ethereum. $\mathbb{I}(\text{Low Network Externality})_{i,t-1}$ and the congestion regime indicators are defined using 90-day and 365-day rolling windows. All specifications include stablecoin fixed effects and controls are the stablecoin basis, Bitcoin return, and Bitcoin volatility. t-statistics are reported in parentheses using robust standard errors clustered by date, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.