Stablecoin Runs and the Centralization of Arbitrage

Yiming Ma
Columbia Business School and NBER, USA

Yao Zeng
The Wharton School, University of Pennsylvania and NBER, USA

Anthony Lee Zhang
University of Chicago Booth School of Business, USA

Abstract

Stablecoins are cryptoassets designed to be pegged to the dollar, but they are backed by imperfectly liquid USD assets. We show that stablecoins feature *concentrated arbitrage*: the largest issuer, Tether, only allows six agents in an average month to redeem stablecoins for cash. We argue that issuers' choice of arbitrage concentration reflects a tradeoff: efficient arbitrage improves stablecoin price stability in secondary markets, but amplifies run risks by reducing investors' price impact from selling stablecoins. Our findings imply that policies designed to improve stablecoin price stability may have the unintended consequence of increasing stablecoin run risks. (*JEL* G01, G23)

We thank the Editor (Andrey Malenko), two anonymous referees, Garth Baughman, Markus Brunnermeier, Chris Cameron, Francesca Carapella, Caitlin Dannhauser, Eduardo Davila, Darrell Duffie, Vincent Glode, Itay Goldstein, Campbell Harvey, Chris Hua, Urban Jermann, Ye Li, Gordon Liao, Cyril Monnet, Cecilia Parlatore, Anna Pavlova, Mila Sherman, Chester Spatt, Dina Tsarapkina, Alexandros Vardoulakis, Dimitri Vayanos, and Ganesh Viswanath-Natraj for detailed comments, and conference and seminar participants at AFA AFFECT, Berkeley Haas, Boston Fed, Central Bank Conference on the Microstructure of Financial Markets, Chicago Booth, CKGSB, Columbia, Department of the Treasury, ECB Money Market Conference, Dartmouth Tuck, Duke/UNC/Milken DeFi Conference, Fed Board, Four Corners Investment Conference, FTG Michigan Meeting, HKUST, NBER Summer Institute Risk of Financial Institutions Meeting, NYU Stern Microstructure Conference, Oxford-Man Institute, Peking University, Purdue FinTech Conference, Texas A&M, Treasury/OFR Rising Scholar Conference, Tsinghua University, UC Irvine, UM Amherst, UT Dallas, University of Washington, Warwick Business School Gillmore Conference, Wharton, and WFA. Ding Ding, Pietro Ducco, Grace Jiarui Fan, Baiyang Han, Junyi Hu, Liming Ning, Yijing Ren, Haichuan Wang, Xinyi Wang, Yuming Yang, Max Yang, and Tianyi Zhang provided excellent research assistance. Send correspondence to Yiming Ma, ym2701@gsb.columbia.edu.

Fiat-backed stablecoins are blockchain assets whose value is claimed to be stable at $1. Such price stability is achieved by promising to back each stablecoin token with at least $1 in US dollar-denominated assets, such as bank deposits, Treasuries, corporate bonds, and loans. The six largest US dollar-backed stablecoins have grown from $5.6 billion in market capitalization at the beginning of 2020 to over $130 billion at the beginning of 2022. The potential for stablecoins to become a widely accepted means of payment that competes with fiat money and bank deposits (e.g., Brunnermeier, James and Landau, 2019; Duffie, 2019) has attracted active discussions about how to mitigate potential risks to financial stability and what the optimal regulatory framework should be.[1]

The ideal stablecoin always trades at a stable $1 and is free from panic runs. However, the determinants of and the relationship between stablecoins' run risk and price stability are far from obvious. Stablecoins hold illiquid assets while promising a fixed $1 redemption value. Unlike an MMF or a commercial bank, the right to redeem at $1 is restricted to a specific set of institutional arbitrageurs. The vast majority of investors can only trade stablecoins on secondary market exchanges, similar to investors trading ETF shares on secondary markets. These investors trade at the secondary market price, which frequently deviates above and below $1 depending on the demand and supply pressures in the market. Since the stablecoin price is determined in equilibrium, when and why would stablecoin investors want to run? How are stablecoin price deviations related to their run-risk?

We answer these questions by developing a framework for the market structure of stablecoins, showing how issuers, arbitrageurs, and customers interact to determine the stability of stablecoin prices and the likelihood of panic runs. Our first contribution is documenting the novel and surprising fact that stablecoins feature *concentrated arbitrage*. For example, the largest stablecoin, USDT, only allows for six arbitrageurs in an average month, whereas all other investors buy and sell stablecoins in competitive secondary markets. Our finding of concentrated arbitrage is surprising because arbitrageurs trade against fluctuations in stablecoin demand and supply: when the stablecoin price falls below $1, arbitrageurs can buy stablecoins from the secondary market and redeem them with the issuer for $1, which pushes the stablecoin price back up. In other words, issuers could simply authorize more arbitrageurs and make arbitrage more efficient to improve price stability in secondary markets.

Second, we show that while limiting arbitrage is harmful to stablecoin price stability, limiting arbitrage can, in fact, reduce the likelihood of panic runs. Panic runs by investors are possible

[1]For example, see G7 Working Group and others, 2019, "Investigating the Impact of Global Stablecoins"; ECB, 2020, "Stablecoins: Implications for monetary policy, financial stability, market infrastructure and payments, and banking supervision in the euro area"; BIS, 2020, "Stablecoins: potential, risks and regulation"; and IMF, 2021, "The Crypto Ecosystem and Financial Stability Challenges".

despite stablecoins trading on competitive exchanges. This is because issuers allow arbitrageurs to redeem stablecoins for $1 in primary markets, but back this promise by holding illiquid reserve assets. If enough stablecoin holders attempt to sell and the issuer cannot meet arbitrageurs' redemption requests through reserve asset fire-sales, it is rational for other stablecoin holders to sell, leading to a self-fulfilling panic run. However, runs are less likely when arbitrage is inefficient because investors' sales decrease secondary-market prices more, discouraging other investors from selling. We thus highlight a new tradeoff faced by stablecoin issuers: by choosing how concentrated arbitrage is, issuers trade off the benefits of arbitrage to price stability with the costs from increased run risk.

Third, our framework has novel implications for the effects of policy interventions proposed in various jurisdictions.[2] Our results imply that although improved price stability and reduced run risk are both desirable traits of stablecoins, they are fundamentally different, and one may come at the expense of the other. For example, proposals that remove issuers' ability to determine arbitrage concentration using the two-layered market structure would directly improve price stability, but may have the unintended consequence of amplifying run risk if not coupled with other measures that reduce asset illiquidity.

Our empirical findings are based on a novel dataset of fiat-backed stablecoins. Each stablecoin creation or redemption involves a stablecoin transaction between an issuer and an arbitrageur on a public blockchain. Thus, to analyze the market structure of the arbitrage sector, we collect transaction-level data on each stablecoin creation and redemption event for the six largest fiat-backed stablecoins: Tether (USDT), Circle USD Coin (USDC), Binance USD (BUSD), Paxos (USDP), TrueUSD (TUSD), and Gemini dollar (GUSD) from the Ethereum, Avalanche, and Tron blockchains. To capture trading activity by investors and arbitrageurs on secondary markets, we also extract trading prices in secondary markets from the main crypto exchanges. Further, we obtain the composition of reserve assets for USDT and USDC, which reported these breakdowns at various points in 2021 and 2022.

We document several stylized facts about stablecoin arbitrage. First, we find that arbitrage is generally concentrated, though the degree of concentration varies significantly across stablecoins. USDT only has six arbitrageurs redeeming stablecoins during an average month, and the largest arbitrageur accounts for 66% of the total redemption activity. In contrast, arbitrage at USDC is more competitive, with 521 redeeming arbitrageurs in an average month. Further, stablecoin trading prices in secondary markets frequently deviate from $1. We note that these price deviations are

[2] Proposed regulatory frameworks for stablecoins in the EU, UK, and US all contain provisions on how stablecoin creations and redemptions should be handled; see European Parliament and Council of the European Union (2023), Financial Conduct Authority (2023), Stablecoin Transparency and Accountability for a Better Ledger Economy (STABLE), and Guiding and Establishing National Innovation for U.S. Stablecoins (GENIUS).

not analogous to MMFs' "breaking the buck" nor are they an indicator of runs. Rather, stablecoins trade below (above) $1 when selling (buying) pressure in secondary markets is not fully absorbed by arbitrage trade, consistent with Lyons and Viswanath-Natraj (2021).

We find that stablecoins with fewer arbitrageurs have larger average price deviations in secondary markets. For example, the median discount at USDT is 11 bps, while the median discount at USDC is less than 1 bps. USDT's average discount is also much larger at 54 bps compared to USDC's at 1 bps. This finding is consistent with the limits to arbitrage literature showing that imperfect arbitrage hurts price efficiency (e.g., Shleifer and Vishny, 1997; Gromb and Vayanos, 2002). However, it also leaves open the question of how stablecoin issuers choose the arbitrage concentration they allow. After all, if approving more arbitrageurs improves price stability in secondary markets, why don't all stablecoin issuers allow for free entry and perfectly efficient arbitrage? At the same time, how is the choice of arbitrage concentration related to the liquidity of reserve assets, given that USDT also has more illiquid assets as part of their reserve assets than USDC?

We develop a tractable model of the two-layered market structure of stablecoins, yielding analytical solutions on both stablecoin price and run risk. In the baseline model, stablecoin investors decide whether to hold stablecoins to capture an exogenous long-term holding benefit. Investors can also prematurely sell their stablecoins in the secondary market to cash out, but only arbitrageurs are allowed to create or redeem stablecoins with the issuer for a fixed $1. Arbitrageurs have trading costs, so that there must be a wedge between secondary market prices and redemption values for arbitrageurs to act.

The issuer meets arbitrageur redemptions at $1 by prematurely liquidating illiquid reserve assets at a discount, which is ultimately why the risk of panic runs by stablecoin investors remains. The conventional view may imply that, like ETFs, stablecoins are not runnable because of exchange trading, where the trading price in secondary markets falls as more investors sell, creating a natural strategic substitutability. In the case of stablecoins, arbitrageurs are promised a fixed in-cash redemption price by the issuer. The issuer's costly sales of illiquid assets to meet arbitrageurs' redemptions at $1 imply that stablecoin investors may end up with less valuable stablecoins in the future. Consequently, stablecoins' fixed primary market price reintroduces strategic complementarity among secondary market investors.

Importantly, what is unique about the two-layered stablecoin market structure is that investors' propensity to run on the stablecoin is influenced by the efficiency of the arbitrage sector. Our core finding is that increasing arbitrage efficiency increases run risk. This is because more efficient arbitrage lowers the price impact for investors who sell in the secondary market. A more favorable selling price implies lower strategic *substitutability* and incentivizes panic selling.

Nevertheless, constraining arbitrage is not without costs. To understand how the stablecoin issuer optimally decides on arbitrage efficiency, we extend our baseline model by adding noise traders, whose trades induce more variance in stablecoin prices when arbitrage is less efficient. Stablecoin issuers thus face a tradeoff: arbitrage efficiency is beneficial for price stability, but also increases run risks. This tradeoff implies that arbitrage efficiency can be a double-edged sword and that some limits to arbitrage may be optimal.

There has been heightened attention on the optimal regulation of stablecoins across several jurisdictions. In the US, for example, the Stablecoin Transparency and Accountability for a Better Ledger Economy (STABLE) and the Guiding and Establishing National Innovation for U.S. Stablecoins (GENIUS) were proposed in the House of Representatives and the Senate, respectively. Regulators and market participants would like stablecoins to both have low run risks and stable prices. Our framework highlights that these two desirable goals are distinct from each other and driven by different economic forces. In particular, we highlight the tradeoff between price stability and financial stability and show that some policies may attain one goal at the expense of the other. In this section, we apply our model predictions to shed light on what the proposed regulation of stablecoins' redemptions, reserves, and interest payments imply for price stability and financial stability.

Our results provide important insights on optimal stablecoin regulation. These insights are especially relevant because different jurisdictions have recently proposed different sets of regulations that do not always agree with each other. First, some proposals require stablecoin issuers to provide unconstrained direct redemptions to all investors, essentially eliminating the current two-layered market structure. Such policies would indeed benefit price stability through more competitive arbitrage. However, to the extent that stablecoins are allowed to continue doing liquidity transformation, our model suggests that more efficient arbitrage would also amplify the risk of panic runs. In contrast, imposing redemption fees on arbitrageurs would reduce run risk through constraining arbitrage, but at the expense of price stability. Second, many policy proposals impose restrictions on how safe and liquid stablecoin issuers' reserve assets can be. Our results suggest that reserve asset policies should be coordinated with policies governing redemptions since these factors jointly determine the dual outcomes of price stability and run risk. Third, we show that allowing for dividend payments further improves stablecoin price stability and may lower run risks. Taken together, we highlight that while price stability and low run risk are both desirable features of stablecoins, they are fundamentally distinct and respond differently to policy interventions.

Finally, we calibrate our global games model for the two largest fiat-backed stablecoins, USDT and USDC, to estimate of their run risk and the extent to which policy interventions would influence this risk. We measure the overall illiquidity of reserve portfolios using collateral haircuts. We then estimate the probability at which the reserve asset payoff does not materialize using CDS spreads. We further proxy for the long-term benefit of holding the stablecoin using the return to lending out the stablecoin. Finally, we choose the slope of investors' stablecoin demand and the cost of price variance to match the slope of investors' demand and the slope of arbitrageurs' demand in the data. Our model estimates imply an economically significant risk of runs at both USDT and USDC. USDT's fragility stems from its higher liquidity transformation, while USDC is vulnerable due to less concentrated arbitrage. Using our calibrated model, we also evaluate the extent to which allowing issuers to make dividend payments could decrease run risks and increase price stability.

Our analysis of stablecoins belongs broadly to the growing literature on digital currencies and their regulations (e.g., Brunnermeier, James and Landau, 2019; Duffie, 2019).[3] While Bitcoin and most other cryptocurrencies exhibit volatile and correlated prices (Hu, Parlour and Rajan, 2019), the defining feature of stablecoins is their relative price stability at $1 and, thus, their potential to become a means of payment. Our overall contribution is to point out how increasing price stability can increase run risks, which has important implications for regulating stablecoins.

We show that the tradeoff between stablecoins' price stability and run risk is determined by arbitrage efficiency: constraints to arbitrage worsen price stability but also reduce financial fragility. The fact that inefficient arbitrage decreases price efficiency has been shown in seminal papers by Shleifer and Vishny (1997) and Gromb and Vayanos (2002). Closely related to us, Gromb and Vayanos (2018) analyze how constrained arbitrageurs exploit price discrepancies across segmented markets, Bryzgalova, Pavlova and Sikorskaya (2023) show that arbitrageurs face entry costs and choose to specialize in some markets, resulting in concentrated arbitrage analogous to our findings, and Davila, Graves and Parlatore (2024) provide a comprehensive study on the social value of closing an arbitrage opportunity. We focus on a two-layer market structure in which we highlight a novel implication of limits to arbitrage: more efficient arbitrage improves price stability, but may unintentionally increase run risks.

Our analysis of stablecoin runs builds on a large literature on panic runs and liquidity transformation (e.g., Diamond and Dybvig, 1983; Allen and Gale, 1998; Morris and Shin, 1998; Bernardo and Welch, 2004; Goldstein and Pauzner, 2005). It has also been shown that MMFs are subject to panic runs because their shares are redeemed by investors at a fixed price (Kacperczyk and Schn-

[3] Also see Harvey, Ramachandran and Santoro (2021); John, Kogan and Saleh (2022) and Makarov and Schoar (2022) for detailed surveys of the market structures of various cryptocurrencies and decentralized finance.

abl, 2013; Sunderam, 2015; Parlatore, 2016; Schmidt, Timmermann and Wermers, 2016), while closed-end funds and ETFs are typically viewed as less runnable because their shares are tradable at market prices (Jacklin, 1987; Allen and Gale, 2004a; Farhi, Golosov and Tsyvinski, 2009; Koont, Ma, Pastor and Zeng, 2021). Our contribution is to incorporate the two-layer market structure of stablecoins and a realistic arbitrage mechanism into a run model, while keeping it tractable and yielding closed-form solutions on both the stablecoin price and run risk.[4] In doing so, our model captures the unique combination of ETFs and MMFs in the design of stablecoins and sheds light on modeling similar financial intermediaries in future work.

In the context of stablecoins, most closely related to us is Lyons and Viswanath-Natraj (2021), who are the first to show that USDT's creation and redemption activity respond to secondary market price deviations. Gorton, Klee, Ross, Ross, and Vardoulakis (2023) show that stablecoins' use in leveraged trading of other crypto-assets helps maintain their price stability. Uhlig (2022) and Liu, Makarov and Schoar (2023) provide comprehensive analysis of runs on algorithmic stablecoins during the Terra-Luna crash in 2022, while Adams and Ibert (2022) analyze earlier algorithmic stablecoins.[5] Aldasoro, Ahmed, and Duley (2023) analyze the effect of disclosure about reserve asset quality on stablecoin runs, while Bertsch (2023) models the effect of stablecoin adoption on fragility. Our contribution to the stablecoins literature is to show how arbitrage concentration within the two-layered market structure shapes price instability and run risks. This aspect of stablecoin design is relatively understudied in the academic literature but has important implications for stablecoin regulation.

More generally, several other papers have explored risks associated with stablecoins other than panic runs. Eichengreen, Nguyen, and Viswanath-Natraj (2023) construct measures of stablecoin devaluation risk using spot and futures prices. Li and Mayer (2021) develop a dynamic model to characterize the endogenous transition between stable and unstable price regimes, focusing on the feedback between debasement and the collapse of demand for stablecoins as money. d'Avernas, Maurin, and Vandeweyer (2022) provide a framework to analyze how price stability can be maintained depending on the issuer's commitment to stablecoin supply. Routledge and Zetlin-Jones (2022) consider the design of exchange rate policies in maintaining price stability. Barthelemy, Gardin and Nguyen (2021), Liao and Caramichael (2022), Flannery (2023), and Kim (2022) analyze the potential impact of fiat-backed stablecoin activities on the real economy, while Baughman and Flemming (2023) argue that the competitive pressure of stablecoins on USD assets is limited.

[4]Technically, our results share similar features with a few other recent developments of run models where there is strategic substitutability from both sides of the action space (e.g., Allen, Carletti, Goldstein and Leonello, 2018; He, Krishnamurthy, and Milbradt, 2019; Kashyap, Tsomocos, and Vardoulakis, 2023) and the solution is guaranteed by the single-crossing property (Athey, 2001).

[5]Also, taking a historical perspective, Frost, Shin, Wierts (2020), Gorton and Zhang (2021), and Gorton, Ross and Ross (2022) compare stablecoins to deposits issued by the banking sector pre-deposit-insurance.

Anadu et al. (2023) show that investors shift from riskier to safer stablecoins during periods of stress similar to the flight-to-safety behavior of MMF investors. Kozhan and Viswanath-Natraj (2021) analyze collateral risk at DAI, which is a stablecoin overcollateralized by risky crypto assets, while Griffin and Shams (2020) show that USDT was used to facilitate bitcoin speculation and likely subject to under-collateralization. Complementary to these papers, we focus on stablecoins as financial intermediaries engaged in liquidity transformation, the arbitrage efficiency between primary and secondary markets, and the resulting relationship between price stability and run risk.

The rest of the paper proceeds as follows. Section 1 describes institutional details of the stablecoin market and Section 2 explains the data we use. Section 3 documents several empirical facts that motivate our model in Section 4. Section 5 shows the policy implications of our results. Section 6 explains the model calibration and Section 7 concludes.

1 Institutional Details

Stablecoins are blockchain assets whose value is claimed to be stable at $1. Blockchain assets can be self-custodial: a user can use crypto wallet software, such as Metamask, to hold, send, and receive stablecoins directly. These stablecoins are not stored with any trusted intermediary; rather, a "private key" – a long numeric code, generally kept only on the user's hardware device – is used to prove to the blockchain network that the user owns her stablecoins and to direct the network to take actions such as transferring stablecoins to other wallets. Others have no access to individuals' private keys so they cannot take funds from individuals' wallets.

Relative to other blockchain assets like bitcoins, the defining feature of stablecoins is (relative) price stability. The largest stablecoin issuers attempt to achieve price stability by promising to back each stablecoin token by at least $1 in off-blockchain US dollar assets. These fiat-backed stablecoins have experienced a rapid expansion over the last few years. Within two years, the total asset size of the six largest fiat-biased stablecoins has grown from $5.6 billion at the beginning of 2020 to exceed $130 billion at the beginning of 2022 (Figure 3). The largest two stablecoins are Tether (USDT) and Circle USD Coin (USDC), which made up more than 50% of the total market size at $76.4 billion in January 2022. Binance USD (BUSD), Paxos (USDP), TrueUSD (TUSD), and Gemini dollar (GUSD) are significantly smaller in size.

The remainder of this section provides an overview of the uses of stablecoins and the stablecoin market structure.

1.1 Uses of Stablecoins

Stablecoins are a reasonably low-cost way to transact and hold US-dollar assets. If a sender in country A sends funds to a receiver in country B, she can purchase stablecoins on a crypto exchange using fiat currency in country A, withdraw these stablecoins to her crypto wallet, and send them to the wallet of the receiver in country B. The receiver can then deposit these funds to a crypto exchange in his country, sell the stablecoins for fiat, and then withdraw the fiat currency. Sending stablecoins from one crypto wallet to another is relatively fast and low-cost.[6] As of January 2023, sending stablecoins on the Ethereum blockchain finalizes in under a minute and costs around $1 USD per transaction, independent of the amount of stablecoins sent. Stablecoins can also be used as a store of value; they can be held in crypto wallets indefinitely at no cost.

As a result, while stablecoins are costlier to use than well-functioning banking services in developed countries, they are competitive when traditional financial infrastructure functions poorly. For example, stablecoins are being used in settings where transactions must cross national borders, capital controls and financial repression are prevalent, inflation is high, or trust in financial intermediaries is low.[7]

Stablecoins are also used to transact with other blockchain smart contracts. For example, market participants can use stablecoin tokens to purchase other blockchain tokens, such as ETH, MKR, or UNI, using an automated market maker protocol such as Uniswap. Market participants can also lend stablecoin tokens on lending and borrowing protocols, such as Aave and Maker, allowing them to receive positive interest rates and also use these assets as collateral to borrow other assets. In a way, stablecoins provide a safe store of value and a medium of exchange for the blockchain ecosystem.

1.2 Market Structure

Stablecoin tokens are created ("minted") or redeemed ("burned") in the primary market with US dollar cash, as shown on the left-hand side of Figure 4. To create a stablecoin token, an arbitrageur

[6] The first and third steps in this process may incur fees and delays from converting fiat to and from crypto using local crypto exchanges, which may vary across exchanges and countries.

[7] Humanitarian organizations have used stablecoins to make cross-border remittance payments, circumventing banking fees and regulatory frictions. See Fortune.com. Some firms in Africa have begun using stablecoins for international payments to suppliers in Asia. See Rest Of World. In settings with high inflation, such as Lebanon and Argentina, individuals have begun storing value and transacting using stablecoins. See Rest Of World for a discussion of the case of Africa, CNBC and Rest Of World for the case of Lebanon, and Coindesk, EconTalk, and Memo for the case of Argentina. Some merchants in these areas have begun accepting stablecoins as a form of payment. For example, the Unicorn Coffee House in Beirut, Lebanon accepts USDT as a form of payment.

sends \$1 to the issuer, and the issuer then sends a stablecoin token into the market participant's crypto wallet. Analogously, to redeem a stablecoin token, for each stablecoin token that the market participant sends to the issuer's crypto wallet, the issuer sends \$1, for example, through a bank transfer, into the market participant's bank account. The primary market for stablecoins resembles a money market fund in the traditional financial system. Please see Appendix A for further details.

Most market participants cannot become arbitrageurs to redeem and create stablecoin tokens and stablecoin issuers differ in their requirements and costs for market participants to access primary markets. According to market participants, USDC allows general businesses to register as arbitrageurs, while USDT requires a lengthy due diligence process and imposes restrictions on where arbitrageurs can be domiciled. Further, USDT imposes a minimum transaction size of \$100,000 and charges the greater of 0.1% and \$1000 per redemption.

The majority of market participants trade existing stablecoins for fiat currencies in secondary markets. Crypto exchanges allow investors to make US dollar deposits and then trade US dollars for stablecoins with other market participants. The price of stablecoin tokens in the secondary market is thus driven by demand and supply. When there is a surge in stablecoin sales, the secondary market price drops but does not induce direct liquidations of reserve assets. In this way, the buying and selling of stablecoins on secondary markets resemble the trading of ETF shares on competitive exchanges.

However, selling pressure in the secondary market can spill over to affect the primary market through arbitrageurs. When secondary market prices drop below \$1, arbitrageurs can profit from purchasing stablecoin tokens in secondary markets, and redeeming them one-for-one for \$1 with the stablecoin issuer in primary markets. Through this arbitrage, the \$1 redemption value of stablecoins in primary markets pulls the trading price of stablecoins towards \$1 in secondary markets. At the same time, this arbitrage process also implies that investor selling pressure in secondary markets eventually triggers sales of reserve assets when stablecoin issuers liquidate reserves to meet arbitrageurs' redemption in cash. These fire sales can be especially costly if illiquid reserve assets are sold at a discount.

2 Data

In this section, we explain our data sources.

Primary market data. The core dataset used in our analysis is data on each stablecoin creation and redemption event for the six largest fiat-backed stablecoins: USDT, USDC, BUSD, USDP, TUSD, and GUSD, on the Ethereum, Avalanche, and Tron blockchains. We obtain this data from each blockchain based on "chain explorer" websites, which process transaction-level blockchain data into a usable format. We use Etherscan for Ethereum, Snowtrace for Avalanche, and Tronscan for Tron. Using our data extraction process, we see, for each stablecoin creation and redemption event, the precise timestamp of the event, the amount of the stablecoin redeemed or created, and the wallet address of the entity involved in stablecoin creation or redemptions. We note that the blockchain only records wallet addresses and the same institution can have multiple wallet addresses. In our data collection process, we combine wallets whose Etherscan labels clearly indicate that they belong to the same institution. However, this process may not be exhaustive, so the degree of arbitrage concentration we find should be viewed as a lower bound to true arbitrage concentration. Appendix B.1 presents further details for the primary markets of stablecoins and the construction of our data. Our baseline analysis uses data from the Ethereum blockchain. We present results for the Tron and Avalanche blockchains in the appendix.

Secondary market data. For each stablecoin, we extract hourly closing prices for direct USD to stablecoin trades from several large exchanges, including Binance, Bitfinex, Bitstamp, Bittrex, Gemini, Kraken, Coinbase, Alterdice, Bequant, and Cexio. We provide further details on why we only use direct USD to stablecoin trades in Appendix B.2. In our main analysis, we calculate daily prices for each stablecoin as the trading-volume-weighted average of hourly closing prices across these exchanges. Differences in stablecoin prices across the main exchanges are generally negligible. Hence the price series are not substantially affected by the weights we put on different exchanges. We winsorize secondary market prices at the 1% level.

Reserves. We use the breakdowns of reserve assets that USDT and USDC self-report at various points in 2021 and 2022 as part of their balance sheets posted online. The other four stablecoins have not released breakdowns of their reserve asset composition but state the broad categories of their reserves. We note that reserve assets are not recorded on the blockchain so we cannot independently verify the reported information. Griffin and Shams (2020), for example, have pointed out that USDT sometimes issues tokens that are insufficiently backed by reserve assets, implying the potential for additional risk or even fraud. We think of the reported reserve asset information as the most optimistic estimate of the actual reserve assets that stablecoins hold. Thus, our estimates of run risk should be interpreted as a best-case scenario or, equivalently, a lower bound.

3 Facts

In this section, we present a set of new facts about stablecoins that informs our model and calibration.

3.1 Secondary Market Prices

Fact 1. *The trading price of stablecoins in the secondary market commonly deviates from $1.*

Figure 5 shows the price at which different stablecoins trade on the secondary market over time. We observe that the secondary market price rarely stays fixed at $1. Instead, stablecoins trade at a discount 27.2% to 41.6% of the time and trade at a premium 57.3% to 72.8% of the time for our sample of stablecoins (see Table 1a). The extent of these price deviations varies by stablecoin. While the average discount at USDT is 54bps, the average discount at USDC is only 1bps. The average discounts of BUSD, TUSD, and USDP are also below that of USDT, while that of GUSD is the highest. The median discounts are generally smaller in magnitude than the average discounts, but the variation in the cross-section remains similar. For example, the median discount at USDT is 11bps, while that at USDC is less than 1bps. The magnitudes also decrease when we consider a common sample period starting from January 2020, when all the six stablecoins were traded, but the variation across coins remains with USDT having a larger average discount than USDC (see Table 1b). The average and median premia also show significant variation in the cross-section.

Trading of stablecoins at a discount has been commonly associated with "breaking the buck" as in the case of money market funds and even as evidence for panic runs.[8] We note that these are misconceptions. Stablecoins' "stable value" of $1 refers to the amount primary market participants receive when they redeem stablecoins with the issuer. "Breaking the buck" thus corresponds to primary market participants not receiving a full $1. The secondary market price is the trading price of stablecoins on exchanges. It is analogous to the share price of an ETF. Just like ETF prices can deviate from the NAV of the underlying portfolio, stablecoin prices can deviate from $1 due to selling pressure in secondary markets and are not a direct indicator of "breaking the buck" or panic runs.

[8]For example, see https://www.nytimes.com/2022/06/17/technology/tether-stablecoin-cryptocurrency.html and https://www.cnbc.com/2022/05/17/tether-usdt-redemptions-fuel-fears-about-stablecoins-backing.html

3.2 Primary Market Concentration

Fact 2. *The redemption and creation of stablecoins in the primary market is performed by a small set of arbitrageurs, whose concentration varies by stablecoin.*

Table 2 shows the characteristics of monthly primary market redemption and creation activity on the Ethereum blockchain for different stablecoins. We observe that in an average month, USDT only has six arbitrageurs engaged in redemptions, whereas USDC has 521. The concentration of arbitrageurs' market shares is generally high but still varies by stablecoin. The largest arbitrageur at USDT performs 66% of all redemption activity, while the largest arbitrageur at USDC performs 45%. In comparison, most other stablecoins lie between USDT and USDC in terms of the number of arbitrageurs and arbitrageur concentration.[9] In terms of transaction volumes, notice that in an average month, the volume of redemptions at USDT is $577 million, while that at USDC is $2976 million. In comparison, the total volume of outstanding tokens at USDT was 1.5 to 2 times that of USDC. Thus, the larger number and lower concentration of arbitrageurs at USDC are correlated with a higher volume of redemptions relative to the total asset size as well. There is a larger volume of creations and relatively more arbitrageurs engaged in creations but the trends across stablecoins and the arbitrage concentration remain similar. In Appendix Tables A.3 and A.4, we repeat the analysis for the Tron and Avalanche blockchains and obtain similar variations in arbitrageur concentration across stablecoins.

3.3 Secondary Market Price and Primary Market Concentration

Fact 3. *Stablecoins with a more concentrated set of arbitrageurs experience more pronounced price deviations in the secondary market.*

We proceed to analyze the relationship between price deviations and arbitrageur concentration. For a given stablecoin, we calculate monthly secondary market price deviations by averaging over the absolute values of daily price deviations from one in a given month, which includes both deviations above and below one. We then average over months to obtain that stablecoin's average price deviation. Similarly, we count the number of unique arbitrageurs that engage in redemptions and/or creations, calculate the market share of the five largest arbitrageurs in each month, and the average over time for each coin. We plot the results in Figure 6a. A clear negative trend emerges: stablecoins with fewer arbitrageurs, like USDT, have higher average price deviations from $1 in their secondary market prices, relative to stablecoins with more arbitrageurs, like USDC. Another

[9] One exception is GUSD, which has the most concentrated arbitrage market for redemptions.

way to capture arbitrageur concentration is through the market share of the largest arbitrageurs. In Figure 6b, we repeat the analysis with the market share of the top 5 arbitrageurs. The relationship is positive. Stablecoins whose top 5 arbitrageurs consistently perform a larger share of total redemptions and creations have higher average price deviations than other stablecoins with lower arbitrageur concentration. In other words, it seems that higher arbitrage competition is associated with reduced price dislocations in secondary markets.

One question arising from Facts 1 to 3 is why some stablecoins choose to have a more concentrated arbitrageur sector. If arbitrageur competition can indeed stabilize secondary market prices, all stablecoins should be incentivized to open up arbitrageur access and encourage the entry of new arbitrageurs. In our model, we show that a counteracting force is the presence of panic runs by investors, which are more likely with a more competitive arbitrageur sector and are fundamentally linked to stablecoin liquidity transformation, which we elaborate on next.

3.4 Liquidity Transformation

Fact 4. *Stablecoins engage in varying degrees of liquidity transformation by investing in illiquid assets.*

Stablecoin issuers hold USD-denominated assets with varying degrees of illiquidity as reserves. Table 3 shows the composition of reserve assets for USDT and USDC on reporting dates. Overall, reserve assets of both USDT and USDC are not fully liquid, with those of USDT being more illiquid.

A significant portion of reserve assets is in deposits and money market instruments, including commercial paper and certificates of deposits. In September 2021, for example, these two asset classes took up 56.2% of reserve assets at USDT, and USDCs' reserve assets were 100% in deposits. Except for deposits in checking accounts, money market instruments and other types of deposits are not fully liquid and experience a discount when demanded or sold before their maturity date. Notice also that deposits are not default-free because their quantities exceed the 250K deposit insurance limit. In fact, USDC was found to be the biggest depositor in Silicon Valley Bank right before its collapse in March 2023.

The remaining reserve assets are comprised of Treasuries and more illiquid assets, including municipal and agency securities, foreign securities, corporate bonds, corporate loans, and other securities. USDT holds a significant portion of reserves in the form of Treasuries, which amounted to 28.1% in September 2021. While Treasuries are relatively liquid and safe, the extent of their liquidity varies by type and over time. For example, Treasury markets were strained in March

2020 following the fire sale by mutual funds and hedge funds. USDT also holds a sizable amount of more illiquid assets.

The other four stablecoins report that their assets are limited to deposits, Treasuries, and money market instruments but, unfortunately, do not provide more detailed breakdowns. That is why our model estimation will focus on USDT and USDC for which reserve asset breakdowns are available.

4 Theoretical Framework

In this section, we develop a theoretical framework to reconcile the facts presented in Section 3 and to analyze the potential for runs on fiat-backed stablecoins. In Section 4.1, we first present a baseline model to formulate the notion of stablecoin runs, pointing to the unique two-layer market structure that differentiates stablecoin runs from bank runs while keeping minimal deviation from an otherwise standard bank run model to highlight our contribution. We show how stablecoin run risk is different from bank run risk and how it is linked to stablecoins' arbitrage concentration. We then extend the model in Section 4.2 to analyze how arbitrage concentration simultaneously affects stablecoins' price stability and solve for the issuers' choice of arbitrage concentration given the tradeoff between run risk and price stability.

4.1 Stablecoin Runs and the Centralization of Arbitrage

The baseline economy builds on Diamond and Dybvig (1983) and has three dates, $t = 1, 2, 3$, with no time discounting. We provide a timeline to illustrate the economy in Figure 1:

Figure 1: Timeline: Baseline Model

ALT TEXT: Graphical representation of the timeline for the baseline model: period 1 in which investors are endowed with stablecoins, period 2 in which investors sell and arbitrageurs bid for stablecoins, and period 3 in which long-term convenience and returns realize.

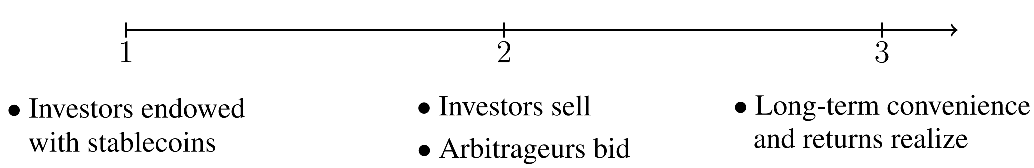

There are two groups of risk-neutral players: 1) a competitive group of stablecoin investors indexed by i, and 2) a sector of n stablecoin arbitrageurs. There are two types of assets: 1) the dollar, which is riskless, liquid, and serves as the numeraire, and 2) an illiquid and potentially productive reserve asset. The investors jointly hold the stablecoin that is initially backed by the reserve asset at $t = 1$. The initial value of the reserve asset is normalized to one dollar. We will formalize investors' participation decisions and the issuer's profit maximization in Section 4.2. Until then, we take n as exogenous and normalize the population of participating investors to one.

At $t = 2$ and $t = 3$, investors decide whether to liquidate their stablecoins early at $t = 2$, potentially triggering runs, or to hold them until maturity at $t = 3$ to capture long-term benefits, which we specify below. Unlike bank depositors, stablecoin investors cannot redeem their holdings directly from the issuer. Instead, they liquidate at $t = 2$ by selling stablecoins in the secondary market, following the mechanism of Jacklin (1987) and Farhi, Golosov and Tsyvinski (2009), to arbitrageurs, who then redeem them for cash from the issuer. As in these models, investors sell stablecoins by independently submitting market orders. This assumption is consistent with empirical evidence showing that retail investors tend to use market orders more, particularly when spreads are narrow (Kelley and Tetlock, 2013), as in the case of stablecoins.[10] We denote by λ the fraction of investors selling their stablecoins at the market price p_2.

There are n symmetric arbitrageurs indexed by j. In any given period, arbitrageurs bid competitively to buy or sell stablecoins from investors and liquidity traders, incur a per-period trading cost if winning the auction, and then create or redeem the stablecoin at the fixed price of $1 if the issuer is solvent. In the baseline model, we assume that arbitrageurs cannot hold net inventory, so they must on net redeem as much on primary markets at $t = 2$ as they purchase in secondary markets. In Appendix C.1, we show that this assumption is consistent with the empirical observation that the vast majority of arbitrageurs hold very small amounts of stablecoins. We further consider in Appendix C.2 an extension model in which arbitrageurs are allowed to hold some stablecoins as reserves where we show that our main results continue to hold. Arbitrageurs face quadratic trading costs: arbitrageur j incurs a cost $\frac{z_j^2}{2\chi}$ for arbitraging z_j units of the stablecoin from secondary to primary markets, where χ can be thought of as capturing arbitrageurs' balance sheet capacity: when χ is higher, trading costs are lower.

The issuer, in turn, meets arbitrageur redemptions in cash by liquidating the illiquid reserve asset at $t = 2$. This involves a liquidation cost of $\phi \in (0, 1]$, i.e., liquidating one unit of the asset yields $1 - \phi$ dollars. Economically, ϕ captures the level of liquidity transformation and the various costs incurred when transacting illiquid assets (see Duffie, 2010, for a review). Note that the issuer

[10] Technically, assuming each investor submits a market order for a single unit of the stablecoin also allows us to model investor behavior as a binary-action global game. It would be much more difficult to solve the model if we assumed investors submitted demand schedules, since this maps less cleanly into the standard literature on runs and global games.

is solvent if and only if $\lambda < 1 - \phi$. When $\lambda \geq 1 - \phi$, the issuer defaults, and arbitrageurs receive the liquidation value of $(1 - \phi)/\lambda$ per stablecoin redeemed.

In deciding whether to liquidate their stablecoins early, investors receive private information at $t = 2$ about the fundamentals of the economy at $t = 3$. Following the global games literature, each investor i obtains a private signal $\theta_i = \theta + \varepsilon_i$ at $t = 2$, where the noise term ε_i are independently and uniformly distributed over $[-\varepsilon, \varepsilon]$. As usual in the literature (e.g., as in Goldstein and Pauzner (2005)), we focus on arbitrarily small noise in the sense that $\varepsilon \to 0$, but the model results also hold beyond the limit case.[11]

Fundamentals θ reflect the level of aggregate risk and determine the stablecoin's long-term value at $t = 3$. With probability $1 - \pi(\theta)$, the economy enters a bad state: the reserve asset fails, and investors do not receive any nominal return nor any long-term benefits from holding the stablecoin backed by assets of no value. With probability $\pi(\theta)$, the economy enters a good state: the reserve asset yields a positive value of $R(\phi) \geq 1$ dollar, which accrues to the issuer. The stablecoin continues to operate, and the remaining $1 - \lambda$ investors consume a long-term benefit $\eta > 0$ per stablecoin and the initial value of 1 per unit of the remaining reserve asset. This long-term benefit η can be motivated by investors' return from lending out the stablecoin, as documented in Gorton, Klee, Ross, Ross, and Vardoulakis (2023).

Following backward induction, we first consider p_2, the price an investor receives when liquidating the stablecoin early. In the main text, we derive the inverse demand function that arises from arbitrageurs bidding competitively. We show in Appendix D that, if arbitrageurs bid strategically, our model resembles the setting of Klemperer and Meyer (1989), which exhibits multiple equilibria. However, our core economic insight that arbitrage capacity influences stablecoin prices remains unchanged under an equilibrium selection rule that chooses the equilibrium bid curve closest to the standard linear solution.

In the main text, we derive the inverse demand function by setting prices equal to marginal costs. We provide a more formal derivation in Appendix E, where we microfound competitive bidding by assuming a measure n of arbitrageurs submit bid curves in a rational expectations equilibrium.

[11]Note that we do not impose any restrictions on the distributions of π, θ, or the increasing function $\pi(\theta)$, which allows us to map the model to any empirical distribution of fundamentals. Also note that the standard assumption in the global games literature that investors obtain a private signal about fundamentals is relatively plausible for the stablecoin market because of its opacity: essentially, no stablecoin issuers disclose asset-level information about their reserves, and investors and arbitrageurs infer stablecoins' value using their private information.

Suppose $\lambda \leq 1-\phi$, implying that the issuer is solvent, an arbitrageur who purchases z_j units in the secondary market at price p_2 and redeems them in the primary market earns

$$\underbrace{z_j 1}_{Redemption\ Value} - \underbrace{z_j p_2}_{Secondary\ Mkt\ Price} - \underbrace{\frac{z_j^2}{2\chi}}_{Trading\ Cost}. \tag{4.1}$$

In any symmetric equilibrium, each arbitrageur absorbs $z_j^* = \frac{\lambda}{n}$ of the total inventory. If issuers take prices as given, in order for z_j^* to be the optimal quantity absorbed, the price must equal the marginal value of absorbing additional stablecoin quantity at z_j^*. That is, differentiating (4.1) with respect to z_j, setting to 0, and solving for p_2, we must have

$$p_2\left(z_j^*\right) = 1 - \frac{z_j^*}{\chi}.$$

In words, the price is just the marginal redemption value 1 less the marginal trading cost $\frac{z_j^*}{\chi}$. Substituting that $z_j^* = \frac{\lambda}{n}$ in symmetric equilibrium, we have

$$p_2(\lambda) = 1 - \frac{\lambda}{n\chi}, \quad \forall \lambda \leq 1-\phi. \tag{4.2}$$

In the insolvent case, $\lambda > 1-\phi$, each unit of the stablecoin can only be redeemed at its liquidation value $\frac{1-\phi}{\lambda}$. Since investors submit market orders, the redemption value does not depend on arbitrageurs' choice of quantity. Arbitrageur profits are thus modified to

$$\underbrace{z_j\left(\frac{1-\phi}{\lambda}\right)}_{Redemption\ Value} - \underbrace{z_j p_2}_{Secondary\ Mkt\ Price} - \underbrace{\frac{z_j^2}{2\chi}}_{Trading\ Cost}. \tag{4.3}$$

Again, in the symmetric equilibrium, the price must equal arbitrageurs' marginal value of absorbing additional quantity. Differentiating (4.3) with respect to z_j, setting to 0, and substituting $z_j^* = \frac{\lambda}{n}$, we have

$$p_2(\lambda) = \frac{1-\phi}{\lambda} - \frac{\lambda}{n\chi}, \quad \forall \lambda > 1-\phi.$$

The only difference from (4.2) is that the redemption value is modified to $\frac{1-\phi}{\lambda}$. We summarize these results in the following lemma.

Lemma 1. *The stablecoin's secondary-market price at $t=2$ is given by*

$$p_2(\lambda) = \begin{cases} 1 - K\lambda & \lambda \leq 1-\phi\,, \\ \dfrac{1-\phi}{\lambda} - K\lambda & \lambda > 1-\phi\,, \end{cases} \tag{4.4}$$

where

$$K \equiv \frac{1}{n\chi}. \tag{4.5}$$

Lemma 1 shows that for any total redemption quantity $\lambda > 0$, p_2 decreases with K and increases with both χ and n. Intuitively, secondary market prices are less affected by investors' sales when arbitrageurs are better capitalized (χ is higher) and more numerous (n is higher). We refer to K as arbitrage concentration, which plays a central role in the analysis as it measures the slope of demand in the secondary market when the issuer remains solvent. Arbitrageurs' bids create a downward-sloping demand curve for the stablecoin, and when n or χ increases, the slope becomes steeper, reducing the price impact of stablecoin sales.

Note also that p_2 is strictly decreasing in λ everywhere: the more investors sell, the lower the price is. This is important because it produces strategic *substitutability* in investors' sale decisions: when many other investors are selling, a given investor anticipates receiving less from selling, thus discouraging her from selling. This force stands in contrast to the strategic complementarity in classic bank run models (e.g., Diamond and Dybvig, 1983), in which depositors get a fixed deposit value from withdrawing.

We then consider v_3, the value an investor may get at $t = 3$ if λ other investors choose to liquidate early. It is given by

$$v_3(\lambda) = \begin{cases} \pi(\theta)\left(\dfrac{1-\phi-\lambda}{(1-\phi)(1-\lambda)} + \eta\right) & \lambda \leq 1-\phi\,, \\ 0 & \lambda > 1-\phi\,. \end{cases} \tag{4.6}$$

To see why this is the case, notice that the issuer needs to liquidate

$$l(\lambda) = \begin{cases} \dfrac{\lambda}{1-\phi} & \lambda \leq 1-\phi\,, \\ 1 & \lambda > 1-\phi\,. \end{cases} \tag{4.7}$$

units of the reserve asset to meet arbitrageur redemptions at $t = 2$, and only $1 - l(\lambda)$ units remain at $t = 3$, whose value will be shared by the remaining $1-\lambda$ late investors. Combining this financial value and the long-term benefit of the stablecoin thus yields (4.6).

An important observation from (4.7) is that more investors selling (i.e., larger λ) and a higher level of liquidity transformation (i.e., larger ϕ) result in more costly liquidations of the reserve asset (i.e., larger $l(\lambda)$). Fundamentally, this arises because the stablecoin issuer, if solvent, has to meet stablecoin redemptions at a fixed cash value of one dollar. As we show shortly below, this force generates strategic complementarities which eventually dominate the strategic substitutability from price impact, thus leading to potential runs.

Investors' incentives to sell stablecoins depend on the sign of the difference between (4.6), the expected utility from holding until date-3, and (4.4), the utility from selling the stablecoin early and receiving the secondary market price. Formally, as a function of the fraction λ of other investors who sell, this difference is:

$$h(\lambda) = v_3(\lambda) - p_2(\lambda) = \begin{cases} \pi(\theta)\left(\dfrac{1-\phi-\lambda}{(1-\phi)(1-\lambda)} + \eta\right) - 1 + K\lambda & \lambda \leq 1-\phi\,, \\ -\dfrac{1-\phi}{\lambda} + K\lambda & \lambda > 1-\phi\,. \end{cases} \tag{4.8}$$

It is easy to see that $h(0) \geq 0$ when $\pi(\theta)$ is sufficiently large while $h(1) < 0$, implying that the model has multiple equilibria when θ is sufficiently large and if θ is common knowledge.

Figure 7 plots the payoff gain function $h(\lambda)$, which first increases, then decreases, and then increases again as λ rises. In the first region, where $h(\lambda)$ increases, strategic substitutability arises from price impact in the secondary market, as captured by (4.4). When few investors sell, falling prices discourage further sales, and investors prefer to wait until $t = 3$ as long as the issuer remains solvent.

As more investors sell, the issuer faces rising liquidation costs, decreasing the value of holding on to stablecoins for remaining investors. Here, $h(\lambda)$ decreases in the second region, reflecting strategic complementarity or first-mover advantage: anticipating high redemptions, investors rush to redeem before reserves are depleted. Note that these two regions differentiate our model from standard global-games bank run models like Goldstein and Pauzner (2005), where illiquidity always leads to complementarity even for small λ. In contrast, due to secondary-market arbitrage, our model shows substitutability at low λ: early sales depress prices, deterring further redemptions. Within $\lambda \leq 1 - \phi$, price impact drives strategic substitutability for small λ, while fire-sale effects drive strategic complementarity for larger λ.

Finally, when $\lambda > 1 - \phi$, all reserves are exhausted, late redemptions yield nothing, and a crowding-out effect emerges: more early redemptions decrease the payoff to early redeemers because each redeemer gets less. This is a mechanical feature shared with Goldstein and Pauzner (2005) and other standard bank run models that does not affect the core economics of runs.[12]

The global games framework allows us to solve for a unique equilibrium for any value of primitives. Formally, we have the following result:

Proposition 1. *There exists a unique threshold equilibrium in which investors sell the stablecoins if they obtain a signal below threshold θ^* and do not sell otherwise.*

Proposition 1 implies that the model with investors' private and noisy signals has a unique threshold equilibrium. An investor's liquidation decision is uniquely determined by her signal:

[12]This matters only for technical conditions of equilibrium selection. As shown in the proof of Proposition 1 and discussed in Goldstein and Pauzner (2005), our model satisfies the single-crossing property, sufficient for threshold equilibria. Goldstein and Pauzner (2005) satisfy the stronger one-sided strategic complementarity, ensuring all equilibria are threshold equilibria. Our model lacks this stronger property, so we explicitly focus on threshold equilibria.

she sells the stablecoin at $t = 2$ if and only if her signal is below a certain threshold. In other words, she is indifferent between selling and holding when her signal is at the threshold. Given the existence of the unique run threshold, we can show that her indifference condition implies the following Laplacian equation:

$$\int_0^{1-\phi} (1 - K\lambda)\, d\lambda + \int_{1-\phi}^{1} \left(\frac{1-\phi}{\lambda} - K\lambda \right) d\lambda = \int_0^{1-\phi} \pi(\theta^*) \left(\frac{1-\phi-\lambda}{(1-\phi)(1-\lambda)} + \eta \right) d\lambda\,. \quad (4.9)$$

Solving the Laplacian equation gives an analytical solution of the run threshold and presents intuitive comparative statics about stablecoin run risk:

Proposition 2. *The run threshold is given by*

$$\pi(\theta^*) = \frac{(1-\phi)(2-2\phi-2(1-\phi)\ln(1-\phi)-K)}{2\left((1+\eta(1-\phi))(1-\phi)+\phi\ln\phi\right)}\,. \quad (4.10)$$

which satisfies the following properties:

i). The run threshold, that is, run risk, is decreasing in K (that is, increasing in n and increasing in χ).

ii). The run threshold, that is, run risk, is increasing in ϕ if and only if

$$g(\phi) \equiv \frac{2(\phi-1)\left(\phi - \ln\phi + \ln(1-\phi)((1+\phi)\ln\phi + 2 - 2\phi) - 1\right)}{1-\phi+\ln\phi} > K\,, \quad (4.11)$$

where $g(\phi)$ is continuous and strictly decreasing in ϕ, and satisfies $\lim_{\phi\to 0} g(\phi) > 0$.

A core theoretical result is part i) of Proposition 2, which shows that more efficient arbitrage, i.e., a smaller value of K, exacerbates run risk. This surprising result is an implication of the way that stablecoin primary and secondary markets are connected. When arbitrage is more efficient, stablecoin sales have a lower price impact, as illustrated by Lemma 1. Thus, investors get higher payoffs from selling early, whereas their payoffs from holding to maturity are unchanged. Investors' incentives to sell early increase, exacerbating run risk. Conversely, when arbitrage is inefficient, sales have more price impact, and investors are discouraged from selling early.[13] Figure 8 illustrates how investors' payoff gain from waiting increases as the secondary market becomes less efficient.

[13]This result assumes rational investors who understand the effect of arbitrage concentration on price stability and run risk. If investors were not fully rational and interpreted price deviations as poor fundamentals, concentrated arbitrage may amplify herding through larger price discounts. However, herding would imply that stablecoins holding more illiquid assets choose less concentrated arbitrage to reduce their price deviations and run risk, which is inconsistent with the empirical findings in Section 3.

Part ii) of Proposition 2 shows that a higher level of stablecoin liquidity transformation leads to a higher run risk when $g(\phi) > K$. This condition is satisfied when ϕ is sufficiently small for a given K. Intuitively, when the stablecoin holds more illiquid reserve assets, the first-mover advantage among investors increases because an investor who chooses not to sell would have to involuntarily bear a higher liquidation cost induced by selling investors. However, when the reserve asset is too illiquid, run risk could be dampened. The intuition can be understood from equation (4.8): investors enjoy the first-mover advantage only when $\lambda \leq 1 - \phi$, that is, only when $h(\lambda)$ takes the value in the first line of (4.8). Otherwise, too high a ϕ shrinks the region in which the first-mover advantage can be realized. Thus, further increasing the level of liquidity transformation when $g(\phi) < K$ reduces run risk.

In Appendix C.2, we further show that if arbitrageurs were allowed to hold some stablecoins as reserves, the run threshold would be strictly higher than the one in Proposition 2. Intuitively, when arbitrageurs use their balance sheet to hold stablecoins, arbitrage becomes more efficient and investors receive higher payoffs from selling early. As a result, the incentive to redeem early increases, exacerbating run risk.[14]

In addition, the analytical solution given in Proposition 2 allows us to calibrate the model to the data to quantify run risk in Section 6. To this end, we translate the run threshold into an ex-ante run probability with the distribution of fundamentals $F(\theta)$. Formally,

Definition 1. *The ex-ante run probability of a stablecoin is given by*

$$\rho = \int_{\pi(\theta)<\pi(\theta^*)} dF(\theta)\,, \tag{4.12}$$

where $\pi(\theta^*)$ is given by (4.10) and $F(\theta)$ is the prior distribution of the fundamentals.

Before proceeding, we make three comments about the notion of stablecoin runs in our framework and highlight our contribution to the literature. We purposefully follow and keep minimal deviation from Diamond and Dybvig (1983) to highlight the unique nature of stablecoin runs. Our contribution is to incorporate the two-layer market structure of stablecoins and a realistic arbitrage mechanism into an otherwise standard run model. Technically, the modeled economy environment does not feature universal or even one-sided strategic complementarity as that in Diamond and Dybvig (1983) and Goldstein and Pauzner (2005). In this aspect, our results share similar features with a few other recent developments of run models where there is strategic substitutability from both sides of the action space (e.g., Allen, Carletti, Goldstein and Leonello, 2018; He, Krishnamurthy, and Milbradt, 2019; Kashyap, Tsomocos, and Vardoulakis, 2023) and the solution is

[14] This result holds except in the extreme case where arbitrageurs are able to absorb all stablecoin holdings at a fixed price and completely eliminate runs.

guaranteed by the single-crossing property (Athey, 2001). Compared to these papers, we highlight the ultimate run incentives from secondary-market investors who face strategic substitutability in selling decisions and who do not directly interact with the stablecoin issuer at all, as opposed to any run incentives among private-market participants, which has been the focus of these existing papers.

Our notion of coordination motives and runs also differs from a few other prominent notions of coordination in the literature. One possible alternative modeling choice is to follow the idea in the new monetarist framework of Kiyotaki and Wright (1989) and Rocheteau and Wright (2005), which shows that an agent adopts a good as a medium of exchange only if other agents adopt and thus accept the same good in transactions. In other words, the value of a medium of exchange becomes higher when more investors adopt it. This approach more explicitly highlights the payment role and network-good feature of stablecoins without capturing liquidity transformation. However, that approach applies to any general form of money or tokens that is not necessarily backed by dollar reserves. Several recent papers that consider general forms of cryptocurrencies and tokens follow this view (e.g., Schilling and Uhlig, 2019; Cong, Li, and Wang, 2021; Li and Mayer, 2021; Baughman and Flemming, 2023; Bertsch, 2023; Sockin, and Xiong, 2023a; Sockin and Xiong, 2023b). Given our focus on reserve-backed stablecoins as a financial intermediary and the financial stability implications for real dollar asset markets, we view Diamond and Dybvig (1983) as the preferred building block for our model. At the same time, we still capture the payment role of stablecoins by modeling its price convenience and linking it to stablecoin price fluctuations in the extended model below.

Finally, despite our focus on secondary-market investors, our notion of coordination and runs also differs from the idea of market runs in Bernardo and Welch (2004). There, if an illiquid secondary asset market features a downward-sloping demand curve, investors fearing future liquidity shocks will be incentivized to front-run each other, fire selling the asset earlier to get a higher price. However, Bernardo and Welch (2004) do not feature an intermediary or liquidity transformation, which is the focus of our paper.

4.2 Price Stability and Optimal Stablecoin Design

Having analyzed the run risk of stablecoins and its relationship with arbitrage concentration, we extend the baseline model by incorporating the stablecoin's price stability at $t = 1$ and the issuer's optimal design choices in a pre-trading game at $t = 0$. The extension serves two purposes. First, it allows us to further formulate the tradeoff between stablecoins' price stability and run risks.

Second, it introduces realistic model ingredients that facilitate the evaluation of run risks and policy proposals in Section 6.

The extended economy has four dates, $t = 0, 1, 2, 3$, with no time discounting. We provide a timeline of the extended model in Figure 2. On top of the baseline model spanning $t = 1, 2, 3$, we introduce one additional group of risk-neutral liquidity traders, who live and trade at $t = 1$, and formulate the issuer's optimal design problem at $t = 0$. At $t = 0$, the issuer designs the primary market. Specifically, the issuer chooses n at $t = 0$, that is, how concentrated its primary market is, to maximize its expected profit. Investors also make participation decisions at $t = 0$. If an investor chooses to participate in the stablecoin market, she incurs a cost of c_i, which follows a distribution function $G(c)$, and receives one stablecoin. An investor participates if her expected utility from participation, which we characterize below, exceeds c_i.

Figure 2: Timeline: Extended Model

ALT TEXT: Graphical representation of the timeline for the extended model: period 0 in which the issuer designs the arbitrageur market structure and investors choose stablecoin holdings, period 1 in which noise traders trade and and arbitrageurs bid for stablecoins, period 2 in which investors sell and arbitrageurs bid for stablecoins, and period 3 in which long-term convenience and returns realize.

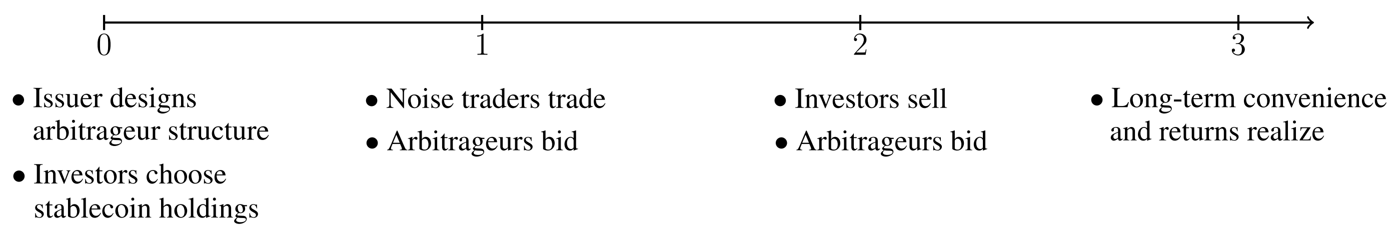

At $t = 1$, liquidity traders trade stablecoins, creating variance in stablecoin prices. With equal probability, liquidity traders either buy a fraction δ of the total stablecoin market cap and then resell them at the end of $t = 1$ or sell short a fraction δ of the stablecoin market cap and then rebuy at the end of the period. Letting ω denote liquidity trader order flow, ω is equal to δ or $-\delta$ with equal probability. Intuitively, we can think of liquidity traders as using stablecoins for remittances: as described in Section 1.1, the remittance process involves buying stablecoins with fiat, sending stablecoins, then selling for fiat.[15] Liquidity traders cannot directly trade with the issuer; instead, they exchange fiat for stablecoins by trading with arbitrageurs in secondary markets. Also, at $t = 1$,

[15] Technically, the specification that liquidity trader order flow perfectly reverts is convenient because, as we will show, it implies that liquidity trading ω affects stablecoin price but does not directly generate fire sales by the issuer. This allows us to focus on the tradeoff between price and financial stability in stablecoin design while ruling out the uninteresting case of liquidity trading itself leading to runs.

n arbitrageurs trade stablecoins in secondary and primary markets to profit from price deviations just like what they do at $t = 2$ in the baseline model.[16] We assume that liquidity-trader-induced price fluctuations lower stablecoin investors' price convenience. Following Gorton and Pennacchi (1990), we let investors enjoy a short-term price convenience of $-\alpha Var(p_1)$ per stablecoin at $t = 1$, where $\alpha > 0$; this captures in reduced-form that stablecoins are less valuable to users as a means of payment when their prices are more volatile.

Consider p_1 and its variance, which determines the price convenience that investors enjoy at $t = 1$. Specifically, we apply the market clearing condition at $t = 1$:

Lemma 2. *The stablecoin's secondary-market price at $t = 1$ is given by*

$$p_1 = \begin{cases} 1 - \delta K & \omega = \delta\,, \\ 1 + \delta K & \omega = -\delta\,, \end{cases} \tag{4.13}$$

where K is given in (4.5). The stablecoin's price convenience at $t = 1$ is thus given by $-\alpha\delta^2 K^2$, which is decreasing in K, that is, increasing in n and χ.

Lemma 2 shows that the stablecoin's price convenience is decreasing in K. This is intuitive because as arbitrage becomes less efficient, the secondary market becomes less elastic and liquidity trading induces larger fluctuations in the secondary market price p_1. Investors thus enjoy a lower convenience, reminiscent of the idea of information sensitivity in Gorton and Pennacchi (1990). The positive relationship between price deviations and arbitrage capacity is consistent with the empirical evidence in Figure 6.

Taken together, Proposition 2 and Lemma 2 point to the tradeoff between price and financial stability of the stablecoin. To formulate this tradeoff, we now consider the stablecoin issuer's design decision at $t = 0$. It involves one key choice variable that determines the elasticity of the stablecoin secondary market: the number of arbitrageurs n who are allowed to perform primary-market redemptions and creations.[17] We suppose that the stablecoin issuer chooses n to maximize its expected revenues at $t = 0$, which depends on how many investors participate at $t = 0$. We also assume that liquidity trading and arbitrageurs' balance sheet capacity are proportional to the population of investors. The issuer's objective function is thus given by

$$\max_n E[\Pi] = \underbrace{G(E[W])}_{\substack{\text{population of} \\ \text{participating investors}}} \underbrace{\int_{\pi(\theta) \geq \pi(\theta^*)} \pi(\theta)(R(\phi) - 1) dF(\theta)}_{\substack{\text{expected issuer revenue per} \\ \text{participating investor}}}\,, \tag{4.14}$$

[16] We note that the separation between $t = 1$ and $t = 2$ is not crucial for the model; it simplifies the model by ruling out the uninteresting case that liquidity trading itself may lead to fire sales or render the stablecoin issuer default. Considering that would complicate the model without new economic insights.

[17] Arbitrage capacity χ also affects arbitrage efficiency, but stablecoin issuers are unlikely to have control over the balance sheet costs and budget constraints of arbitrageurs, which is why we let the issuer choose n for a given χ.

where each investor's expected utility of participation is

$$E[W] = \underbrace{-\alpha\delta^2 K^2}_{\text{short-term convenience}} + \underbrace{\int_{\pi(\theta)<\pi(\theta^*)} (1-\phi-K)\, dF(\theta)}_{\text{short-term payoff if runs}} + \underbrace{\int_{\pi(\theta)\geq\pi(\theta^*)} \pi(\theta)(1+\eta)\, dF(\theta)}_{\text{long-term payoff if no runs}}, \tag{4.15}$$

in which $\pi(\theta^*)$ is given by (4.10) in Proposition 2.

The stablecoin issuer's objective function (4.14) captures its revenue base. Absent a panic run, the issuer obtains the expected net long-term return of the remaining reserve asset. At the same time, a larger population of participating investors allows the issuer to scale up its investment in reserve assets. Investors' participation is, in turn, driven by their expected utility $E[W]$, which is comprised of three components, as shown in (4.15). The first term denotes investors' expected price convenience loss due to stablecoin price fluctuations. The second term denotes their expected payoff when a panic run happens, while the third term corresponds to their expected payoff without a run.

Solving the stablecoin issuer's problem (4.14), we have the following result about the stablecoin issuer's optimal choice of arbitrageur concentration:

Proposition 3. *Suppose G is linear, and ϕ is small enough that (4.11) holds. Then, the issuer's optimal choice of K is increasing in ϕ: if reserves are more illiquid, then the issuer optimally chooses a more concentrated arbitrage sector.*

Proposition 3 stems from the tradeoff between price stability and financial stability. The stablecoin issuer chooses arbitrage concentration K to trade off its benefits from decreasing run risks with its costs from decreasing price stability. When asset illiquidity ϕ is higher, run risk is increased. The issuer should then be more willing to sacrifice price stability to limit run risk, leading to a higher optimal value of K.[18]

In the main text, we think of ϕ as exogenous. In practice, differences in ϕ across issuers can be driven by differences in the sets of investments available. An issuer like USDT may not have the same access to safe and liquid US dollar assets as US-based USDC. USDT may therefore hold higher-ϕ assets and choose more concentrated arbitrage via Proposition 3. In Appendix F, we

[18] While economically intuitive, we state Proposition 3 under two technical conditions to rule out mechanically driven or economically uninteresting cases. First, we require that ϕ is sufficiently small to satisfy condition (4.11), excluding cases where liquidation becomes extremely costly, mechanically making early sales unprofitable, reducing run incentives, and complicating the role of arbitrageur concentration. This condition is also commonly seen in the banking literature since Allen and Gale (1998). Second, we restrict the form of consumer demand G. Since the issuer's first-order condition has a "markup" interpretation, changes in ϕ can otherwise affect demand elasticity and lead the issuer to either increase or decrease arbitrageur concentration. Assuming linearity in G keeps demand elasticity constant and rules out such mechanical effects.

analyze an extension of the model where issuers jointly optimize asset liquidity ϕ and arbitrage concentration K. We formally show that issuers who have access to assets with a higher illiquidity premium optimally choose higher values of ϕ despite the potentially higher run risk. Consequently, they will also tend to choose lower values of n to mitigate run risk. We leave a full-fledged endogenous determination of ϕ for future research because information on stablecoins' accessible asset pools, business models, and regulatory risks is currently not readily available.

5 Policy Implications

There has been heightened attention on the optimal regulation of stablecoins across several jurisdictions. Regulators and market participants would like stablecoins to have both low run risks and stable prices. Our framework highlights that these two desirable goals are distinct from each other and driven by different economic forces. In particular, we highlight the tradeoff between price stability and financial stability and show that some policies may attain one goal at the expense of the other. In this section, we apply our model predictions to shed light on what the proposed regulation of stablecoins' redemptions, reserves, and interest payments imply for price stability and financial stability.

5.1 Redemptions and Primary Market Access

Regulators in many jurisdictions are considering how to regulate stablecoin issuers' redemption policies. In the EU, the proposed policy requires unconstrained access for all stablecoin holders to redeem their stablecoins in cash without delay.[19] In the UK, the draft regulation provides a more nuanced view by incorporating the possibility of restricting redemptions to reduce run risk, stating that "by temporarily suspending the 'next UK business day redemption requirement' we hope that regulated stablecoin issuers are enabled to better deal with the exceptional circumstance and reduce the chance of failure of the regulated stablecoin issuer."[20]

What are the implications of these proposed measures? In our model, allowing for unconstrained redemptions by stablecoin holders would effectively improve arbitrage efficiency, i.e., a lower K. This improvement in arbitrage efficiency affects stablecoin run risk and price stability differently, as the following Corollary shows.

[19] Specifically, Article 49 of the European Parliament and Council of the European Union (2023) states that "Upon request by a holder of an e-money token, the issuer of that e-money token shall redeem it, at any time and at par value, by paying in funds, other than electronic money, the monetary value of the e-money token held to the holder of the e-money token."

[20] See Sec. 3.48 of Financial Conduct Authority (2023).

Corollary 1. *When arbitrage efficiency increases (that is, K decreases) holding other primitives fixed, price variance $Var[p_1^*]$ decreases, but run risk increases (that is, the run threshold $\pi(\theta^*)$ increases).*

The first part of Corollary 1, which follows from Lemma 2, suggests that unconstrained redemptions would unambiguously benefit price stability. Intuitively, since more market participants would have the option to quickly redeem stablecoins for dollars, stablecoin prices in secondary markets would remain more stable at $1 than with a more constrained set of arbitrageurs.[21]

However, the second part of Corollary 1, which follows from Proposition 2, suggests that, unrestricted redemptions would have the perverse effect of amplifying run risk, all else equal. Without constraints to the set of arbitrageurs, arbitrage becomes more efficient and increases investors' tendency to panic sell because the price impact is lower. We thereby highlight higher run risk as an important side-effect of requiring unconstrained access to redemptions.

Beyond policy proposals allowing all customers to redeem stablecoins, our conclusions generalize to several other possible policies affecting the structure of stablecoin creations and redemptions. In Appendix G, we formally show in Proposition 6 that imposing fees or gates on redemptions would have a similar effect to allowing more concentrated arbitrage: secondary-market price stability would decrease, but run risk would also decrease, consistent with Corollary 1.

Taken together, our model highlights that regulations for stablecoin redemptions inherently trade off the effects of price stability and run risk. Policies that require unconstrained arbitrage and redemptions improve price stability but at the expense of worsening run risk; policies restricting redemptions through fees and gates lower run risk at the cost of price stability. While price stability is observable on a daily basis and run risk only materializes in tail events, both are essential considerations for the regulation of stablecoins.

5.2 Reserves and Liquidity Transformation

The proposed regulations on stablecoin reserves also generally attempt to limit the illiquidity and credit risk of reserve assets, that is, the level of liquidity transformation by stablecoins. For example, draft regulation in the EU requires 30% of reserves to be deposited in credit institutions, and the rest to be invested in a set of assets classified as secure, low-risk, highly liquid, with low

[21] There could still be requirements that redemptions are only available for market participants who pass anti-money-laundering and other checks, as the Financial Conduct Authority (2023) proposal suggests. In this case, investors who pass AML/CTF checks can be thought of as arbitrageurs for customers who do not pass these checks. However, this is a broad enough set of customers so that price stability is most likely still improved in such a scenario.

market, credit, and concentration risk. In the UK, the proposal is to restrict reserves to direct holdings of government debt and short-term cash, while in the US, the STABLE Act and the GENIUS Act limit stablecoin reserve assets to currency, deposits, repo, and other high-quality liquid assets (HQLAs).

Under our theoretical framework, more liquid reserve assets reduce run risk for a given arbitrage concentration. Recall from part ii) of Proposition 2 that when the stablecoin holds less illiquid reserve assets, the first-mover advantage among investors decreases. In this sense, limiting the illiquidity of stablecoin reserve assets would be beneficial for financial stability.[22]

Nevertheless, it is important to jointly design policies that govern reserve asset liquidity and policies that govern arbitrage concentration. This is because when reserve assets are more liquid, stablecoin issuers are incentivized to reduce the constraints on arbitrage as pointed out in Proposition 3. This increased arbitrage efficiency benefits price stability but may curtail the reduction in run risk from more liquid reserve asset holdings. Therefore, our results imply that the interaction between reserve assets and redemption policies is essential to consider in regulating stablecoins.

5.3 Dividend Issuance

For fiat-backed stablecoins, returns from reserve assets are fully accrued to the issuer, and no dividends are issued to investors holding stablecoins. In the US, neither the STABLE Act nor the GENIUS Act designate stablecoins as securities, which effectively prohibits dividend payments. Further, neither the STABLE Act nor the GENIUS Act mention the issuance of interest payments. The EU has proposed to explicitly ban interest payments to stablecoin holders. The UK similarly proposes to ban income or interest payments, but notes that "this may be perceived as unfair to consumers".[23]

To provide a sense of the asset returns that stablecoin investors are currently deprived of, we show the percentage return and revenue of USDT and USDC in Table 4. At the end of the sample period in March 2022, for example, USDT earned almost 2% return on its assets, amounting to $1.6 trillion. None of this income was passed on to USDT investors. These forgone returns increase with the monetary policy rate that increases the nominal return on portfolio assets. In other words,

[22]In practice, we note that limiting reserve asset liquidity may reduce but not fully eliminate run risk. This is because very few assets are 100% liquid and safe. For example, deposits at FDIC-insured banks may not be fully insured if the account balance exceeds the deposit insurance limit. Deposits may also not be immediately demandable without cost. For US Treasuries, secondary markets tend to be generally liquid, but strains may arise in stressed market conditions like in March 2020.

[23]See article 50 of European Parliament and Council of the European Union (2023) for the EU and article 3.12 of Financial Conduct Authority (2023) for the UK.

the absence of distributions to stablecoin investors has particularly large repercussions in high interest rate environments.

In terms of the implications for price stability and run risk, our framework shows that issuing dividends to investors can improve price stability and may, under some conditions, reduce run risk. Formally, we model dividends by assuming that, in the good state of the world, the stablecoin issuer is forced to pay τ per unit stablecoin to its long-term investors at $t = 3$. Each investor's value at $t = 3$ thus becomes:

$$v_3(\lambda;\tau) = \begin{cases} \pi(\theta)\left(\dfrac{1-\phi-\lambda}{(1-\phi)(1-\lambda)}(1+\tau)+\eta\right) & \lambda \leq 1-\phi\,, \\ 0 & \lambda > 1-\phi\,, \end{cases} \tag{5.1}$$

Compared to (4.6), there is an additional τ term that can be collected when the stablecoin is solvent. Accordingly, the stablecoin issuer's objective function becomes:

$$\max_{n}\; E[\Pi] = G(E[W])\int_{\pi(\theta)\geq\pi(\theta^*)} \pi(\theta)\left(R(\phi)-1-\tau\right) dF(\theta)\,, \tag{5.2}$$

which nests (4.14) as a special case. We have the following result:

Proposition 4. *Suppose ϕ is small in the sense that (4.11) holds, when the stablecoin issuer distributes a positive dividend τ to its long-term investors, the run risk of the stablecoin decreases before the issuer re-optimizes n. In equilibrium, the stablecoin issuer optimally designs a less concentrated arbitrageur sector, that is, $n^*_\tau > n^*$, resulting in higher price stability of the stablecoin while the change in run risk is ambiguous.*

Intuitively, Proposition 4 states that, if we hold arbitrage efficiency fixed, dividends lower run risk since they increase investors' incentives to hold the stablecoin until the final period. In response to lowered run risks, issuers have an incentive to increase n, increasing price stability. Regarding run risk, it is in theory possible for the issuer re-optimization effect to dominate and increase run risk relative to the no-dividends case. This is because the issuer's expected revenue per participating investor decreases after distributing dividends, which reduces the issuer's incentive to prevent runs.

We acknowledge that there may be other effects of dividend issuance. For example, dividend issuance may intensify price competition among stablecoin issuers, which could encourage entry and improve allocative efficiency. Stablecoins that distribute dividends would also likely be classified as a security under US securities law and exposed to regulatory risk.[24] We leave the analysis of these and other forces to future work.

6 Model Calibration: Quantifying Run Risks and Policies

In this section, we use our framework to assess the run risk of major stablecoins and the extent to which various policy interventions would influence their run risk and price stability, following the quantitative exercises in recent papers such as Davila and Goldstein (2023). We calibrate our model to empirical moments for the two largest fiat-backed stablecoins, USDT and USDC, for which reserve asset breakdowns are available.

6.1 Empirical Moments ϕ, $p(\theta)$, and η

We first estimate asset illiquidity ϕ, the distribution of $p(\theta)$, and the long-term benefit η directly from the data.

Asset Illiquidity ϕ. We proxy asset illiquidity ϕ, that is, the per unit cash value lost when selling an illiquid asset at short notice, with repo haircuts. This approach follows Bai, Krishnamurthy and Weymuller (2018), which relies on the idea that one minus the haircut in a repo transaction directly captures how much cash can be obtained against an illiquid asset at short notice.[25] Economically, repo haircuts serve as a proxy for liquidation costs because illiquid assets such as bonds and loans are traded over-the-counter through dealers, with prices shaped by dealers' balance sheet and funding conditions. In particular, since repos are commonly used to finance dealers' bond positions, repo haircuts are closely linked to the costs of liquidating these assets. (Macchiavelli and Zhou, 2022).

To measure the overall illiquidity of USDT and USDC's reserve portfolios, we calculate the average discounts of their reserve assets weighted by their portfolio weights. One challenge is that we do not know the exact liquidity of their deposits, which include demandable deposits, time deposits, and certificates of deposits (CDs). In the baseline estimate, we assume that one-quarter

[24] US regulators have deemed many programs that take funds from users and return funds with dividend or interest payments, to be securities that fall under the SEC's jurisdiction. For example, the June 2023 SEC case against Coinbase argued that Coinbase's Staking Program is a security. The June 2023 SEC case against Binance argued that Binance's BNB Vault and Simple Earn programs, and the BAM Trading Staking program, constituted securities under US law. In our conversations with market participants, many believed that a stablecoin that offered to pay accrued interest on reserves as dividends would be classified as securities.

[25] The New York Fed publishes haircuts on different securities when pledged as collateral in repos at https://www.newyorkfed.org/data-and-statistics/data-visualization/tri-party-repo#interactive/margins.

of the deposits are fully liquid while the remainder is subject to the lowest money market discount. The results are shown in Table 5.

Distribution of $p(\theta)$. To estimate the distribution of $p(\theta)$, the signal of how likely the risky asset will pay nothing, we use historical CDS prices to evaluate the daily recovery value of each portfolio component and then take a weighted average to obtain the daily expected recovery value of the reserve portfolio. Finally, we multiply the recovery value by one plus the amount of over-collateralization because some stablecoins are backed by slightly more than \$1 per coin before asset liquidations. Using daily data from 2008 to 2022 from Markit, we then fit a beta distribution to match the mean and variance of daily expected recovery values. Appendix J contains further details of this procedure. The means and variances as well as the fitted beta distribution parameters are shown in Appendix Table A.5.

Long-term Benefit η. To proxy for investors' long-term benefit from holding and using the stablecoin, we follow Gorton, Klee, Ross, Ross, and Vardoulakis (2023) to use investors' return from lending out the stablecoin. Specifically, we focus on Aave, which is a smart contract lending platform that allows market participants to lend cryptoassets for interest, overcollateralized by other cryptoassets. This lending rate captures the compensation to the investor for not being able to use the stablecoin herself while it is on loan to another investor. Our data on lending rates is from aavescan.com. Table 5 shows the annual return from lending out USDT and USDC in each reporting period.

6.2 Estimating $\alpha\delta^2$ and $G(\cdot)$ using K and $\frac{\partial \log G(E[W])}{\partial \eta}$

The remaining model parameters are $\alpha\delta^2$, the cost of liquidity trading to investors, and $G(\cdot)$, investors' demand function for the stablecoin. We will estimate the product $\alpha\delta^2$ as a single parameter; our approach does not separately identify risk aversion α and the size of liquidity trading shocks δ. We parametrize $G(\cdot)$ as:

$$G(EW) = \max\left[1 - \gamma\left(1 - EW\right), 0\right].$$

That is, γ is simply the slope of investor demand: the issuer has a unit mass of consumers if she produces $EW = 1$, and loses γ customers for any gap between 1 and EW, until the point where she loses all consumers. We allow the demand slopes for USDC and USDT to differ, calling them γ_{Circle} and γ_{Tether}, respectively, accounting for their different investor bases.

We then estimate $\alpha\delta^2$ and $G(\cdot)$ through moment matching. For each choice of $\alpha\delta^2$, γ_{Circle}, γ_{Tether} and each coin-month combination in our data, we calculate the optimal value of K, by solving the issuer's optimization problem (4.14). At the optimal K, we then numerically compute the partial elasticity of investors' demand with respect to η in the model:

$$\frac{\partial \log G\left(E\left[W\right]\right)}{\partial \eta}.$$

For each choice of $\alpha\delta^2$ and γ, this procedure gives us a model-predicted value of K and $\frac{\partial \log G(E[W])}{\partial \eta}$ for each month. We then choose parameters to minimize the sum of squared distances between model-predicted log values of K and $\frac{\partial \log G(E[W])}{\partial \eta}$, averaged across months for each coin, and their counterparts in the data.

Recall from Lemma 1 that K is the slope of demand in the secondary market when the issuer remains solvent. Thus, to obtain K from the data, we regress daily price deviations against daily redemption or creation volume for each stablecoin:

$$Deviation_t = \beta Redemption/Creation_t + FE_y, \tag{6.1}$$

where $Deviation_t$ is one minus the lowest observed secondary market price on redemption days and the highest observed secondary market price minus one on creation days, $Redemption/Creation_t$ is the volume of redemptions or creations divided by the total outstanding volume of stablecoins on day t. We use the lowest and highest secondary market prices on each day to capture the extent of price dislocations that demand arbitrage rather than the price dislocations resulting from arbitrage. We normalize the volume of redemptions and creations by the total outstanding volume of stablecoins to consider the difference in market sizes across stablecoins. Finally, we include a year fixed effect to capture potential structural shifts in the arbitrageur sector for each stablecoin. For example, the number and constraints of arbitrageurs may evolve after some time with the growth of stablecoins.

From the results in Table 6, we observe that the estimated K for USDT is larger in absolute magnitude as for USDC, which is consistent with the higher arbitrageur concentration of USDT constraining redemption volume to be less sensitive to price dislocations. That is, a larger price dislocation is required to induce the same amount of redemptions for USDT than for USDC. Magnitude-wise, a 10 percentage point higher redemption/creation volume as a fraction of the total volume outstanding corresponds to a 2.1 cent larger price deviation USDT and a 1.6 cent larger price deviation at USDC. For the detailed estimation results with respect to K, please refer to Appendix Table A.6.

To obtain $\frac{\partial \log G(E[W])}{\partial \eta}$ from the data, we regress the monthly log change in the number of shares outstanding against the beginning-of-month long-term benefit, i.e., the lending rate. The results in Table 6 show that the demand for USDC is more responsive to a given change in the long-term benefit than the demand for USDT.

The parameter estimates are shown in the first two columns of Table 5. We estimate $\alpha\delta^2$, γ_{Tether}, and γ_{Circle} to be 7.06, 0.54, and 1.10. As is standard in structural models, both parameters contribute to variation in both moments; however, the intuition behind the identification of model parameters is as follows. When $\alpha\delta^2$ is high, the cost of price variance is high. Thus, issuers will tend to choose lower values of K, trading off slightly increased run probabilities for lower price variance and thus lower costs of liquidity trading. Hence, the level of K in the data, relative to fundamentals, contributes to identifying $\alpha\delta^2$. The parameter γ controls investors' elasticity of demand; when γ is higher, the stablecoin market size will increase more for any given increase in η.

The fit of our model to the targeted moments is shown in Table 6. The model-predicted arbitrageur demand slopes K are in the same range but slightly higher than those in the data.[26] Note that we can match the stylized fact that the optimal K is higher for USDT than USDC, with approximately the same magnitude as in the data. In terms of the second moment, we can match the elasticity of investors' demand for stablecoins fairly well, on average over time within coins. The mapping from moments to parameters is intuitive: we estimate investors' demand elasticity to be somewhat higher for USDC than USDT, which is why we find that γ is slightly higher for USDC.

6.3 Run Probability

Table 5 shows the implied run probabilities. Notice that the run risk of USDC remains substantial even without holding illiquid assets like corporate bonds and corporate loans as USDT. For example, the run probabilities for USDT and USDC were 2.495% and 2.134% in September 2021, respectively. This is because of USDC's concentrated exposure to bank deposits, which incur a higher default risk than Treasuries in the case of uninsured deposits and retain some illiquidity in the case of time deposits. Over time, there was a decline in run probabilities from 3.188% in May 2021 to 1.828% in October 2021 for USDC because of ϕ declining and the long-term benefit η trending up. For USDT, both illiquidity ϕ and the long-term benefit η display less variation over time, resulting in relatively stable run risk over the reporting period from 2.590% in June 2021 to 1.664% in March 2022.

Our estimates of stablecoin run probability complement the empirical literature of bank runs (e.g., Demirguc-Kunt and Detragiache, 2002; Calomiris and Mason, 2003; Iyer and Puri, 2012), and particularly, the findings in Egan, Hortacsu and Matvos (2017) and Albertazzi, Burlon, Jankauskas, and Pavanini (2022), who build dynamic structural models to estimate the run probability of com-

[26]Technically, this mismatch is because, under our estimates, K values in the data would imply overly high run probabilities for USDC, which could not be consistent with issuer optimization under any parameter settings.

mercial banks. Their focus is on the feedback loop between a bank's credit risk and uninsured depositor outflows. We estimate run probabilities derived from a global games model that captures the unique interaction between the primary and secondary markets of stablecoins. In this sense, our approach provides a complementary way to quantify the run risk of tradable assets that are also involved in liquidity transformation.

Our estimation suggests that regulators should monitor the market structure of the stablecoin arbitrage sector.[27] Regulators could track the number of arbitrageurs and concentration metrics such as top-1 or top-5 shares of arbitrage activity. Taking our model more seriously, specification (6.1) suggests measuring the coefficient from regressing price deviations on net redemption volumes. These quantities can be readily estimated without imposing extra reporting requirements on stablecoin issuers, because the wallet identifiers and volumes of all primary market transactions are recorded in real-time on public blockchains, and can be freely downloaded by regulators and academics alike. Regulators could use these metrics to evaluate the extent to which arbitrage concentration limits run risks.

6.4 The Effect of Dividend Issuance and Redemption Fees

Having estimated run risk, we show the extent to which different policies could mitigate this risk. We first evaluate how run risk would change with dividend issuance by stablecoins. Recall from Section 5.3 that issuing dividends to investors can improve price stability and may, under some conditions, reduce run risk. If arbitrage efficiency is held constant, dividends lower run risk since they increase investors' incentives to hold the stablecoin until the final period to get the dividend. However, this reduction in run incentive, in turn, allows the issuer to choose more arbitrageurs for maximum price stability, which counteracts the initial reduction in run risk. Which of these two forces dominates is thus an empirical question.

Using our calibrated model, we find that issuing dividends leads to a net reduction in run risk for USDT and USDC in our sample period. The results are shown in Figure 10 for the September 2021 reporting period.[28] Quantitatively, as dividend issuance increases from 0 to 4%, the run probabilities of USDC and USDT are lowered by 1.34% and 0.80%, respectively, as panel (c) shows. Consistent with our model predictions, issuers optimally choose a lower K to make arbitrage more efficient (panel (a)), and the cost of price variance $\alpha\delta^2K^2$ decreases relative $\tau = 0$ (panel (b)). Although a lower K contributes to higher run risk, the direct effect of dividends dominates in our sample period so run risk is reduced.

[27] Regulators already impose similar reporting requirements on ETFs with respect to their authorized participants through the N-CEN filings.

[28] Results for other reporting periods follow a similar trend and are shown in Appendix Figure A.6.

Finally, we estimate the effect of redemption fees on run risk for USDT and USDC. In Appendix G, we show that redemption fees make redemptions more costly and increase the constraint to arbitrage. Imposing redemption fees thus leads to the same tradeoff between price stability and run risk. Magnitude-wise, we evaluate how much run risk would change with different values of ν, holding the arbitrageur demand slope at the optimal values of K. We find that as redemption fees increase from 0 to 50bps, USDC and USDT run probabilities decrease by 2.01% and 2.38% for the September 2021 reporting period of USDT and USDC (Figure 9).[29] Overall, redemption fees would be quite effective at decreasing run risk.

7 Conclusion

In this paper, we analyzed the tradeoff between stablecoin run risk and price stability. At a high level, stablecoin runs arise from liquidity transformation. Stablecoin issuers hold illiquid assets while offering arbitrageurs the option to redeem stablecoins for a fixed \$1 in the primary market. This liquidity mismatch spills over from the primary market to trigger the possibility of runs among investors on the secondary market despite exchange trading.

Importantly, we show that stablecoin run risk is mediated by the market structure of the arbitrageur sector, which serves as a "firewall" between the secondary and primary markets. When the arbitrageur sector is more efficient, shocks in the secondary market are transmitted more effectively to the primary market. The price stability of stablecoins is thus improved, but the first-mover advantage for sellers is also higher, increasing run risk. If the arbitrageur sector is less efficient, shocks in secondary markets transmit less effectively. Price stability suffers, but run risk decreases, as the price impact of stablecoin trades in secondary markets discourages market participants from panic selling.

Our results have important policy implications. We show that although regulators and market participants would like stablecoins to both have low run risks and stable prices, these two desirable goals are distinct from each other and driven by different economic forces. In particular, allowing unconstrained redemptions increases arbitrage efficiency, which improves price stability but may come at the expense of higher run risk. It is also imperative to consider the dampening effect of dividend issuance on run risk in the discussion on whether stablecoins should be allowed to pass

[29]Note that we hold the arbitrageur demand slope K fixed rather than allowing issuers to re-optimize because it is not possible to quantify the effects of redemption fees on price stability within our estimated model. The effect of redemption fees depends on risk aversion α; our estimation identifies only the product $\alpha\delta^2$ of risk aversion and the size of liquidity trading shocks δ, so we cannot quantitatively evaluate how costly redemption fees are to price stability from consumers' perspective. Results for other reporting periods follow a similar trend and are shown in Appendix Figure A.7.

through income to investors in the form of dividends. Overall, the joint consideration of price stability and run risk in designing regulation is essential for a well-functioning and safe stablecoin sector going forward.

References

ADAMS, AUSTIN and MARKUS IBERT (2022). "Runs on Algorithmic Stablecoins: Evidence from Iron, Titan, and Steel." Working Paper.

ALBERTAZZI, BURLON, LORENZO BURLON, TOMAS JANKAUSKASA, and NICOLA PAVANINI (2022). "The Shadow Value of Unconventional Monetary Policy." Working Paper.

ALDASORO, IÑAKI, AHMED, RASHAD, and DULEY, CHANELLE (2023). "Par for the course: Public information and stablecoin runs." Available at SSRN.

ALLEN, FRANKLIN and DOUGLAS GALE (1998). "Optimal Financial Crises." *Journal of Finance*, 53: 1245-1284.

ALLEN, FRANKLIN and DOUGLAS GALE (2004). "Financial Intermediaries and Markets." *Econometrica*, 72: 1023-1061.

ALLEN, FRANKLIN and DOUGLAS GALE (2004). "Competition and financial stability." *Journal of money, credit and banking*, 453–480.

ANADU, KENECHUKWU, AZAR, PABLO, HUANG, CATHERINE, CIPRIANI, MARCO, EISENBACH, THOMAS M, LA SPADA, GABRIELE, LANDONI, MATTIA, MACCHIAVELLI, MARCO, MALFROY-CAMINE, ANTOINE, and WANG, J CHRISTINA (2023). "Runs and Flights to Safety: Are Stablecoins the New Money Market Funds?" FRB of Boston Working Paper.

ATHEY, SUSAN (2001). "Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information." *Econometrica*, 69(4): 861-889.

AUER, RAPHAEL, JON FROST, LEONARDO GAMBACORTA, CYRIL MONNET, TARA RICE, and HYUN SONG SHIN (2022). "Central Bank Digital Currencies: Motives, Economic Implications, and the Research Frontier." *Annual Review of Economics*, 14: 697-721.

ALLEN, FRANKLIN, ELENA CARLETTI, ITAY GOLDSTEIN and AGNESE LEONELLO (2018). "Government Guarantees and Financial Stability." *Journal of Economic Theory*, 177: 518-557.

BAI, JENNIE, ARVIND KRISHNAMURTHY and CHARLES-HENRI WEYMULLER (2018). "Measuring Liquidity Mismatch in the Banking Sector." *Journal of Finance*, 73: 51-93.

BARTHELEMY, JEAN, PAUL GARDIN and BENOÎT NGUYEN (2018). "Stablecoins and the real economy." Working Paper.

BAUGHMAN, GARTH and JEAN FLEMMING (2023). "Global Demand for Basket-Backed Stablecoins." *Review of Economic Dynamics*, 48: 244-264.

BERNARDO, ANTONIO and IVO WELCH (2004). "Liquidity and Financial Market Runs." *Quarterly Journal of Economics*, 119: 135-158.

BERTSCH, CHRISTOPH (2023). "Stablecoins: Adoption and Fragility." Working Paper.

Brunnermeier, Markus, Harold James, and Jean-Pierre Landau. 2019. "The Digitalization of Money." Working Paper.

Bryzgalova, Svetlana, Anna Pavlova, and Taisiya Sikorskaya. 2023. "Strategic Arbitrage in Segmented Markets." Working Paper.

CALOMIRIS, CHARLES W. and JOSEPH R. MASON (2003). "Fundamentals, Panics, and Bank Distress during the Depression." *The American Economic Review*, 93(5): 1615-1647.

CLARITY FOR PAYMENT STABLECOINS ACT OF 2023 (2023). Bill introduced in the U.S. Senate. 118th Congress, 2nd Session.

CONG, LIN WILLIAM, YE LI, and NENG WANG (2021). "Tokenomics: Dynamic Adoption and Valuation." *The Review of Financial Studies*, 34: 1105-1155.

DAVILA, EDUARDO and ITAY GOLDSTEIN (2023). "Optimal Deposit Insurance." *Journal of Political Economy*, forthcoming.

DAVILA, EDUARDO, DANIEL GRAVES and CECILIA PARLATORE (2024). "The Value of Arbitrage." *Journal of Political Economy*, forthcoming.

D'AVERNAS, ADRIEN, VINCENT MAURIN, and QUENTIN VANDEWEYER (2022). "Can Stablecoins be Stable?" Working Paper.

DEMIRGUC-KUNT, ASLI and ENRICA DETRAGIACHE (2002). "Does Deposit Insurance Increase Banking System Stability? An Empirical Investigation." *Journal of Monetary Economics*, 49: 1373-1406.

DIAMOND, DOUGLAS and PHILIP DYBVIG (1983). "Bank Runs, Deposit Insurance, and Liquidity." *Journal of Political Economy*, 91: 401-419.

DU, SONGZI and HAOXIANG ZHU (2017). "What is the Optimal Trading Frequency in Financial Markets?" *Review of Economic Studies*, 84: 1606-1651.

DUFFIE, DARRELL (2010). "Presidential Address: Asset Price Dynamics with Slow-Moving Capital." *Journal of Finance*, 65: 1237-1267.

DUFFIE, DARRELL (2019). "Digital Currencies and Fast Payment Systems: Disruption is Coming." Asian Monetary Policy Forum.

EICHENGREEN, BARRY, T NGUYEN, MY, and VISWANATH-NATRAJ, GANESH (2023). "Stablecoin Devaluation Risk." *WBS Finance Group Research Paper*.

EGAN, MARK, ALI HORTACSU and GREGOR MATVOS (2017). "Deposit Competition and Financial Fragility: Evidence from the US Banking Sector." *American Economic Review*, 107: 169-216.

EUROPEAN PARLIAMENT AND COUNCIL OF THE EUROPEAN UNION (2023). "Regulation (EU) 2023/1114 on markets in crypto-assets, and amending Regulations (EU) No 1093/2010 and (EU) No 1095/2010 and Directives 2013/36/EU and (EU) 2019/1937." Official Journal of the European Union, L 150, 40-205. https://eur-lex.europa.eu/eli/reg/2023/1114/oj.

FARHI, EMMANUEL, MIKHAIL GOLOSOV and ALEH TSYVINSKI (2009). "A Theory of Liquidity and Regulation of Financial Intermediation." *Review of Economic Studies*, 76: 973-992.

FINANCIAL CONDUCT AUTHORITY (2023). "Regulating Cryptoassets Phase 1: Stablecoins." Discussion Paper, DP23/4. Available at FCA: https://www.fca.org.uk/publication/discussion/dp23-4.pdf.

FLANNERY, MARK J (2023). "How the Rise of Stablecoins Could Threaten Traditional Banks' Lending Capacity." Available at SSRN 4320945.

FROST, JON, HYUN SONG SHIN, and PETER WIERTS 2020. "An Early Stablecoin? The Bank of Amsterdam and the Governance of Money." Working Paper.

GORTON, GARY, ELIZABETH C. KLEE, CHASE ROSS, SHARON ROSS, and ALEXANDROS P. VARDOULAKIS (2023). “Leverage and Stablecoin Pegs.” Working Paper.

GORTON, GARY, CHASE ROSS and SHARON ROSS (2022). “Making Money.” Working Paper.

GORTON, GARY and JEFFERY ZHANG (2021). “Taming Wildcat Stablecoins.” Working Paper.

GORTON, GARY and GEORGE PENNACCHI (1990). “Financial Intermediaries and Liquidity Creation.” *Journal of Finance*, 45: 49-71.

GOLDSTEIN, ITAY and ADY PAUZNER (2005). “Demand Deposit Contracts and the Probability of Bank Runs.” *Journal of Finance*, 60: 1293-1328.

GRIFFIN, JOHN and AMIN SHAMS (2020). “Is Bitcoin Really Untethered?” *Journal of Finance*, 75: 1913-1964.

GROMB, DENIS and DIMITRI VAYANOS (2002). “Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs.” *Journal of Financial Economics*, 66: 361-407.

GROMB, DENIS and DIMITRI VAYANOS (2018). “The Dynamics of Financially Constrained Arbitrage.” *The Journal of Finance*, 73(4): 1713-1750.

GENIUS ACT (2025). Bill introduced in the U.S. Senate. 119th Congress, 1st Session. https://www.hagerty.senate.gov/wp-content/uploads/2025/02/GENIUS-Act.pdf.

HARVEY, CAMPBELL, ASHWIN RAMACHANDRAN, and JOEY SANTORO (2021). *DeFi and the Future of Finance*, 2021, John Wiley & Sons.

HE, ZHIGUO, ARVIND KRISHNAMURTHY and KONSTANTIN MILBRADT (2019). “A Model of Safe Asset Determination.” *American Economic Review*, 109: 1230-1262.

HE, ZHIGUO and WEI XIONG (2012). “Rollover Risk and Credit Risk.” *Journal of Finance*, 67: 391-430.

HOLMSTROM, BENGT (2015). “Understanding the role of debt in the financial system.” Working Paper.

HU, ALBERT, CHRISTINE PARLOUR and UDAY RAJAN (2019). “Cryptocurrencies: Stylized Facts on a New Investible Instrument.” *Financial Management*, 48.4: 1049-1068.

IYER, RAJKAMAL and MANJU PURI (2012). ”Understanding Bank Runs: The Importance of Depositor-Bank Relationships and Networks.” *American Economic Review*, 102: 1414-1445.

JACKLIN, CHARLES (1987). “Demand Deposits, Trading Restrictions and Risk Sharing.” in: Prescott, E. D., Wallace, N. (Eds.), *Contractual Arrangements for Intertemporal Trade*. Minnesota: University of Minnesota Press.

JOHN, KOSE, LEONID KOGAN and FAHAD SALEH (2022). “Smart Contracts and Decentralized Finance.” Working Paper.

JIN, DUNHONG, MARCIN KACPERCZYK, BIGE KAHRAMAN and FELIX SUNTHEIM (2022). “Swing Pricing and Fragility in Open-end Mutual Funds.” *Review of Financial Studies*, 35: 1-50.

KACPERCZYK, MARCIN and PHILIPP SCHNABL (2013). “How Safe Are Money Market Funds?” *Quarterly Journal of Economics*, 128: 1073-1122.

KASHYAP, ANIL, DIMITRIOS TSOMOCOS, and ALEXANDROS VARDOULAKIS (2023). “Optimal Bank Regulation in the Presence of Credit and Run-Risk.” *Journal of Political Economy*, forthcoming.

KELLEY, ERIC and PAUL TETLOCK (2013). "How Wise Are Crowds? Insights from Retail Orders and Stock Returns." *Journal of Finance*, 68: 1229-1265.

KIYOTAKI, NOBUHIRO and RANDALL WRIGHT (1983). "On Money as A Medium of Exchange." *Journal of Political Economy*, 97: 927-954.

KIM, SANG RAE (2022). "How the Cryptocurrency Market is Connected to the Financial Market." Working Paper.

KLEMPERER, PAUL and MARGARET MEYER (1989). "Supply Function Equilibria in Oligopoly under Uncertainty." *Econometrica*, 57: 1243-1277.

KOONT, NAZ, YIMING MA, LUBOS PASTOR and YAO ZENG (2022). "Steering a Ship in Illiquid Waters: Active Management of Passive Funds." Working Paper.

KOZHAN, ROMAN and GANESH VISWANATH-NATRAJ (2021). "Decentralized Stablecoins and Collateral Risk." Working Paper.

KYLE, ALBERT (1989). "Informed Speculation with Imperfect Competition." *Review of Economic Studies*, 56: 317-355.

LI, YE and SIMON MAYER (2021). "Money Creation in Decentralized Finance: A Dynamic Model of Stablecoin and Crypto Shadow Banking." Working Paper.

LIU, JIAGENG, IGOR MAKAROV and ANTOINETTE SCHOAR (2023). "Anatomy of a Run: The Terra Luna Crash." Working Paper.

LIAO, GORDON and JOHN CARAMICHAEL (2021). "Stablecoins: Growth Potential and Impact on Banking." Working Paper.

LYONS, RICHARD and GANESH VISWANATH-NATRAJ (2021). "What Keeps Stablecoins Stable?" Working Paper.

MA, YIMING, KAIRONG XIAO and YAO ZENG (2021). "Mutual Fund Liquidity Transformation and Reverse Flight to Liquidity." *Review of Financial Studies*, forthcoming.

MACCHIAVELLI, MARCO, and XING ZHOU (2022). "Funding Liquidity and Market Liquidity: The Broker-Dealer Perspective." *Management Science*, 68: 3379-3398.

MAKAROV, IGOR and ANTOINETTE SCHOAR (2022). "Cryptocurrencies and Decentralized Finance (DeFi)." *Brookings Papers on Economic Activity*, Spring 2022: 1-71.

MORRIS, STEPHEN and HYUN SHIN (1998). "Unique Equilibrium in a Model of Self-fulfilling Currency Attacks." *American Economic Review*, 88: 587-597.

MORRIS, STEPHEN and HYUN SHIN (2003). "Global Games: Theory and Applications." in Mathias Dewatripont, Lars Peter Hansen and Stephen J. Turnovsky (eds), *Advances in Economics and Econometrics (Proceedings of the Eighth Congress of the Econometric Society)*, 3: 56-114.

PARLATORE, CECILIA (2016). "Fragility in Money Market Funds: Sponsor Support and Regulation." *Journal of Financial Economics*, 121: 595-623.

RANALDO, ANGELO, GANESH VISWANATH-NATRAJ, and JUNXUAN WANG (2024). "Blockchain Currency Markets." Working Paper.

ROCHETEAU, GUILLAUME and RANDALL WRIGHT (2005). "Money in Search Equilibrium, in Competitive Equilibrium, and in Competitive Search Equilibrium." *Econometrica*, 73: 175–202.

ROUTLEDGE, BRYAN and ARIEL ZETLIN-JONES (2022). "Currency Stability Using Blockchain Technology." *Journal of Economic Dynamics and Control*, 142: 104-155.

SCHILLING, LINDA and HARALD UHLIG (2019). "Some Simple Bitcoin Economics." *Journal of Monetary Economics*, 106: 16-26.

SCHMIDT, LAWRENCE, ALLAN TIMMERMANN and RUSS WERMERS (2016). "Runs on Money Market Mutual Funds." *American Economic Review*, 106: 2625-2657.

SHLEIFER, ANDREI and ROBERT VISHNY (1997). "The Limits of Arbitrage." *Journal of Finance*, 52: 35-55.

SOCKIN, MICHAEL and WEI XIONG (2023). "Decentralization Through Tokenization." *Journal of Finance*, 78: 247-299.

SOCKIN, MICHAEL and WEI XIONG (2023). "A Model of Cryptocurrencies." *Management Science*, forthcoming.

STABLE ACT (2025). Bill introduced in the U.S. Congress 119th Congress, 1st Session. https://www.congress.gov/bill/119th-congress/house-bill/2392/text/ih?overview=closed&format=xml.

SUNDERAM, ADI (2015). "Money Creation and the Shadow Banking System." *Review of Financial Studies*, 28: 939-977.

UHLIG, HARALD (2022). "A Lunatic Stablecoin Crash." Working Paper.

ZHANG, ANTHONY LEE (2022). "Competition and Manipulation in Derivative Contract Markets." *Journal of Financial Economics*, 144: 396-413.

Figure 3: Asset Size of Fiat-backed Stablecoins

This figure shows the asset size of the six largest fiat-backed stablecoins over time.

ALT TEXT: Line chart depicting how the asset sizes of the six largest fiat-backed stablecoins, including USDT, USDC, BUSD, USDP, TUSD, and GUSD, evolve over time.

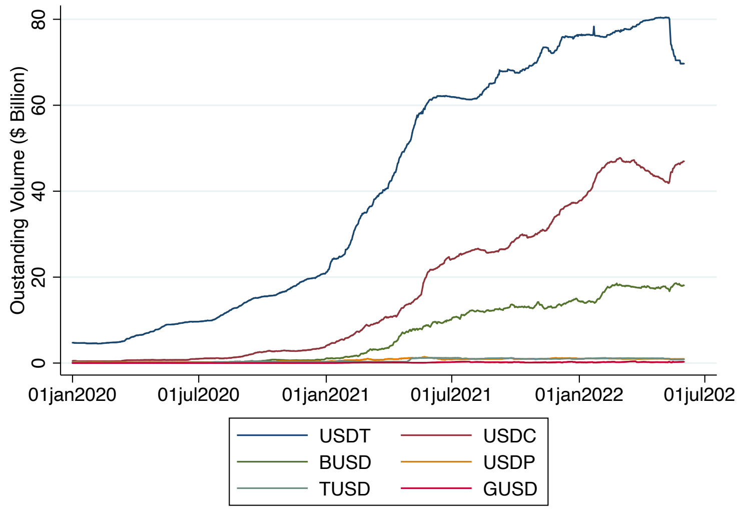

Figure 4: The Design of Fiat-backed Stablecoins

This figure illustrates the design of fiat-backed stablecoins.

ALT TEXT: Graphical representation of how arbitrageurs connect the primary and secondary markets for stablecoins. Depicts the trading of stablecoins between buyers and sellers in the secondary market, arbitrageurs' purchase and sale of stablecoins in the secondary market, arbitrageurs' redemption and creation of stablecoins for $1 with the stablecoin issuer in the primary market, and the stablecoin issuer's purchase and sale of securities.

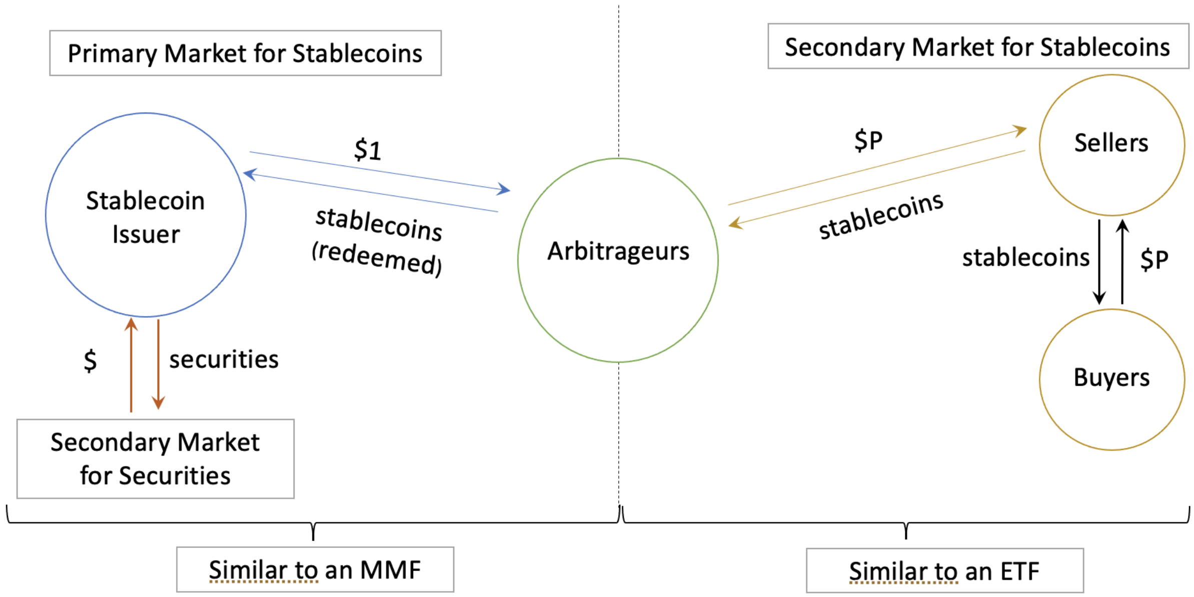

Figure 5: Secondary Market Trading Price

Panels (a) to (f) show the daily secondary market trading price of USDT, USDC, BUSD, USDP, TUSD, and GUSD, respectively. Secondary market prices are volume-weighted averages of trading prices from the exchanges listed in Section 1.

ALT TEXT: Six line graphs depicting the daily secondary market trading price of USDT, USDC, BUSD, USDP, TUSD, and GUSD.

(a) USDT

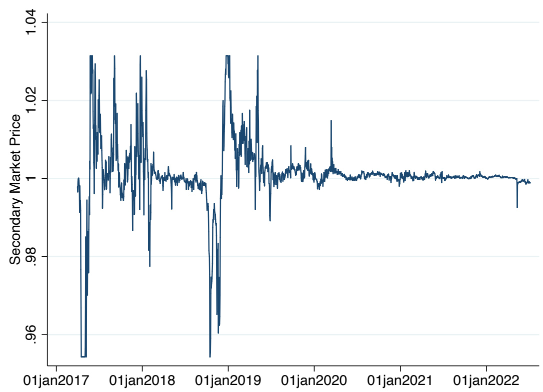

(b) USDC

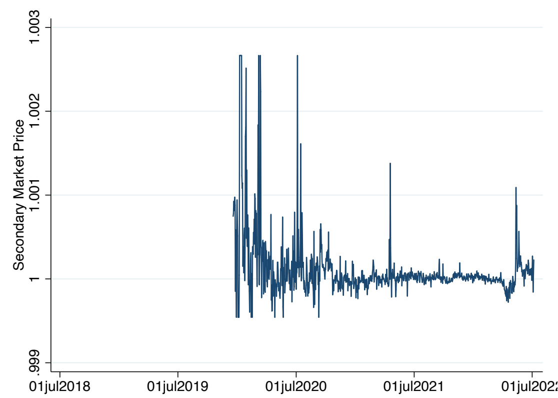

(c) BUSD

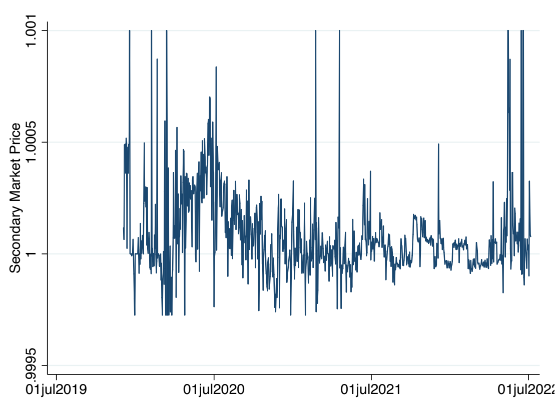

(d) USDP

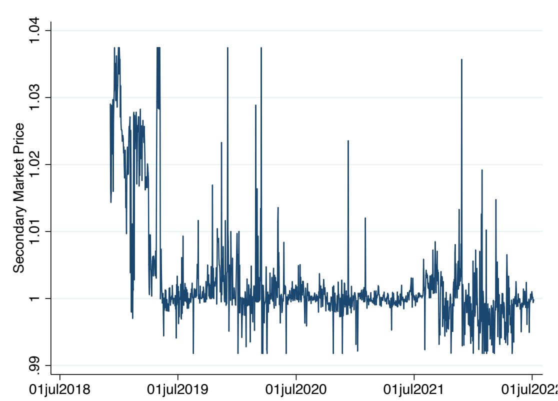

(e) TUSD

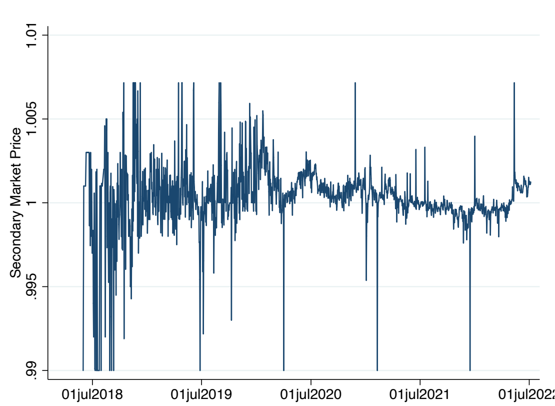

(f) GUSD

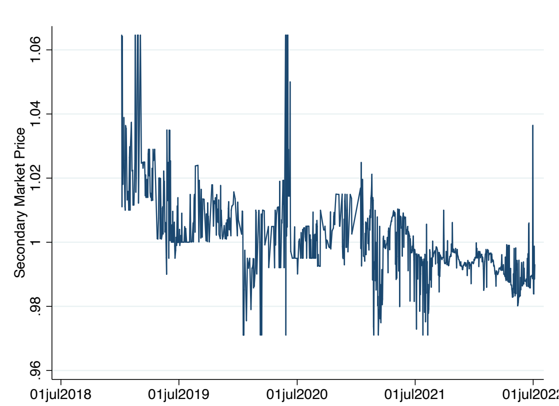

Figure 6: Secondary Market Price Dislocations and Primary Market Structure

This figure shows the relationship between secondary market price dislocations and primary market structure. In panel (a), each dot indicates the average secondary market price deviation and the average number of arbitrageurs in a month for a given stablecoin. In panel (b), each dot indicates the average secondary market price deviation and the average market share of the top 5 arbitrageurs in a month for a given stablecoin. We first calculate monthly secondary market price deviations for a given stablecoin by averaging over the absolute values of daily price deviations from one in a given month, which includes both deviations above and below one. We then average over months to obtain that stablecoin's average secondary market price deviation. Similarly, we count the number of unique arbitrageurs that engage in redemptions and/or creations, calculate the market share of the largest five arbitrageurs in each month, and then average over time for each coin. For ease of presentation, we take the number of arbitrageurs for USDC, which exceeds 5000, to be 500.

ALT TEXT: Two scatter plots depicting the relationship between average secondary market price deviations and the number of arbitrageurs and the relationship between average secondary market price deviations and the market share of the largest five arbitrageurs. Each observation corresponds to one stablecoin.

(a) Number of Arbitrageurs **(b)** Market Share of Top 5 Arbitrageurs

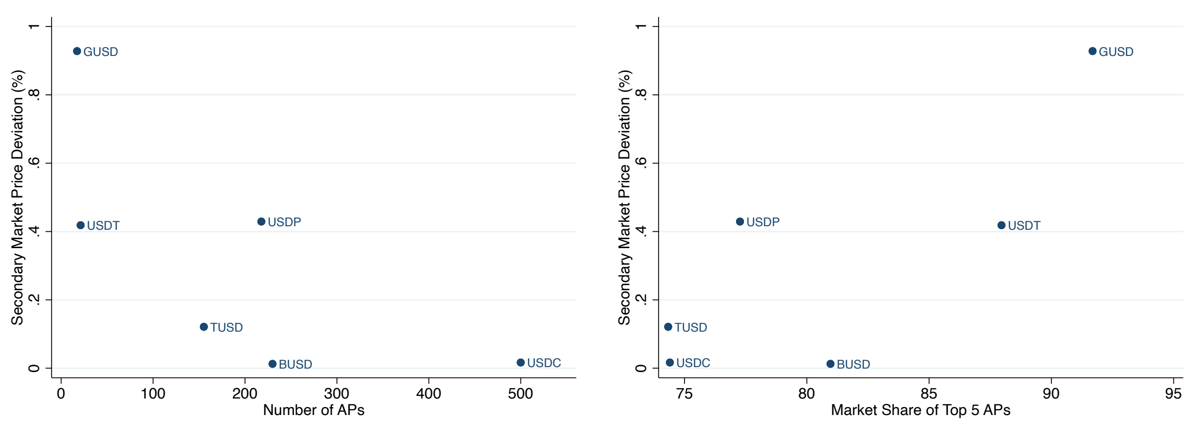

Figure 7: Investors' Payoff Gain from Waiting versus Selling Early

This figure shows an investor's payoff gain from waiting until $t = 3$ relative to selling early at $t = 2$ against the population of investors selling early at $t = 2$. Parameters used are $\pi(\theta) = 0.97, \eta = 0.2$, $\phi = 0.05$, and $K = 0.3$.

ALT TEXT: Line graph depicting an investor's payoff gain from waiting until $t = 3$ relative to selling early at $t = 2$ against the population of investors selling early at $t = 2$.

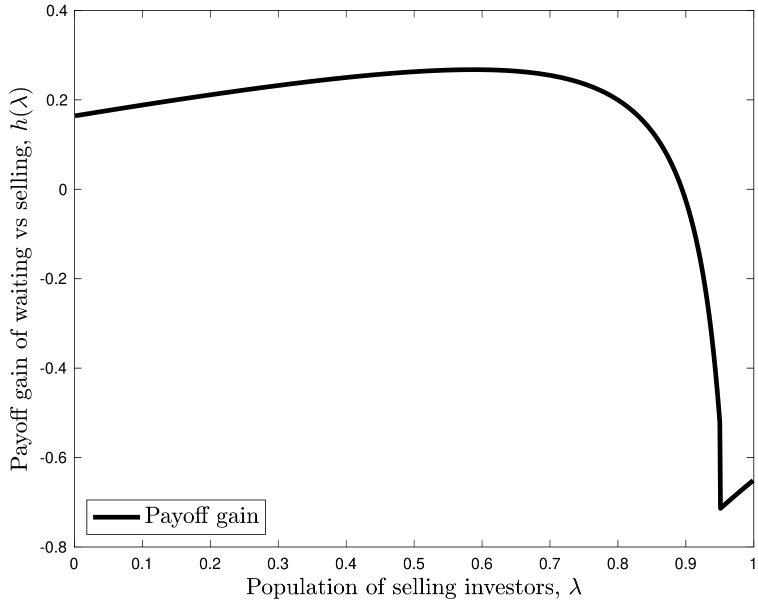

Figure 8: Investors' Payoff Gain from Waiting versus Selling Early: Comparative Statics with Respect to K

This figure shows an investor's payoff gain from waiting until $t = 3$ relative to selling early at $t = 2$ against the population of investors selling early at $t = 2$ under different values of arbitrage concentration K. Parameters used are $\pi(\theta) = 0.97, \eta = 0.2$, and $\phi = 0.05$.

ALT TEXT: Line graph depicting an investor's payoff gain from waiting until $t = 3$ relative to selling early at $t = 2$ against the population of investors selling early at $t = 2$ under two different values of arbitrage concentration K.

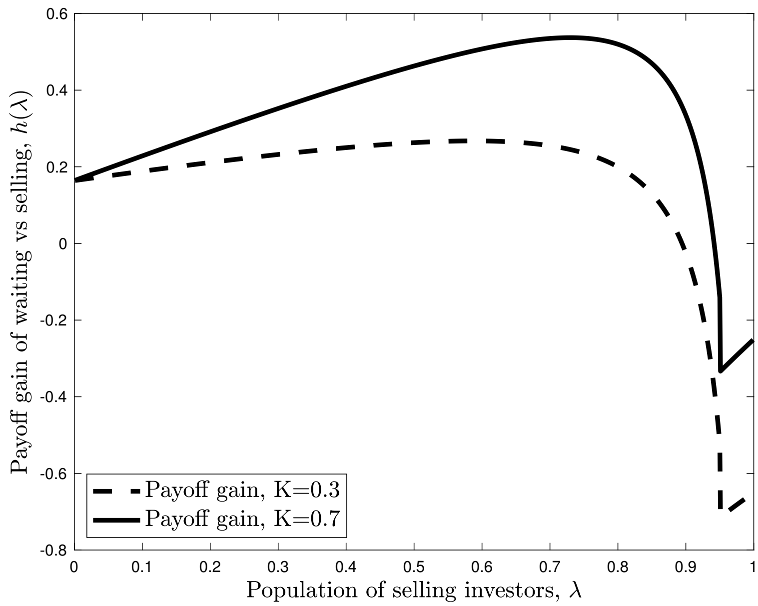

Figure 9: Effect of Redemption Fees

This figure shows the predicted effect of redemption fees ν on run probabilities. Throughout the exercise, we hold K equal to the model-predicted optimal value of K, in the absence of redemption fees.

ALT TEXT: Line graph depicting the relationship between run probabilities and redemption fees for USDC and USDT.

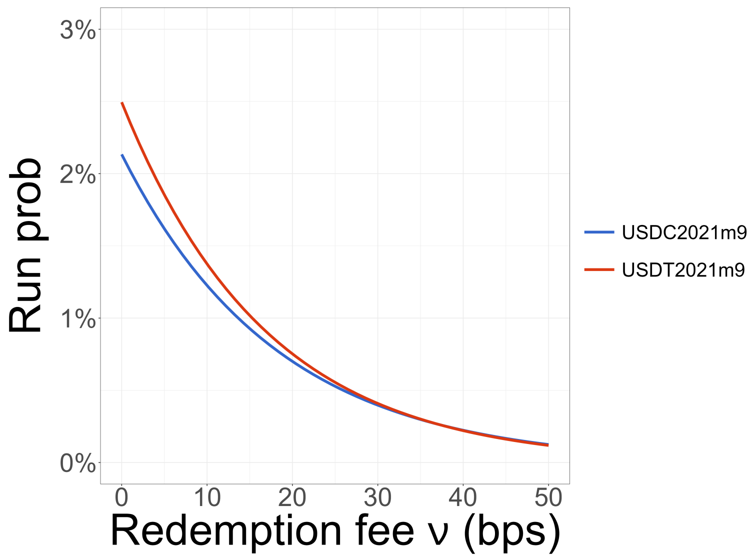

Figure 10: Effect of Dividend Payments

This figure shows the predicted effect of dividend payments to investors on the issuer's choice of K, the cost of price variance $K\alpha\delta^2$, and run probability.

ALT TEXT: Three line graphs depicting the relationship between dividend payments and arbitrage concentration K, the cost of price variance $K\alpha\delta^2$, and run probability for USDC and USDT.

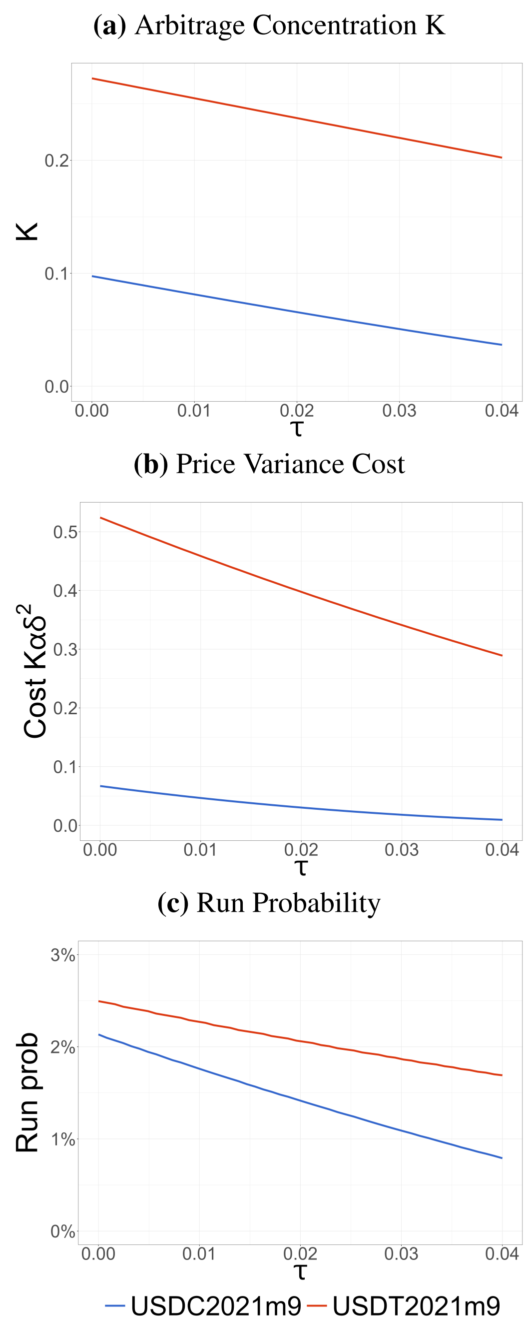

Table 1: Secondary Market Price and Volume

This table provides statistics about secondary market trading, including the average daily trading volume, the proportion of days with discounts and premiums, the average discount and premium, and the median discount and premium. Table 1a is based on the full sample period, whereas Table 1b is based on the sample period starting in January 2020.

(a) Full Sample

	USDT	USDC	BUSD	TUSD	USDP	GUSD
Average Daily Volume	16.4	15.4	13.5	11.4	10.5	7.3
Proportion of Discount Days (%)	30.5	27.2	34.9	38.2	41.6	39.7
Proportion of Premium Days (%)	69.5	72.8	64.4	61.4	57.3	58.9
Average Discount (%)	0.54	0.01	0.01	0.11	0.18	0.78
Average Premium (%)	0.36	0.02	0.02	0.13	0.64	1.17
Median Discount (%)	0.11	0.00	0.00	0.05	0.09	0.63
Median Premium (%)	0.11	0.01	0.01	0.10	0.18	0.82

(b) Sample starting from January 2020

	USDT	USDC	BUSD	TUSD	USDP	GUSD
Average Daily Volume	18.3	15.5	13.6	13.0	11.1	7.6
Proportion of Discount Days (%)	21.8	40.7	37.5	38.3	53.5	58.0
Proportion of Premium Days (%)	78.2	59.3	62.1	61.7	45.9	41.9
Average Discount (%)	0.06	0.01	0.01	0.05	0.19	0.81
Average Premium (%)	0.07	0.02	0.02	0.10	0.20	0.81
Median Discount	0.05	0.00	0.00	0.04	0.09	0.64
Median Premium (%)	0.05	0.01	0.01	0.08	0.10	0.65

Table 2: Primary Market Monthly Redemption and Creation Activity

Panels (a) to (f) provide statistics about monthly primary market redemption and creation activity on the Ethereum blockchain, including the number of arbitrageurs, the market share of the top 1 and top 5 arbitrageurs, and the transaction volume. For each variable, we show the average, 25^{th} percentile, 50^{th} percentile, and 75^{th} percentile values across months in our sample.

(a) USDT

	mean	p25	p50	p75
RD AP Num	6	3	6	8
RD Top 1 Share	66	42	61	89
RD Top 5 Share	97	98	100	100
RD Vol (mil)	577	46	123	763
CR AP Num	18	9	17	26
CR Top 1 Share	59	35	57	77
CR Top 5 Share	90	84	93	99
CR Vol (mil)	1271	101	470	1800

(b) USDC

	mean	p25	p50	p75
RD AP Num	521	114	168	262
RD Top 1 Share	45	38	49	50
RD Top 5 Share	85	81	85	90
RD Vol (mil)	2976	160	460	4965
CR AP Num	5067	284	406	13112
CR Top 1 Share	45	31	44	51
CR Top 5 Share	81	70	84	92
CR Vol (mil)	3953	184	680	7448

(c) BUSD

	mean	p25	p50	p75
RD AP Num	214	157	202	274
RD Top 1 Share	48	30	50	62
RD Top 5 Share	81	74	82	87
RD Vol (mil)	1596	233	1498	2720
CR AP Num	16	8	11	19
CR Top 1 Share	65	53	68	82
CR Top 5 Share	98	97	99	100
CR Vol (mil)	2116	290	1628	3739

(d) USDP

	mean	p25	p50	p75
RD AP Num	178	71	174	284
RD Top 1 Share	41	24	37	54
RD Top 5 Share	74	62	77	88
RD Vol (mil)	260	94	174	262
CR AP Num	41	5	8	67
CR Top 1 Share	58	48	61	70
CR Top 5 Share	93	94	99	100
CR Vol (mil)	279	107	170	341

(e) TUSD

	mean	p25	p50	p75
RD AP Num	66	49	74	85
RD Top 1 Share	50	36	46	64
RD Top 5 Share	86	79	91	94
RD Vol (mil)	154	31	85	260
CR AP Num	92	53	106	130
CR Top 1 Share	50	33	46	65
CR Top 5 Share	87	83	87	92
CR Vol (mil)	164	30	77	259

(f) GUSD

	mean	p25	p50	p75
RD AP Num	1	1	1	1
RD Top 1 Share	100	100	100	100
RD Top 5 Share	100	100	100	100
RD Vol (mil)	113	7	17	164
CR AP Num	17	1	12	19
CR Top 1 Share	55	29	40	100
CR Top 5 Share	85	72	82	100
CR Vol (mil)	117	4	13	155

Table 3: Asset Composition

This table shows the breakdown of reserves by asset class for USDT and USDC. Data are available for the dates on which USDT and USDC publish reserve breakdowns. For USDT, the "Deposit" category includes bank deposits, while for USDC, the "Deposit" category includes US dollar deposits at banks and short-term, highly liquid investments.

(a) USDT

	Deposits	Treas	Muni	MM	Corp	Loans	Others
2021/06	10.0	24.3	0.0	50.7	7.7	4.0	3.3
2021/09	10.5	28.1	0.0	45.7	5.2	5.0	5.5
2021/12	5.3	43.9	0.0	34.5	4.6	5.3	6.4
2022/03	5.0	47.6	0.0	32.8	4.5	3.8	6.4

(b) USDC

	Deposits	Treas	Muni	MM	Corp	Loans	Others
2021/05	60.4	12.2	0.5	22.1	5.0	0.0	0.0
2021/06	46.4	13.1	0.4	24.2	15.9	0.0	0.0
2021/07	47.4	12.4	0.7	23.0	16.4	0.0	0.0
2021/08	92.0	0.0	0.0	6.5	1.5	0.0	0.0
2021/09	100.0	0.0	0.0	0.0	0.0	0.0	0.0
2021/10	100.0	0.0	0.0	0.0	0.0	0.0	0.0

Table 4: Asset Returns

This table shows the return and revenue of asset portfolios for USDT and USDC by date. Returns are in % and revenue is in million dollars. Data are available for the dates on which USDT and USDC publish reserve breakdowns. For each date, we also list the Fed funds rate in that month.

(a) USDT

	Fed Funds Rate (%)	Return (%)	Revenue (Million $)
2021/06	0.10	0.59	370.4
2021/09	0.06	0.66	458.6
2021/12	0.07	0.97	759.7
2022/03	0.33	1.98	1634.2

(b) USDC

	Fed Funds Rate (%)	Return (%)	Revenue (Million $)
2021/05	0.06	0.25	56.6
2021/06	0.10	0.51	127.6
2021/07	0.10	0.48	130.9
2021/08	0.08	0.18	50.2
2021/09	0.06	0.14	44.4
2021/10	0.08	0.16	52.8

Table 5: Parameter Estimates

Parameters for asset illiquidity ϕ and the long-term benefit η are estimated as described in Section 6.1. Parameters for the price variance cost $\alpha\delta^2$ and the elasticity of demand γ, η, ϕ are estimated as described in Section 6.2. Run prob is the run probability at the issuer's optimal choice of K.

Coin	Month	$\alpha\sigma_\epsilon^2$	γ	η	ϕ	Run Prob
USDC	2021m5	7.06	1.10	0.0301	0.0250	3.188%
USDC	2021m6			0.0198	0.0296	3.893%
USDC	2021m7			0.0221	0.0293	3.737%
USDC	2021m8			0.0575	0.0178	1.883%
USDC	2021m9			0.0443	0.0150	2.134%
USDC	2021m10			0.0525	0.0150	1.828%
USDT	2021m6		0.54	0.0301	0.0431	2.590%
USDT	2021m9			0.0292	0.0436	2.495%
USDT	2021m12			0.0250	0.0413	2.040%
USDT	2022m3			0.0365	0.0395	1.664%

Table 6: Model Fit

Target K is the slope of arbitrageur demand for the stablecoin, estimated from the data using (6.1). Model K is the model-predicted slope of arbitrageur demand. Target elas. is the partial elasticity of investors' demand for the stablecoin with respect to the long-term benefit η, as described in Subsection (6.2). Model elas. is the model partial elasticity of investors' demand for the stablecoin with respect to η.

Coin	Month	Target K	Model K	Target elas.	Model elas.
USDC	2021m5	0.156	0.166	2.486	2.638
USDC	2021m6	0.156	0.206	2.486	4.080
USDC	2021m7	0.156	0.201	2.486	3.822
USDC	2021m8	0.156	0.090	2.486	1.444
USDC	2021m9	0.156	0.097	2.486	1.501
USDC	2021m10	0.156	0.084	2.486	1.387
USDT	2021m6	0.209	0.269	1.600	1.687
USDT	2021m9	0.209	0.273	1.600	1.737
USDT	2021m12	0.209	0.267	1.600	1.662
USDT	2022m3	0.209	0.236	1.600	1.332

Internet Appendix for

Stablecoin Runs and the Centralization of Arbitrage

Yiming Ma Yao Zeng Anthony Lee Zhang

A Additional Institutional Details

A.1 Minting of Stablecoins

Technically, stablecoins on Ethereum are ERC-20 tokens, and stablecoins on other blockchains are implemented as similar token "smart contracts." The stablecoin "smart contract," that is, the blockchain code that governs the behavior of the stablecoin, gives the stablecoin issuer the arbitrary right to create, or "mint", new stablecoin tokens into arbitrary wallet addresses. Stablecoin issuers adopt technically slightly different strategies to issue and redeem stablecoins in primary markets. Some, like USDC, directly "mint" new coins using the token smart contract into customers' wallets. Others, like USDT, occasionally mint large amounts of stablecoin tokens to "treasury" wallets under their control and then issue stablecoins in primary markets by sending tokens from the "treasury" address to customers' wallets.[30]

A.2 Trading on Crypto Exchanges

There are several ways individuals can purchase stablecoins with local fiat currency. One method is to deposit fiat on a custodial centralized crypto exchange (CEX), such as Binance or Coinbase. Centralized exchanges, like stock brokerages, keep custody of fiat and crypto assets on behalf of users, and allow users to purchase or sell crypto assets using fiat currencies. After purchasing stablecoins on a CEX, the user can then "withdraw" the stablecoins, instructing the CEX to send her stablecoins to a wallet address of her choosing, to self-custody the purchased stablecoins. Another approach is to use peer-to-peer exchanges, such as Paxful. On these platforms, users list offers to buy or sell stablecoins or other crypto tokens for other forms of payment. Accepted forms

[30] Treasury address tokens technically count towards the market cap of any given stablecoin, but they are not economically meaningful as part of the market cap, since USDT does not have to hold US dollar assets against tokens it holds in its treasury. Thus, we will not count tokens held in treasury addresses as part of the stablecoin supply in circulation.

of payment in the US include Zelle, Paypal, Western Union, ApplePay, and many others. The exchange platform plays an escrow, insurance, and mediation role in these transactions. When a user buys a stablecoin, she sends funds to the exchange's escrow account and the stablecoin seller sends stablecoins to an address of the buyer's choosing. Once the buyer confirms receipt of the stablecoins, the exchange sends funds from the escrow account to the seller's account. In this process, purchased stablecoins are sent directly to the user's self-custodial wallet.

B Further Information regarding the Data

B.1 Primary Market Data

As mentioned in Appendix A, there are two ways that stablecoin tokens can be minted or redeemed. First, the stablecoin's "mint" or "burn" functions can be called directly to the primary market participant's wallet. To capture this category of actions, we query Etherscan for all cases in which the "mint" and "burn" functions are called for each stablecoin. Second, the stablecoin issuer can send or receive stablecoins from their "treasury" address. To capture this category, we identify the treasury address or addresses for each stablecoin, and then query Etherscan for every send or receive transaction involving the treasury address. Logistically, some issuers, such as USDT, tend to mint a large number of stablecoin tokens into "treasury" addresses they control, then issue tokens to market participants simply by transferring tokens out of their treasury wallet; whereas other issuers, such as TrueUSD, occasionally directly mint stablecoin tokens into the wallet addresses of market participants. On the other hand, most issuers handle redemptions by having market participants send tokens to a treasury wallet address. If the treasury wallet has a large balance of redeemed stablecoins, the issuer will occasionally "burn" quantities of the stablecoin, removing them from the technical outstanding balance of the token.[31]

Different wallet addresses that we observe minting and burning stablecoins could in principle be controlled by the same entity. We gather wallet "labels" provided by Etherscan, and group wallets that are labeled as belonging to the same entity. For example, Etherscan labels the addresses 0x345d8e3a1f62ee6b1d483890976fd66168e390f2 as "Binance 23", and 0x21a31ee1afc51d94c2efccaa2092ad1028285549 as "Binance 15"; we group all arbitrage transactions of these and other Binance-related wallets, and treat this group of wallets as one economic entity, for the purpose of our analysis. Not all wallets are labeled in Etherscan, so it is conceivable that wallets

[31]The exception to this rule is that TrueUSD occasionally handles redemptions by "burning" tokens directly from market participants' wallets, rather than the treasury.

we treat as separate in fact are operated by the same economic entity. Our results may thus slightly underestimate the degree of arbitrage concentration in practice.

We calculate the total issued market capitalization of a given stablecoin at any point in time as the total technical market capitalization of the stablecoin minus the amount of the stablecoin held in "treasury" addresses. This is because tokens held in treasury wallets need not be backed one-to-one by US dollars and thus should not count as part of the total market capitalization of stablecoins in circulation.

B.2 Secondary Market Price Data

Our sample uses direct USD to stablecoin trading pairs to calculate USD stablecoin prices because using a larger set of trading pairs to back out the USD price of stablecoins, as is done by CoinGecko, for example, introduces several issues.

First, using USDC/USDT trading pairs in calculating USDC (or USDT) prices introduces some complications for stablecoins. For example, when the USDC/USDT trading pair depegs, it becomes unclear whether this deviation is driven by USDC or USDT. Thus, we avoid using such pairs in calculating stablecoin prices.

Second, when CoinGecko includes all stablecoin/other non-stablecoin cryptocurrency pairs in its calculations, it presumably converts these back to stablecoin prices by dividing by some cryptocurrency/USD metric; we think that this strategy overlooks the dispersion in other cryptocurrency prices across exchanges due to factors like demand shocks and incomplete markets (e.g., the "kimchi premium" effect), which can be much more salient for other cryptocurrencies than for stablecoins. This could result in variations in stablecoin prices that reflect fluctuations in other cryptocurrency demand across exchanges, rather than in stablecoins per se.

As an illustration, we plot the secondary market price for our sample, CoinGecko, and two major exchanges, Kraken and Binance, in Figure A.1 below. Observe that the CoinGecko data shows significantly more price deviations through the whole sample period for USDC than that uncovered in our data. Also, the price deviations in our data resemble the USD/USDC data on Kraken and Binance much better. This is consistent with CoinGecko's being influenced by noise and rate dispersion across exchanges on which different crypto trading pairs trade.

Finally, the direct stablecoin to USD trading pairs are conceptually more suitable for our research question. We hope to analyze the "on/off ramp" behavior of consumers where they are

Figure A.1: Secondary Market Trading Price

Panels (a) to (d) show the secondary market trading price of USDC in our data, on Coingecko, Binance, and Kraken, respectively.

(a) Our Data

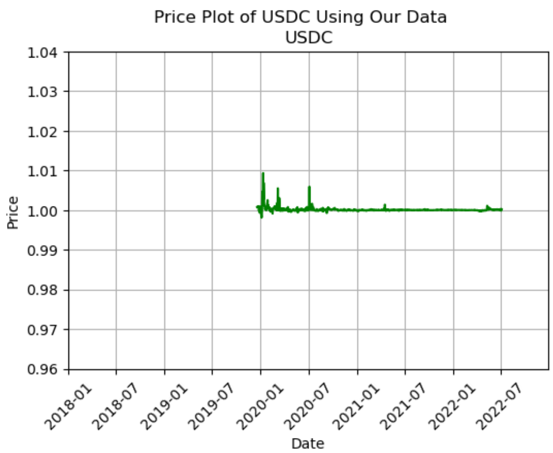

(b) Coingecko Data

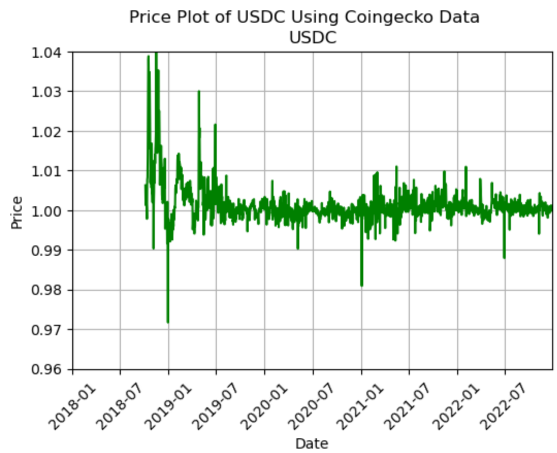

(c) Binance Data

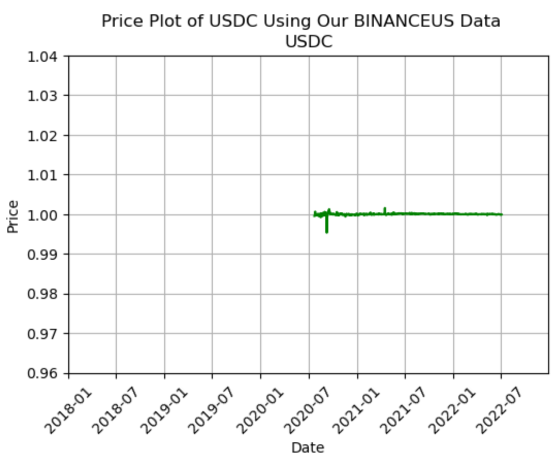

(d) Kraken Data

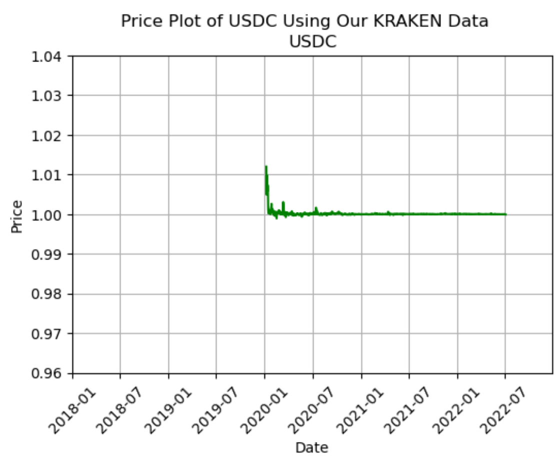

primarily concerned with entering or exiting the market in terms of USD. The CoinGecko data may be more suitable for studying trading between stablecoins and other crypto assets.

C Arbitrageur Inventory Holding

In the baseline model, we assume that arbitrageurs have zero net position in stablecoins so they cannot hold stablecoins on their balance sheet between periods. In this appendix, we first examine arbitrageurs' stablecoin holdings in the data. Then, we extend the model to allow arbitrageurs to maintain a positive balance of stablecoins.

C.1 Arbitrageurs' Stablecoin Ownership in the Data

To calculate arbitrageurs' stablecoin positions in the data, we take the following steps. First, using Etherscan, Snowtrace, and Tronscan, we gather all send and receive transactions of stablecoins involving each arbitrageur wallet address in our dataset. Having in hand all creation, redemption, send, and receive transactions, we can construct a time series of each arbitrageur's holdings of each token. Using this data, for each coin and chain, we calculate each arbitrageur wallet's monthly balance of stablecoin holdings and divide it by the total supply of that stablecoin.

In Figures A.2, A.3, and A.4, we plot the average stablecoin ownership of arbitrageurs on the Ethereum, Tron, and Avalanche blockchains, respectively. The average positions are non-zero and fluctuate over time but they remain well below 1.5% across all chains and coins. In Table A.1, we further show the variation across arbitrageurs' stablecoin ownership for the full sample of arbitrageurs. These positions are minimal. The 99^{th} percentile of stablecoin holdings remain below 1% except USDT arbitrageurs on Tron that have a 99^{th} percentile holding of 5.32%. When we further limit the sample to arbitrageurs with active redemptions or creations in that month, almost all positions at the 99^{th} percentile remain minimal and well below 2%. The only exceptions are GUSD on Ethereum and USDT on Tron.

Overall, our findings suggest that the vast majority of stablecoin arbitrageurs hold very small amounts of stablecoins and do not make up a substantial fraction of total holdings. This is consistent with our conversations with market participants. Arbitrageurs specialize in profiting from differences between secondary-market prices and primary-market redemption values. This is a sensible strategy for profit-seeking institutions because unlike arbitrage between primary and secondary markets, holding stablecoins per se does not generate returns because stablecoins currently do not pay out dividends.

Nevertheless, we acknowledge that our calculations may to some extent underestimate the concentration of stablecoin holdings by arbitrageurs. First, multiple wallets belonging to the same arbitrageur on the same chain may not be fully linked despite our best efforts. Second, we are unable to link wallets on different chains that belong to the same arbitrageur. Although arbitrageurs would need to own a very large number of wallets within and across chains for their holdings to be economically significant given our empirical findings, we further extend our model to explore the effect of arbitrageurs' allowing stablecoin holdings on our model predictions.

Figure A.2: Average Share of Stablecoins Held by Arbitrageurs (Ethereum)

Panels (a) to (f) show the average stablecoin ownership of arbitrageurs on the Ethereum blockchain. For each stablecoin and chain, we show the monthly average stablecoin ownership for arbitrageurs in our sample. Stablecoin ownership is expressed as a percentage of the total supply of stablecoins.

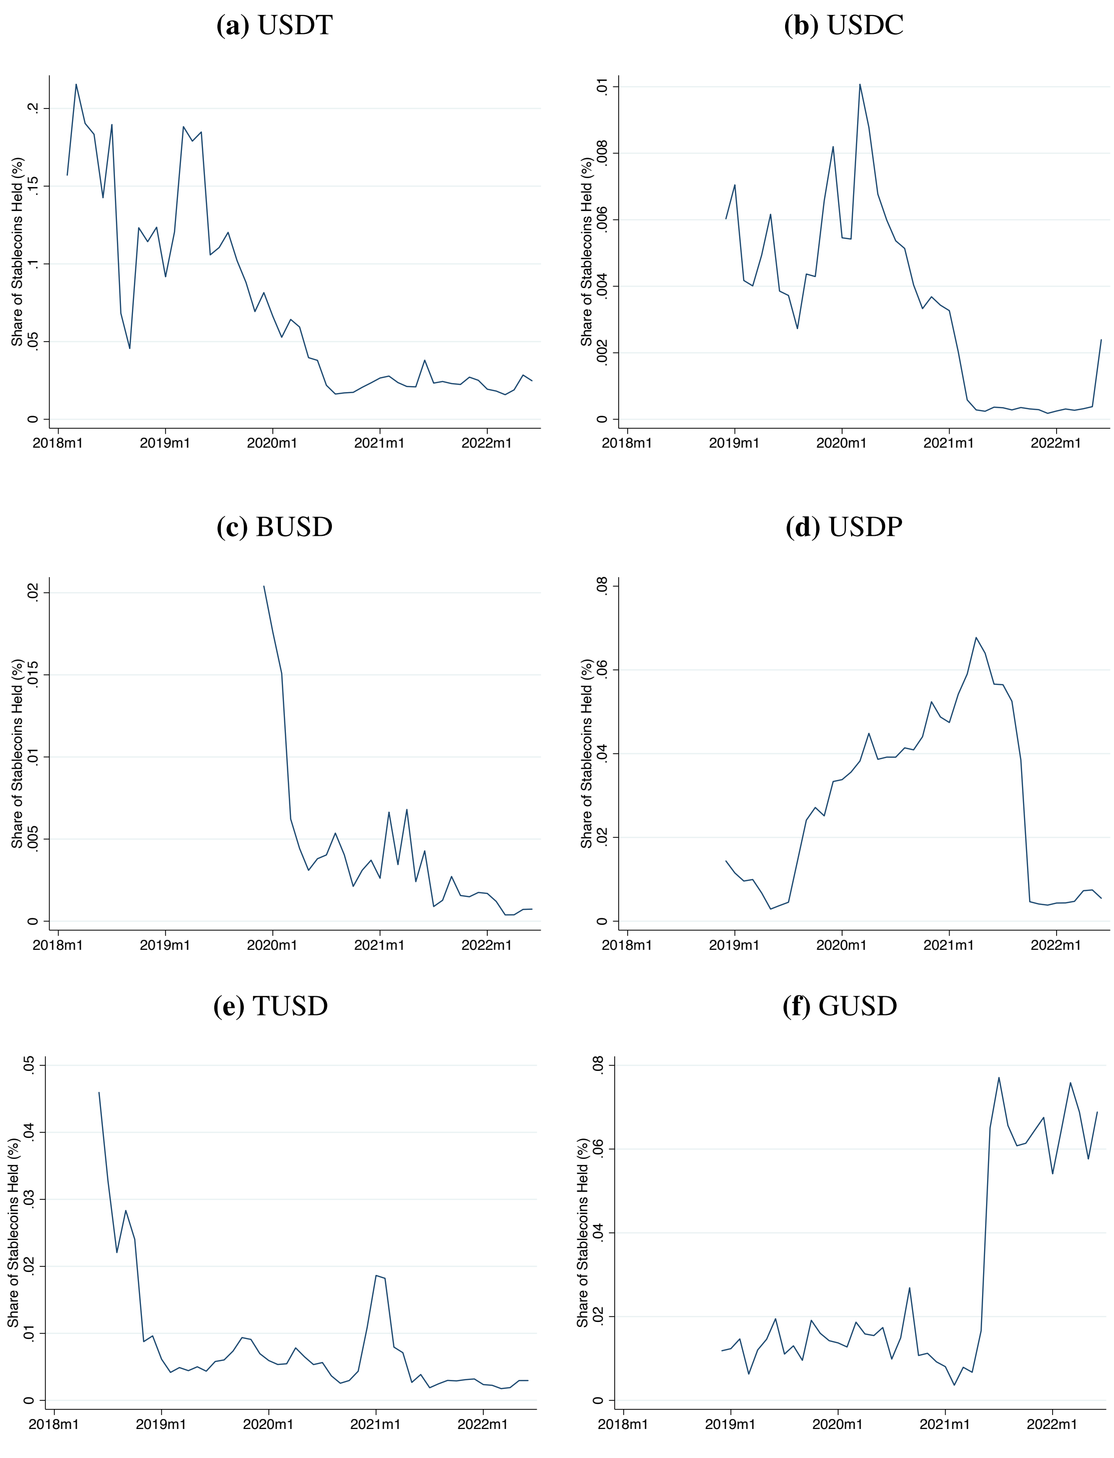

Figure A.3: Average Share of Stablecoins Held by Arbitrageurs (Tron)

Panels (a) to (f) show the average stablecoin ownership of arbitrageurs on the Tron blockchain. For each stablecoin and chain, we show the monthly average stablecoin ownership for arbitrageurs in our sample. Stablecoin ownership is expressed as a percentage of the total supply of stablecoins.

(a) USDT

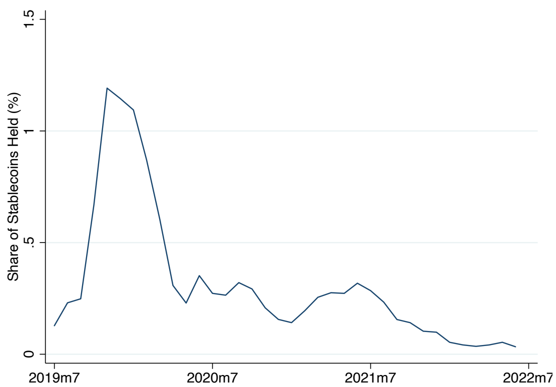

(b) USDC

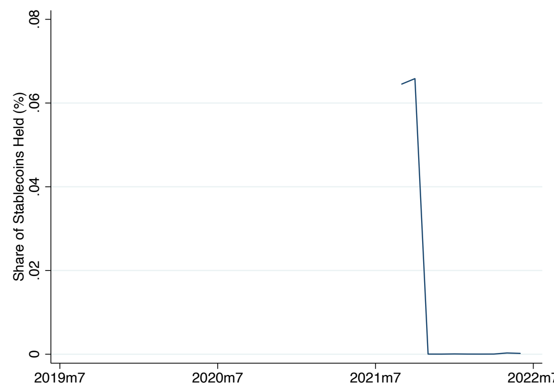

(c) TUSD

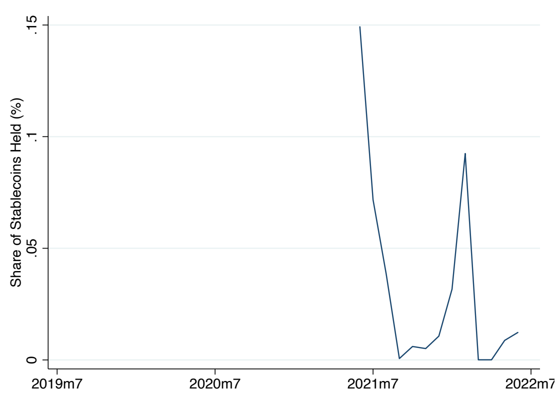

Figure A.4: Average Share of Stablecoins Held by Arbitrageurs (Avalanche)

Panels (a) to (f) show the average stablecoin ownership of arbitrageurs on the Avalanche blockchain. For each stablecoin and chain, we show the monthly average stablecoin ownership for arbitrageurs in our sample. Stablecoin ownership is expressed as a percentage of the total supply of stablecoins.

(a) USDT

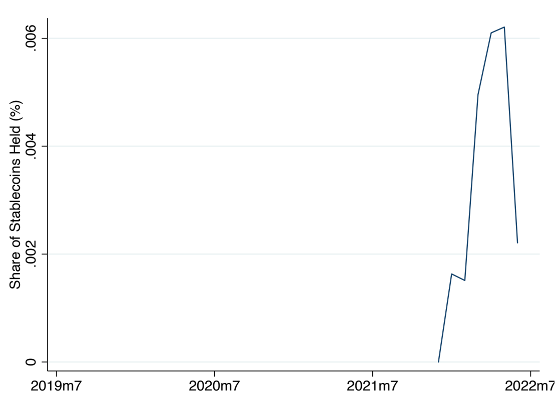

(b) USDC

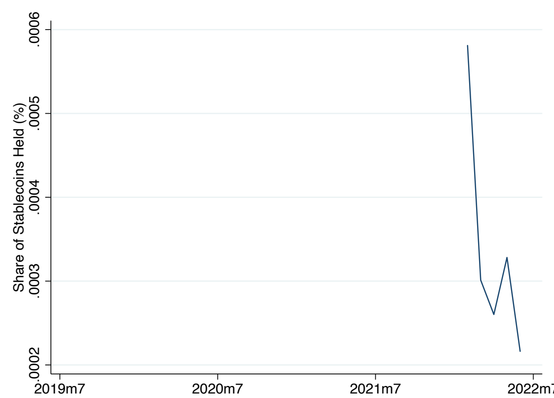

(c) TUSD

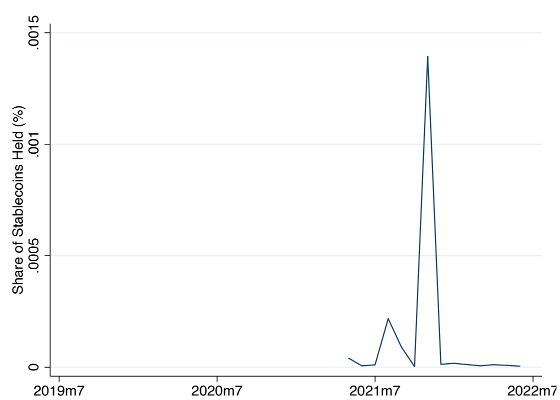

Table A.1: Share of Stablecoins Held by Arbitrageurs

Panels (a) to (c) provide statistics about the share of stablecoins owned by active arbitrageurs on the Ethereum, Tron, and Avalance blockchains. For each stablecoin and chain, we show the average and the 25^{th}, 50^{th}, 75^{th}, 90^{th}, 95^{th}, and 99^{th} percentile of stablecoin ownership for arbitrageur-month observations in our sample. Stablecoin ownership is expressed as a percentage of the total supply of stablecoins.

(a) Ethereum

	mean	p25	p50	p75	p90	p95	p99
USDC	0.00	0.00	0.00	0.00	0.00	0.00	0.00
USDT	0.07	0.00	0.00	0.00	0.00	0.01	0.70
BUSD	0.00	0.00	0.00	0.00	0.00	0.00	0.02
GUSD	0.03	0.00	0.00	0.00	0.00	0.01	0.13
TUSD	0.01	0.00	0.00	0.00	0.00	0.00	0.02
USDP	0.03	0.00	0.00	0.00	0.00	0.00	0.04

(b) Tron

	mean	p25	p50	p75	p90	p95	p99
USDC	0.00	0.00	0.00	0.00	0.00	0.00	0.00
USDT	0.18	0.00	0.00	0.00	0.04	0.47	5.32
TUSD	0.03	0.00	0.00	0.00	0.00	0.00	0.77

(c) Avalanche

	mean	p25	p50	p75	p90	p95	p99
USDC	0.00	0.00	0.00	0.00	0.00	0.00	0.00
USDT	0.00	0.00	0.00	0.00	0.01	0.03	0.04
TUSD	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Table A.2: Share of Stablecoins Held by Arbitrageurs (Active Arbitrageurs)

Panels (a) to (c) provide statistics about the share of stablecoins owned by active arbitrageurs on the Ethereum, Tron, and Avalance blockchains. For each stablecoin and chain, we show the average and the 25^{th}, 50^{th}, 75^{th}, 90^{th}, and 95^{th}, and 99^{th} percentile of stablecoin ownership for all arbitrageur-month observations with active redemptions or creations. Stablecoin ownership is expressed as a percentage of the total supply of stablecoins.

(a) Ethereum

	mean	p25	p50	p75	p90	p95	p99
USDC	0.00	0.00	0.00	0.00	0.00	0.00	0.00
USDT	0.09	0.00	0.00	0.01	0.06	0.24	1.81
BUSD	0.02	0.00	0.00	0.00	0.00	0.00	0.22
GUSD	0.67	0.00	0.00	0.00	0.13	1.03	36.17
TUSD	0.02	0.00	0.00	0.00	0.00	0.01	0.11
USDP	0.12	0.00	0.00	0.00	0.00	0.01	0.33

(b) Tron

	mean	p25	p50	p75	p90	p95	p99
USDC	0.00	0.00	0.00	0.00	0.00	0.00	0.00
USDT	0.40	0.00	0.00	0.01	0.62	1.09	12.11
TUSD	0.04	0.00	0.00	0.00	0.00	0.35	0.88

(c) Avalanche

	mean	p25	p50	p75	p90	p95	p99
USDC	0.03	0.00	0.03	0.06	0.06	0.07	0.07
USDT	0.01	0.00	0.01	0.02	0.03	0.04	0.04
TUSD	0.00	0.00	0.00	0.00	0.00	0.00	0.01

C.2 Model Extension Allowing for Arbitrageurs' Stablecoin Holdings

We now consider a stripped-down extension of the model in which arbitrageurs can absorb stablecoins at the initial price of \$1 and hold some inventories without redeeming.

In an extreme case in which arbitrageurs are willing to absorb all stablecoin holdings at a fixed price, arbitrageurs can completely eliminate runs, serving as a form of insurance for investors. However, this would require arbitrageurs to absorb much more inventory than they hold in practice, as we show in the previous subsection. Interestingly, in an intermediate case where arbitrageurs are willing to absorb small amounts of inventory at fixed prices, we will show that arbitrageurs can actually exacerbate run risk because first-mover advantage is greater when price impact is lower.

Suppose each of the n arbitrageurs is willing to purchase up to $z_j = h$ units of stablecoin in the secondary market at the initially fixed price of 1 and hold them in inventory. Let $H = nh$ denote the aggregate inventory holding capacity of the arbitrageur sector. This setup can be interpreted as capturing, for instance, implicit or explicit agreements between arbitrageurs and exchanges to provide a minimum quantity of liquidity at a negligible bid-ask spread. Within this inventory range, we model that arbitrageurs simply hold stablecoins in inventory and offer to buy from investors at a fixed price of 1. For arbitrageur purchase quantities exceeding $z_j > h$, arbitrageurs redeem the entire purchase quantity with the issuer, incurring the usual cost $\frac{z_j^2}{2\chi}$ as in the baseline model. We also impose a mild and plausible assumption that $H < 1-\phi$, which ensures that arbitrageurs cannot accumulate a sufficiently large balance sheet to effectively take over the role of the stablecoin issuer in maintaining the stablecoin price when the issuer would have become insolvent; in other words, arbitrageurs will still have to turn to the issuer to redeem if their inventory balance becomes too high despite the capacity and willingness to actively absorb some stablecoin sales.

Under this model extension, it is straightforward to characterize the equilibrium secondary-market price of the stablecoin. If $\lambda \leq H$, arbitrageurs absorb all stablecoins sold by investors without redeeming, and the secondary market price remains fixed at 1. If $\lambda > H$, arbitrageurs begin redeeming their purchases with the issuer, and the pricing dynamics revert to those of the baseline model. The secondary market price of the stablecoin is thus given by:

$$p_2^{\text{ArbHolding}}(\lambda) = \begin{cases} 1 & \lambda \leq H\,, \\ 1 - K\lambda & H < \lambda \leq 1 - \phi\,, \\ \dfrac{1-\phi}{\lambda} - K\lambda & \lambda > 1 - \phi\,. \end{cases} \tag{C.1}$$

Note that the price function (C.1) lies above the baseline price function (4.4) for $\lambda \leq H$, and coincides with it for $\lambda > H$.

Accordingly, the payoff gain from waiting until $t = 3$ rather than selling at $t = 2$ is now given by:

$$h(\lambda) = v_3(\lambda) - p_2^{\text{ArbHolding}}(\lambda) = \begin{cases} \pi(\theta)\left(\dfrac{1-\phi-\lambda}{(1-\phi)(1-\lambda)} + \eta\right) - 1 & \lambda \leq H\,, \\ \pi(\theta)\left(\dfrac{1-\phi-\lambda}{(1-\phi)(1-\lambda)} + \eta\right) - 1 + K\lambda & H < \lambda \leq 1-\phi\,, \\ -\dfrac{1-\phi}{\lambda} + K\lambda & \lambda > 1-\phi\,. \end{cases} \tag{C.2}$$

Similarly, (C.2) lies below the baseline payoff gain function (4.8) for $\lambda \leq H$, and coincides with it for $\lambda > H$.

Intuitively, the arbitrage sector's capacity to absorb redemptions at a fixed price of 1 reduces price impact and temporarily stabilizes the market. However, this stabilization also weakens the disincentive to sell early, thereby reducing investors' incentive to wait until $t = 3$. In effect, the ability of arbitrageurs to buffer modest redemption flows can inadvertently raise run risk by muting early price signals that would otherwise discourage preemptive withdrawals.

To show this result formally, we follow the same logic as in the main text to derive the run threshold under this extension and compare it to the baseline model. As before, each investor's liquidation decision is uniquely determined by her private signal: she sells the stablecoin at $t = 2$ if and only if her signal falls below a threshold. In other words, an investor is indifferent between selling and holding when her signal is exactly at the threshold. Given the existence of a unique run threshold, the investor's indifference condition implies the following revised Laplacian equation:

$$\int_0^H 1\, d\lambda + \int_H^{1-\phi} (1-K\lambda)\, d\lambda + \int_{1-\phi}^1 \left(\frac{1-\phi}{\lambda} - K\lambda\right) d\lambda = \int_0^{1-\phi} \pi(\theta^{**}) \left(\frac{1-\phi-\lambda}{(1-\phi)(1-\lambda)} + \eta\right) d\lambda\,, \tag{C.3}$$

where θ^{**} denotes the new run threshold under the inventory-holding extension.

Comparing the left-hand side of (C.3) with that of (4.9), it follows that the run threshold under the extended model is strictly higher: $\theta^{**} > \theta^*$. This implies that the run risk is greater when arbitrageurs absorb some initial stablecoin sales into inventory at a fixed price without redeeming them, even though this behavior reduces secondary-market price volatility.

The intuition behind this result echoes the main tradeoff highlighted in the baseline model between secondary market price stability and financial stability. When arbitrageurs use their balance sheets to support prices without demanding a discount, arbitrage becomes more efficient and price impact declines. However, this also means investors receive higher payoffs from selling early, while their payoff from holding until maturity remains unchanged. As a result, the incentive to redeem early increases, exacerbating run risk. Thus, somewhat counterintuitively, allowing arbitrageurs to stabilize prices through more active inventory holdings can worsen financial fragility, consistent with the core message of the baseline model that greater price stability can come at the cost of increased run risks.

D Strategic Bidding

In this appendix, we extend our baseline model to incorporate imperfect competition between and strategic bidding by arbitrageurs. The framework closely follows Klemperer and Meyer (1989), with three key differences: (i) demand is perfectly inelastic, (ii) payoffs are slightly modified by the possibility of runs, and (iii) we consider the case of $n > 2$ arbitrageurs, which is briefly analyzed in Proposition 8a of Klemperer and Meyer (1989). We show that under imperfect competition and strategic bidding, there may exist multiple equilibrium prices for any given set of primitives. Despite these technical complexities, the core economic insight that arbitrage capacity influences secondary stablecoin prices remains unchanged after imposing a reasonable equilibrium selection mechanism. Thus, for tractability, we assume competitive bidding in the main text without loss of generality.

To start, as in the main text, we assume there are n arbitrageurs, who each face a cost

$$\frac{z_j^2}{2\chi}$$

for arbitraging z_j units of the stablecoin. As in the main text, the redemption price p_r that arbitrageurs receive per unit of the stablecoin redeemed depends on the total amount λ redeemed due to the illiquidity of the reserve asset:

$$p_r = \begin{cases} 1 & \lambda \leq 1 - \phi\,, \\ \frac{1-\phi}{\lambda} & \lambda > 1 - \phi\,. \end{cases}$$

From the perspective of arbitrageurs, we view the demand shock λ as a random variable. We assume λ has full support on $[0, 1]$. While in the limit of $\varepsilon_i \to 0$, consumers will perfectly coordinate on running or not running, on the path to the limit, λ has positive support everywhere on $[0, 1]$. Since investors make market orders, the demand shock is perfectly inelastic.

In contrast to the main text, we now assume arbitrageurs bid strategically, taking into account the influence of their bids on equilibrium prices. We search for a symmetric equilibrium bid curve $z(p)$, where $z(p)$ is a strictly decreasing function defining the quantity arbitrageurs are willing to arbitrage at price p. We will sometimes work with the inverse bid curve, which we will write as

$p(z)$, defined as:

$$p(\tilde{z}) = \{p : z(p) = \tilde{z}\} \; .$$

Since the maximum value of λ is 1, with n arbitrageurs, we require $p(z)$ to be defined only on the range $z \in \left[0, \frac{1}{n}\right]$, since in a symmetric equilibrium each arbitrageur absorbs at most a quantity $\frac{1}{n}$. As in Klemperer and Meyer (1989), we require that each arbitrageur finds $z(p)$ to be an ex-post best response conditional on any realization of demand λ.

We next characterize the first-order conditions that define the equilibrium bid curves. Suppose the equilibrium bid curve is $z(p)$. Let p^* represent the equilibrium market clearing price assuming symmetric bids, which is related to λ according to:

$$nz(p^*) = \lambda \, . \tag{D.1}$$

Assume all arbitrageurs bid according to $z(\cdot)$, and fix some value of p^*; this is equivalent to fixing a value of the demand shock λ, through (D.1). For this value p^*, suppose arbitrageur i considers deviating from the symmetric equilibrium bid curves $z(p)$. If $n-1$ arbitrageurs bid according to $z(p)$, given p^*, the last arbitrageur i faces a residual supply curve – a menu of pairs of prices p and quantities, which we will write as $\tilde{z}_i(p;p^*)$ – which satisfies:

$$\tilde{z}_i(p;p^*) = nz(p^*) - (n-1)\,z(p) \; . \tag{D.2}$$

That is, the quantity available to i if the price is p is total demand λ, which by the definition of p^* in (D.1) is equal to $nz(p^*)$, minus total quantity purchased by other arbitrageurs, $(n-1)\,z(p)$.

Now, first, suppose the aggregate demand shock $\lambda < 1-\phi$, so the issuer is solvent, and stablecoins are redeemed for a dollar each. Taking into account arbitrage costs, i's profit if the price is p is:

$$(1-p)\,\tilde{z}_i(p;p^*) - \frac{\tilde{z}_i(p;p^*)^2}{2\chi} \, . \tag{D.3}$$

That is, i makes $(1-p)$ per unit times her quantity $\tilde{z}_i(p;p^*)$, less the arbitrage cost $\frac{\tilde{z}_i(p;p^*)^2}{2\chi}$. In a symmetric equilibrium, for any p^*, arbitrageur i must find it optimal to absorb exactly $\frac{1}{n}$ of total inventory, which is only possible from (D.2) if $p = p^*$. Thus, in order for $z(p)$ to form a symmetric equilibrium, we must have:

$$p^* = \arg\max_p \,(1-p)\,\tilde{z}_i(p;p^*) - \frac{\tilde{z}_i(p;p^*)^2}{2\chi} \, .$$

Differentiating with respect to p, and requiring the derivative to equal 0 at $p = p^*$, we have:

$$z'(p) = -\frac{z}{(n-1)\left(1 - \frac{z}{\chi} - p^*\right)} \, , \tag{D.4}$$

where we used (D.2) to differentiate $\tilde{z}_i(p;p^*)$. Applying the inverse function theorem, this implies that:

$$p'(z) = -(n-1)\left(\frac{1}{z} - \frac{p}{z} - \frac{1}{\chi}\right) \;\forall z < \frac{1-\phi}{n}\,. \tag{D.5}$$

Then, suppose $\lambda > 1-\phi$, so that the issuer is insolvent. The redemption value of each unit of the coin is then $\frac{1-\phi}{\lambda}$. The residual supply curve is still (D.2). However, (D.3), the value to the arbitrageur if the price is p, now becomes:

$$\left(\frac{1-\phi}{nz(p^*)} - p\right)\tilde{z}_i(p;p^*) - \frac{\tilde{z}_i(p;p^*)^2}{2\chi}\,.$$

Relative to (D.3), the redemption value is modified to only $\frac{1-\phi}{nz(p^*)}$ per unit of the coin redeemed. Differentiating with respect to p, and setting to 0 at p^*, the equivalent of (D.4) is:

$$-z(p^*) - (n-1)\,z'(p^*)\left(\frac{1-\phi}{nz(p^*)} - p^* - \frac{z(p^*)}{\chi}\right) = 0\,.$$

Rearranging and inverting:

$$p'(z) = -(n-1)\left(\frac{1-\phi}{nz^2} - \frac{p}{z} - \frac{1}{\chi}\right) \;\forall z \geq \frac{1-\phi}{n}\,. \tag{D.6}$$

Writing conditions (D.5) and (D.6) together, we find first-order necessary conditions for inverse demand curves $p(z)$ to constitute a symmetric equilibrium:[32]

$$p'(z) = \begin{cases} -(n-1)\left(\frac{1}{z} - \frac{p}{z} - \frac{1}{\chi}\right) & z < \frac{1-\phi}{n}\,, \\ -(n-1)\left(\frac{1-\phi}{nz^2} - \frac{p}{z} - \frac{1}{\chi}\right) & z \geq \frac{1-\phi}{n}\,. \end{cases} \tag{D.7}$$

Expression (D.7) is simply a first-order differential equation, specifying the derivative $p'(z)$ as a function of z as well as p. This implies that there is a simple way to find candidate equilibria: we can pick any point (p_0, z_0), and solve (D.7) to find a function which satisfies the differential equation on $z = \left[0, \frac{1}{n}\right]$ and passes through (p_0, z_0). The top-left panel of Figure A.5 depicts the aggregate supply curves generated by a number of such equilibria, in green, for fixed values of n, χ, ϕ, where we plot $\lambda = nz$ on the x-axis. There is a range of different equilibria with different levels of prices.

[32] We have not verified that the SOCs are guaranteed to hold; however, in a very similar model to ours, Klemperer and Meyer (1989) show that the FOCs holding and the bid slope $p'(z) > 0$ are sufficient conditions for the second-order conditions to hold globally. The equilibrium inverse demand functions $p(z)$, substituting $\lambda = nz$, define an aggregate arbitrage supply curve facing consumers. If consumers demand λ, then each arbitrageur must absorb quantity $z = \frac{\lambda}{n}$ in symmetric equilibrium, meaning that the price will be $p\left(\frac{\lambda}{n}\right)$.

Following Klemperer and Meyer (1989), we can characterize the set of equilibria as follows. Since all equilibria must satisfy that $p'(z) < 0$ throughout, we can find an upper bound to all equilibrium price curves by finding a function which satisfies $p'(z) = 0$, that is, which exactly satisfies the second-order condition when the demand shock $\lambda = 1$. This function is illustrated as the dashed gray line in the top left panel of Figure A.5. We can further solve for a "highest" possible equilibrium, which can be characterized by finding the equilibrium, which exactly reaches $p'\left(\frac{1}{n}\right) = 0$; this is the red line in the top left panel of Figure A.5. All other valid equilibrium bid curves lie below this one, as the figure illustrates.

Figure A.5

These plots illustrate solutions to the FOCs in (D.7), in various settings. The x axis shows demand $\lambda = nz(p)$, and the y-axis shows the inverse aggregate supply function $p\left(\frac{\lambda}{n}\right)$. The top left plot shows, for fixed parameters n, χ, ϕ, a number of different equilibria in green; the highest possible equilibrium, satisfying $p'\left(\frac{1}{n}\right) = 0$, in red; and the linear equilibrium conjecture from (D.8), which fails to be an equilibrium, in blue. The top right plot illustrates equilibrium for different values of χ, passing through a single (p, λ) pair. The bottom plot shows the highest possible equilibria for different values of χ.

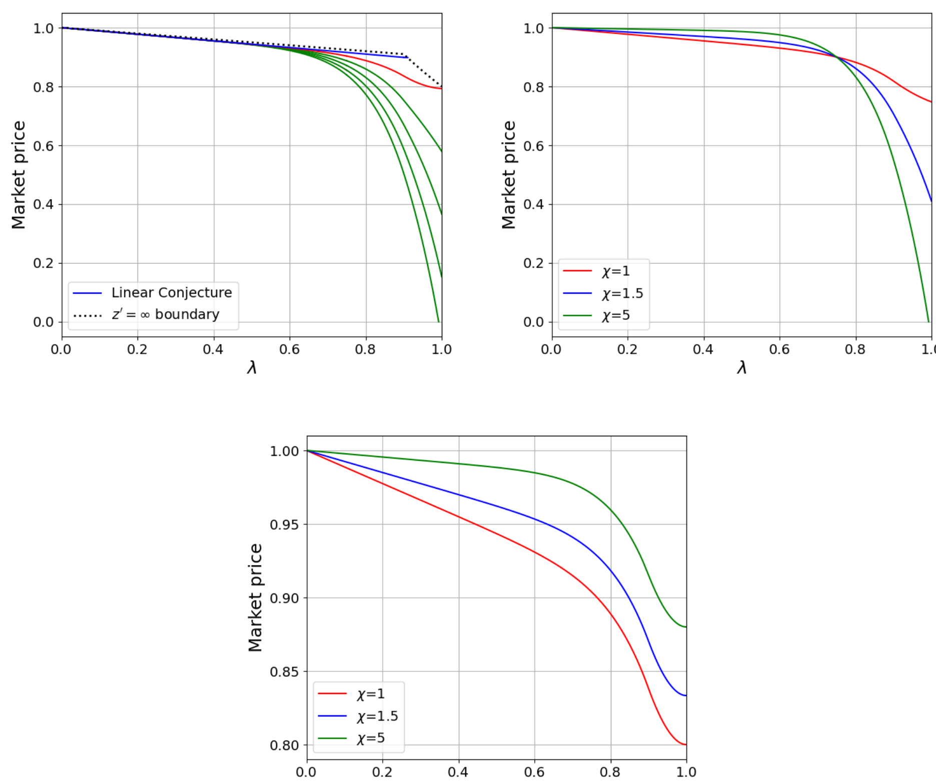

Note that a natural conjecture for the equilibrium bid curve – that $p(z)$ follows the well-known linear equilibrium bids under linear marginal costs for $z(p) < \frac{1-\phi}{n}$, that is,

$$p(\lambda) = 1 - \frac{n-1}{n-2} \cdot \frac{\lambda}{n\chi} \tag{D.8}$$

and is nonlinear thereafter, is in general *not* an equilibrium in this setting. The reason for this is that, in the Klemperer and Meyer (1989) setting, there are interestingly non-local effects of marginal costs for large values of $p(z)$, on equilibria for small values of $p(z)$. Essentially, an equilibrium bid curve satisfies second-order conditions when its slope does not become 0 or ∞. We can solve the differential equation for $p(z)$ beginning at the linear solution (D.8), and it is linear until $z = \frac{1-\phi}{n}$. However, for higher values of z, this differential equation reaches slope $p'(z) = 0$ before $z = \frac{1}{n}$, thus violating the second-order conditions for equilibria. In the figure, this potential equilibrium is illustrated as the blue line; which intersects the dotted boundary well before $\lambda = 1$. This argument also illustrates that the standard equilibrium selection criteria that equilibrium bid curves are linear, imposed in much of the applied finance literature, is not, in fact, innocuous when the demand shock has finite support. In this setting, there is actually a multiplicity of valid nonlinear bidding equilibria even when marginal costs are fully linear, illustrated by the various trajectories to the left of $z = \frac{1-\phi}{n}$ in the figure.

Comparative statics are generally difficult to obtain due to the multiplicity of equilibria under strategic bidding. However, we can derive a counterpart to our main result that lower values of χ lead to less aggressive bidding and greater price impact under an equilibrium selection approach commonly used in the auction literature. Since the linear equilibrium commonly used in applied settings is not always valid in our framework, a natural selection criterion is to choose the linear equilibrium if it exists and does not reach $p'(z) = 0$ before $z = \frac{1}{n}$. Otherwise, we select the *closest possible* equilibrium to the linear equilibrium. This approach is straightforward to characterize: if the linear equilibrium fails to exist, it does so because it lies "above" the set of feasible equilibria, as illustrated in the top-left panel of Figure A.5. Thus, whenever the linear equilibrium is invalid, the closest alternative, measured, for example, by the L^p norm on functions for $1 \leq p \leq \infty$ over the domain where the linear equilibrium is valid is simply the *highest possible* equilibrium, depicted by the red curve in the top-left panel. We refer to this as the "closest-to-linear" equilibrium selection.

Under this equilibrium selection approach, we can cleanly rank equilibria for different values of χ to facilitate comparative statics. Consider two values $\chi_1 > \chi_2$, such that the linear equilibrium is not valid in both cases. We argued that the upper bound of all equilibrium bid curves has exactly $p'(z) = 0$ at $z = \frac{1}{n}$; we can explicitly solve for the price which implements this equilibrium by setting (D.6) to 0 for $z = \frac{1}{n}$ and solving:

$$\frac{1-\phi}{n\left(\frac{1}{n}\right)^2} - \frac{p}{\left(\frac{1}{n}\right)} - \frac{1}{\chi} = 0 \tag{D.9}$$

The price that solves (D.9) is clearly monotone in χ. This shows that, when we apply the "closest-to-linear" equilibrium selection, the equilibrium price at $z = \frac{1}{n}$ must be higher under χ_1 than χ_2.

We further show that the equilibrium price is weakly greater under χ_1 than χ_2 for *any* quantity $z \in \left[0, \frac{1}{n}\right]$. To prove this, assume for contradiction that there exists some $z \in \left(0, \frac{1}{n}\right)$ such that the χ_1-equilibrium price is lower than the χ_2-equilibrium price. Since the χ_1-equilibrium price is also lower at $z = \frac{1}{n}$, and both equilibrium price curves are continuous, this implies that the χ_1-equilibrium price curve must cross the χ_2-equilibrium price curve from below at some point. However, this is impossible due to a single-crossing property of the differential equations characterizing equilibrium. Specifically, Equations (D.5) and (D.6) are strictly monotone in χ, holding p and z fixed. Thus, at any given pair (p, z), the slope $p'(p, z)$ must be more negative for larger χ. Consequently, if $\chi_1 > \chi_2$, the function $p(z; \chi_1)$ can only cross $p(z; \chi_2)$ from above as z increases. This single-crossing property is illustrated in the top-right panel of Figure A.5, where equilibrium bid curves passing through a single point exhibit steeper slopes for higher values of χ, consistently crossing lower-χ curves from above. This contradiction implies that for $\chi_1 > \chi_2$, the closest-to-linear equilibrium under χ_1 must yield a higher price than the closest-to-linear equilibrium under χ_2 for all $z \in \left[0, \frac{1}{n}\right]$. We illustrate this result in the bottom panel of Figure A.5, where the closest-to-linear equilibria for different values of χ are plotted. The curves become completely flat at $z = \frac{1}{n}$, with higher prices corresponding to higher values of χ, and they never cross each other.

To summarize, we have shown that under a "closest-to-linear" equilibrium selection mechanism inspired by the double auction literature, lower values of χ lead to less aggressive bidding and greater price impact. This finding suggests that, the core economic insight that arbitrage capacity influences secondary stablecoin prices remains intact with an appropriate equilibrium selection mechanism. Specifically, our equilibrium selection approach attempts to stay as close as possible to the linear equilibrium implicitly selected in the double-auctions literature.

E Microfoundation of Competitive Bidding Benchmark

In the main text, we derived the market inverse demand function in Lemma 1 by arguing that the price must equal investors' marginal cost of absorbing additional quantity while holding total redemption quantity fixed, without explicitly solving for equilibrium in a demand function submission game. In this appendix, we explicitly derive this demand curve as a simple rational expectations equilibrium in a demand function submission game with competitive arbitrageurs.

Suppose there is a measure n of competitive arbitrageurs indexed by $j \in [0, n]$. As in the main text, arbitrageurs incur a cost of $\frac{z_j^2}{2\chi}$ when purchasing and redeeming z_j units of the asset. Each arbitrageur submits a demand schedule $z_j(p_2)$, representing the quantity they are willing to purchase and redeem at secondary market price p_2. In a symmetric equilibrium where all arbitrageurs bid the same demand curve $z(p_2)$, the aggregate market demand curve is given by:

$$Z(p_2) = \int_0^n z(p_2)\, dj = nz(p_2) \ .$$

In equilibrium, the inelastic quantity of investors' market-order redemption demand, λ, must equal $Z(p_2)$. Here, we write $Z(p_2)$ to conceptually distinguish arbitrageurs' supply from investors' redemption demand.

In purely private-valued models, the competitive bidding benchmark simply involves arbitrageurs bidding up to the point where the price equals their marginal utility for the asset. Our model follows this structure in the range where $Z(p_2) \leq 1 - \phi$ and the issuer remains solvent. However, the setting becomes more complex when $Z(p_2) > 1 - \phi$, as the issuer is then insolvent, meaning the redemption value per stablecoin depends on the total quantity redeemed. Consequently, investors can anticipate the redemption value based on the price.

To address this, we seek a rational expectations equilibrium in which arbitrageurs' demand schedules, $z(p_2)$, are best responses to the induced aggregate demand schedule, $Z(p_2) = nz(p_2)$, while correctly inferring the redemption value of the stablecoin from the price p_2.

First, consider values of p_2 such that $z_j(p_2) \leq \frac{1-\phi}{n}$, ensuring that total redemption quantity satisfies $Z(p_2) \leq 1 - \phi$. In this range, the issuer remains solvent, so each arbitrageur j correctly perceives that they can redeem each unit of stablecoin for one dollar with the issuer. The arbitrageur's problem is thus purely private-valued and straightforward to solve.

Each arbitrageur's demand schedule must maximize utility, given other arbitrageurs' demand schedules and a correct forecast of redemption values. Suppose the market price is p_2. Since arbitrageurs are competitive, they take prices as given and do not account for their impact on the market. The utility of an arbitrageur purchasing and redeeming quantity $\tilde{z}$ at a fixed price p_2, when the redemption value is 1, is:

$$\tilde{z}(1 - p_2) - \frac{\tilde{z}^2}{2\chi}.$$

Differentiating with respect to $\tilde{z}$ and setting the derivative to zero, the demand schedule $z(p_2)$ must satisfy:

$$1 - p_2 - \frac{z(p_2)}{\chi} = 0,$$

which implies:

$$z_j(p_2) = \chi(1 - p_2).$$

This is the standard demand curve under private values, where the arbitrageur bids the inverse of her marginal utility function. The aggregate demand curve is then:

$$Z(p_2) = nz_j(p_2) = n\chi(1 - p_2).$$

Inverting, in a rational expectations equilibrium, the price $p_2(Z)$ satisfies:

$$p_2(Z) = 1 - \frac{Z}{n\chi}, \quad \forall Z \leq 1 - \phi. \tag{E.1}$$

The situation becomes more complex when $z(p_2) > \frac{1-\phi}{n}$, so that $Z(p_2) > 1 - \phi$. In this range, the stablecoin's redemption value to an arbitrageur depends on the behavior of other arbitrageurs, since a higher total redemption quantity lowers the per-stablecoin redemption value. In a rational expectations equilibrium, each arbitrageur correctly infers the redemption value from the price, anticipating that if others bid $z(p_2)$ and the price is p_2, then the total redemption quantity is $nz(p_2)$, and the per-stablecoin redemption value is:

$$\frac{1 - \phi}{Z(p_2)} = \frac{1 - \phi}{nz(p_2)}.$$

Under this redemption value, the arbitrageur's utility from purchasing quantity $\tilde{z}$ at price p_2 is:

$$\tilde{z}\left(\frac{1 - \phi}{nz(p_2)} - p_2\right) - \frac{\tilde{z}^2}{2\chi}.$$

A rational expectations equilibrium requires that each arbitrageur bidding $z(p_2)$ is playing a best response for all p_2, meaning:

$$z(p_2) = \arg\max_{\tilde{z}} \tilde{z}\left(\frac{1 - \phi}{nz(p_2)} - p_2\right) - \frac{\tilde{z}^2}{2\chi}, \tag{E.2}$$

whenever $z(p_2) > \frac{1-\phi}{n}$. Differentiating the right-hand side of (E.2) with respect to $\tilde{z}$, substituting $z(p_2)$ for $\tilde{z}$, and setting the derivative to zero, we obtain:

$$\frac{1 - \phi}{nz(p_2)} - p_2 - \frac{z(p_2)}{\chi} = 0. \tag{E.3}$$

Thus, an rational expectations equilibrium demand curve $z(p_2)$ must satisfy (E.3) whenever $Z > 1 - \phi$, i.e., $z(p_2) > \frac{1-\phi}{n}$. The equation (E.3) defines an inverse demand curve:

$$p_2(z) = \frac{1 - \phi}{nz} - \frac{z}{\chi}, \quad \forall z > \frac{1 - \phi}{n}.$$

Substituting $Z = nz$, we obtain:

$$p_2(Z) = \frac{1 - \phi}{Z} - \frac{Z}{n\chi}, \quad \forall Z > 1 - \phi. \tag{E.4}$$

Equations (E.1) and (E.4) together constitute (4.4) in Lemma 1.

F Joint Design of Liquidity and Arbitrage Concentration

In this appendix, we explore a model extension where the stablecoin issuer concurrently determines the levels of liquidity transformation ϕ and arbitrage concentration n, both at $t = 0$. We show that issuers' optimal choice of ϕ is monotone in a specific ordering on the function $R(\phi)$, which determines the returns from holding illiquid assets: when the illiquidity premium becomes higher, in a sense we will define, the gains from increasing ϕ are larger, so issuers optimally choose higher values of ϕ.

Note that, without other market participants taking any action at $t = 0$, this joint optimization problem is equivalent to a sequential decision problem in which the issuer first decides the optimal level of liquidity transformation ϕ, and then decides the optimal arbitrage concentration n as analyzed in the baseline model.

Hence, in the extended joint optimization problem, the issuer's objective can be written as, from (4.14):

$$\max_{n} E[\Pi] = G(E[W]) \int_{\pi(\theta)\geq\pi(\theta^*)} \pi(\theta) \left(R(\phi) - 1 - \tau\right) dF(\theta) \,.$$

Factoring, and ignoring τ, we have:

$$\max_{\phi} \max_{n} \left(R(\phi) - 1\right) G(E[W]) \int_{\pi(\theta)\geq\pi(\theta^*)} \pi(\theta) dF(\theta)$$

$$= \max_{\phi} \left[\left(R(\phi) - 1\right) \max_{n} \left[G\left(EW\left(\phi, n\right)\right) \int_{\pi(\theta)\geq\pi(\theta^*(n,\phi))} \pi(\theta) dF(\theta) \right] \right] \,.$$

Define the function:

$$F\left(\phi, n\right) \equiv G\left(EW\left(\phi, n\right)\right) \int_{\pi(\theta)\geq\pi(\theta^*(n,\phi))} \pi(\theta) dF(\theta) \,.$$

Note that there is no dependence on $R\left(\phi\right)$. Also define:

$$Q\left(\phi\right) \equiv R\left(\phi - 1\right) \,.$$

The objective function is then:

$$\max_{\phi} \left[\left(R(\phi) - 1\right) \max_{n} \left[G\left(EW\left(\phi, n\right)\right) \int_{\pi(\theta)\geq\pi(\theta^*(n,\phi))} \pi(\theta) dF(\theta) \right] \right]$$

$$= \max_{\phi} \left[Q\left(\phi\right) \max_{n} F\left(\phi, n\right) \right] \,,$$

which can be restated by taking logs of the objective function:

$$\arg\max_{\phi} \left[\log\left(Q\left(\phi\right)\right) + \log \max_{n} F\left(\phi, n\right) \right] \,.$$

We have the following formal result about the optimal choice of liquidity transformation with respect to the revenue function R, which can be understood as capturing the asset market that the stablecoin issuer has access to. We highlight that the result in Proposition 3 still holds, that is, the optimal arbitrage concentration can be sequentially determined as the optimal liquidity transformation is pinned down.

Proposition 5. *Suppose there are stablecoins issuers 1 and 2, with different revenue functions R_1 and R_2, with the "monotone increasing ratios" property as stated below. For any $\phi' > \phi$, suppose:*

$$\frac{R_1(\phi') - 1}{R_1(\phi) - 1} \geq \frac{R_2(\phi') - 1}{R_2(\phi) - 1}. \tag{F.1}$$

Then, the optimal ϕ^* is always greater for issuer 1.

Intuitively, Proposition 5 implies that the observed variations in liquidity transformation across different stablecoins can be justified by the access some issuers have to asset markets that command higher illiquidity premiums. In Proposition 5, stablecoin issuer 1 could be understood as USDT, while issuer 2 USDC. Condition (F.1) in Proposition 5 means that the asset market that USDT has access to pocesses a higher illiquidity premium in that a more illiquid asset holding offers a higher expected return of the asset. Proposition 5 then predicts that USDT optimally chooses to hold more illiquid assets, that is, transforms more liquidity. To accomodate a larger amount of endogenous liquidity transformation, Proposition 3 further implies that USDT admits a more concentrated arbitrage sector, as we highlighted above.

An issuer's optimal choice of ϕ is also affected by other model parameters, such as the long-term benefit η and the demand function $G(\cdot)$. However, we could not able to prove clean monotonicity results regarding the relationship between the optimal ϕ and these parameters. Technically, the issue is that the expected welfare function $EW(\phi, n)$ is not monotone in ϕ, meaning that the monotone comparative statics approach we use to prove Proposition 5 cannot be applied.

G Redemption Fees and Gates

Consistent with the idea of opening up primary market redemptions to more stablecoin holders, we show in this appendix that imposing redemption fees can reduce stablecoin run risk but faces the same tradeoff between price stability and run risk. Redemption fees render redemptions more costly and thereby increase the constraints to arbitrage. Less efficient arbitrage helps to reduce run risk by raising the price impact from selling stablecoins but hampers price stability at the same time. In the model, we can think of redemption fees as the stablecoin issuer charging a fee of ν per stablecoin created or redeemed. The following result characterizes the effect of the redemption fee on the run risk of the stablecoin.

Proposition 6. *Suppose ϕ is small in the sense that (4.11) holds. Suppose we impose an exogenous creation/redemption fee ν per coin created/redeemed, which is paid by arbitrageurs to the issuer, holding fixed arbitrage capacity K. There exists a unique threshold equilibrium in which the run threshold $\pi(\theta^*; \nu)$ decreases in ν, implying that the run risk uniformly decreases as the redemption fee increases.*

The effect of redemption gates resembles that of redemption fees. In our model, redemption gates can be effectively captured by the stablecoin issuer's choice of arbitrage concentration n in our model. To see this logic, note that a full redemption gate implies that effectively no arbitrageur can redeem the stablecoin, that is, $n = 0$. Consequently, the use of a full redemption gate implies that run risk would be eliminated at the cost of an exceedingly volatile secondary market. If the issuer adopts a selective redemption gate, we can think of an effective arbitrage concentration $n' \leq n$. As Proposition 2 shows, a more selective redemption gate leads to a lower run risk at $t = 2$, but at the expense of worse price stability.

Taken together, our model highlights that regulations for stablecoin redemptions should carefully trade off the effects of price stability and run risk. Policies that require unconstrained arbitrage and redemptions improve price stability but at the expense of worsening run risk; policies restricting redemptions through fees and gates improve run risk at the cost of price stability. While price stability is observable on a daily basis and run risk only materializes in tail events, both are essential considerations for the optimal regulation of stablecoins.

H Omitted Proofs

Proof of Proposition 1. Denote the run threshold as θ', that is, if investor i observes a private signal $s_i < \theta'$ she sells her stablecoin at $t = 2$; otherwise she waits until $t = 3$. Then the population of investors who run, λ, can be written as

$$\lambda(\theta, \theta') = \begin{cases} 1 & \text{if } \theta \leq \theta' - \varepsilon \\ \frac{\theta' - \theta + \varepsilon}{2\varepsilon} & \text{if } \theta' - \varepsilon < \theta \leq \theta' + \varepsilon \\ 0 & \text{if } \theta > \theta' + \varepsilon \end{cases} . \tag{H.1}$$

Let $h(\theta, \lambda)$ be the payoff gain from waiting until $t = 3$ versus selling at $t = 2$. It is straightforward that

$$h(\theta, \lambda) = v_3(\theta, \lambda) - p_2(\theta, \lambda) = \begin{cases} \pi(\theta) \left(\dfrac{1 - \phi - \lambda}{(1 - \phi)(1 - \lambda)} + \eta \right) - 1 + K\lambda & \lambda \leq 1 - \phi \, , \\ -\dfrac{1 - \phi}{\lambda} + K\lambda & \lambda > 1 - \phi \, . \end{cases}$$

Notice that $h(\theta, \lambda)$ is concave in λ over $(0, 1 - \phi)$ because

$$\frac{\partial^2 h(\theta, \lambda)}{\partial \lambda^2} = -\frac{2\pi(\theta)\phi}{(1 - \lambda)^3(1 - \phi)} < 0 \, .$$

If investor i observes signal s_i, given that other households use the threshold strategy, she will sell her stablecoin if

$$\int_{s_i - \varepsilon}^{s_i + \varepsilon} h\left(\theta, \lambda\left(\theta, \theta'\right)\right) d\theta < 0 \, ,$$

or stay otherwise. To prove that there exists a unique run threshold θ^*, we need to prove that there is a unique θ^* such that if $\theta' = \theta^*$, the investor who observes signal $s_i = \theta' = \theta^*$ is indifferent between selling and waiting. That is,

$$V\left(\theta^*\right) \equiv \int_{\theta^* - \varepsilon}^{\theta^* + \varepsilon} h\left(\theta, \lambda\left(\theta, \theta^*\right)\right) d\theta = 0.$$

According to Morris and Shin (2003) and Goldstein and Pauzner (2005), it then suffices to show that $h(\lambda)$ crosses 0 only once, that is, satisfies the single-crossing property when the upper dominance region exists. To show this, first note that $h(1) = -1 + \phi + K < 0$ and that $h(\lambda)$ increases in $(1-\phi, 1)$. It then must be that $h(1-\phi) < 0$. On the other hand, note that $h(0) > 0$ when θ, and thus $\pi(\theta)$, are sufficiently large. Because $h(\lambda)$ is continuous and concave in $(0, 1 - \phi)$, it then immediately follows that $h(\lambda)$ must cross 0 once and only once in $(0, 1-\phi)$. Since $h(\lambda)$ does not cross 0 in $(1 - \phi, 1)$, this implies that $h(\lambda)$ crosses 0 once and only once in $(0, 1)$, concluding the proof. ∎

Proof of Proposition 2. Based on Proposition 1, we first compute the run threshold $\pi(\theta^*)$ directly. By construction, an investor with signal θ^* must be indifferent between selling her stablecoin at $t = 2$ and waiting until $t = 3$. This investor's posterior belief of θ is uniform over the interval $[\theta^* - \varepsilon, \theta^* + \varepsilon]$. On the other hand, she understands that the proportion of investors who sell at $t = 2$, as a function of θ, is $\lambda(\theta, \theta^*)$, where the function $\lambda(\theta, \theta')$ is given by (H.1) in the proof of Proposition 1. Therefore, her posterior belief of λ is also uniform over $(0, 1)$. At the limit, this gives the indifference condition as the Laplace condition:

$$\int_0^{1-\phi} (1 - K\lambda) \, d\lambda + \int_{1-\phi}^{1} \left(\frac{1 - \phi}{\lambda} - K\lambda \right) d\lambda = \int_0^{1-\phi} \pi(\theta^*) \left(\frac{1 - \phi - \lambda}{(1 - \phi)(1 - \lambda)} + \eta \right) d\lambda \, , \quad \text{(H.2)}$$

which we also give in the main text as (4.9). Solving this Laplace condition (H.2) yields the run threshold (4.10).

We then perform comparative statics about the run threshold $\pi(\theta^*)$. With respect to ϕ, we have

$$\frac{\partial \pi(\theta^*)}{\partial \phi} = \frac{(2-2\phi-K)((\phi-1)(\eta(\phi-1)+1)-\ln\phi)-2(\phi-1)\ln(1-\phi)(-2\phi+(\phi+1)\ln\phi+2)}{2\left((1-\phi)(1+\eta(1-\phi))+\phi\ln\phi\right)^2}, \tag{H.3}$$

whose denominator is positive. Thus, (H.3) is positive if its numerator is positive. This holds when

$$g(\phi) \equiv \frac{2(\phi-1)\left(\phi-\ln\phi+\ln(1-\phi)((1+\phi)\ln\phi+2-2\phi)-1\right)}{1-\phi+\ln\phi} > K\,, \tag{H.4}$$

where $g(\phi)$ is continuous and strictly decreasing in ϕ, and it satisfies $\lim_{\phi\to 0} g(\phi) = 2 > 0$. Thus, conditions (H.3) and (H.4) hold when ϕ is sufficiently small for any given $K \leq 2$, and then the equilibrium run threshold $\pi(\theta^*)$ increases in ϕ.

With respect to K, we have

$$\frac{\partial \pi(\theta^*)}{\partial K} = \frac{\phi-1}{2\left((1-\phi)(1+\eta(1-\phi))+\phi\ln\phi\right)} < 0\,. \tag{H.5}$$

To see why (H.5) holds, notice that its numerator is negative. On the other hand, define its denominator as

$$\zeta(\phi) \equiv 2\left((1-\phi)(1+\eta(1-\phi))+\phi\ln\phi\right)\,.$$

It is straightforward to show that $\zeta(\phi)$ strictly decreases in ϕ while $\lim_{\phi\to 1}\zeta(\phi) = 0$ when $\eta = 0$. Thus, the denominator of (H.5) is positive. This concludes the proof. ∎

Proof of Proposition 3. Recall that the issuer's objective function is

$$\max_n\; E[\Pi] = \underbrace{G(E[W])}_{\substack{\text{population of}\\ \text{participating investors}}} \underbrace{\int_{\pi(\theta)\geq\pi(\theta^*)} \pi(\theta)(R(\phi)-1)dF(\theta)}_{\substack{\text{expected issuer revenue per}\\ \text{participating investor}}}\,,$$

as defined in (4.14). Because $R(\phi) > 1$ and at the same time $R(\phi)$ does not affect $G(E[W])$, the issuer effectively considers the reduced problem below:

$$\max_n\; E[\Pi] = G(E[W]) \int_{\pi(\theta)\geq\pi(\theta^*)} \pi(\theta)dF(\theta)\,, \tag{H.6}$$

We now consider the first-order condition (FOC) for the issuer's problem (H.6) that determines the optimal K, the slope of arbitrageurs' demand. When $G(\cdot)$ is linear, the FOC is:

$$0 = \frac{\partial E[\Pi]}{\partial K} = \underbrace{\frac{\partial E[W]}{\partial K} \int_{\pi(\theta) \geq \pi(\theta^*)} \pi(\theta) dF(\theta)}_{\substack{\text{marginal cost from} \\ \text{reduced investor participation}}} - \underbrace{E[W] \frac{\partial \pi(\theta^*)}{\partial K} (f(\theta^*)\pi(\theta^*))}_{\substack{\text{marginal benefit from} \\ \text{reduced run risk}}}, \tag{H.7}$$

where according to (4.15),

$$\frac{\partial E[W]}{\partial K} = \underbrace{-2\alpha\delta^2 K}_{\substack{\text{marginal utility cost from} \\ \text{decreasing price stability}}} + \underbrace{\frac{\partial \pi(\theta^*)}{\partial K} (f(\theta^*)(1 - \phi - K - \pi(\theta^*)(1 + \eta)))}_{\substack{\text{marginal utility benefit from} \\ \text{increasing financial stability}}} - \int_{\pi(\theta) < \pi(\theta^*)} dF(\theta) . \tag{H.8}$$

This first-order condition reveals the various channels through which increasing K affects the stablecoin issuer's expected revenue. The first part of (H.7) captures the marginal effect of changing the population of participating investors, which in turn depends on each investor's expected utility from participating. The second part of (H.7) captures the marginal benefit that directly results from the reduced run risk on issuer revenue (since the issuer captures the revenue only if a run is avoided). Furthermore, (H.8) captures the marginal effects of increasing K on an investor's expected utility: the first term of (H.8) is the marginal cost that results from higher price fluctuations, while the second term is the marginal benefit from the reduced run risk on investor utility. Notice that this last marginal benefit then indirectly affects the issuer's expected revenue. In equilibrium, the issuer cares about run risk both directly and indirectly, which are captured by the second term of (H.7) and the second term of (H.8), respectively.

Now, suppose condition (H.4) holds, that is, ϕ is sufficiently small. Under this condition, we know from condition (H.3) in the proof of Proposition 2 that the equilibrium run threshold $\pi(\theta^*)$ increases in ϕ, that is $\partial \pi(\theta^*)/\partial \phi > 0$. We can now compute $dK^*/d\phi$ and sign the respective derivatives below. Specifically, using the FOC (H.7) above:

$$\begin{aligned}
\frac{\partial FOC_K(K, \phi)}{\partial \phi} = & \underbrace{\frac{\partial^2 E[W]}{\partial K \partial \phi}}_{+} \underbrace{\int_{\pi(\theta) \geq \pi(\theta^*)} \pi(\theta) dF(\theta)}_{+} \\
& - \underbrace{\pi(\theta^*) f(\theta^*)}_{+} \underbrace{\left(\frac{\partial E[W]}{\partial K} \frac{\partial \pi(\theta^*)}{\partial \phi} + \frac{\partial E[W]}{\partial \phi} \frac{\partial \pi(\theta^*)}{\partial K} \right)}_{-} \\
& - \underbrace{E[W] \pi(\theta^*) f(\theta^*)}_{+} \underbrace{\left(\frac{\partial^2 \pi(\theta^*)}{\partial K \partial \phi} \pi(\theta^*) + \frac{\partial \pi(\theta^*)}{\partial K} \frac{\partial \pi(\theta^*)}{\partial \phi} \right)}_{-} \\
& > 0 .
\end{aligned}$$

On the other hand, because K^* is an interior solution, we have the second-order condition:

$$\frac{\partial FOC_K(K,\phi)}{\partial K} < 0\,.$$

Applying the implicit function theorem thus yields:

$$\frac{dK^*}{d\phi} = -\frac{\dfrac{\partial FOC_K(K,\phi)}{\partial \phi}}{\dfrac{\partial FOC_K(K,\phi)}{\partial K}} > 0\,,$$

which immediately implies that $dn^*/d\phi < 0$. This concludes the proof. ∎

Proof of Proposition 4. Following the proofs of Propositions 1 and 2, the run threshold when the stablecoin issuer pays dividend τ can be re-derived as:

$$\pi(\theta^*;\tau) = \frac{(1-\phi)(2-2\phi-2(1-\phi)\ln(1-\phi)-K)}{2\left((1+\tau+\eta(1-\phi))(1-\phi)+(1+\tau)\phi\ln\phi\right)}\,, \tag{H.9}$$

where $\partial\pi(\theta^*;\tau)/\partial K < 0$ still holds.

It is obvious that $\pi(\theta^*;\tau)$ is decreasing in τ, implying that the run risk decreases as the issuer pays dividends with n and other model parameters fixed.

We next consider the issuer's optimization problem with respect to n. With dividend τ, the issuer's objective function changes to

$$\max_K \; E_\tau[\Pi] = \underbrace{G(E_\tau[W])}_{\substack{\text{population of}\\ \text{participating investors}}} \underbrace{\int_{\pi(\theta)\geq\pi(\theta^*;\tau)} \pi(\theta)(R(\phi)-1-\tau)dF(\theta)}_{\substack{\text{expected issuer revenue per}\\ \text{participating investor}}}\,,$$

where each investor's expected utility of participation changes to

$$E_\tau[W] = \underbrace{-\alpha\delta^2K^2}_{\text{short-term convenience}} + \underbrace{(1-\phi-K)\int_{\pi(\theta)<\pi(\theta^*;\tau)} dF(\theta)}_{\text{short-term payoff if runs}} + \underbrace{\int_{\pi(\theta)\geq\pi(\theta^*;\tau)} \pi(\theta)\left(1+\eta+\tau\right)dF(\theta)}_{\text{long-term payoff if no runs}}\,,$$

in which $\pi(\theta^*;\tau)$ is given by (4.10) in Proposition 2.

Similarly, we consider the FOC with respect to K:

$$0 = \frac{\partial E_\tau[\Pi]}{\partial K} = \underbrace{\frac{\partial E_\tau[W]}{\partial K}\int_{\pi(\theta)\geq\pi(\theta^*;\tau)} \pi(\theta)(R-1-\tau)dF(\theta)}_{\substack{\text{marginal cost from}\\ \text{reduced investor participation}}}$$

$$
\underbrace{-E_\tau[W]\frac{\partial \pi(\theta^*;\tau)}{\partial K}\left(f(\theta^*)\pi(\theta^*;\tau)(R-1-\tau)\right)}_{\substack{\text{marginal benefit from}\\ \text{reduced run risk}}}, \tag{H.10}
$$

where

$$
\frac{\partial E_\tau[W]}{\partial K} = \underbrace{-2\alpha\delta^2 K}_{\substack{\text{marginal utility cost from}\\ \text{decreasing price stability}}} + \underbrace{\frac{\partial \pi(\theta^*;\tau)}{\partial K}\left(f(\theta^*)(1-\phi-K-\pi(\theta^*;\tau)(1+\eta+\tau))\right)}_{\substack{\text{marginal utility benefit from}\\ \text{increasing financial stability}}} - \int_{\pi(\theta)<\pi(\theta^*;\tau)} dF(\theta)\,.
$$

We first use (H.9) to calculate that

$$
\frac{\partial^2 \pi(\theta^*;\tau)}{\partial K \partial \tau} = \frac{(1-\phi)(1-\phi+\phi\ln\phi)}{2\left((1-\phi)(1+\tau+\eta(1-\phi))+(1+\tau)\phi\ln\phi\right)^2} > 0\,,
$$

and also

$$
\begin{aligned}
\frac{\partial\left[(\pi(\theta^*;\tau)(1+\eta+\tau))\right]}{\partial \tau} &= \frac{\eta(1-\phi)\phi(1-\phi+\ln\phi)(K+2\phi+2(1-\phi)\ln(1-\phi)-2)}{2\left((1-\phi)(1+\tau+\eta(1-\phi))+(1+\tau)\phi\ln\phi\right)^2} \\
&> 0\,,
\end{aligned}
$$

when (H.4) holds. Thus, for $\tau > 0$ we have

$$
\begin{aligned}
&\left.\frac{\partial \pi(\theta^*;\tau)}{\partial K}\right|_{K=K^*}\left(f(\theta^*)(1-\phi-K^*-\pi(\theta^*;\tau)(1+\eta+\tau))\right) \\
<&\left.\frac{\partial \pi(\theta^*)}{\partial K}\right|_{K=K^*}\left(f(\theta^*)(1-\phi-K^*-\pi(\theta^*)(1+\eta))\right)\,.
\end{aligned} \tag{H.11}
$$

On the other hand, similar calculation yields:

$$
E_\tau[W]|_{K=K^*}\left.\frac{\partial \pi(\theta^*;\tau)}{\partial K}\right|_{K=K^*}\pi(\theta^*;\tau) < E[W]|_{K=K^*}\left.\frac{\partial \pi(\theta^*)}{\partial K}\right|_{K=K^*}\pi(\theta^*)\,. \tag{H.12}
$$

Because $R-1-\tau < R-1$, conditions (H.11) and (H.12) thus jointly imply that the new FOC (H.10) also evaluated at K^* is smaller than the old FOC (H.7) evaluated at K^*, which is zero. This immediately implies that $K^*_\tau < K^*$, and hence $n^*_\tau > n^*$. This concludes the proof. ∎

Proof of Proposition 5. Note that, without other market participants taking any action at $t = 0$, this joint optimization problem is equivalent to a sequential decision problem in which the issuer first decides the optimal level of liquidity transformation ϕ, and then decides the optimal arbitrage concentration n as analyzed in the baseline model.

In the extended joint optimization problem, the issuer's objective can be written as, from (4.14):

$$
\max_n\ E[\Pi] = G(E[W])\int_{\pi(\theta)\geq\pi(\theta^*)} \pi(\theta)\left(R(\phi)-1-\tau\right)dF(\theta)\,.
$$

Factoring, and ignoring τ, we have:

$$\max_{\phi} \max_{n} (R(\phi) - 1) G(E[W]) \int_{\pi(\theta) \geq \pi(\theta^*)} \pi(\theta) dF(\theta)$$

$$= \max_{\phi} \left[(R(\phi) - 1) \max_{n} \left[G(EW(\phi, n)) \int_{\pi(\theta) \geq \pi(\theta^*(n,\phi))} \pi(\theta) dF(\theta) \right] \right] .$$

Define the function:

$$F(\phi, n) \equiv G(EW(\phi, n)) \int_{\pi(\theta) \geq \pi(\theta^*(n,\phi))} \pi(\theta) dF(\theta) .$$

Note that there is no dependence on $R(\phi)$. Also define:

$$Q(\phi) \equiv R(\phi - 1) .$$

The objective function is then:

$$\max_{\phi} \left[(R(\phi) - 1) \max_{n} \left[G(EW(\phi, n)) \int_{\pi(\theta) \geq \pi(\theta^*(n,\phi))} \pi(\theta) dF(\theta) \right] \right]$$

$$= \max_{\phi} \left[Q(\phi) \max_{n} F(\phi, n) \right] ,$$

which can be restated by taking logs of the objective function:

$$\arg\max_{\phi} \left[\log(Q(\phi)) + \log \max_{n} F(\phi, n) \right] .$$

By contradiction suppose that $\phi_1^* < \phi_2^*$. The optimality for issuer 2 implies:

$$\log(Q_2(\phi_2^*)) + \log \max_{n} F(\phi_2^*, n) \geq \log(Q_2(\phi_1^*)) + \log \max_{n} F(\phi_1^*, n) ,$$

and

$$\log(Q_2(\phi_2^*)) - \log(Q_2(\phi_1^*)) \geq \log \max_{n} F(\phi_1^*, n) - \log \max_{n} F(\phi_2^*, n) .$$

But (F.1) implies:

$$\log(Q_1(\phi_2^*)) - \log(Q_1(\phi_1^*)) \geq \log(Q_2(\phi_2^*)) - \log(Q_2(\phi_1^*)) \geq \log \max_{n} F(\phi_1^*, n) - \log \max_{n} F(\phi_2^*, n) ,$$

which further implies that:

$$\log(Q_1(\phi_2^*)) - \log \max_{n} F(\phi_2^*, n) \geq \log(Q_1(\phi_1^*)) - \log \max_{n} F(\phi_1^*, n) ,$$

contradicting the optimality of issuer ϕ_1^*. This concludes the proof. ∎

Proof of Proposition 6. Under the redemption fee, the stablecoin's secondary-market price at $t = 2$ is given by

$$p_2(\lambda;\nu) = \begin{cases} 1-\nu-K\lambda & \lambda \le \frac{1-\phi}{1-\nu}\,, \\ \dfrac{1-\phi}{\lambda} - K\lambda & \lambda > \frac{1-\phi}{1-\nu}\,, \end{cases} \tag{H.13}$$

and an investor's value at $t = 3$ thus becomes:

$$v_3(\lambda;\nu) = \begin{cases} \pi(\theta)\left(\dfrac{1-\phi-\lambda(1-\nu)}{(1-\phi)(1-\lambda)} + \eta\right) & \lambda \le \frac{1-\phi}{1-\nu}\,, \\ 0 & \lambda > \frac{1-\phi}{1-\nu}\,. \end{cases} \tag{H.14}$$

Note that, beyond a lower stablecoin price at $t = 2$ due to the redemption fee and a higher remaining value at $t = 3$, the solvency threshold for the stablecoin also changes from $\lambda = 1 - \phi$ in the baseline model to $\lambda = \frac{1-\phi}{1-\nu} > 1 - \phi$ because the stablecoin can use the collected redemption fees to meet redemption needs, thus becomes more resilient to redemption shocks.

Following the proofs of Propositions 1 and 2, (H.13) and (H.14) imply that the investor's indifference condition at $t = 2$ is now given by:

$$\int_0^{\frac{1-\phi}{1-\nu}} (1-\nu-K\lambda)\,d\lambda + \int_{\frac{1-\phi}{1-\nu}}^1 \left(\frac{1-\phi}{\lambda} - K\lambda\right) d\lambda = \int_0^{\frac{1-\phi}{1-\nu}} \pi(\theta^*)\left(\frac{1-\phi-\lambda(1-\nu)}{(1-\phi)(1-\lambda)} + \eta\right) d\lambda\,. \tag{H.15}$$

Solving (H.15) then yields:

$$\pi(\theta^*;\nu) = \frac{(1-\nu)(1-\phi)\left(2-2\phi-2(1-\phi)\ln\left(\frac{1-\phi}{1-\nu}\right)-K\right)}{2(1-\nu+\eta(1-\phi))(1-\phi)+2(1-\nu)(\phi-\nu)\ln\left(\frac{\phi-\nu}{1-\nu}\right)}\,, \tag{H.16}$$

which is decreasing in ν. This concludes the proof. ∎

I Additional Empirical Results

Table A.3: Primary Market Monthly Redemption and Creation Activity (Tron)

Panels (a) to (f) provide statistics about monthly primary market redemption and creation activity on the Tron blockchain, including the number of arbitrageurs, the market share of the top 1 and top 5 arbitrageurs, and the transaction volume. For each variable, we show the average, 25^{th} percentile, 50^{th} percentile, and 75^{th} percentile values across months in our sample.

(a) USDT

	mean	p25	p50	p75
RD AP Num	5	2	4	6
RD Top 1 Share	72	53	68	94
RD Top 5 Share	100	100	100	100
RD Vol (mil)	4625	651	3575	7515
CR AP Num	11	2	12	14
CR Top 1 Share	65	46	54	96
CR Top 5 Share	98	96	99	100
CR Vol (mil)	4991	628	3515	7475

(b) USDC

	mean	p25	p50	p75
RD AP Num	446	11	317	391
RD Top 1 Share	58	33	51	81
RD Top 5 Share	84	78	85	100
RD Vol (mil)	41	3	24	70
CR AP Num	442	8	493	655
CR Top 1 Share	77	56	92	98
CR Top 5 Share	94	97	99	100
CR Vol (mil)	259	11	70	153

(c) TUSD

	mean	p25	p50	p75
RD AP Num	4	2	3	7
RD Top 1 Share	87	69	95	100
RD Top 5 Share	100	100	100	100
RD Vol (mil)	61	0	21	32
CR AP Num	3	1	2	3
CR Top 1 Share	95	98	100	100
CR Top 5 Share	100	100	100	100
CR Vol (mil)	85	0	24	80

Table A.4: Primary Market Monthly Redemption and Creation Activity (Avalanche)

Panels (a) to (f) provide statistics about monthly primary market redemption and creation activity on the Avalanche blockchain, including the number of arbitrageurs, the market share of the top 1 and top 5 arbitrageurs, and the transaction volume. For each variable, we show the average, 25^{th} percentile, 50^{th} percentile, and 75^{th} percentile values across months in our sample.

(a) USDT

	mean	p25	p50	p75
RD AP Num	1	1	1	1
RD Top 1 Share	100	100	100	100
RD Top 5 Share	100	100	100	100
RD Vol (mil)	50	1	10	120
CR AP Num	1	1	1	2
CR Top 1 Share	88	93	100	100
CR Top 5 Share	100	100	100	100
CR Vol (mil)	84	1	45	140

(b) USDC

	mean	p25	p50	p75
RD AP Num	34	18	32	47
RD Top 1 Share	49	31	42	60
RD Top 5 Share	94	87	96	99
RD Vol (mil)	111	3	16	219
CR AP Num	44	34	44	60
CR Top 1 Share	54	43	49	64
CR Top 5 Share	89	83	86	96
CR Vol (mil)	287	20	267	524

(c) BUSD

	mean	p25	p50	p75
RD AP Num	22	10	18	30
RD Top 1 Share	37	30	40	42
RD Top 5 Share	83	73	82	94
RD Vol (mil)	0	0	0	0
CR AP Num	33	11	18	43
CR Top 1 Share	41	34	38	50
CR Top 5 Share	87	82	94	98
CR Vol (mil)	0	0	0	0

(d) TUSD

	mean	p25	p50	p75
RD AP Num	66	49	74	85
RD Top 1 Share	50	36	46	64
RD Top 5 Share	86	79	91	94
RD Vol (mil)	154	31	85	260
CR AP Num	92	53	106	130
CR Top 1 Share	50	33	46	65
CR Top 5 Share	87	83	87	92
CR Vol (mil)	164	30	77	259

J Additional Calibration Details and Results

The CDS spread s_c on an asset class $c \in \{1 \dots C\}$ can be thought of as the probability of default under a recovery rate of 0. Since we assume 0 recovery rates in our model, for a single asset, s_c maps exactly to p in our model. Now, suppose the issuer holds a fraction q_c of her portfolio in asset class c. If each asset pays off 1 with probability s_c and 0 with probability $(1 - s_c)$, the portfolio as a whole has an expected recovery value:

$$\sum_{c=1}^{C} s_c q_c$$

We add an adjustment factor to account for the fact that stablecoin issuers tend to be overcollateralized. If the issuer holds $1 + \xi$ in assets times the total number of stablecoin issued, then the expected recovery value of assets, for each unit of stablecoin issued, is:

$$p = (1 + \xi) \sum_{c=1}^{C} (1 - s_c) \, q_c \tag{J.1}$$

Since p in the model is equal to the expected recovery value of assets per unit stablecoin issued, we will use (J.1) on each date we observe CDS spreads as one realization of p. We can think of (J.1) as the price of a composite security, which averages across CDS spreads of different components of a stablecoin issuer's portfolio, and accounts for the fact that issuers are slightly overcollateralized. With any set of CDS spreads on a given day, we can calculate a value of p using (J.1). By plugging CDS spreads from different dates into (J.1), we can calculate a distribution of signals p. Note that, when we plug CDS spreads into (J.1), we use spreads from a single day; hence, this method accounts for correlations between CDS prices of different asset classes.

We implement (J.1). We choose the historical CDS series from Markit that is liquid and that best fits each reported asset category. For deposits, we assign the average CDS of unsecured debt at the top six US banks to capture the riskiness of the banking sector.[33] We note that despite stablecoin issuers' claim that deposits are riskless in FDIC-insured institutions, they are not riskless or fully insured because deposit accounts exceeding 250K are not covered by deposit insurance, as evident from the recent Silicon Valley Bank episode. For Treasuries, we assign the CDS spreads on 3-year US treasuries. For money market instruments, we use CDX spreads on 1-year investment-grade corporate debt. For USDC's corporate bonds, we assign the 10-year investment-grade corporate CDX because they are stated to be of at least a BBB+ rating. For USDT's corporate bonds, we assign the average 10-year corporate CDX. The remaining categories, "foreign" and "other", do not have a clear mapping to the existing CDS series. For USDT, for example, assets in the "other"

[33] These include Bank of America, Wells Fargo, JP Morgan Chase, Citigroup, Goldman Sachs, and Morgan Stanley.

category include cryptocurrency, which could potentially be very risky. In our baseline results, we use the emerging market CDX spread as a proxy. We use the 10-year high-yield CDX spread as a robustness check. Our sample period is from 2008 to 2022.

Using the daily portfolio-level CDS spreads as observations, we fit a beta distribution for each coin-month by choosing the two beta distribution parameters to match the mean and variance of the empirical distribution of signals p. We then use this beta distribution as the distribution of $p(\theta)$ in the model. Appendix Table A.5 shows the parameters of the beta distributions we estimate.

Table A.5: Distribution of $p(\theta)$

This table shows the fitted beta distributions for $p(\theta)$, for each stablecoin and month in our data. α and β are respectively the two beta distribution parameters. Mean $p(\theta)$ and SD $p(\theta)$ are the mean and SD of the estimated beta distributions for $p(\theta)$.

Coin	Month	α	β	Mean $p(\theta)$	SD $p(\theta)$
USDT	2021m6	156.24	1.16	0.9926	0.0068
USDT	2021m9	170.15	1.33	0.9922	0.0067
USDT	2021m12	211.54	1.60	0.9925	0.0059
USDT	2022m3	213.25	1.42	0.9934	0.0055
USDC	2021m5	127.59	0.57	0.9955	0.0059
USDC	2021m6	137.00	0.57	0.9959	0.0054
USDC	2021m7	138.22	0.58	0.9958	0.0055
USDC	2021m8	122.20	0.83	0.9933	0.0073
USDC	2021m9	121.81	0.89	0.9928	0.0076
USDC	2021m10	121.81	0.89	0.9928	0.0076

Table A.6: Secondary Market Price Deviation versus Redemptions/Creations

This table shows the results from regressing daily secondary market price deviations against the daily volume of redemptions/creations for USDT and USDC. For redemptions, price deviation is one minus the lowest hourly secondary market price on that day. For creations, price deviation is the highest hourly secondary market price on that day minus one. The daily volumes of redemptions and creations are expressed as a proportion of the total outstanding volume of each stablecoin. We include a year fixed effect to account for structural shifts over time.

	USDT	USDC
	(1)	(2)
Redemption/Creation	0.21***	0.16***
	(0.06)	(0.02)
Year FE	Yes	Yes
Observations	1225	1792
Adjusted R2	0.01	0.05

Figure A.6: Effect of Dividend Payments (Full Sample Period)

This figure shows the predicted effect of dividend payments to investors on the issuer's choice of K, the cost of price variance $K\alpha\delta^2$, and run probability, for all periods in our sample.

(a) Elasticity K

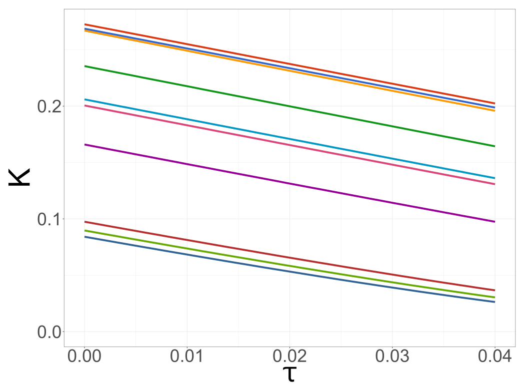

(b) Price Variance Cost

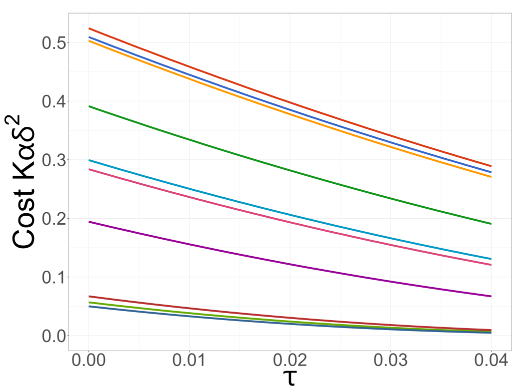

(c) Run Probability

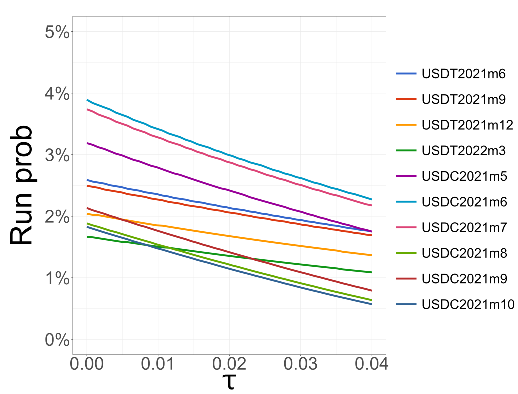

Figure A.7: Effect of Redemption Fees (Full Sample Period)

This figure shows the predicted effect of redemption fees ν on run probabilities. Throughout the exercise, we hold K equal to the model-predicted optimal value of K, in the absence of redemption fees, for all periods in our sample.