Stablecoin Runs and the Centralization of Arbitrage*

Yiming Ma† — Columbia Business School

Yao Zeng‡ — Wharton

Anthony Lee Zhang§ — Chicago Booth

This version: December 2023

Click **here** for the latest version

*We thank Garth Baughman, Chris Cameron, Francesca Carapella (discussant), Caitlin Dannhauser (discussant), Eduardo Davila (discussant), Campbell Harvey (discussant), Chris Hua, Urban Jermann, Ye Li (discussant), Gordon Liao, Igor Makarov, Cyril Monnet (discussant), Cecilia Parlatore, Anna Pavlova, Mila Sherman, Dina Tsarapkina, Alexandros Vardoulakis, and Ganesh Viswanath Natraj (discussant) 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 SI 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, and Wharton. Ding Ding, Grace Jiarui Fan, Baiyang Han, Junyi Hu, Liming Ning, Haichuan Wang, Yuming Yang, and Max Yang provided excellent research assistance. All errors are ours.

†Finance Division, Columbia Business School, 3022 Broadway, Uris Hall 820, New York, NY 10027. Email: ym2701@gsb.columbia.edu.

‡Wharton School, University of Pennsylvania, Steinberg-Dietrich Hall Suite 2300, Philadelphia, PA 19104. Email: yaozeng@wharton.upenn.edu.

§Booth School of Business, University of Chicago, 5807 S. Woodlawn Ave. Chicago, IL 60637. Email: anthony.zhang@chicagobooth.edu.

Stablecoin Runs and the Centralization of Arbitrage

Click here for the latest version

Abstract

Stablecoins are cryptoassets which are designed to be pegged to the dollar, but are backed by imperfectly liquid USD assets. We show that stablecoins feature *concentrated arbitrage*: the largest issuer, Tether, only allows 6 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 increases run risks by reducing investors' price impact from selling stablecoins. We show that run risk could be reduced by imposing redemption fees or issuing dividends to investors.

1 Introduction

Stablecoins are blockchain assets whose value is claimed to be stable at $1. The largest stablecoins are fiat-backed. They attempt to achieve price stability by promising to back each stablecoin token with at least $1 in US dollar-denominated assets, which range from bank deposits and Treasuries to corporate bonds and loans. The potential for stablecoins to become a safe asset and a means of payment has contributed to their meteoric rise. The six largest US dollar-backed stablecoins have grown from $5.6 billion in asset size at the beginning of 2020 to exceed $130 billion at the beginning of 2022.

The rapid expansion of fiat-backed stablecoins has raised financial stability concerns because of the potential spillover effects on the traditional financial system.[1] Distinct from defaults of other crypto assets, a run on fiat-backed stablecoins may strain important asset markets for deposit funding, Treasuries, and corporate bonds. For example, if the largest USD stablecoin, USDT, sold its Treasury positions in a run, the sale would amount to one-sixth of the Treasuries liquidated by mutual funds in March 2020. These ramifications have led to widespread discussions about how stablecoin runs can be mitigated.

In this paper, we analyze the run risk of US dollar fiat-backed stablecoins. Our contributions are three-fold. First, we document the novel and surprising fact that stablecoins feature *concentrated arbitrage*. Only a small number of arbitrageurs are allowed to redeem and create stablecoins with the stablecoin issuer for $1 in cash in primary markets; whereas the vast majority of investors buy and sell stablecoins in competitive secondary markets, analogous to how investors trade ETF shares on exchanges. Arbitrageurs trade against fluctuations in stablecoin demand. For example, 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 this sense, our finding of concentrated arbitrage is surprising because more efficient arbitrage should help to improve price stability on secondary markets, as also pointed out by Lyons and Viswanath-Natraj (2021). This raises the question of why issuers do not simply authorize more arbitrageurs to improve stablecoins' price stability.

Second, we reconcile arbitrage concentration by showing that more competitive arbitrage increases the risk of stablecoin runs. Stablecoins are subject to panic runs because of illiquidity in their assets

[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".

and the fixed $1 redemption value. These kinds of runs are more likely when arbitrage is more efficient because investors have less price impact, which makes them more incentivized to sell. 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. Our theory rationalizes, for example, why USDT, the largest stablecoin issuer, chooses to concentrate arbitrage to a greater degree than USDC, the second largest. USDT has more illiquid reserve assets than USDC, and thus faces higher run risks, implying that USDT may optimally choose more concentrated arbitrage to limit the risk of runs, at the cost of decreased price stability.

Finally, we quantify the run risks of the two largest stablecoins and show how regulatory interventions could influence stablecoin run risks and price stability. Stablecoin issuers currently do not distribute dividends to investors, in part because of how securities regulation is applied to stablecoins. We show that allowing positive dividends could both increase price stability and decrease run risk. This is because dividends increase the opportunity cost for investors to sell their stablecoin. We also show that imposing restrictions on redemptions could reduce run risk, but comes at the cost of decreased price stability.

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 6 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 fairly concentrated, though the degree of concentration varies significantly across stablecoins. USDT only has 6 arbitrageurs redeeming tokens during the 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. We also find that stablecoin trading prices in secondary markets frequently deviate from $1. We note that these price deviations are not analogous to MMFs' "breaking the buck" nor are they an indicator of runs. We posit that stablecoins trade below (above) $1 when selling (buying) pressure in secondary markets is not fully absorbed by arbitrage trade. In this sense, stablecoin price fluctuations resemble ETF shares trading at a premium or discount to their NAVs.[2]

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. The gap between the average discounts is larger at 54 bps for USDT and 1 bps for USDC. This finding is consistent with the limits to arbitrage literature showing that imperfect arbitrage hurts price efficiency (e.g., Gromb and Vayanos, 2002, Oehmke, 2010, Du and Zhu, 2017, Davila, Graves and Parlatore, 2022). 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 model to show how the issuer's choice of arbitrage efficiency reflects a tradeoff between financial stability and price stability. There are four types of agents: stablecoin investors, arbitrageurs, noise traders, and a stablecoin issuer. Investors decide whether to hold stablecoins by comparing their expected benefits and costs. Holding stablecoins becomes less attractive when price fluctuations induced by noise traders are larger, and when there is a greater probability of runs that destroy the long-term benefit and recovery value of stablecoins. Investors can also prematurely sell the stablecoins they hold in the secondary market, but they cannot directly redeem stablecoins from the issuer. Instead, they sell to arbitrageurs, who can redeem and create stablecoins with the issuer in the primary market at a fixed price of $1. Arbitrageurs trade stablecoins in the direction of equalizing prices in primary and secondary markets, but leave a wedge depending on their inventory costs. The issuer meets arbitrageur redemptions by prematurely liquidating illiquid reserve assets at a discount, until the point at which the issuer defaults.

Our model shows that panic runs by investors on stablecoins can occur even though stablecoin investors are selling at market prices in exchange-traded markets like investors selling ETF shares. The conventional view may imply that like ETFs, stablecoins are not runnable because of exchange

[2]In contrast, the parallel to "breaking the buck" at money market funds would be a failure by stablecoin issuers to honor the $1 redemption value in primary markets, which has not yet materialized thus far.

trading, where the trading price in secondary markets falls as more investors sell, creating a natural strategic substitutability. In the case of stablecoins, however, arbitrageurs are promised a fixed in-cash redemption price by the issuer. As a result, investors who hold stablecoins may end up with less valuable stablecoins in the future due to the costs from the issuer's forced sales of illiquid assets to meet arbitrageurs' redemptions at \$1. Consequently, stablecoins' fixed primary market price reintroduces strategic complementarity among secondary market investors.

We endogenize the probability of runs using a global games approach to arrive at our core result: increasing the efficiency of arbitrage can increase 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 incentivizes selling and amplifies investors' first-mover advantage when they expect others to sell. In contrast, more concentrated arbitrage increases price impact in secondary markets discourages panic selling, and mitigates run risk. Nevertheless, this mitigation of run risk comes at the expense of secondary market price stability, which suggests that arbitrage concentration is a double-edged sword. Since stablecoin investors care about both run risk and price stability, the stablecoin issuer optimally chooses arbitrage efficiency to trade off its benefits for price stability with its costs for run risk.

We then calibrate our model to quantify the run risk of the two largest stablecoins, USDT and USDC. We first measure the overall illiquidity of USDT and USDC's reserve portfolios using collateral haircuts. On average, the reserve assets of USDT are more illiquid than those of USDC. 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. This lending rate captures the compensation to investors for not being able to use the stablecoin while it is on loan. Finally, we obtain the remaining two parameters using moment matching. Specifically, we choose the slope of investors' stablecoin demand and the cost of price variance to most closely match the slope of investors' demand and the slope of arbitrageurs' demand, K, in the data. Intuitively, when the cost of price instability is high, the stablecoin issuer chooses a more efficient arbitrage sector, i.e., a lower K, to better arbitrage away price fluctuations. In the data, the K estimate for USDT is larger than that for USDC, consistent with USDT maintaining a more concentrated arbitrage sector.

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. USDC is also exposed to default risk in the banking sector because of its concentrated deposit holdings. In September 2021, the run risk at USDT and USDC remain elevated at 2.495% and 2.134% respectively.

Our results have several policy implications. A direct implication of our findings is that regulators should also pay close attention to stablecoin arbitrage capacity since this also has important effects on run risk. Stablecoin primary market structure is also straightforward to measure because creation and redemption activity leaves records on public blockchains. One way for issuers and regulators to influence arbitrage capacity and reduce run risk at stablecoins is to impose redemption fees on arbitrageurs. Quantitatively, our calibrated model shows that moderate redemption fees can have fairly large effects on stablecoin run risk.

Our calibrated model further shows that the run risk at both USDT and USDC could be meaningfully reduced if they started paying dividends to investors. Most stablecoins currently do not pay out dividends to investors, in part because paying dividends would likely lead stablecoins to be classified as securities, imposing uncertain regulatory costs from securities regulation for cryptocurrencies. In our model, dividend payments to investors could both decrease run risk and increase price stability: the increased payoffs from holding stablecoins discourage investors from running, and issuers may opt to increase arbitrage efficiency and thus price stability in response. As dividends increase from 0% to 4%, for example, we estimate that run probabilities would be lowered by 0.80% and 1.34% for USDT and USDC, respectively. Thus, changes in regulation that incentivize or require stablecoin issuers to pay out dividends to investors could contribute to improving the financial stability and price stability of the two largest fiat-backed stablecoins going forward.

Our paper contributes to a large literature on bank runs and liquidity transformation (e.g., Diamond and Dybvig, 1983, Allen and Gale, 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 Schnabl, 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). We model the unique combination of ETFs and MMFs in the design of stablecoins, and we show that panic runs may still happen despite stablecoins' trading on competitive secondary markets and investors' inability to access primary markets.

We also contribute to the fast-growing stablecoin literature by shedding light on arbitrage concentration at US dollar reserve-backed stablecoins and analyzing its effect on stablecoin run risk. Most closely related to us is Lyons and Viswanath-Natraj (2021), who 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. Aldasoro, Ahmed, and Duley (2023) analyzes the effect of disclosure about reserve asset quality on stablecoin runs, while Bertsch (2023) models the effect of stablecoin adoption on fragility. 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 stablecoin. 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.

Several other papers have explored the 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. 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 risk of 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 run risks.

Our paper also fits more broadly into the literature on cryptocurrencies and decentralized finance, discussed and surveyed in Harvey, Ramachandran and Santoro (2021), John, Kogan and Saleh (2022), and Makarov and Schoar (2022).

The rest of the paper proceeds as follows. Section 2 describes institutional details of the stablecoin market and Section 3 explains the data we use. Section 4 documents several empirical facts that motivate our model in Section 5. Section 6 explains the model calibration and results. Section 7 shows two policy counterfactual results of issuing dividends to investors and imposing redemption fees. Finally, Section 8 concludes.

2 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 tokens 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 tokens, and to direct the network to take actions such as transferring tokens 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 1). 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 (PUSD), TrueUSD (TUSD), and Gemini dollar (GUSD) are significantly smaller in size.

In the remainder of this section, we provide an overview of the uses of stablecoins and the stablecoin market structure.

2.1 Uses of Stablecoins

Stablecoins are a fairly 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.[3] As of January 2023, sending tokens 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.[4]

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.

---

[3] 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.

[4] 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 (Tether) as a form of payment.

2.2 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 2. To create a stablecoin token, an arbitrageur 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 how easily and costly market participants can 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.

3 Data

In this section, we explain our main data sources.

Primary market data. The core dataset used in our analysis is data on each stablecoin creation and redemption event for the 6 largest fiat-backed stablecoins: USDT, USDC, 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. Appendix B.1 presents further details for the primary markets of stablecoins and the construction of our data.

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 weighted average of hourly closing prices across these exchanges, where the weights are by trading volume. 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.

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 at times issues tokens 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.

4 Facts

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

4.1 Secondary Market Prices

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

Figure 3 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. Rather, 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 6 coins 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.

The 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.[5] We note that these are misconceptions. Stablecoins' "stable value" of $1 refers to the amount that 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, which has not yet occurred at any of the stablecoins in our sample. The secondary market price is the trading price of stablecoins on exchanges. It is essentially the share price of a closed-end fund and 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 is not a direct indicator of "breaking the buck" or panic runs.

[5] 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

4.2 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 6 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.[6] In terms of transaction volumes, notice that in the 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.2 and A.3, we repeat the analysis for the Tron and Avalanche blockchains and obtain similar variations in arbitrageur concentration across stablecoins.

4.3 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 the average price deviation of that stablecoin. Similarly, we count the number of unique arbitrageurs that engage in redemptions and/or creations and calculate the market share of the largest five arbitrageurs in each month and the average

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

over time for each coin. We plot the results in Figure 4a. A clear negative trend emerges: stablecoins with fewer arbitrageurs, like USDT, have higher average price deviations from one in their secondary market prices than stablecoins with more arbitrageurs, like USDC. Another way to capture arbitrageur concentration is through the market share of the largest arbitrageurs. In Figure 4b, 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.

4.4 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 the form of 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.

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 firesale 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 details breakdowns. That is why our model estimation will focus on USDT and USDC for which reserve asset breakdowns are available.

5 Model

In this section, we develop a model to reconcile the facts presented in Section 4 and to analyze the potential for runs at fiat-backed stablecoins. We first show how run risk is linked to stablecoins' level of liquidity transformation and arbitrage concentration. We then analyze how arbitrage concentration simultaneously affects price stability. Finally, we solve for the issuers' choice of arbitrage efficiency given the tradeoff between run risk and price stability.

5.1 Setup

The economy has four dates, $t = 0, 1, 2, 3$, with no time discounting. There are four groups of risk-neutral players: 1) a competitive group of infinitesimal investors indexed by i, 2) noise traders, 3) a sector of n arbitrageurs, and 4) a stablecoin issuer. 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.

At $t = 0$, the stablecoin 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. The issuer also holds the stablecoin that is initially backed by the reserve asset. The initial value of the reserve asset is normalized to one dollar.

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. We will formalize investors' participation decisions and issuer's profit maximization in Section 5.3. Until then, we take n as exogenous and normalize the population of participating investors to one.

At $t = 1$, noise traders trade stablecoins, creating a variance in stablecoin prices. With equal probability, noise 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 stablecoin market cap, and then rebuy at the end of the period. Letting ω denote noise trader order flow, ω is equal to δ or $-\delta$ with equal probability. Intuitively, we can think of noise traders as using stablecoins for remittances: as we describe in Subsection 2.1, the remittance process involves buying stablecoins with fiat, sending stablecoins, then selling for fiat.[7] Noise traders cannot directly trade with the issuer; instead, they exchange fiat for stablecoins by trading with arbitrageurs in secondary markets.

Also at $t = 1$, n arbitragers trade stablecoins in secondary and primary markets to profit from price deviations. Arbitrageurs cannot hold net inventory, so they must on net redeem as much on primary markets as they purchase in secondary markets. This is consistent with the empirical observation that arbitrageurs hold negligible amounts of stablecoins (see Appendix B.3). Arbitrageurs face quadratic inventory 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 scales proportionally with the stablecoin's market cap. χ can be thought of as capturing arbitrageurs' balance sheet capacity: when χ is higher, inventory costs are lower. Arbitrageurs submit competitive bids in secondary markets to trade stablecoins in a uniform-price multi-unit double auction. We characterize the solution to the double auction in Appendix C. Essentially, arbitrageurs' bids produce a downwards-sloping demand curve in secondary markets for the stablecoin; when the number of arbitrageurs n is larger, and when each arbitrageur's balance sheet capacity χ is larger. the slope of demand is higher, so stablecoin sales have less price impact.

We assume that noise trader-induced price fluctuations lower stablecoin investors' convenience yields. Following Gorton and Pennacchi (1990), we let investors enjoy a short-term 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, both as a transaction medium and as a store of value, when their prices are more volatile.

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

In stages $t = 2$ and $t = 3$, investors decide whether to liquidate stablecoins early, potentially leading to runs, or hold stablecoins to maturity to capture benefits from using stablecoins. At $t = 2$, investors can choose to sell their stablecoins in the manner of Jacklin (1987) and Farhi, Golosov and Tsyvinski (2009). We use λ to denote the fraction of investors that sell their stablecoins at the market price p_2. Like noise traders, investors cannot directly redeem stablecoins from the issuer; instead, they liquidate by selling stablecoins to arbitrageurs in the secondary market, who redeem stablecoins for cash from the issuer. Arbitrageurs, considering the amount they expect to be able to redeem from the issuer, bid in a double auction to determine the price p_2 at which investors' sales of λ stablecoins occur.[8]

The issuer, in turn, meets arbitrageur redemptions in cash by liquidating the illiquid reserve asset. 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 as well as the various costs incurred when transacting illiquid assets (see Duffie, 2010, for a review). Note that the issuer 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.[9]

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

[8]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 noise trading itself may lead to fire sales or render the stablecoin issuer default. Considering that would complicate the model without new economic insights.

[9]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.

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 parameter η can be understood as the convenience and benefits derived from holding and using the stablecoin in the long run.

We solve the model by backward induction: we first analyze investors' run decisions at $t = 2$, and link the run risk to the concentration of arbitrage in Section 5.2. Then in Section 5.3, we analyze how the choice of arbitrage concentration at $t = 0$ in turn affects stablecoin price and investors' convenience at $t = 1$.

5.2 Stablecoin Runs and the Centralization of Arbitrage

To start, we first consider p_2, the price an investor receives from liquidating the stablecoin early:

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

$$p_2(\lambda) = \begin{cases} 1 - K\lambda & \lambda \leq 1 - \phi\,, \\ \dfrac{1-\phi}{\lambda} - K\lambda & \lambda > 1 - \phi\,, \end{cases} \tag{5.1}$$

where

$$K = \frac{1}{\chi n}\,. \tag{5.2}$$

Lemma 1 shows that, for any total redemption quantity $\lambda > 0$, p_2 is decreasing in K and thereby increasing in χ and n. Intuitively, investors' sales decrease secondary market prices less when each arbitrageur is better capitalized, so χ is higher, and when there are more arbitrageurs, so n is higher.

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 long-term value an investor may get at $t = 3$ if λ other investors choose to liquidate early. It is given by

$$v_3(\lambda) = \begin{cases} \pi(\theta)\left(\dfrac{1-\phi-\lambda}{(1-\phi)(1-\lambda)} + \eta\right) & \lambda \leq 1-\phi\,, \\ 0 & \lambda > 1-\phi\,. \end{cases} \tag{5.3}$$

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

$$l(\lambda) = \begin{cases} \dfrac{\lambda}{1-\phi} & \lambda \leq 1-\phi\,, \\ 1 & \lambda > 1-\phi\,. \end{cases} \tag{5.4}$$

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 (5.3).

An important observation from (5.4) 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 from the fact that 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 tokens depend on the sign of the difference between (5.3), the expected utility from holding until date-3, and (5.1), the return 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:

$$h(\lambda) = v_3(\lambda) - p_2(\lambda) = \begin{cases} \pi(\theta)\left(\dfrac{1-\phi-\lambda}{(1-\phi)(1-\lambda)} + \eta\right) - 1 + K\lambda & \lambda \leq 1-\phi\,, \\ -\dfrac{1-\phi}{\lambda} + K\lambda & \lambda > 1-\phi\,. \end{cases} \tag{5.5}$$

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 5 plots the payoff gain function $h(\lambda)$. Observe that $h(\lambda)$ first increases in λ, then decreases, and then increases in λ again. The first region where $h(\lambda)$ increases reflects strategic substitutability arising from the price impact on the secondary market of stablecoins, as shown in (5.1). Selling pressure decreases secondary-market prices. Thus, if a small number of other investors are selling, but not enough to make the issuer insolvent, any given investor has a stronger preference to hold until $t = 3$. However, when a larger number of investors sell, the issuer incurs increasing liquidation costs until there are insufficient assets left to meet redemptions at one dollar at $t = 3$. This strategic complementarity, or "first-mover advantage", force become dominant when investors expect λ to be very large because they want to sell before other investors deplete the issuer's reserves.

The global games framework we use 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: she sells the stablecoin at $t = 2$ if and only if her signal is below a certain threshold. Given the existence of the unique run threshold, we can show that it satisfies the following Laplace equation:

$$\int_0^{1-\phi} (1 - K\lambda)\, d\lambda + \int_{1-\phi}^{1} \left(\frac{1-\phi}{\lambda} - K\lambda\right) d\lambda = \int_0^{1-\phi} \pi(\theta^*) \left(\frac{1-\phi-\lambda}{(1-\phi)(1-\lambda)} + \eta\right) d\lambda\,. \tag{5.6}$$

Solving the Laplace 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*

$$\pi(\theta^*) = \frac{(1-\phi)(2 - 2\phi - 2(1-\phi)\ln(1-\phi) - K)}{2\left((1+\eta(1-\phi))(1-\phi) + \phi\ln\phi\right)}\,. \tag{5.7}$$

which satisfies the following properties:

i). The run threshold, that is, run risk, is increasing in ϕ if and only if $g(\phi) > K$, where $g(\phi)$ is continuous and strictly decreasing in ϕ, and satisfies $\lim_{\phi\to 0} g(\phi) > 0$.[10]

[10]The function $g(\phi)$ can be solved in closed form and is given in the proof.

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

Part i) of Proposition 2 shows that a higher level of stablecoin liquidity transformation leads to a higher run risk when $g(\phi) > K$. This condition may be satisfied when ϕ is not too large 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 (5.5): 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 (5.5); 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.

A core theoretical result is part ii) of Proposition 2, which shows that more efficient arbitrage – that is, a larger 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. 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 to sell early. Imperfect arbitrage thus reduces run risk. Figure 6 illustrates how investors' payoff gain from waiting increases as the secondary market becomes less efficient.

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*

$$\rho = \int_{\pi(\theta)<\pi(\theta^*)} dF(\theta)\,, \tag{5.8}$$

where $\pi(\theta^*)$ is given by (5.7) and $F(\theta)$ is the prior distribution of the fundamentals.

Before proceeding, we make two comments about the notion of stablecoin runs in our framework. We purposefully follow the framework of Diamond and Dybvig (1983) to focus on liquidity transformation and the resulting first-mover advantage and coordination failure in liquidation. Other conceptions of coordination motives, and thus other modeling choices, are possible. One possibility is to follow the idea in the new monetarism 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. In fact, 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). In contrast, given our focus on reserve-backed stablecoins as a financial intermediary, as well as 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 also capture the payment role of stablecoins by modeling its convenience and linking it to stablecoin price fluctuations, consistent with the micro-foundation provided by Gorton and Pennacchi (1990).

Another possibility for modeling coordination and runs is to follow the idea of market runs in Bernardo and Welch (2004). There, if an illiquid asset market features a downward-sloping demand curve, investors fearing future liquidity shocks will have an incentive to front-run each other, fire selling the asset earlier to get a higher price. Similarly, Bernardo and Welch (2004) do not feature an intermediary or liquidity transformation, which is the focus of our paper.

5.3 Price Stability and Optimal Stablecoin Design

Having analyzed the run risk of stablecoins and its relationship with arbitrage concentration, we analyze how arbitrage affects the stablecoin's price stability at $t = 1$ and the issuer's optimal design choices at $t = 0$.

Consider p_1 and its variance, which determines the convenience that investors enjoy at $t = 1$:

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

$$p_1 = \begin{cases} 1 - \delta K & \omega = \delta\,, \\ 1 + \delta K & \omega = -\delta\,, \end{cases} \tag{5.9}$$

where K *is given in* (5.2)*. The stablecoin's 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 convenience is decreasing in K. This is intuitive because as arbitrage becomes less efficient, the secondary market becomes less elastic and noise 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).

Taken together, Proposition 2 and Lemma 2 point to the trade-off between price and financial stability of the stablecoin. To formulate this trade-off, 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.[11] We suppose that the stablecoin issuer chooses n to maximize its expected revenues at $t = 0$, which in turn depends on how many investors participate at $t = 0$. We also assume that noise trading and arbitrageurs' balance sheet capacity are proportional to the population of investors. The issuer's objective function is thus given by

$$\max_n \; E[\Pi] = \underbrace{G(E[W])}_{\substack{\text{population of} \\ \text{participating investors}}} \underbrace{\int_{\pi(\theta)\geq\pi(\theta^*)} \pi(\theta)(R(\phi)-1)dF(\theta)}_{\substack{\text{expected issuer revenue per} \\ \text{participating investor}}}\,, \tag{5.10}$$

where each investor's expected utility of participation is

$$E[W] = \underbrace{-\alpha\delta^2 K^2}_{\text{short-term convenience}} + \underbrace{\int_{\pi(\theta)<\pi(\theta^*)} (1-\phi-K)\,dF(\theta)}_{\text{short-term payoff if runs}} + \underbrace{\int_{\pi(\theta)\geq\pi(\theta^*)} \pi(\theta)\,(1+\eta)\,dF(\theta)}_{\text{long-term payoff if no runs}}\,, \tag{5.11}$$

in which $\pi(\theta^*)$ is given by (5.7) in Proposition 2.

[11] 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.

The stablecoin issuer's objective function (5.10) 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 (5.11). The first term denotes investors' expected 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 (5.10), we have the following result about the stablecoin issuer's optimal choice of arbitrageur concentration:

Proposition 3. *When the stablecoin engages in a higher level of liquidity transformation, the stablecoin issuer optimally designs a more concentrated arbitrageur sector, that is, n^* decreases in ϕ when ϕ is not too large and the cumulative distribution function G is close to linear.*

Proposition 3 stems from the trade-off between price stability and financial stability of stablecoins. Intuitively, when an investor values the convenience of the stablecoin, she prefers smaller price fluctuations as captured by the first term in (5.11). Further, she would like to receive a higher price when she decides to liquidate the stablecoin for dollars. For both reasons, the stablecoin issuer would like to maintain an efficient and elastic secondary market to keep the price stable. However, a more elastic secondary market at the same time leads to higher run risks, as captured by the second and third terms in (5.11). This hurts investors' expected utility and thus their participation incentive as captured by the first term in (5.10). Further, a higher run risk also cuts into the stablecoin issuer's expected revenue per participating investor because it only obtains the net long-term return of the reserve asset when no run happens, as captured by the second term in (5.10). Thus, the issuer accepts some level of price fluctuations to avoid runs. In particular, when asset illiquidity makes runs more likely, the issuer optimally chooses a more concentrated arbitrageur sector to reduce the first-mover advantage among investors, as illustrated in Figure 6.

In the main text, we focus on stablecoin issuers' choice of the degree of arbitrage concentration, and how it depends on illiquidity ϕ, taking ϕ as exogeneous. In the background, we can think of stablecoins as adopting different choices of ϕ because of differences in their business models, which influence the returns to purchasing more illiquid assets. For example, since Circle is in the US and Tether is offshore, Circle may have access to more liquid assets with a higher liquidity premium. In Appendix D, we analyze an extension of the model in which issuers optimally choose both ϕ and arbitrage concentration n concurrently. We show that issuers who face uniformly higher returns from increasing ϕ optimally choose higher values of ϕ, and consequently, lower values of n. This is consistent with the empirical fact that Tether tends to hold more illiquid assets than Circle and imposes more restrictions on redemptions.

6 Model Calibration and Results

In this section, we calibrate our model to estimate run probability as defined in Definition 1. We focus our analysis on the largest two fiat-backed stablecoins, USDT and USDC, because of the availability of their reserve asset breakdowns. We first estimate asset illiquidity ϕ, the distribution of $p(\theta)$, and the long-term benefit η from the data. Using these parameters, we can calculate run thresholds for any given value of K, and then investor welfare and issuers' profits given investor demand $G(\cdot)$. Second, we choose the investor risk parameter $\alpha\delta^2$ and investor demand elasticity to match the model-predicted K and investor demand elasticity in the data.

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

Asset Illiquidity ϕ. We proxy asset illiquidity with haircuts following Bai, Krishnamurthy and Weymuller (2018) and Ma, Xiao and Zeng (2021). These haircuts measure the discount incurred when illiquid assets are converted into cash at short notice. More liquid assets are more readily pledged to obtain cash while more illiquid assets incur a higher discount. [12]

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 demandable deposits, time deposits, and certificates of deposits (CDs). In the baseline estimate, we assume that one-quarter of the deposits are fully liquid while the remainder is subject to the lowest money market discount. The results are shown in Table 4.

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

[12] 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.

component, and then take a weighted average, adjusting for over-collateralization, to obtain the daily expected recovery value of the reserve portfolio. 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 G contains further details of this procedure. The means and variances, as well as the fitted beta distribution parameters, are shown in Appendix Table A.4.

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, which allows market participants to lend cryptoassets for interest, overcollateralized by other cryptoassets. Intuitively, 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 4 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 noise 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 noise 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 $\alpha\delta^2, \gamma_{Circle}, \gamma_{Tether}$ and each coin-month combination in our data, we calculate the optimal value of K, by solving the issuer's optimization problem (5.10). 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.

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 tokens 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 tokens 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 5, we observe that the estimated K for USDT is larger in absolute magnitude than 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.5.

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 5 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 4. 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 noise 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 5. The model-predicted arbitrageur demand slopes K are in the same range but slightly higher than those in the data.[13] 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 4 shows that 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 findings in Egan, Hortacsu and Matvos (2017) and Albertazzi, Burlon, Jankauskas, and Pavanini (2022), who build dynamic structual models

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

to estimate the run probability of commercial 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 framework suggests that there are two core measures that regulators should track to monitor stablecoin run risks going forward. First, regulators should monitor the market structure of the stablecoin arbitrage sector.[14] 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 could limit run risks. Second, regulators should monitor issuers' asset holdings because stablecoin run risks are greater when issuers hold risky and illiquid backing assets.

7 Policy Implications

We discuss two sets of policy implications of our results in this section. We first consider the effect of issuing dividends to investors and then analyze the impact of redemption restrictions.

7.1 Effect of 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. One potential reason for this is that a stablecoin that distributed dividend payments would likely be classified as a security under US securities law.[15] A token

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

[15]The authors are not lawyers; however, US regulators have deemed many programs which 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

being classified as a security potentially exposes the issuer, and other parties such as exchanges that trade the token, to substantial regulatory risk if the issuer does not register the token as a security.[16] However, regulatory treatment of cryptocurrencies is also sufficiently uncertain that very few large crypto projects have successfully registered their tokens as securities with the SEC. Increased regulatory clarity, which decreases the barrier to registering crypto assets as securities, could in principle lower issuers' perceived costs of paying dividends out to investors.

Using our calibrated model, we can show that issuing dividends to investors can both lower run risk and increase price stability. 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:

$$v_3(\lambda;\tau) = \begin{cases} \pi(\theta)\left(\dfrac{1-\phi-\lambda}{(1-\phi)(1-\lambda)}(1+\tau)+\eta\right) & \lambda \leq 1-\phi\,, \\ 0 & \lambda > 1-\phi\,, \end{cases} \tag{7.1}$$

Compared to (5.3), there is an additional τ term that can be collected when the stablecoin is solvent. Accordingly, the stablecoin issuer's objective function becomes:

$$\max_n \; E[\Pi] = G(E[W]) \int_{\pi(\theta)\geq\pi(\theta^*)} \pi(\theta)\,(R(\phi)-1-\tau)\,dF(\theta)\,, \tag{7.2}$$

which nests (5.10) as a special case. We have the following result:

Proposition 4. *Suppose ϕ is not too large, 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.*

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. A number of online sources also concur with this opinion. Web3 University states: "If you want to organize payouts to your token holders, it is safer to structure them as bonus points instead of paying out what could be considered dividends." Unblock states: "In the U.S. dividend-bearing cryptocurrencies are classified as securities."

[16] For example, the June 2023 SEC case against Coinbase accused Coinbase of operating as an exchange for trading unregistered securities.

Intuitively, Proposition 4 states that, if we hold fixed arbitrage efficiency, 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 decrease n, increasing price stability. It is in principle 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.

However, our calibrated model shows that issuing dividends would lead to a net reduction in run risk for USDT and USDC. The results are shown in Figure 7 for the September 2021 reporting period.[17] 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 $\alpha\delta^2 K^2$ decreases relative $\tau = 0$ (panel (b)).

Our findings on dividend issuance shed light on the broader question of how the design of financial intermediaries engaged in liquidity transformation can improve their stability. For mitigating bank runs, Davila and Goldstein (2023) and Kashyap, Tsomocos, and Vardoulakis (2023) explore the optimal design of deposit insurance and banking regulation, respectively. Empirically, Demirguc-Kunt and Detragiache (2002) and Iyer and Puri (2012) show that deposit insurance indeed mitigates run risks by changing the behavior of banks and depositors. In the context of non-banks, Jin, Kacperczyk, Kahraman and Suntheim (2022) and Ma, Xiao and Zeng (2021) show that swing pricing can prevent panic-driven runs at open-ended mutual funds. Our results complement these papers by showing that issuing dividends can also reduce fragility in the context of stablecoins. We note 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. We leave the analysis of these and other forces to future work.

7.2 Redemption Fees

Finally, we examine policies that restrict redemptions when arbitrageurs seek to redeem stablecoins from the issuer. One way is through raising the cost of redemptions by imposing redemption fees.

[17] Results for other reporting periods follow a similar trend and are shown in Appendix Figure A.2.

For instance, USDT charges a 0.1% fee for fiat withdrawals over $1,000, meaning that USDT can be redeemed at $0.999, with a minimum fiat withdrawal or deposit requirement of $100,000. In the model, suppose the stablecoin issuer charges 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 5. *Suppose ϕ is not too large. 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.*

Intuitively, redemption fees play a similar role to constrained arbitrage: investors' sales have a larger price impact, discouraging selling. Moreover, issuers also collect additional profits from creation and redemption fees, lowering the likelihood they have to fire sale assets in order to meet redemptions. We explicitly give the run threshold in the proof of Proposition 5, which allows us to evaluate the run risk under counterfactual scenarios with different redemption fees.

In our estimated model, we evaluate how run risk would change with different values of ν, holding the arbitrageur demand slope at the optimal values of K. The results are shown in Figure 8, for the September 2021 reporting period of USDT and USDC.[18] As redemption fees increase from 0 to 50bps, USDC and USDT run probabilities decrease by 2.01% and 2.38%. Overall, redemption fees would be quite effective at decreasing run risk.

Why do issuers not charge higher redemption fees to limit run risk in practice? One possibility is that fees would be too costly for price stability since prices depeg by an amount independent of redemption size: at a 20bp fee, arbitrarily small redemption amounts cause prices to jump to 20bps below par. Using arbitrage concentration to create price impact has the benefit of creating larger price movements when there are more redemptions. In principle it is possible to charge redemption fees contingent on the number of redeeming customers, effectively creating an issuer token supply curve, but issuers may not be willing or able to commit to this kind of redemption fee schedule.

[18]Note that we hold the arbitrageur demand slope K fixed rather than allowing issuers to reoptimize 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 noise 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.3.

Regarding the use of redemption gates rather than fees, note that 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 arbitrager 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 $n' \leq n$.[19] Then, 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.

8 Conclusion

In this paper, we analyzed the possibility of panic runs on stablecoins. 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.

We show, however, 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. Calibrating the model to data, we find that the two leading fiat-backed stablecoins by market cap, USDT, and USDC, have significant run risk. We also showed how requiring stablecoin issuers to pay dividends to token holders could simultaneously decrease run risk and increase price stability.

Our results have implications for understanding stablecoin issuers' behavior. Some industry participants currently view the difficulty of becoming a stablecoin arbitrageur as essentially a bureaucratic oversight on behalf of issuers. We posit instead that issuers may be purposefully limiting the efficiency

[19] For example, USDC does not typically impose such redemption restrictions but temporarily halted redemptions of USDC from arbitrageurs on the evening of March 10, 2023, as investors became aware that $3.3 billion of its cash deposits remained uninsured at Silicon Valley Bank.

of primary-secondary market arbitrage, in response to the tension between price stability and run risks inherent in the design of fiat-backed stablecoins. Our results also have implications for the measures policymakers should track to monitor stablecoin run risks, and for how policy changes around dividend issuance and redemption fees have the potential to influence price stability and financial stability in the market for fiat-backed stablecoins.

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.

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.

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.

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 (2022). "The Value of Arbitrage." Working Paper.

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." Working Paper.

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.

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

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.

Harvey, Campbell, Ashwin Ramachandran, and Joey Santoro (2021). *DeFi and the Future of Finance*, 2021, John Wiley & Sons.

He, Zhiguo and Wei Xiong (2012). “Rollover Risk and Credit Risk.” *Journal of Finance*, 67: 391-430.

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.

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

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.

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

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

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.

OEHMKE, MARTIN (2010). "Gradual Arbitrage." Working Paper.

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

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.

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.

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 1: Asset Size of Fiat-backed Stablecoins

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

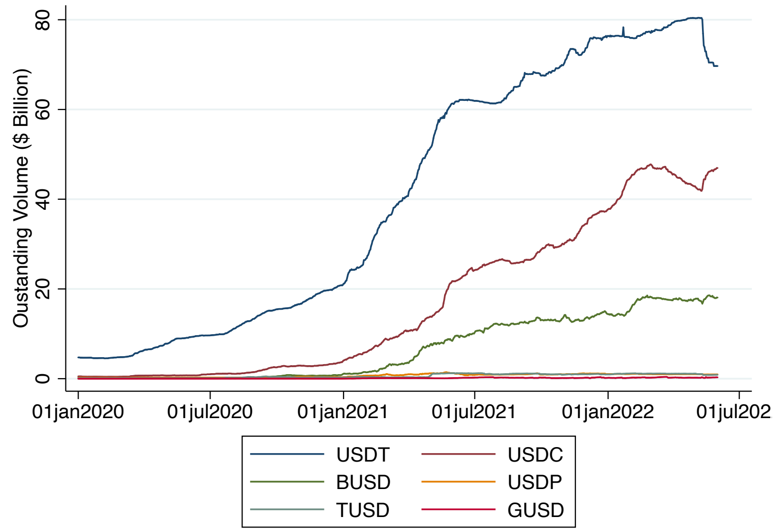

Figure 2: The Design of Fiat-backed Stablecoins

This figure illustrates the design of fiat-backed stablecoins.

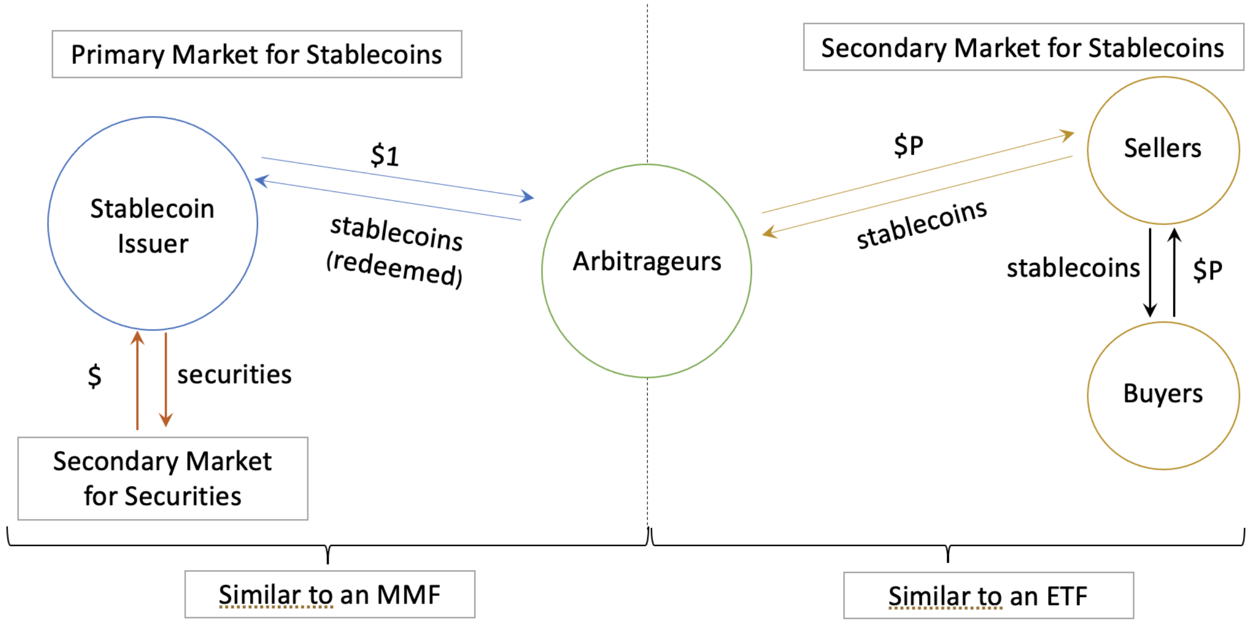

Figure 3: Secondary Market Trading Price

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 2.

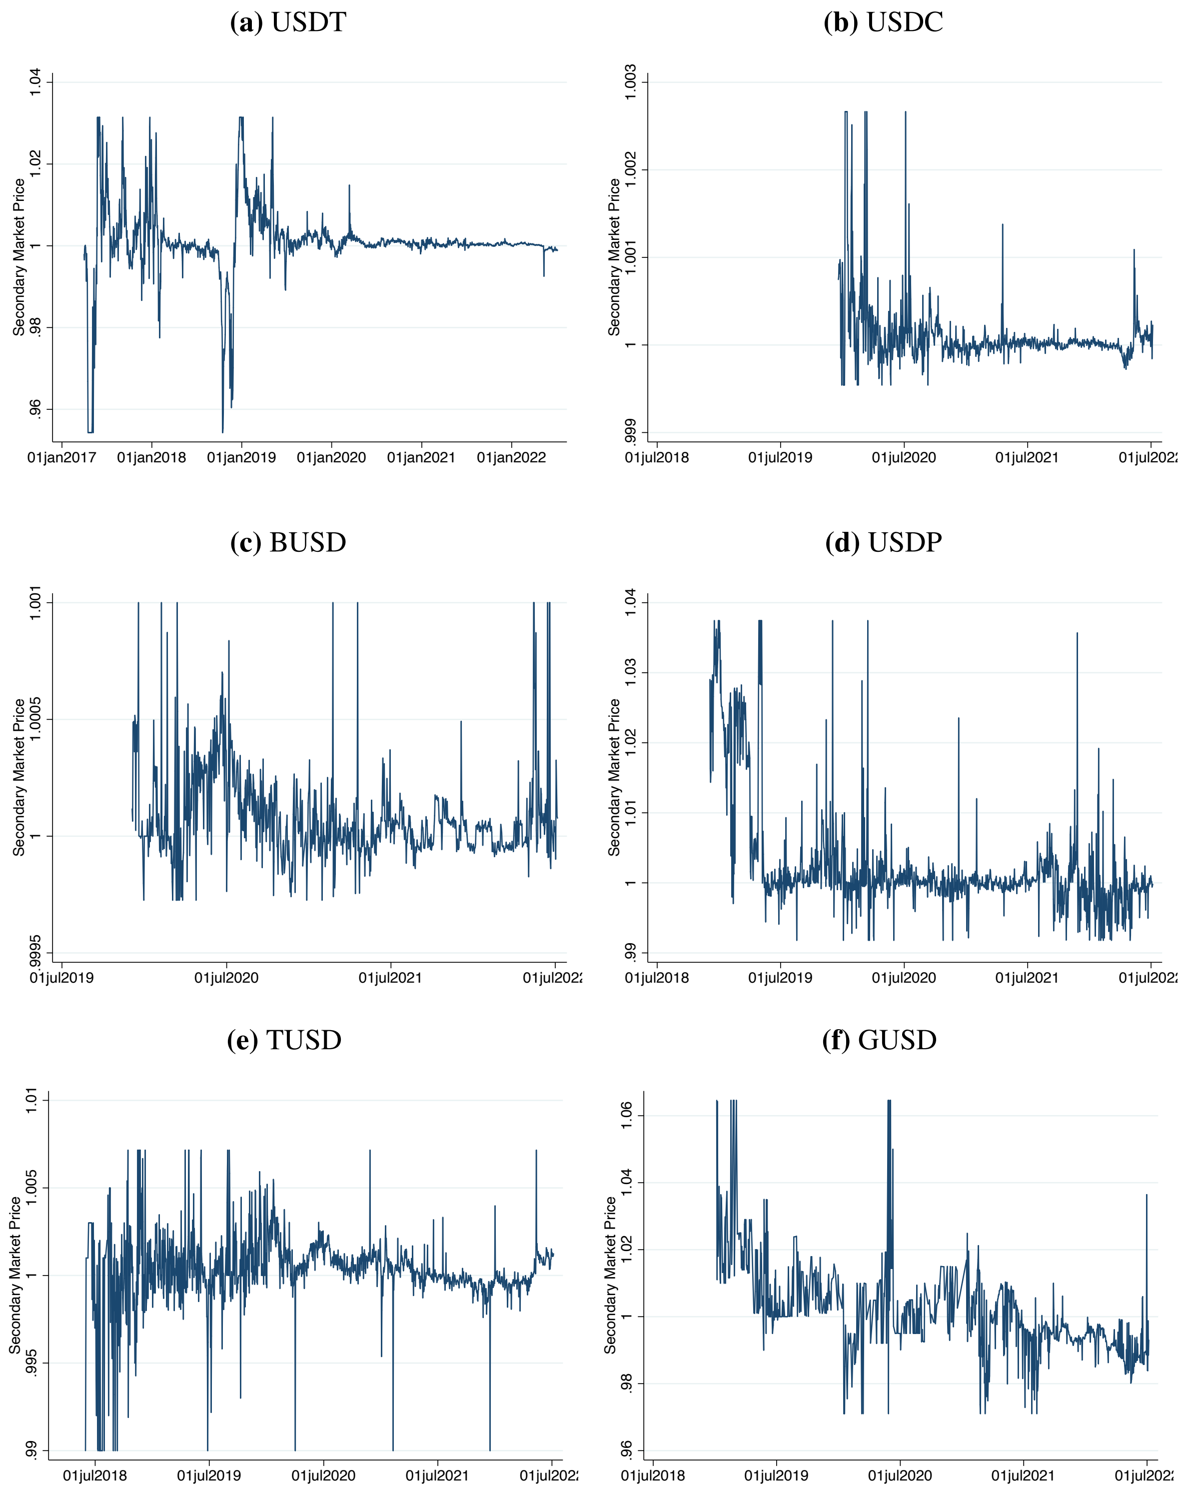

Figure 4: Secondary Market Price Dislocations and Primary Market Structure

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 five 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 the average secondary market price deviation for that stablecoin. Similarly, we count the number of unique arbitrageurs that engage in redemptions and/or creations and 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.

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

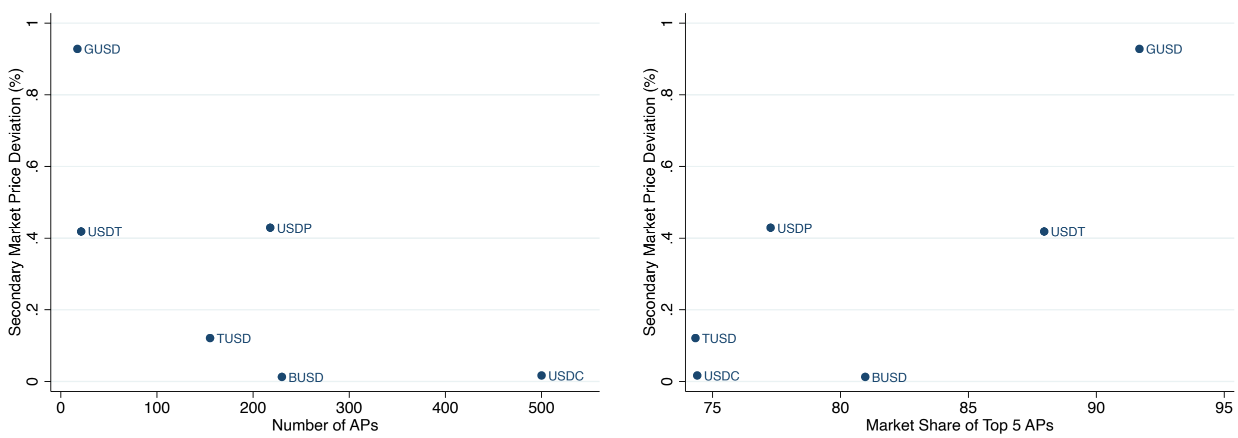

Figure 5: Investors' Payoff Gain from Waiting versus Selling Early

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

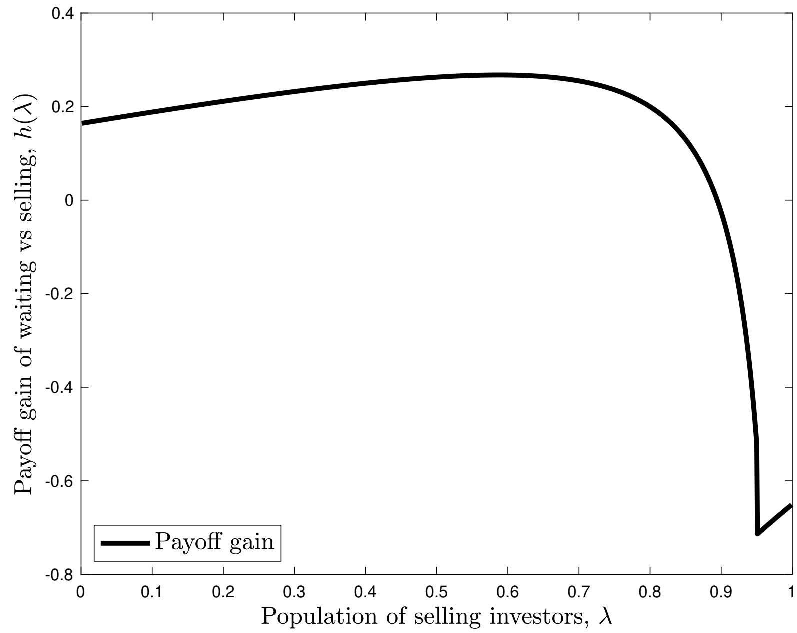

Figure 6: Investors' Payoff Gain from Waiting versus Selling Early: Comparative Statics with respect to K

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

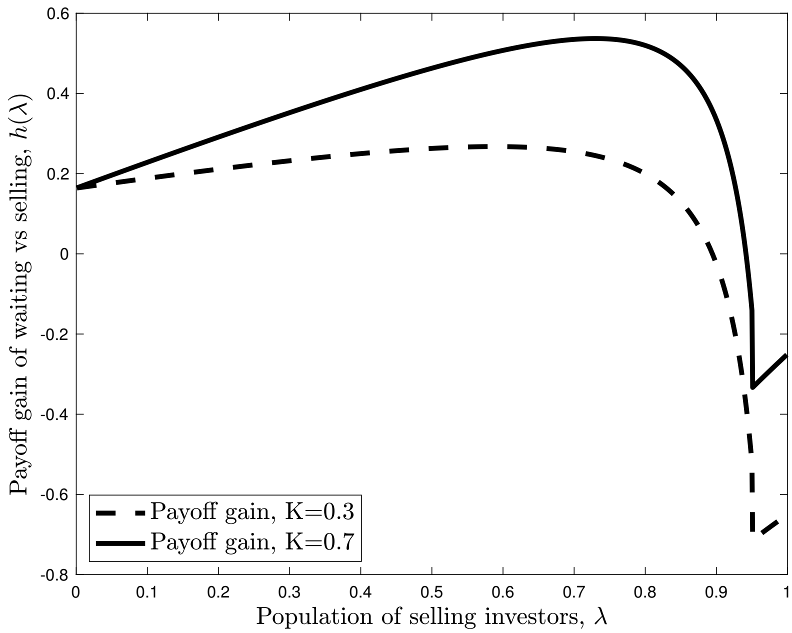

Figure 7: Effect of Dividend Payments

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.

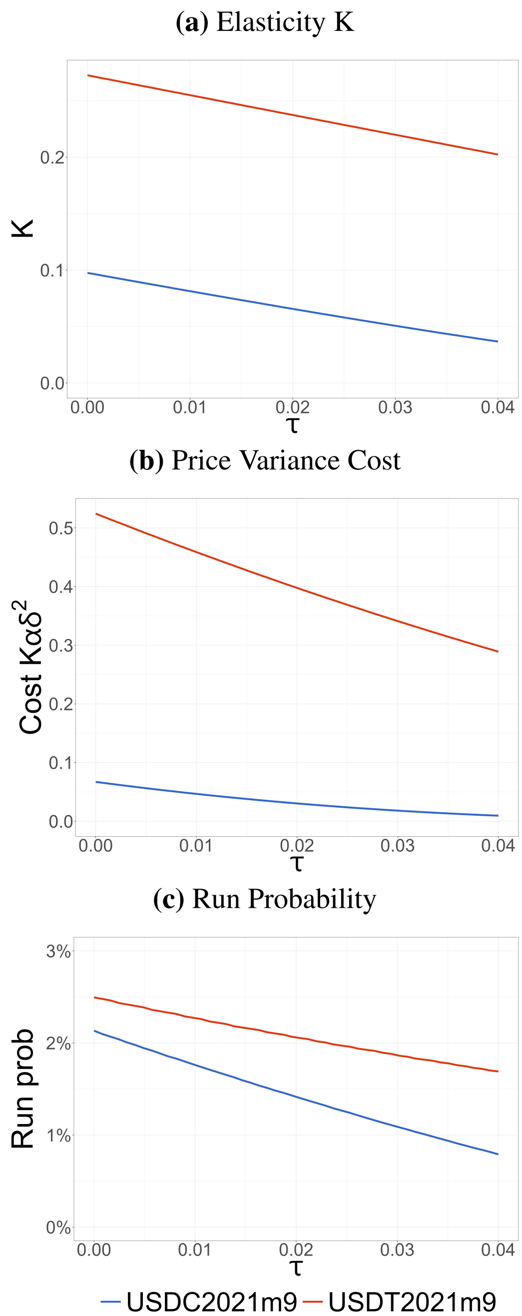

Figure 8: Effect of Redemption Fees

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.

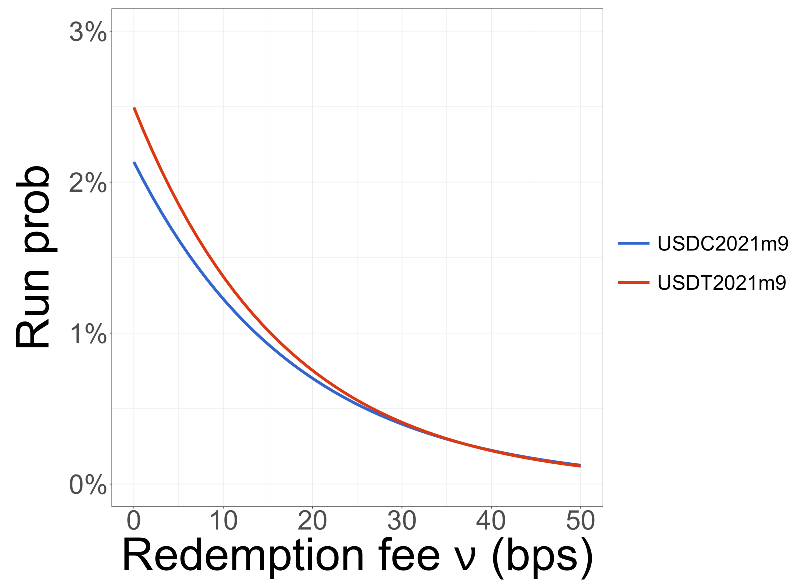

Table 1: Secondary Market Price and Volume

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 of 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 reserve breakdowns are published by USDT and USDC. 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: 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 5: Model Fit

Target K is the slope of arbitrageur demand for the stablecoin, estimated from the data, from (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 Tether, 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.[20]

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

[20]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 Tether 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.

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 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 Tether, 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.[21]

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

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

"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. Observe also that 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 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.

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

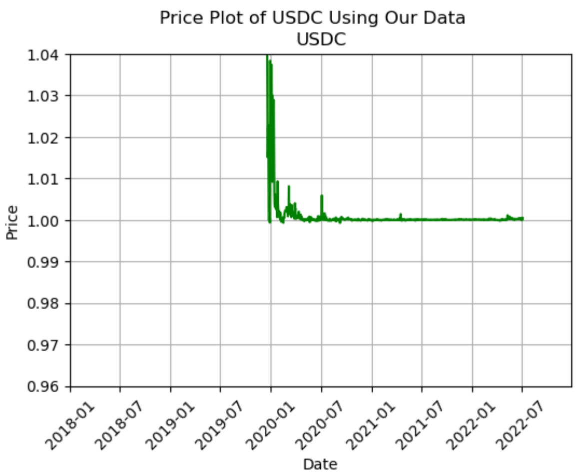

(b) Coingecko Data

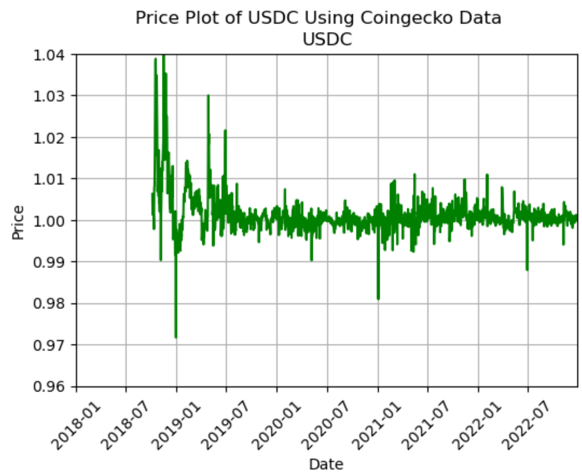

(c) Binance Data

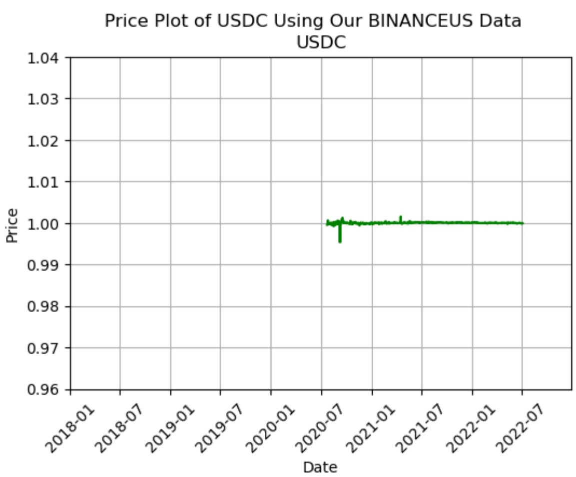

(d) Kraken Data

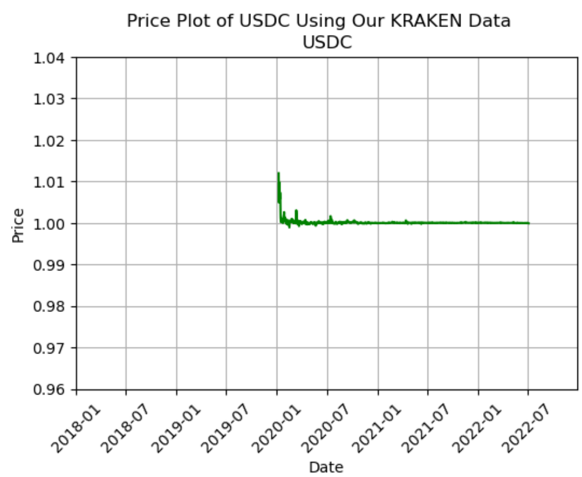

B.3 Arbitrageurs' Stablecoin Ownership

In the model, we assume that arbitrageurs must maintain zero net position, so cannot hold stablecoins on their balance sheet between periods. This is a reasonable approximation to reality: we find that arbitrageurs engaged in redemptions and creations own negligible amounts of stablecoins in their wallets, and never make up a substantial fraction of total holdings. To do this, we use Etherscan to calculate the stablecoin holdings of the largest 15 arbitrageurs for each coin as of the end of our sample period, June 2022; we show summary statistics of arbitrageurs' holdings in Table A.1. This is consistent with our conversations with market participants. Arbitrageurs specialize in profiting from differences between secondary-market prices and primary-market redemption values; in the language of our model, arbitrageurs likely do not attain a substantial convenience yield from holding stablecoins, so they derive little value from holding more of the stablecoin than they need to conduct arbitrage at any point in time.

Table A.1: Share of Stablecoins Held by Arbitrageurs

Panels (a) to (f) provide statistics about the share of stablecoins owned by the largest arbitrageurs on the Ethereum blockchain. For each stablecoin, we show the average, 25^{th} percentile, 50^{th} percentile, and 75^{th} percentile of stablecoin ownership for the largest 15 arbitrageurs engaged in creations and redemptions at the end of our sample period in June 2022. Stablecoin ownership is expressed as a percentage of the total supply of stablecoins.

(a) USDT

	mean	p25	p50	p75
Redemption	0.05	0.00	0.00	0.05
Creation	0.02	0.00	0.00	0.00

(b) USDC

	mean	p25	p50	p75
Redemption	0.00	0.00	0.00	0.00
Creation	0.14	0.00	0.00	0.01

(c) BUSD

	mean	p25	p50	p75
Redemption	0.00	0.00	0.00	0.00
Creation	0.04	0.00	0.00	0.00

(d) USDP

	mean	p25	p50	p75
Redemption	0.00	0.00	0.00	0.00
Creation	0.07	0.00	0.00	0.00

(e) TUSD

	mean	p25	p50	p75
Redemption	0.00	0.00	0.00	0.00
Creation	0.02	0.00	0.00	0.00

(f) GUSD

	mean	p25	p50	p75
Redemption	0.03	0.03	0.03	0.03
Creation	0.00	0.00	0.00	0.00

C Double Auction

In this appendix, we follow Kyle (1989) and Du and Zhu (2017) to model the arbitrage sector as a double auction. At $t = 1$ and $t = 2$, there are n symmetric arbitrageurs indexed by j. At any given period, arbitrageurs bid to buy or sell stablecoins from investors and noise traders, incur a per-period inventory cost if winning the auction, and then create or redeem the stablecoin at the fixed price of one dollar if the issuer is solvent. Thus, arbitrageurs always hold zero inventory at the beginning and the end of each period. Specifically, at any given period, the winning arbitrageurs incur a per-period inventory cost $z_j^2/2\chi$ of arbitraging z_j of the stablecoin, where χ is a parameter. Hence, the expected profit, conditional on winning the auction, for arbitraging z_j of the stablecoin at price p is:

$$z_j\left(1-p\right)-\frac{z_j^2}{2\chi}\,.$$

In the auction, arbitrageurs submit a competitive bid, z_{Bj}, of the amount of stablecoins they are willing to arbitrage if the secondary-market price is p. Taking first-order condition yields:

$$z_{Bj}\left(p\right)=\chi\left(1-p\right)\ . \tag{C.1}$$

Aggregating (C.1) then yields the arbitrageurs' market demand curve:

$$\sum_j z_{Bj}\left(p\right)=\chi n\left(1-p\right)\ , \tag{C.2}$$

and market clearing finally requires:

$$\sum_j z_{Bj}\left(p\right)=m\,, \tag{C.3}$$

where m is the amount of stablecoins supplied by investors or noise traders.

If the stablecoin issuer facing redemptions is insolvent at $t=2$, a similar derivation yields the adjusted market demand curve:

$$\sum_j z_{Bj}\left(p\right)=\chi n\left(\frac{1-\phi}{m}-p\right)\ .$$

Lemmas 1 and 2 immediately follow from applying the market clearing condition (C.3) at $t=2$ and $t=1$.

D 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 (5.10):

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

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:

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

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

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

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

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 6. *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:*

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

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

Intuitively, Proposition 6 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 6, stablecoin issuer 1 could be understood as Tether, while issuer 2 Circle. Condition (D.1) in Proposition 6 means that the asset market that Techer has access to processes higher illiquidity premium in that a more illiquid asset holding offers a higher expected return of the asset. Proposition 6 then predicts that Tether optimally designs USDT's asset holdings in that it holds more illiquid assets, that is, transforming more liquidity. To handle such a greater level 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 convenience yield η and the demand function $G(\cdot)$. However, we were 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 6 cannot be applied.

E 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

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

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

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

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,

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

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 (E.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:

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

which we also give in the main text as (5.6). Solving this Laplace condition (E.2) yields the run threshold (5.7).

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

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

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

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

where $g(\phi)$ is continuous and strictly decreasing in ϕ, and it satisfies $\lim_{\phi\to 0} g(\phi) = 2 > 0$. Thus, conditions (E.3) and (E.4) hold when ϕ is not too large for any given $K \leq 2$, and then the equilibrium run threshold $\pi(\theta^*)$ increases in ϕ.

With respect to K, we have

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

To see why (E.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 (E.5) is positive. This concludes the proof. ∎

Proof of Proposition 3. Suppose condition (E.4) holds, that is, ϕ is not too large. Under this condition, we know from condition (E.3) in the proof of Proposition 2 that the equilibrium run threshold $\pi(\theta^*)$ increases in ϕ. We also consider the limit case of $R'(\phi) = 0$ and the general case of $R'(\phi) > 0$ follows by continuity when ϕ is not too large.

We now consider the first-order condition (FOC) for the issuer's problem (5.10) that determines the optimal K, the slope of arbitrageurs' demand. When $G(\cdot)$ is linear, the FOC is:

$$0 = \frac{\partial E[\Pi]}{\partial K} = \underbrace{\frac{\partial E[W]}{\partial K}\int_{\pi(\theta)\geq\pi(\theta^*)} \pi(\theta)(R-1)dF(\theta)}_{\substack{\text{marginal cost from} \\ \text{reduced investor participation}}} - \underbrace{E[W]\frac{\partial \pi(\theta^*)}{\partial K}\left(f(\theta^*)\pi(\theta^*)(R-1)\right)}_{\substack{\text{marginal benefit from} \\ \text{reduced run risk}}}, \tag{E.6}$$

where according to (5.11),

$$\frac{\partial E[W]}{\partial K} = \underbrace{-2\alpha\delta^2 K}_{\substack{\text{marginal utility cost from} \\ \text{decreasing price stability}}} + \underbrace{\frac{\partial \pi(\theta^*)}{\partial K}\left(f(\theta^*)(1-\phi-K-\pi(\theta^*)(1+\eta))\right)}_{\substack{\text{marginal utility benefit from} \\ \text{increasing financial stability}}} - \int_{\pi(\theta)<\pi(\theta^*)} dF(\theta)\,. \tag{E.7}$$

This first-order condition reveals the various channels through which increasing K affects the stablecoin issuer's expected revenue. The first part of (E.6) 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 (E.6) 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, (E.7) captures the marginal effects of increasing K on an investor's expected utility: the first term of (E.7) is the marginal cost that results from a lower convenience due to 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 (E.6) and the second term of (E.7), respectively.

We now compute $dK^*/d\phi$. Using the FOC (E.6) above:

$$\begin{aligned}\frac{\partial FOC_K(K,\phi)}{\partial \phi} &= \underbrace{\frac{\partial^2 E[W]}{\partial K \partial \phi}}_{+} \underbrace{\int_{\pi(\theta)\geq\pi(\theta^*)} \pi(\theta)(R-1)dF(\theta)}_{+} \\ &- \underbrace{\pi(\theta^*)f(\theta^*)(R-1)}_{+} \underbrace{\left(\frac{\partial E[W]}{\partial K}\frac{\partial \pi(\theta^*)}{\partial \phi} + \frac{\partial E[W]}{\partial \phi}\frac{\partial \pi(\theta^*)}{\partial K}\right)}_{-} \\ &- \underbrace{E[W]\pi(\theta^*)f(\theta^*)(R-1)}_{+} \underbrace{\left(\frac{\partial^2 \pi(\theta^*)}{\partial K \partial \phi}\pi(\theta^*) + \frac{\partial \pi(\theta^*)}{\partial K}\frac{\partial \pi(\theta^*)}{\partial \phi}\right)}_{-} \\ &> 0\,.\end{aligned}$$

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:

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

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:

$$\pi(\theta^*;\tau) = \frac{(1-\phi)(2-2\phi-2(1-\phi)\ln(1-\phi)-K)}{2\left((1+\tau+\eta(1-\phi))(1-\phi)+(1+\tau)\phi\ln\phi\right)}, \tag{E.8}$$

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 (5.7) in Proposition 2.

Similarly, we consider the FOC with respect to K:

$$\begin{aligned}0 = \frac{\partial E_\tau[\Pi]}{\partial K} &= \underbrace{\frac{\partial E_\tau[W]}{\partial K}\int_{\pi(\theta)\geq\pi(\theta^*;\tau)} \pi(\theta)(R-1-\tau)dF(\theta)}_{\substack{\text{marginal cost from}\\ \text{reduced investor participation}}} \\ &\underbrace{- E_\tau[W]\frac{\partial\pi(\theta^*;\tau)}{\partial K}\left(f(\theta^*)\pi(\theta^*;\tau)(R-1-\tau)\right)}_{\substack{\text{marginal benefit from}\\ \text{reduced run risk}}},\end{aligned} \tag{E.9}$$

where

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

We first use (E.8) 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 ϕ is sufficiently small. Thus, for $\tau > 0$ we have

$$\begin{aligned}&\left.\frac{\partial \pi(\theta^*;\tau)}{\partial K}\right|_{K=K^*}\left(f(\theta^*)(1-\phi-K^*-\pi(\theta^*;\tau)(1+\eta+\tau))\right) \\ <&\left.\frac{\partial \pi(\theta^*)}{\partial K}\right|_{K=K^*}\left(f(\theta^*)(1-\phi-K^*-\pi(\theta^*)(1+\eta))\right)\,. \end{aligned} \tag{E.10}$$

On the other hand, similar calculation yields:

$$E_\tau[W]|_{K=K^*}\left.\frac{\partial \pi(\theta^*;\tau)}{\partial K}\right|_{K=K^*}\pi(\theta^*;\tau) < E[W]|_{K=K^*}\left.\frac{\partial \pi(\theta^*)}{\partial K}\right|_{K=K^*}\pi(\theta^*)\,. \tag{E.11}$$

Because $R-1-\tau < R-1$, conditions (E.10) and (E.11) thus jointly imply that the new FOC (E.9) also evaluated at K^* is smaller than the old FOC (E.6) 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. Following the proofs of Propositions 1 and 2, the run threshold when the stablecoin issuer imposes redemption fee ν can be re-derived as:

$$\pi(\theta^*;\nu) = \frac{(1-\nu)(1-\phi)\left(2-2\phi-2(1-\phi)\ln\left(\frac{1-\phi}{1-\nu}\right)-K\right)}{2(1-\nu+\eta(1-\phi))(1-\phi)+2(1-\nu)(\phi-\nu)\ln\left(\frac{\phi-\nu}{1-\nu}\right)}, \tag{E.12}$$

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

Proof of Proposition 6. By contradiction suppose that $\phi_1^* < \phi_2^*$. The optimality for issuer 2 implies:

$$\log\left(Q_2\left(\phi_2^*\right)\right) + \log\max_n F\left(\phi_2^*, n\right) \geq \log\left(Q_2\left(\phi_1^*\right)\right) + \log\max_n F\left(\phi_1^*, n\right) ,$$

and

$$\log\left(Q_2\left(\phi_2^*\right)\right) - \log\left(Q_2\left(\phi_1^*\right)\right) \geq \log\max_n F\left(\phi_1^*, n\right) - \log\max_n F\left(\phi_2^*, n\right) .$$

But (D.1) implies:

$$\log\left(Q_1\left(\phi_2^*\right)\right) - \log\left(Q_1\left(\phi_1^*\right)\right) \geq \log\left(Q_2\left(\phi_2^*\right)\right) - \log\left(Q_2\left(\phi_1^*\right)\right) \geq \log\max_n F\left(\phi_1^*, n\right) - \log\max_n F\left(\phi_2^*, n\right) ,$$

which further implies that:

$$\log\left(Q_1\left(\phi_2^*\right)\right) - \log\max_n F\left(\phi_2^*, n\right) \geq \log\left(Q_1\left(\phi_1^*\right)\right) - \log\max_n F\left(\phi_1^*, n\right) ,$$

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

F Additional Empirical Results

Table A.2: 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 of 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.3: 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 of 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

G Additional Calibration Details and Results

The CDS spread s_c on an asset class $c \in \{1 \ldots C\}$ 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:

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

Since p in the model is equal to the expected recovery value of assets per unit stablecoin issued, we will use (G.1) on each date we observe CDS spreads as one realization of p. We can think of (G.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 (G.1). By plugging CDS spreads from different dates into (G.1), we can calculate a distribution of signals p. Note that, when we plug CDS spreads into (G.1), we use spreads from a single day; hence, this method accounts for correlations between CDS prices of different asset classes.

We implement (G.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 6 US banks to capture the riskiness of the banking sector.[22] 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

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

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" 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.4 shows the parameters of the beta distributions we estimate.

Table A.4: 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.5: 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.2: 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

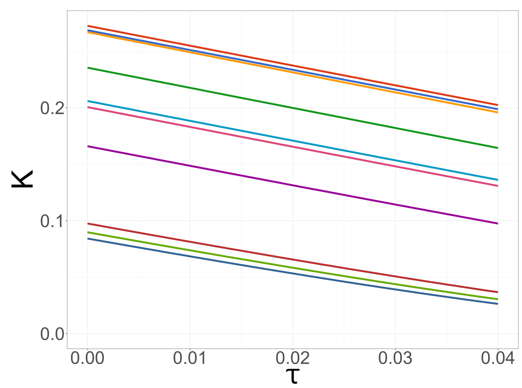

(b) Price Variance Cost

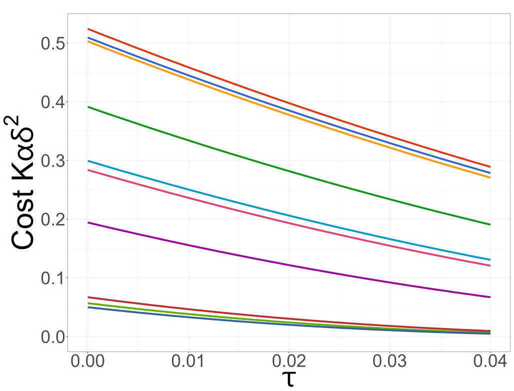

(c) Run Probability

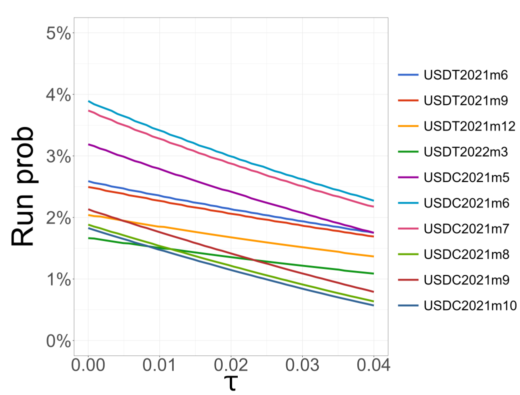

Figure A.3: Effect of Redemption Fees (Full Sample Period)

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.