Stablecoin Runs and the Centralization of Arbitrage*

Yiming Ma†
Columbia Business School

Yao Zeng‡
Wharton

Anthony Lee Zhang§
Chicago Booth

This version: March 2023

Click **here** for the latest version

*We thank Campbell Harvey, Xavier Giroud, Urban Jermann, Nellie Liang, Anna Pavlova, and Mila Sherman for detailed comments, and conference and seminar participants at AFA AFFECT, Duke/UNC/Milken DeFi Conference, Texas A&M, UT Dallas, and Wharton. 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

We analyze the run risk of USD-backed stablecoins and uncover a dilemma between stablecoins' price stability and financial stability. Stablecoin runs bear important financial stability implications through the fire sale of US dollar assets like bank deposits, Treasuries, and corporate bonds. We show that panic runs exist even though general investors only trade stablecoins in secondary markets with flexible prices. Run incentives are reinstated by stablecoin issuers' liquidity transformation and the fixed $1 at which arbitrageurs redeem stablecoins for cash in the primary market. We discover that more efficient arbitrage amplifies run risk. This explains why stablecoin issuers only authorize a small set of arbitragers even though it comes at the expense of maintaining a stable secondary price. In other words, the centralization of arbitrage embeds an inherent tradeoff between run risk and price stability. Our findings are based on a model and a novel dataset on stablecoin redemptions, trading, and reserve assets. Calibrating our model, we find a higher run risk for USDT, the largest stablecoin, compared to USDC, the second-largest stablecoin. However, even USDC bears significant run risk due to its less concentrated arbitrage and more concentrated deposit holdings.

1 Introduction

Stablecoins are blockchain assets whose value is claimed to be stable at $1. The main type of stablecoins, fiat-backed stablecoins, 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 the safe asset for the blockchain ecosystem as well as a means of payment for real purchases 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 stablecoins has also raised financial stability concerns.[1] Terra USD, one of the largest algorithmic stablecoins, experienced a sharp run in May 2022, which led to its collapse in a week (Liu, Makarov and Schoar, 2023).[2] In March 2023, Circle's USDC, the second largest fiat-backed stablecoin, also experienced a run amid the collapse of Silicon Valley Bank with its price plummeting by more than 15% within a few hours. These financial stability concerns have also been a major driving force behind efforts to introduce central bank digital currencies (Brunnermeier, James and Landau, 2019, Duffie, 2019, Auer, Frost, Gambacorta, Monnet, Rice and Shin, 2022, Makarov and Schoar, 2022).

Unlike other crypto assets, fiat-backed stablecoins are directly connected to the traditional financial system through their US dollar asset holdings. A run on them would not only lead to losses for stablecoin investors but could also contract bank deposit funding, strain US Treasury markets, and induce the fire sales of illiquid assets like corporate bonds. These ramifications may become even more pronounced going forward as stablecoins potentially become a more widely adopted means of payment and an increasingly important holder of financial assets. Thus, it is essential to understand whether runs could materialize in the future and what design features of stablecoins could affect their occurrence.

[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".

[2]Different from fiat-backed stablecoins, algorithmic stablecoins use a different pegging mechanism without physically holding a pool of reserve assets.

In this paper, we analyze the economics of US dollar fiat-backed stablecoins and shed light on the possibility and probability of stablecoin runs. Stablecoins are uniquely designed with features of both exchange-traded funds (ETFs) and money market funds (MMFs). The majority of investors trade stablecoins in competitive secondary markets. Fluctuations in the secondary market price reflect changes in demand and supply but do not involve any direct fire sale of assets similar to fluctuations in the price of ETF shares. Asset sales only occur when the stablecoin issuer meets redemption requests in the primary market. The issuer liquidates some of its assets to pay \$1 in cash for each stablecoin redeemed similar to MMFs, but redemption requests can only be submitted by a limited number of arbitrageurs approved by the stablecoin issuer. These arbitrageurs buy stablecoins trading below \$1 in secondary markets to redeem them for \$1 in primary markets, which allows them to pocket arbitrage profits while providing liquidity to investors.

Despite stablecoins' unique design and their tradability in competitive secondary markets, we show that they remain vulnerable to panic runs by investors in the spirit of Diamond and Dybvig (1983). This is because the fixed \$1 redemption price in the primary market reinstates run incentives among secondary market investors, who fear that APs will retract from providing liquidity to them if the stablecoin issuer can no longer honor the \$1 redemption value.

Interestingly, the concentration of arbitrageurs embeds an inherent tradeoff between run risk and price stability. If issuers only approve a small number of arbitrageurs to redeem tokens for cash in primary markets, arbitrage is less efficient and the same selling pressure would depress prices more in secondary markets. However, precisely because sellers in secondary markets would receive lower prices, their "first-mover advantage" from selling stablecoins in a run decreases. In other words, approving more arbitrageurs for more efficient arbitrage would exacerbate run risk and be counterproductive for financial stability.

More specifically, our dataset of fiat-backed stablecoins is constructed as follows. 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), on the Ethereum, Avalanche, and Tron blockchains. We obtain this data from each blockchain by converting transaction-level blockchain data into a usable format.

For each stablecoin, we also extract average trading prices in secondary markets from the main exchanges. Further, we obtain the composition of reserve assets for USDT and USDC, which reported these breakdowns at various points in 2021 and 2022.

From our novel data, we observe that the concentration of arbitrageurs in the primary market, where stablecoins are directly redeemed for cash with issuers, varies across stablecoins. For example, USDT only has 6 arbitrageurs redeeming shares during the average month and the largest arbitrageur accounts for 64% of the total redemption activity. In contrast, the arbitrage market at USDC is more competitive with 521 active arbitrageurs in an average month. We also find that trading prices in the secondary market for stablecoins frequently deviate from zero with discounts occurring 27.2% to 41.6% of the time and premia occurring 57.3% to 72.8% of the time. We note that these price deviations are not analogous to money market funds "breaking the buck" nor are they an indicator of runs. Rather, stablecoin prices fall below \$1 when secondary market investors' selling pressure is not fully absorbed by the arbitrageurs, who purchase stablecoins in secondary markets and redeem them for \$1 each in primary markets. In this sense, stablecoins trading below \$1 is similar to ETF shares trading at a discount to their NAVs.[3]

We further observe that stablecoins with fewer arbitrageurs have higher average discounts in secondary markets. For example, the average discount at USDT is 55bps, while the average discount at USDC is only 1bps. At the same time, USDT also has more illiquid assets, like corporate bonds and loans, as part of their reserve assets than USDC. These observations leave open the question of how stablecoins choose the concentration of their arbitrageurs and how the choice relates to their asset illiquidity. After all, if approving more arbitrageurs can minimize secondary price deviations, then why wouldn't all stablecoin issuers simply allow for a competitive arbitrage market?

We develop a model to rationalize our empirical observations, assess the potential for stablecoin runs, and analyze the effect of market structure. Our theory applies a Diamond-Dybvig-style model to stablecoins and characterizes its unique design with features of both ETFs and MMFs. There are three types of agents: investors, arbitrageurs, and a stablecoin issuer. Specifically, investors are endowed

[3] 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.

with stablecoins that are aimed at providing a fixed value and backed by an illiquid reserve asset. They may sell stablecoins to arbitrageurs in the secondary market but they cannot directly redeem them from the issuer, similar to the case of ETFs. Arbitrageurs bid in a double auction to absorb any residual selling pressure from investors, and can then redeem stablecoins with the issuer in the primary market for one dollar, which resembles the redemption of MMFs shares. To honor the fixed redemption price of one dollar, the issuer liquidates its illiquid reserve asset pre-maturely until it defaults, after which only the liquidation value will be paid to redeeming arbitrageurs.

Our model shows that panic runs by investors on stablecoins can happen despite investors only being able to sell stablecoins in the secondary exchange at the market price. The conventional view is that exchange-traded claims like stocks and ETF shares are less runnable than bank deposits because the trading price falls as more investors sell, which creates a natural strategic substitutability. In the context of stablecoins, however, arbitrageurs are promised a fixed redemption price by the issuer. Hence, investors who choose to hold stablecoins may end up getting a less valuable stablecoin in the future because they bear the costs induced by the issuer's firesales of illiquid assets while meeting arbitrageurs' redemptions at $1. In this way, stablecoins' fixed primary market price re-introduces strategic complementarity among secondary market investors.

We endogenize the run probability using global games to evaluate the effect of arbitrage efficiency. Surprisingly, we find that the run risk of stablecoins decreases in the concentration and increases in the balance sheet capacity of the arbitrage market. To understand this result, note that investors compare the benefit of selling their shares in the secondary market, i.e., the secondary market price, to the benefit of remaining invested in the stablecoin in the long run. When arbitrage is more efficient, price stability in the secondary market is higher because arbitrageurs are more willing to absorb selling pressure from investors. This higher trading price increases the benefit of selling stablecoins for investors and thereby amplifies their first-mover advantage when they expect other investors to sell. In contrast, when arbitrage is less efficient, small quantities of investor sales can have a substantial impact on stablecoin prices so the risk of secondary price discounts is higher. This price impact of stablecoin sales discourages panic selling and thereby mitigates run risk. Therefore, arbitrageurs act as a firewall between the primary and secondary markets for stablecoins, which induces a trade-off between stablecoins' price stability and run risk.

Stablecoin issuers optimally design the structure of their arbitrage sectors to trade off price stability and run risk. Recall that USDT holds more illiquid assets than USDC while also approving fewer arbitrageurs than USDC. This is consistent with USDT restricting entry into their arbitrage market to partially offset the increased run-risk from holding more illiquid reserve assets.

Our model further provides an analytical solution for stablecoins' run probability, which we calibrate to quantify the run risk of the two largest stablecoins, Tether (USDT) and USD Coin (USDC). Our first input is the elasticity of redemptions in the primary market. Based on the model, redemption volumes should be more responsive to deviations in the secondary market price when there is a larger number of arbitrageurs. Empirically, we regress daily discounts against daily redemption volumes normalized by the total outstanding volume for each stablecoin. We find that the coefficient for USDT is larger in absolute magnitude than for USDC, which is consistent with the higher AP concentration of USDT constraining redemption volume to be less sensitive to price dislocations. Magnitude-wise, a 10 percentage point higher redemption volume corresponds to a 3.0 cent larger discount at USDT and a 1.3 cent larger discount at USDC.

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. We follow Bai, Krishnamurthy and Weymuller (2018) to proxy asset discounts with collateral haircuts by asset class. Intuitively, more liquid assets are more readily pledged to obtain cash at short notice while more illiquid assets incur a higher discount. On average, the reserve assets of USDT are more illiquid than those of USDC, but both of them shift towards holding more liquid assets over the sample period.

We estimate the distribution of the probability at which the risky asset payoff does not materialize. We use CDS spreads to evaluate the expected recovery value of each portfolio component and then calculate how the expected recovery value of the stablecoin issuer's overall reserve portfolio varies over time using historical data. The resulting empirical distribution is close to but not only concentrated at 1, consistent with USDT and USDC holding mostly but not exclusively safe assets.

Finally, we use our model to quantify the run risk at the two largest US dollar stablecoins. Tether and Circle, which make up the bulk of the market at $76.4 billion and $37.7 billion in January 2022. To calibrate our model parameters, we construct a novel dataset comprising of stablecoins' primary market transactions, secondary market trades, and reserve asset composition. Our estimates imply a higher run risk for USDT, the largest stablecoin, compared to USDC, the second-largest stablecoin, due to higher liquidity transformation. However, USDC also processes significant run risk due to less concentrated arbitrage and more concentrated deposit holdings.

Our paper contributes to a large literature on 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, 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 without direct liquidation of the underlying assets (Jacklin, 1987, Allen and Gale, 2004, Koont, Ma, Pastor and Zeng, 2021). By carefully modeling the unique combination of ETFs and MMFs in the design of stablecoins, we show that panic runs may still happen despite their trading on competitive secondary markets and investors' inability to access primary markets.

Methodologically, our estimation of run risks is enabled by the use of the global games approach to derive a unique run threshold. Goldstein and Pauzner (2005) shows that in the classic Diamond and Dybvig (1983) bank setting global strategic complementarity fails and thus the standard global games approach (e.g., Morris and Shin, 1998) does not directly apply, but a unique threshold run equilibrium exists as long as there is one-sided strategic complementary. In the stablecoin setting where the secondary and primary markets are separated, even one-sided strategic complementarity may not hold because selling the first unit of stablecoin in the secondary market generates a first-order price impact. However, we are able to show that a unique threshold run equilibrium still exists, which provides a foundation for our calibration to quantify stablecoin run risks. Relatedly, Egan, Hortacsu and Matvos (2017) build a structural model to quantify bank instability, highlighting the feedback between endogenous bank default and deposit withdrawals.

We also contribute to the emerging stablecoin literature by analyzing and quantifying the run risk of US dollar stablecoins. Several recent papers explore runs on algorithmic stablecoins (e.g., Adams and Ibert, 2022, Uhlig, 2022) after the Terra-Luna crash in 2022. Closely related to our work is Liu, Makarov and Schoar (2023), who examine the dynamics of the Terra USD run, focusing on how investor characteristics affect run behavior and the financial inclusion implications. On fiat-backed stablecoins, Barthelemy, Gardin and Nguyen (2021) and Liao and Caramichael (2022) analyze the potential impact of fiat-backed stablecoin activities on the real economy. Frost, Shin, Wierts (2020), Gorton and Zhang (2021), and Gorton, Ross and Ross (2022) compare stablecoins to the banking sector pre-deposit-insurance. Griffin and Shams (2020) suggest that, prior to 2020, Tether was used to manipulate Bitcoin prices. Lyons and Viswanath-Natraj (2021) show that USDT's creation and redemption respond to price deviations and who relate stablecoin price stability to defending exchange rate pegs. Kim (2022) finds that increases in the issuance of USDT and USDC lead to decreases in Treasury and commercial paper yields. Kozhan and Viswanath-Natraj (2021) analyze DAI, which is a stablecoin overcollateralized by risky non-USD assets. Li and Mayer (2021), d'Avernas, Maurin, and Vandeweyer (2022) and Routledge and Zetlin-Jones (2022) are theoretical papers on the mechanisms stablecoins use to maintain peg stability, encompassing algorithmic and collateral-backed stablecoins as well as fiat-backed stables. We provide a complementary yet distinct perspective of stablecoins as financial intermediaries engaged in liquidity transformation. Through this lens, we highlight the possibility of panic runs and relate run risk to the design of the primary market.

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. Finally, Section 7 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, 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 transfer tokens to other wallets. Others have no access to individuals' private keys so they have no ability to take funds from individuals' wallets. Stablecoins are thus a useful way to hold US dollar assets in settings where there is a lack of trusted financial intermediaries that can be relied on to custody US dollar assets on behalf of market participants.

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' time, 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 2). The largest stablecoin is Tether (USDT), which made up more than 50% of the total market size at $76.4 billion in January 2022. Circle USD Coin (USDC) and Binance USD (BUSD) are second and third at $37.7 and $14.4 billion. Paxos, (PUSD), TrueUSD (TUSD), and Gemini dollar (GUSD) are significantly smaller with a market size of around or below $1 billion. The asset size of fiat-backed stablecoins has experienced ups and downs in 2022 but remains high at $136 billion in June 2022.

We proceed to explain the design of stablecoins and how they attempt to achieve price stability. A diagram illustrating the primary and secondary market for stablecoins is shown in Figure 1.

2.1 The Primary Market

Stablecoin tokens are created/minted and redeemed/destroyed in the primary market with US dollar cash as shown on the left-hand side of Figure 1. To create a stablecoin token, an arbitrageur sends $1 US dollar to the issuer, through a bank transfer or other means; 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 US dollar, 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.

Importantly, not all market participants can freely become arbitrageurs to participate in the redemption and creation of stablecoin tokens in the primary market. Stablecoin issuers differ in how easily and costly market participants can access primary markets. For example, while USDC allows general businesses to register as arbitrageurs and charges no fees for redemptions and creations, USDT restricts AP registration, imposes a minimum transaction size of $100,000, and charges the greater of 0.1% and $1000 per redemption. USDT also requires a lengthy due-diligence process and imposes restrictions on where arbitrageurs can be domiciled.

2.2 The Secondary Market

The majority of market participants trade existing stablecoins for fiat currencies in secondary markets, as shown on the right-hand side of Figure 1. Crypto-exchanges like Binance allows customers to make US dollar deposits, and then trade US dollars for USDT, USDC, or BUSD with other market participants.[4] The price of stablecoin tokens in the secondary market is thus driven by the demand from stablecoin buyers and the supply from stablecoin sellers. When there is a surge in stablecoin sellers on the secondary market, the secondary market price would drop but the closed-end nature implies that there are no direct liquidations of any reserve assets involved. In this way, the buying and selling of stablecoins on secondary markets resemble the trading of ETF shares on the exchange.

[4] Please see Appendix A for details regarding the use of different crypto exchanges.

2.3 Shock Transmission from the Secondary to the Primary Market

Nevertheless, selling pressure in the secondary market for stablecoins can spill over to affect the primary market through arbitrageurs. When investor selling pressure in the secondary market depresses stablecoin prices to be below \$1, arbitrageurs can profit from purchasing stablecoin tokens for below \$1 in secondary markets, and redeeming them one-for-one for \$1 with the stablecoin issuer in primary markets as long as the issuer does not default. Analogously, if positive demand shocks in secondary markets caused stablecoins to trade above \$1, arbitrageurs could profit from creating stablecoin tokens one-for-one for dollars in primary markets and then selling them at higher prices in secondary markets. Thus, the \$1 redemption value of stablecoins in primary markets pulls the trading price of stablecoins towards \$1 in secondary markets through the trading incentive of arbitrageurs.

At the same time, this arbitrage process implies that investor selling pressure in secondary markets can eventually trigger fire sales of assets when stablecoin issuers liquidate reserves to meet arbitrageurs' \$1 redemption in cash. These fire sales can become especially costly if large amounts of redemptions occur in a short period of time and if illiquid reserve assets can only be converted to cash at a discount. If redemptions and discounts are large enough, the issuer may fail to pay the promised \$1 for each stablecoin token redeemed, and the stablecoin defaults.

2.4 Uses of Stablecoins

Stablecoins have a number of uses. First, they are a fairly low-cost way to transact in US-dollar assets. As of January 2023, sending tokens on the Ethereum costs around \$1 per transaction, and transactions finalize in under a minute.

Stablecoins are also being used as a store of value and medium of exchange in settings where inflation is high, capital controls and financial repression are prevalent, and trust in intermediaries is low. For example, humanitarian organizations have used stablecoins to circumvent banking fees and regulatory frictions.[5] Some firms in Africa have begun using stablecoins for international payments to

[5] See Fortune.com and Rest Of World.

suppliers in Asia.[6] In settings with high inflation, such as Lebanon and Argentina, individuals have begun storing value and transacting using stablecoins.[7] Some merchants in these areas have begun accepting stablecoins as a form of payment.[8]

Finally, stablecoins are used with other smart contracts within the space of "decentralized finance." 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 to use these assets as collateral to borrow other assets. In a way, stablecoins provide a safe store of value and medium of exchange resemble for the blockchain ecosystem similar to the role of deposits and money market fund shares in the traditional financial system.

3 Data

We compile a novel and comprehensive dataset that sheds light on stablecoins' on-chain primary market activity, secondary market prices, and reserve assets.

3.1 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: Tether (USDT), Circle USD Coin (USDC), Binance USD (BUSD), Paxos (USDP), TrueUSD (TUSD), and Gemini dollar (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.

[6] See Rest Of World.

[7] See CNBC and Rest Of World for a discussion of the Lebanon case, and Coindesk and EconTalk for a discussion of Argentina.

[8] For example, the Unicorn Coffee House in Beirut, Lebanon accepts USDT (Tether) as a form of payment.

As described in Section 2, 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 quantity 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.[9]

Using our data extraction process, we see, for each stablecoin creation and redemption event, the precise timestamp of the event; the amount of the stablecoin redeemed or created; and the wallet address of the entity involved in stablecoin creation or redemptions. We also observe the "gas" fee – that is, the transaction fee paid to Ethereum miners for including the transaction in the blockchain – paid for each transaction. Some wallet addresses are tagged on Etherscan, as they are known to belong to large entities such as crypto exchanges. Using Etherscan wallet tags, we are able to group some wallets that are known to belong to the same economic entity.

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

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

3.2 Secondary Market Data

For each of the 6 stablecoins in our data, we extract the hourly closing prices for trades from the main exchanges, including Binance, Bitfinex, Bitstamp, Bittrex, Gemini, Kraken, Coinbase, Alterdice, Bequant, and Cexio. 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.

3.3 Reserves

Stablecoins' reserve assets are not recorded on the blockchain. However, USDT and USDC reported breakdowns of their reserve assets at various points in 2021 and 2022 as part of their balance sheets. We obtain these breakdowns for USDT and USDC. The other four stablecoins have not released breakdowns of their reserve asset composition but state the broad categories of their reserves. We obtain and discuss these asset types in the next section.

4 Stylized Facts

In this section, we present a set of new stylized facts about stablecoins, which informs our model and calibration to quantify the run risk of stablecoins.

4.1 Secondary Market Price

Fact 1. *The trading price of stablecoins in the secondary market commonly deviates from $1. This price deviation per se does not constitute a run by investors.*

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 to \$1 27.2% to 41.6% of the time and trade at a premium to \$1 57.3% to 72.8% of the time for our sample of stablecoins (see Table 2).

The extent of these price deviations varies by stablecoin. While the average discount at USDT is 55bps, the average discount at USDC is only 1bps. The average discount of BUSD, TUSD, and USDP are also below that of USDT at 1bps, 11bps, and 18bps, respectively, while that of GUSD is the highest at 78bps. The median discounts are generally smaller in magnitude than the average discounts, but the variation in the cross-section remains similar. The average and median premia also show significant variation in the cross-section.

The trading of stablecoins at a discount to \$1 has been commonly associated with "breaking the buck" as in the case of money market funds and even as evidence for panic runs.[10] We note that these are misconceptions. Stablecoins maintaining a "stable value" of \$1 refers to the amount that primary market participants receive or pay when they redeem existing stablecoins or create new stablecoins with the stablecoin issuer. The notion of "breaking the buck" thus corresponds to primary market participants not receiving a full \$1. This scenario has not yet occurred at any of the stablecoins in our sample despite their secondary market price frequently deviating below \$1. 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. This stablecoin price falling below \$1 simply captures the selling pressure of stablecoins in the secondary market and is not a direct indicator of "breaking the buck" or panic runs.

4.2 Primary Market Concentration

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

[10]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

Table 3 shows the characteristics of daily primary market redemption activity on the Ethereum blockchain for different stablecoins. We observe that on an average day, USDT only has 2 APs engaged in redemptions, whereas USDC has 33. The concentration of APs' market shares also varies. The largest AP at USDT performs 93% of all redemption activity, while the largest AP at USDC performs 54%. BUSD, USDP, and TUSD lie in between USDT and USDC in terms of the number of redeeming APs and AP concentration. GUSD has the most concentrated AP market with one AP essentially being in charge of all redemptions.

We repeat the analysis at the monthly level in Table 4. The monthly snapshot may better capture the market structure of the primary market than the daily snapshot if not all APs are active every day. Indeed, we observe that the number of APs redeeming stablecoins is larger at the monthly level. However, the AP market remains highly concentrated for USDT, with only 6 APs redeeming shares during the average month and the largest AP accounting for 64% of the total redemption activity. In contrast, USDC has 521 active APs in an average month but the top 1 and top 5 APs make up 45% and 85% of all redemption activity. As before, USDP, and TUSD lie in between USDT and USDC in terms of the number of redeeming APs and AP concentration. GUSD has the most concentrated AP market at the monthly level as well.

Further, notice that in the average month, the volume of redemptions at USDT is \$615 million, while that at USDC is \$2976 million. In comparison, the total volume of outstanding tokens at USDT was 1.5 to 2 times of that of USDC. Thus, the larger number and lower concentration of APs at USDC is correlated with a higher volume of redemptions relative to the total asset size as well.

In the appendix, we repeat Tables 3 and 4 for the Tron and Avalanche blockchains and obtain similar variations in AP 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 discounts in the secondary market.*

We proceed to analyze the potential effects of AP concentration. We calculate the average monthly discount and the average number of redeeming APs for each stablecoin and plot them in Figure 4a. A clear negative trend emerges: stablecoins with fewer APs, like USDT and GUSD, have higher average discounts in their secondary market prices than stablecoins with more APs, like USDC and BUSD. Another way to capture AP concentration is through the market share of the largest APs. In Figure 4b, we repeat the analysis with the market share of the top 5 APs. The relationship is positive. Stablecoins whose top 5 APs consistently perform a larger share of total redemptions, like USDT and GUSD, have higher average discounts than other stablecoins with lower AP concentration. In other words, it seems that higher AP competition is associated with reduced price dislocations in secondary markets.

One question arising from this trend is why some stablecoins choose to have a more concentrated AP sector. If AP competition can indeed stabilize secondary market prices, all stablecoins should be incentivized to open up AP access and encourage the entry of new APs. 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 AP sector. We show that the probability of panic runs is especially pronounced if the reserve assets are more illiquid, which makes AP concentration even costlier. In the next subsection, we illustrate that USDT indeed also has more illiquid reserve assets.

4.4 Liquidity Transformation

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

Stablecoins are not literally backed by US dollars in the form of cash. Rather, they hold USD-denominated assets with varying degrees of illiquidity as reserves. Table 1 shows the composition of reserve assets for USDT and USDC on reporting dates. Overall, reserve assets of both USDT and USDC are far from being 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. In June 2021, these two asset classes took up 60.7% and 59.5% of reserve assets at USDT and USDC, respectively. Money market instruments include commercial paper and certificates of deposits. For USDT, deposits include "cash deposits at financial institutions and call deposits, i.e., deposits that may be withdrawn with two days' notice or less; fiduciary deposits, i.e., deposits made by banks on behalf of and for the benefit of members of the consolidated group; and, term deposits, i.e., deposits placed by members of the consolidated group at its banks for a fixed term." For USDC, deposits include "US dollar deposits at banks and short-term, highly liquid investments that are readily convertible to known amounts of cash and have a maturity of less than or equal to 90 days from purchase." Thus, except for deposits in checking accounts, money market instruments and other types of deposits are not fully liquid, i.e., they can not be freely converted to cash at short notice. For example, time deposits and certificates of deposit experience a discount when demanded before their maturity date.

USDT also holds a significant portion of reserves in the form of Treasury bills, which increased from 24.3% in June to 47.6% in March 2022. In contrast, USDC reduced its Treasury holdings from 15.0% in July 2021 to 0% in August 2021. USDC states that their Treasuries include "US government treasury bills, notes and bonds with a maximum maturity of 3 years". While Treasuries are generally liquid and safe security, the extent of their liquidity varies by type and over time. For example, on-the-run Treasuries and Treasury bills are much more liquid than off-the-run Treasuries and non-bills.

The remaining reserve assets are comprised of more illiquid assets, including municipal and agency securities, foreign securities, corporate bonds, corporate loans, and other securities. USDT still held a sizable amount of these illiquid assets in March 2022, with 4.5%, 3.8%, and 6.0% in corporate bonds, corporate loans, and other assets, respectively. While the exact identity of other assets is not disclosed, it is mentioned that they do include crypto investments. In June 2021, USDC held 0.4%, 15.9%, and 9.5% in municipal and agency securities, foreign securities, and corporate bonds respectively. USDC's holding of these assets is reported to have dropped to zero starting in September 2021, with all assets held in the form of the deposits described above.

The other four stablecoins do not publish reserve breakdowns, but they report that their assets are limited to deposits, Treasuries, and money market instruments. For example, a statement issued by BUSD and USDP in July 2021 claims that they hold 96% of cash equivalents and 4% of Treasury bills. GUSD states that their reserves are "held and maintained at State Street Bank and Trust Company, Signature Bank, and within a money market fund managed by Goldman Sachs Asset Management, invested only in U.S. Treasury Obligations." TUSD also claims that their US dollar balance is held by

"U.S. depository institutions and Hong Kong depository institutions" and that they "include US dollar cash and cash equivalents that include short-term, highly liquid investments of sufficient credit quality that are readily convertible to know amounts of cash."

5 Model

In this section, we build a model to analyze the potential for stablecoin runs. The model aims for achieving three goals. First, the model formulates the notion of runs on the primary market of stablecoins and explicitly derives the run probability, linking it to the level of stablecoin liquidity transformation and the concentration of arbitragers. Second, the model formulates the stablecoin issuer's optimal design of its primary market structure. Finally, the model allows us to quantify the run risks for a number of major stablecoins.

5.1 Setting

The economy has three dates, $t = 0, 1, 2$, with no time discount. There are three groups of risk-neutral players, 1) a competitive fringe of identical, infinitesimal investors indexed by i, 2) a sector of $n \geq 3$ arbitrageurs or APs, and 3) a stablecoin issuer.

At $t = 0$, investors are born; each investor would incur a cost of c_i, which follows a distribution function $G(c)$, to participate in the stablecoin market. Once participated, each investor is endowed with one stablecoin. An investor participates only when its participation cost is smaller than its expected utility from participation, which will be determined in equilibrium. There are two types of assets: the dollar, which is also the consumption good and serves as the numeraire, and an illiquid and potentially productive reserve asset. The stablecoin is initially backed by the reserve asset held by the issuer. The initial value of the reserve asset is normalized to one dollar. We introduce the features of the two assets shortly below. The stablecoin issuer may also choose n at $t = 0$, that is, design the structure of its primary market, to maximize its expected profit, which we introduce in Section 5.4. Before that, we take n as exogenous and focus on investors' equilibrium behaviors after they have participated, and as such, we also normalize the population of participating investors to one.

Participating investors are subject to idiosyncratic liquidity shocks at $t = 1$ as in Diamond and Dybvig (1983). Each investor is uncertain about her preferences over consumption at $t = 1$ and $t = 2$. At the beginning of $t = 1$, an investor learns her preferences privately: with probability $\pi > 0$ she is an early-type and gets utility from date-1 consumption only, while with probability $1 - \pi$ she is a late-type and gets utility from consumption from both dates.

At $t = 1$, a total of $\lambda \geq \pi$ investors decide to sell stablecoins to APs in a secondary market at the market price p in exchange for a consumption good called dollar, where both λ and p will be determined in equilibrium. Dollar is riskless, liquid, and it serves as the numeraire. Households cannot directly redeem the stablecoin for dollar from the issuer, but APs are able to redeem the stablecoins from the issuer in a primary market, getting a fixed price of one dollar per stablecoin if the issuer is solvent. To raise dollars to meet AP redemptions at $t = 1$, the issuer has to liquidate the illiquid reserve asset prematurely at a liquidation cost of $\phi \in (0, 1]$, that is, liquidating one unit of asset yields $1 - \phi$ dollar only. Hence, the issuer is solvent if and only if $\lambda < 1 - \phi$. We assume $\pi < 1 - \phi$ to rule the uninteresting case that early investors alone render the issuer default. When $\lambda \geq 1 - \phi$, the issuer defaults, and the redeeming APs will get the liquidation value per total stablecoins redeemed, that is $(1 - \phi)/\lambda$. Expecting the amount of dollars to be redeemed from the issuer, APs bid in a double auction (e.g., in the manner of Kyle (1989) and Du and Zhu (2017)) to buy the stablecoins from λ selling investors. Denote the AP sector's total balance sheet capacity in the auction by S. The auction determines the secondary-market price p, the magnitude of which reflects the de-pegging risk of the stablecoin.

At $t = 2$, the economy entails aggregate risk. With probability $p(\theta)$, the economy enters a good state: the reserve asset matures and yields a value of $R(\phi) \geq 1$ dollar. With probability $1 - p(\theta)$, the economy enters a bad state: the reserve asset fails and yields zero. Here, $R(\phi)$ is decreasing in ϕ and $p(\theta)$ is increasing in $\theta \in \Theta$. We call θ the fundamentals of the economy which captures the level of aggregate risk, which is unknown to investors, APs, or the stablecoin issuer before $t = 2$. Intuitively, the reserve asset is more likely to profit as the fundamentals are better, and its long-term maturing value $R(\phi)$ increases in its illiquidity, capturing a notion of liquidity risk premium.

In the good state of the economy, participating investors and the stablecoin issuer share the value of the reserve asset based on the following rule. Unlike a security, a stablecoin never pays dividends. Thus, the net maturing gain of the reserve asset, that is, $R(\phi) - 1$, only gets accrued to the stablecoin issuer but not the investors. Rather, the remaining $1 - \lambda$ investors consume the initial value 1 per unit of the remaining reserve asset, plus a convenience value $\eta > 0$ per stablecoin at $t = 2$. Beyond the stablecoin, investors cannot access the underlying asset market or any other investment technology to transfer wealth across time.

To endogenize investors' stablecoin selling decisions and hence the stablecoin's run risk, we follow the global games literature to assume that each investor i obtains a private signal $\theta_i = \theta + \varepsilon_i$ at $t = 1$, 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. An investor's selling decision depends on the signal that she obtains. Note that we do not impose any restrictions on the distributions of p, θ, or the increasing function $p(\theta)$, which would conveniently allow us to map the model to any empirical distribution of fundamentals.

5.2 Discussion of Model Specification

Before proceeding, it is useful to discuss several important modeling choices to highlight the economics underlying the model. The discussion also highlights in what sense our model parsimoniously captures the most important features of the stablecoin markets.

First, our model purposefully features the separate but connected primary and secondary markets of stablecoins, as discussed in Section 2 and observed in Section 4. This separation is important for us to separately define the de-pegging and run risks of stablecoins and to analyze the relationship between these two types of risks. In reality, most retail investors of stablecoins cannot directly participate in the stablecoin creation and redemption process with the issuers, which is captured by investors only being accessible to the secondary market. Given our focus on stablecoin selling and redemptions, any excess supply for stablecoins by investors in the secondary market then must be met by AP redemptions with the issuer in the primary market. Our modeling of the AP's activity as a double auction with a fixed redemption price from the issuer closely mirrors the real-world redemption and destroy process of stablecoins, which is only available to the small set of APs.

It is also instructive to draw connections to ETFs and MMFs to further highlight the uniqueness of the stablecoin market and how our model captures this uniqueness. Like stablecoins, ETFs also feature the separation of the primary and secondary markets in that only APs can access the primary market and any excess demand or supply of ETF shares from investors in the secondary market must be met by APs (e.g., see Koont, Ma, Pastor and Zeng (2021) for a model of ETFs highlighting these features). However, ETFs notably differ from stablecoins in that AP creations and redemptions are predominantly performed in-kind with the issuer, that is, APs are delivered the underlying assets rather than cash when redeeming ETF shares. In contrast, a stablecoin AP gets a fixed amount of one dollar when redeeming one unit of stablecoin with the issuer provided the stablecoin issuer is solvent. This key difference thus resembles MMFs before the 2016 Money Market Reform in that MMF investors also get a fixed amount of one dollar in redemption provided the issuer is solvent (e.g., see Parlatore (2016) for a model of MMFs). Note that, however, MMF shares are not tradable in any secondary markets. Hence, stablecoins uniquely combine the two-layer market structure of ETFs and the fixed-value in-cash redemption feature of MMFs, and our model parsimoniously captures this combination.

Our model captures liquidity transformation and concentration of APs as the two most important economic sources of variation across different stablecoins, as documented in Section 4. Liquidity transformation is captured by the illiquidity cost parameter ϕ. Below, we call ϕ as either liquidation cost or haircut. Economically, ϕ captures that the liquidation of loans and bonds in secondary markets can depress their prices (see Duffie, 2010, for a review). It may also capture negative real impacts when liquidations of loans and bonds affect the capacity of governments and corporates rolling over their debt (e.g. He and Xiong, 2012). The concentration of APs is captured by the parameter n, holding the AP sector's total balance sheet capacity in bidding fixed. This specification helps us separately consider the effects of AP concentration and AP balance sheet costs on de-pegging and run risks.

The model also parsimoniously captures how the stablecoin issuer and coin-holding investors share the long-term value of the stablecoin. As of now, stablecoins do not pay any dividends to investors and thus are not regulated as securities by the U.S. SEC. Hence, the net return of the underlying reserve assets is only accrued to the issuer. However, by holding stablecoins in the long term (rather than selling them to meet immediate liquidity demand), participating investors enjoy a convenience yield that is currently specific to the use of stablecoins as a payment method in crypto investments and decentralized finance contexts, and potentially beyond as a widely adopted means of payment going forward. As we further specify in Section 5.4, the expected revenue of the stablecoin issuer hinges on this specific form of value sharing between the issuer and investors.

Finally, our model follows the global games approach to endogenize the run risk of stablecoins, which is the focus of this paper. One key assumption of the global games approach is the information structure: the fundamentals are unobservable but each agent obtains a private signal about them. We view this assumption to be plausible for the stablecoin market because of its opacity: essentially no stablecoin issuers disclose asset-level information about their reserves. On the other hand, investors in the stablecoin market are likely to be more sophisticated than those in more traditional financial markets, justifying their ability to obtain private and heterogeneous signals about the fundamentals.

5.3 Equilibrium Analysis

We first solve for the equilibrium secondary-market stablecoin price at $t = 1$, p, when λ investors choose to sell. Define

$$K = \frac{1}{S}\frac{n-1}{n-2}\,, \tag{5.1}$$

and impose the following parametric assumption to ensure that the secondary-market price is positive:

$$1 - \phi - K > 0\,. \tag{5.2}$$

We have the following result:

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

$$q(\lambda) = \begin{cases} 1 - K\lambda & \lambda \leq 1 - \phi\,, \\ \dfrac{1-\phi}{\lambda} - K\lambda & \lambda > 1 - \phi\,. \end{cases} \tag{5.3}$$

Proposition 1 shows that the stablecoin's secondary-market price depends on selling pressure λ, and further, the level of liquidity transformation ϕ, AP balance sheet capacity S, and AP concentration n. Specifically, q is decreasing in λ and ϕ while increasing in S and n. All these comparative statics are intuitive. A higher selling pressure λ depresses stablecoin price due to the standard excess supply effect, leading to higher de-pegging. A higher level of liquidity transformation ϕ does not affect the stablecoin price when the issuer is solvent but translates to a lower stablecoin price when the issuer defaults because of the lower liquidation value. Indeed, a higher ϕ also makes the issuer more likely to default. A higher AP balance sheet capacity implies that APs are more willing to bid to absorb the selling pressure, supporting a higher secondary-market price. Finally, a less concentrated AP sector, that is, a higher n implies that APs bid more competitively, also leading to a higher secondary-market price. Looking forward, we will show that these features play an important role in determining the relationship between secondary-market de-pegging risk and primary-market run risk of stablecoins.

Viewing the stablecoin's secondary-market price q as a function of λ specifically, we highlight two important features of $q(\lambda)$. First, it is strictly decreasing in λ everywhere. Second, it features a kink at $\lambda = 1 - \phi$, that is, when the stablecoin issuer just defaults. The first feature points to the standard notion of strategic substitutability usually present in many financial markets including the ETF market: the more investors sell, the lower the price is, making an investor less likely to sell. However, we show that the second feature, that is, the jump in price due to the issuer's inability to keep the fixed redemption price as it defaults, which resembles MMFs, may eventually give rise to a strong enough first-mover advantage in selling, as we analyze later.

Now we consider the late investors' decision of selling the stablecoin at $t = 1$ or not (recall that early investors always sell). In making this decision, a late investor compares the secondary-market price $q(\lambda)$ she may get by selling the stablecoin at $t = 1$ to the return she may get at $t = 2$ if she does not sell, which is given by

$$v(\lambda) = \begin{cases} p(\theta)\left(\dfrac{1-\phi-\lambda}{(1-\phi)(1-\lambda)} + \eta\right) & \lambda \le 1-\phi\,, \\ 0 & \lambda > 1-\phi\,. \end{cases} \tag{5.4}$$

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

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

unit of the reserve asset to meet AP redemptions at $t = 1$, and only $1 - l(\lambda)$ unit remains at $t = 2$, whose value will be shared by the remaining $1 - \lambda$ late investors.

It is useful to compare the date-2 stablecoin value (5.4) to the date-1 secondary-market stablecoin price (5.3) and define a late investor's payoff gain of waiting until $t = 2$ versus selling at $t = 1$:

$$h(\lambda) = v(\lambda) - q(\lambda) = \begin{cases} p(\theta)\left(\dfrac{1-\phi-\lambda}{(1-\phi)(1-\lambda)} + \eta\right) - 1 + K\lambda & \lambda \le 1-\phi\,, \\ -\dfrac{1-\phi}{\lambda} + K\lambda & \lambda > 1-\phi\,. \end{cases} \tag{5.5}$$

It is easy to see that $h(0) \ge 0$ when $p(\theta)$ 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.

The intuition behind a stablecoin run can be further illustrated by Figure 5, which plots the function $h(\lambda)$. It is clear from Figure 5 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 secondary market of stablecoins. Because selling investors generate a price impact on the secondary market and depress the secondary-market stablecoin price, a late investor may find it less appealing to sell if other late investors sell if the price impact is sufficiently large. This force works against the standard first-mover advantage that is typically present in a bank run model like Diamond and Dybvig (1983) and acts to prevent a run from happening. However, because APs redeem stablecoins from the issuer at a fixed price of \$1, the cost that waiting investors have to bear will increase as more and more late investors choose to sell. This force will ultimately dominate, offsetting the secondary-market strategic substitutability and leading to a decreasing $h(\lambda)$ in the second region. Eventually, the second force dominates as λ becomes sufficiently large, pushing $h(\lambda)$ to be negative, which reinstalls the first-mover advantage and leads to runs in equilibrium.

Under the global games framework, we have the following result:

Proposition 2. *There exists a unique threshold equilibrium in which late investors sell the stablecoins if they obtain a signal below threshold θ^* and do not sell otherwise.*

Proposition 2 implies that the model with investors' private and noisy signals has a unique threshold equilibrium. A late investor's selling decision is uniquely determined by her signal: she sells the stablecoin at $t = 1$ 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:

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

Solving the Laplace equation gives the following result:

Proposition 3. *The run threshold is given by*

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

which satisfies the following properties:

i). The run threshold, that is, run risk, is increasing in ϕ if and only if $g(\phi) > K$,[11] *where $g(\phi)$ is continuous and strictly decreasing in ϕ over $(0, 1-\pi)$, and satisfies $\lim_{\phi\to 0} g(\phi) > 0$.*

[11] The explicit form of $g(\phi)$ is given by

$$g(\phi) = \frac{2\left((\pi+\phi-1)((\phi-1)(\eta(\phi-1)+1)+\pi-\ln\phi) + (\phi-1)\ln(1-\phi)(\pi(\eta(\phi-1)-2) - 2\phi + (\phi+1)\ln\phi + 2) + \ln(1-\pi)\left(\pi - \left(\phi^2-1\right)\ln(1-\phi) + \phi - 1\right)\right)}{\left(1-\pi^2\right)\left((1-\phi)(1-\eta(1-\phi)) - \pi + \ln\phi - \ln(1-\pi)\right)} \,.$$

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

Proposition 3 gives an analytical solution of the run threshold and presents intuitive comparative statics about the stablecoin's run risk with respect to the level of liquidity transformation and the organization of the AP sector.

Part i) of Proposition 3 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 engages in a higher level of liquidity transformation in the sense that it holds less liquid reserve asset, the first-mover advantage among investors becomes larger because an investor who chooses not to sell would have to involuntarily bear a higher cost of liquidation induced by selling investors. This leads to a higher run risk. However, when the reserve asset is too illiquid, the first-mover advantage could be dampened because too few investors can enjoy it. This intuition can be understood from equation (5.5): investors enjoy the first-mover advantage only when $\lambda \leq 1 - \phi$, that is, $h(\lambda)$ takes the value in the first line of (5.5); 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 $g(\phi) < K$ will reduce the run risk. Looking forward, we confirm empirically in Section 6 that $g(\phi) > K$ indeed holds for the major stablecoins, suggesting that further increasing liquidity transformation will likely increase their run risks.

Part ii) of Proposition 3 shows that a more efficient AP sector in terms of less AP concentration and higher AP balance sheet capacity leads to higher run risk. To understand this more surprising result, note that the connection between stablecoins' secondary and primary markets implies a trade-off between de-pegging and run risks. A more efficient AP sector implies a lower de-pegging risk because the APs are more willing to absorb selling pressure from investors and thus more able to support a stable secondary market trading price. However, this means that APs will support a higher trading price to selling investors and subsequently redeem more stablecoins in the primary market, leading to a larger first-mover advantage and higher run risk. In contrast, to reduce run risk, the stablecoin issuer has to bear a higher de-pegging risk with a less stable secondary market price. In other words, when the AP sector is more efficient, the de-pegging risk is lower, because APs are more willing to absorb selling pressure from investors and thus more able to support a stable secondary market trading price. However, run risk actually increases, because APs support a higher trading price to selling investors, increasing the first-mover advantage for stablecoin sellers. When the AP sector is less efficient, the de-pegging risk is higher, since small quantities of stablecoin sales can have a substantial impact on stablecoin prices. However, the price impact of stablecoin sales in fact decreases first-mover advantage and discourages "panic selling", contributing to decreasing run risk. In this sense, the AP sector acts as a firewall between stablecoins' secondary and primary markets, and the stablecoin issuer optimally designs the structure of its AP sector to trade off between these two risks.

The analytical solution given in Proposition 3 allows us to calibrate the model and quantify the run risks of the stablecoins in reality. To this end, we can easily translate the run threshold into an ex-ante run probability, with the additional input of the fundamental distribution $F(\theta)$. The following definition gives us a formal notion of run risk, which we use in the calibration exercise in Section 6.

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

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

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

5.4 Optimal Design of the Stablecoin Primary Market

To further illustrate the idea of APs act as a firewall between stablecoins' secondary and primary markets and the trade-off between de-pegging and run risks, we further study the optimal design of the stablecoin primary market. Specifically, we focus on the optimal concentration of APs.

Given the potential for a panic run, the stablecoin issuer's design decision at $t = 0$ involves one key choice variable: n, that is, how many APs are allowed to perform primary-market redemptions. As described in Section 2, stablecoin issuers indeed consider the number of APs as one of the most important market design choices. 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 $t = 0$. The issuer's objective function is given by

$$\max_{n} E[\Pi] = \underbrace{G(E[W])}_{\substack{\text{population of} \\ \text{participating investors}}} \underbrace{\int_{p(\theta)\geq p(\theta^*)} \left(p(\theta)(R(\phi)-1)\frac{1-\phi-\pi}{1-\phi} \right) dF(\theta)}_{\substack{\text{expected issuer revenue per} \\ \text{participating investor}}}, \tag{5.9}$$

where each investor's expected utility of participation is given by

$$E[W] = \int_{p(\theta)<p(\theta^*)} q(1)dF(\theta) + \int_{p(\theta)\geq p(\theta^*)} (\pi q(\pi) + (1-\pi)v(\pi))\, dF(\theta), \tag{5.10}$$

in which $q(\cdot)$ and $v(\cdot)$ are given by (5.3) and (5.4), which is in turn a function of θ, and $p(\theta^*)$ is given by (5.7) in Proposition 3.

The stablecoin issuer's objective function (5.9) intuitively captures its revenue base: it enjoys the net long-term return of the remaining reserve asset if no panic run happens (after possible liquidation to meet redemptions driven by early investors), and more participating investors allow the issuer to start with investing in more reserve assets. Turning to participating investors' expected utility $E[W]$, the first term in (5.10) denotes the expected welfare of all investors when a panic run happens, while the second term corresponds to the expected investor welfare when a run does not happen.

Solving the stablecoin issuer's problem (5.9), we have the following result about the stablecoin issuer's optimal choice of AP concentration:

Proposition 4. *When the stablecoin engages in a higher level of liquidity transformation, the stablecoin issuer optimally designs a more concentrated AP sector, that is, n^* decreases in ϕ when ϕ is not too large.*

Proposition 4 stems fundamentally from the trade-off between the de-pegging and run risks of stablecoins. Intuitively, when investors are subject to idiosyncratic liquidity risks, they do not know ex-ante whether they have to consume early or late. To attract more investor participation, the stablecoin allows investors to share their idiosyncratic risks by them jointly holding a pool of reserve assets and offering the ability to sell the stablecoin in the secondary market at a price potentially higher than what an investor would have gotten by holding the reserve assets herself. However, because of the run risks, risk sharing may not always be achieved because everyone would just get the autarky outcome in a run scenario, hurting investors' expected utility and thus their participation as captured by the first term in (5.9). Further, a higher run risk also directly hurt the stablecoin issuer's expected revenue per participating investor as captured by the second term in (5.9), because the issuer would only enjoy the net long-term return of the reserve asset when no run happens. Thus, the issuer optimally accepts some level of de-pegging risk, that is, some deviation of the secondary-market price from its peg to avoid runs. This limits the ability of the stablecoin to provide immediate liquidity to early investors but would avoid a run. To achieve so, the issuer optimally chooses a concentrated AP sector to reduce the first-mover advantage among investors.

6 Model Calibration and Results

In this section, we calibrate our model to estimate run probability as defined in Definition 1. We start with a simple benchmark case of $\pi = 0$, which implies that the idiosyncratic shock of buying and selling stablecoins is mean zero and thus does not directly drive runs. This benchmark allows us to focus on panic runs and relate the run risk to the key stablecoin design features that we highlight. 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 explain our estimation of redemption elasticity, K, asset liquidity, ϕ, and the distribution of $p(\theta)$, before reporting the estimation results.

6.1 Elasticity of Redemptions in the Primary Market K

To estimate how responsive the volume of redemptions is to price discounts, we regress daily discounts against daily redemption volume for each stablecoin:

$$Discount_t = \beta Redemptions_t + FE_y, \tag{6.1}$$

where $Discount_t$ is the lowest observed secondary market price minus 1 on day t and $Redemptions_t$ is the volume of redemptions divided by the total outstanding volume of tokens on day t. We use the lowest secondary market price to better capture the extent of price dislocations that demand AP arbitrage rather than the price dislocations resulting from AP arbitrage. We normalize the volume of redemptions 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 AP sector for each stablecoin. For example, the number and constraints of APs may evolve after some time with the growth of stablecoins.

Table 5 shows the results. We observe that the regression coefficients are negative for both USDT and USDC, which is consistent with larger redemption volumes on days with steeper discounts, i.e., more negative secondary market prices. Further, the coefficient for USDT is larger in absolute magnitude than for USDC, which is consistent with the higher AP 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 volume as a fraction of the total volume outstanding corresponds to a 3.0 cent larger discount at USDT and a 1.3 cent larger discount at USDC.

6.2 Asset Illiquidity ϕ

We proxy asset illiquidity with haircuts following Bai, Krishnamurthy and Weymuller (2018) and Ma, Xiao and Zeng (2021). These haircuts proxy for the discount incurred when illiquid assets are converted into cash at short notice.[12] In other words, one minus the haircut is the amount of cash that stablecoin issuers can provide to APs redeeming at short notice by borrowing against the asset. More liquid assets are more readily pledged to obtain cash while more illiquid assets incur a higher discount. Figure 6 shows the median of asset discounts over time. In comparison, Treasuries are generally the most liquid, while corporate loans are the most illiquid.

[12] The New York Fed publishes haircuts on different securities when pledged as collateral in repo loans.

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. The results are shown in Figures 7a, and 7b. One challenge is that we do not know the liquidity of their deposits. As discussed in Section 4, these deposits include time deposits and CDs for which an early withdrawal penalty is incurred. These penalties generally range from half-years to two years' worth of interest rates, depending on the financial institution and contract length. We set the discount on the early withdrawal of deposits to be 0.5%. This is a relatively conservative measure given that the lowest asset discounts are at 2%. Further, 0.5% would have been the approximate one-year penalty rate on 5-year CDs in the latter half of 2021, which is the period for which asset breakdowns are available.

Overall, the reserve assets of USDT are more illiquid than those of USDC, but both of them shift towards holding more liquid assets over the sample period. The discount on USDT reserve assets decreased from 4.3% in September 2021 to 4.0% in March 2022. In comparison, the discount on USDC reserve assets drops from 0.9% in August 2021 to 0.5 % in September 2021. We use these estimates for the asset illiquidity parameter, ϕ, in our model.

6.3 Distribution of $p(\theta)$

Finally, our model requires us to take a stance on the distribution of $p(\theta)$, which is the signal of how likely the risky asset held in the issuer's portfolio is to pay nothing. To estimate p empirically, we use historical CDS prices to evaluate the extent to which the value of each portfolio component varies over time, allowing us to calculate a synthetic measure for how much the expected recovery value of the reserve portfolio is likely to fluctuate over time.

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 expected recovery value:

$$\sum_{c=1}^{C} (1 - s_c)\, q_c$$

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:

$$p = (1+\xi)\sum_{c=1}^{C}(1-s_c)\,q_c \tag{6.2}$$

Since p in the model is equal to the expected recovery value of assets per unit stablecoin issued, we will use (6.2) on each date we observe CDS spreads as one realization of p. We can think of (6.2) 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 (6.2). By plugging CDS spreads from different dates into (6.2), we can calculate a distribution of signals p. Note that, when we plug CDS spreads into (6.2), we use spreads from a single day; hence, this method accounts for correlations between CDS prices of different asset classes.

We implement (6.2) using historical CDS spread data from Markit, from 2008 to 2022.

For deposits, we assign the average CDS of unsecured debt at the top 6 US banks to capture the riskiness of the banking sector.[13] 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. For Treasuries, we assign the CDS spreads on 3-year US treasuries. For money market instruments, we use CDX spreads on 1-year investment-grade corporate debt. For USDC's corporate bonds, we assign the 10-year investment-grade corporate CDX because they are stated to be of at least a BBB+ rating. For USDT's corporate bonds, we assign the average 10-year corporate CDX. The remaining categories, "foreign" and "other", do not have a clear mapping to the existing CDS series. For USDT, for example, assets in the "other" 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.

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

Table 6 shows the distributions of p for USDT and USDC on dates with reported balance sheets. The distributions of p are fairly concentrated near 1, with a narrow range from roughly 97% to 99.5%. In comparison, the distribution of p's for USDC is slightly worse than USDT, which arises from USDT's large holdings of Treasuries that have lower CDS spreads than bank deposits, which are the bulk of USDC's portfolio.

6.4 Calibration Results

Combining our estimates of the redemption elasticity, K and the asset illiquidity, ϕ, calculate run cutoffs according to (5.7) in Proposition 3. Then, we can infer run probabilities for each stablecoin in each time period based on the corresponding empirical distribution of the signal $p(\theta)$ following (5.8) in Definition 1.

The results for run probabilities are shown in Table 7. Overall, run probabilities are substantial. USDT's run probability was 3.45% in March 2022 and USDC's run probability was 0.14% in October 2021.

7 Conclusion

In this paper, we analyzed the possibility of panic runs on stablecoins. At a high level, stablecoins holders engage in liquidity transformation, offering APs the option to redeem stablecoins for cash dollars, while holding partially illiquid portfolios of assets. This creates the possibility for runs, where market participants sell tokens in secondary markets, leading APs to buy and redeem stablecoins for dollars with the issuer. We show, however, that stablecoin run risk is mediated by the market structure of the AP sector, which serves as a "firewall" between the secondary and primary markets. When the AP sector is more efficient, shocks in the secondary market transmit more effectively to the primary market; peg stability of stablecoins is thus improved, but the first-mover advantage for sellers is also higher, increasing run risk. If the AP sector is less efficient, shocks in secondary markets transmit less effectively; peg stability suffers, but run risk actually decreases, as the price impact of stablecoin trades in secondary markets discourages market participants from panic selling. Calibrating the model to data, we quantified run risk for the two leading fiat-backed stablecoins by market cap.

References

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

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*, 72.4: 1023-1061.

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.

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

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

COOPER, RUSSELL and THOMAS ROSS (1998). "Bank Runs: Liquidity Costs and Investment Distortions." *Journal of Financial Economics*, 41: 27-38.

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

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.

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.

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, CHASE ROSS and SHARON ROSS (2022). "Making Money." Working Paper.

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

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.

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.

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.

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

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

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

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

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.

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

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

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

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

Figure 1: The Design of Fiat-backed Stablecoins

This figure illustrates the design of fiat-backed stablecoins.

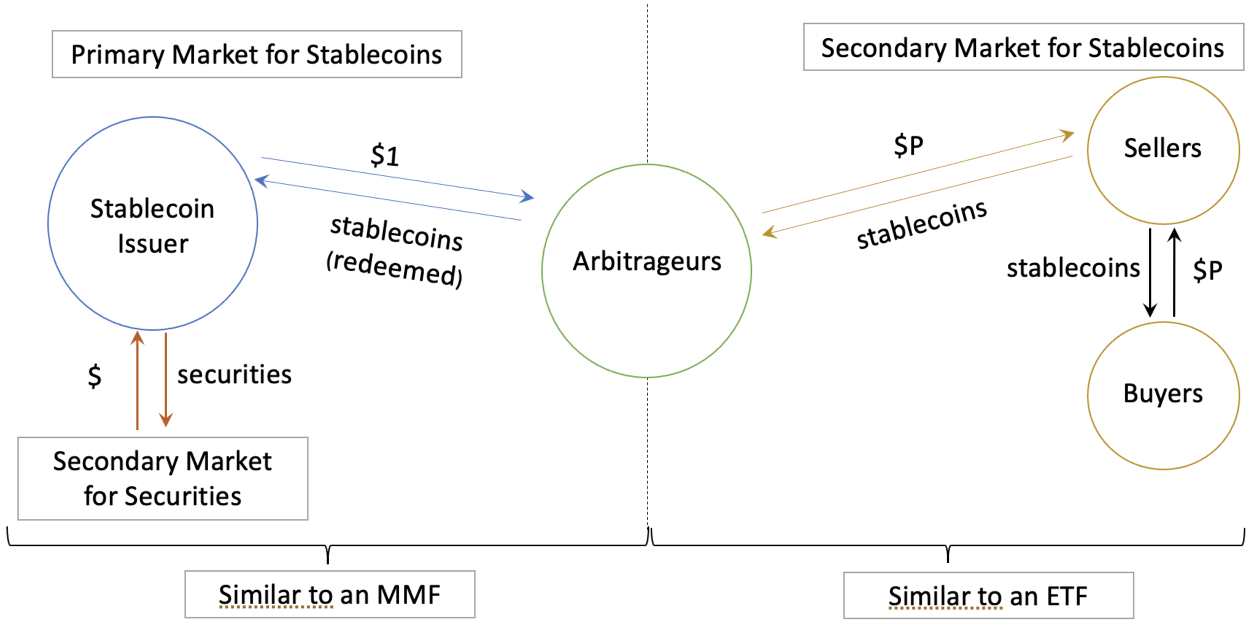

Figure 2: Asset Size of Fiat-backed Stablecoins

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

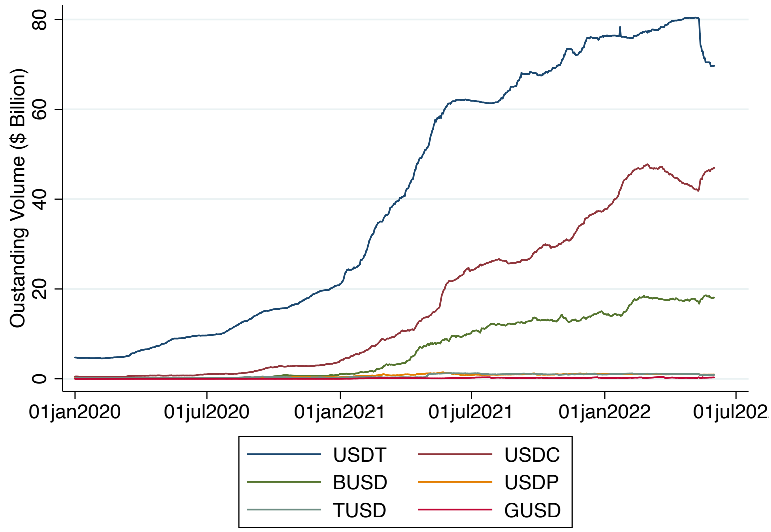

Figure 3: Secondary Market Trading Price

Panels (a) to (f) show the the daily secondary market trading price of USDT, USDC, BUSD, USDP, TUSD, and GUSD, respectively. Secondary market prices are volume-weighted average of trading prices from the exchanges listed in Section 2.

(a) USDT

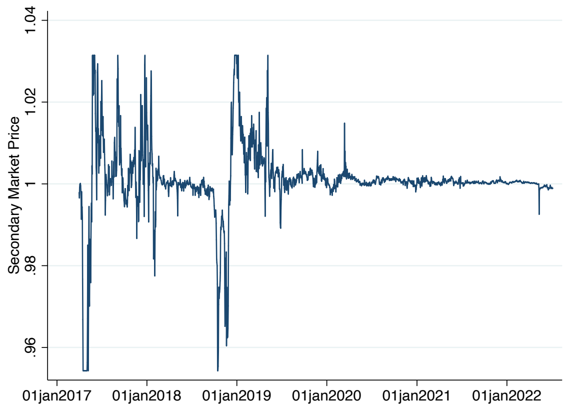

(b) USDC

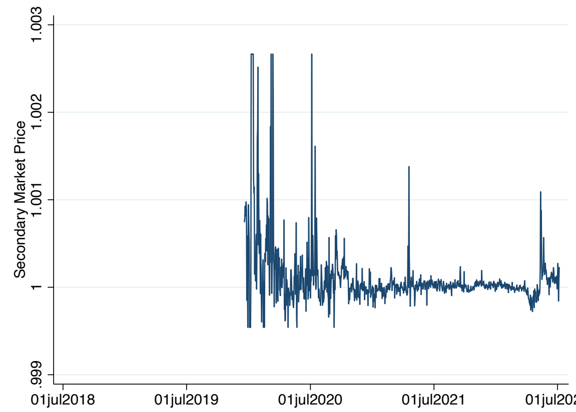

(c) BUSD

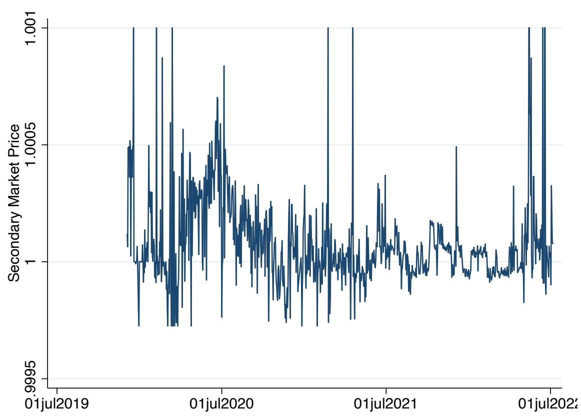

(d) USDP

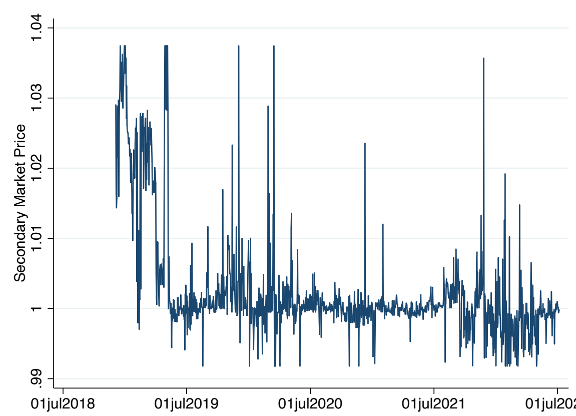

(e) TUSD

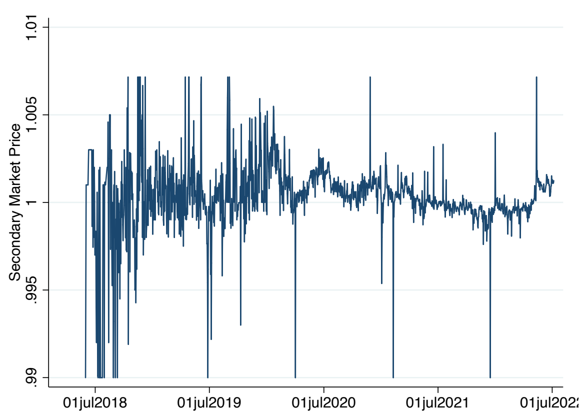

(f) GUSD

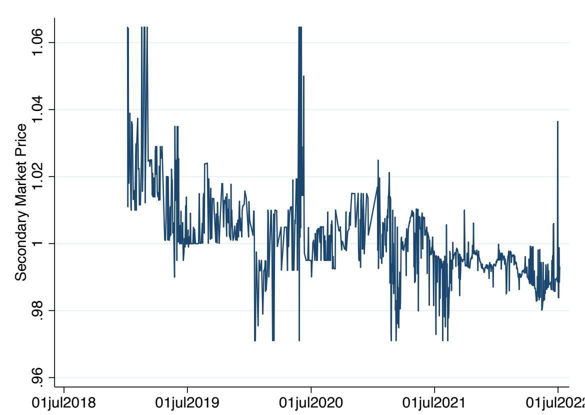

Figure 4: Secondary Market Discount and Primary Market Structure

This figure shows the relationship between secondary market price dislocations and primary market structure. In panel (a), each dot indicates the average secondary market discount and the average number of redeeming APs in a month for a given stablecoin. In panel (b), each dot indicates the average secondary market discount and the market share of the top give redeeming APs in a month for a given stablecoin.

(a) Number of APs

(b) Market Share of Top 5 APs

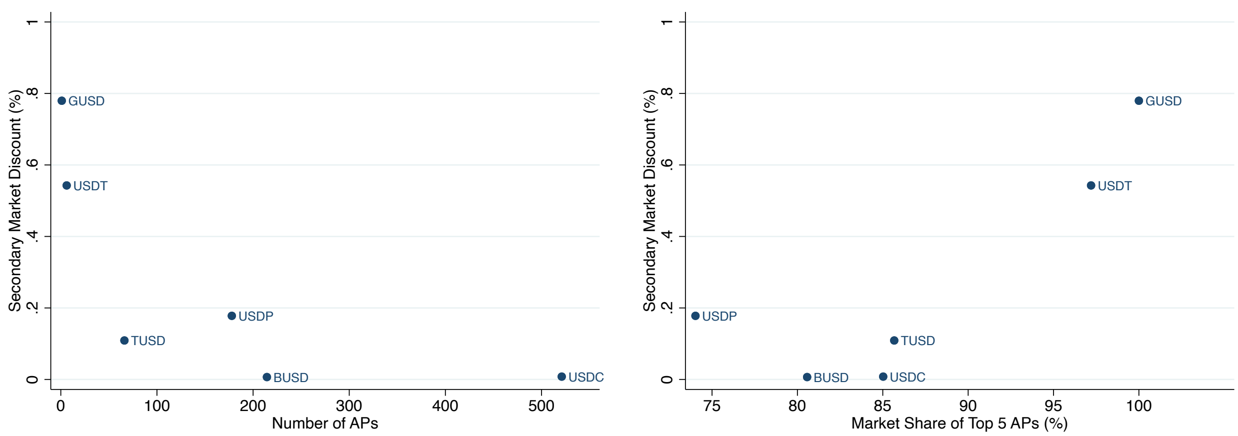

Figure 5: Payoff Gain of Late Investors

This figure shows a late investor's payoff gain between waiting until $t = 2$ versus selling at $t = 1$. Parameters: $p(\theta) = 0.97, \eta = 0.2, \phi = 0.05, K = 0.3$.

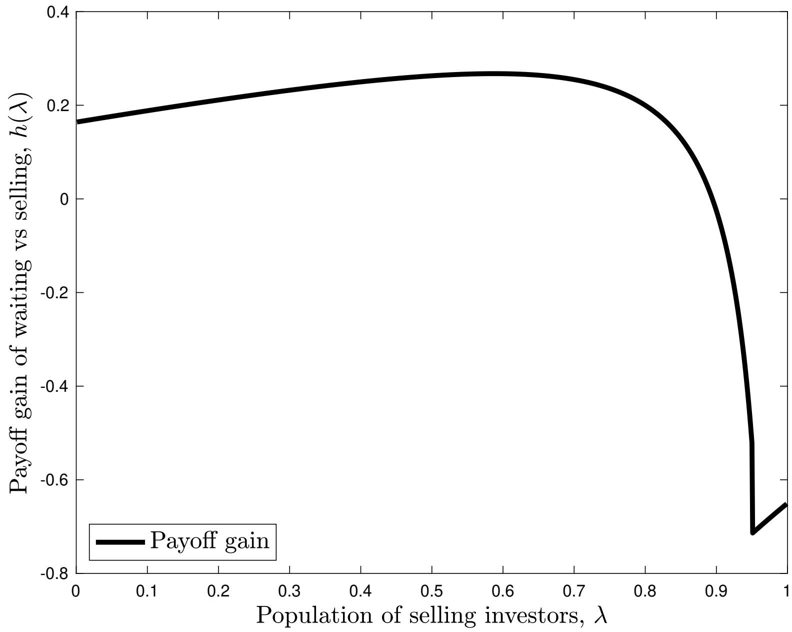

Figure 6: Liquidation Discounts

This figure shows median haircuts by collateral type. Data is from the New York Fed

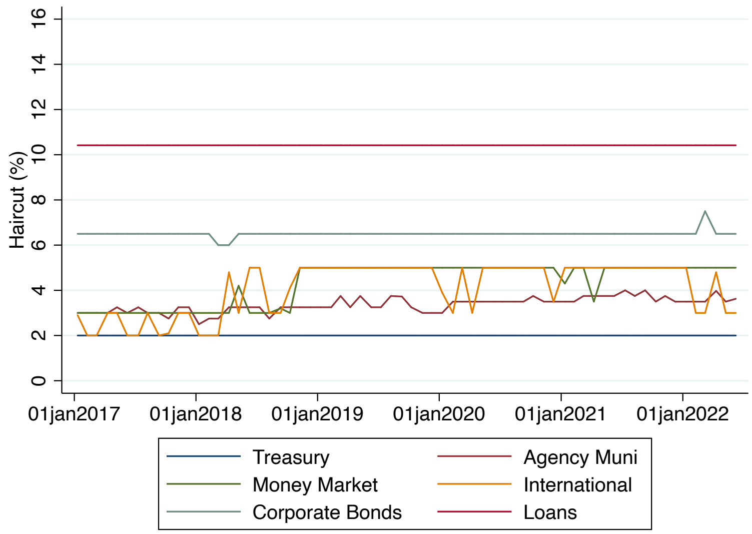

Figure 7: Asset Illiquidity

Panels (a) and (b) show the liquidation discount for USDT's and USDC's reserves. The sample period covers the dates for which a breakdown of reserve holdings for USDT and USDC overlapped.

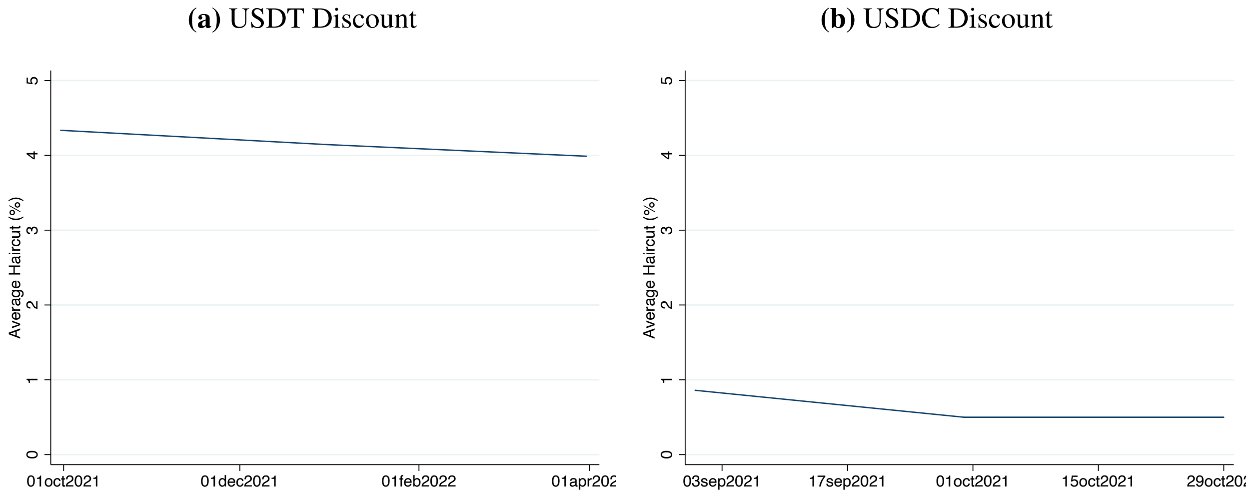

Table 1: Asset Composition

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 2: 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.

	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

Table 3: Primary Market Daily Redemption Activity

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

(a) USDT

	mean	p25	p50	p75
AP Num	1	1	1	2
Top 1 Share	94	100	100	100
Top 5 Share	100	100	100	100
Vol (mil)	57	2	12	60

(b) USDC

	mean	p25	p50	p75
AP Num	33	8	14	28
Top 1 Share	54	45	50	59
Top 5 Share	96	95	98	100
Vol (mil)	103	2	15	134

(c) BUSD

	mean	p25	p50	p75
AP Num	21	8	15	28
Top 1 Share	59	40	56	76
Top 5 Share	94	90	96	100
Vol (mil)	62	8	27	82

(d) USDP

	mean	p25	p50	p75
AP Num	18	8	17	27
Top 1 Share	55	37	52	73
Top 5 Share	90	85	95	100
Vol (mil)	12	3	6	13

(e) TUSD

	mean	p25	p50	p75
AP Num	6	3	6	8
Top 1 Share	72	54	73	91
Top 5 Share	99	99	100	100
Vol (mil)	6	1	2	5

(f) GUSD

	mean	p25	p50	p75
AP Num	1	1	1	1
Top 1 Share	100	100	100	100
Top 5 Share	100	100	100	100
Vol (mil)	6	0	1	3

Table 4: Primary Market Monthly Redemption Activity

Panels (a) to (f) provide statistics about monthly primary market redemption activity on the ethereum blockchain, including the number of APs, the market share of the top 1 and top 5 APs, and the volume of redemptions. 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
AP Num	6	3	6	8
Top 1 Share	66	42	61	89
Top 5 Share	97	98	100	100
Vol (mil)	577	46	123	763

(b) USDC

	mean	p25	p50	p75
AP Num	521	114	168	262
Top 1 Share	45	38	49	50
Top 5 Share	85	81	85	90
Vol (mil)	2976	160	460	4965

(c) BUSD

	mean	p25	p50	p75
AP Num	214	157	202	274
Top 1 Share	48	30	50	62
Top 5 Share	81	74	82	87
Vol (mil)	1596	233	1498	2720

(d) USDP

	mean	p25	p50	p75
AP Num	178	71	174	284
Top 1 Share	41	24	37	54
Top 5 Share	74	62	77	88
Vol (mil)	260	94	174	262

(e) TUSD

	mean	p25	p50	p75
AP Num	66	49	74	85
Top 1 Share	50	36	46	64
Top 5 Share	86	79	91	94
Vol (mil)	154	31	85	260

(f) GUSD

	mean	p25	p50	p75
AP Num	1	1	1	1
Top 1 Share	100	100	100	100
Top 5 Share	100	100	100	100
Vol (mil)	113	7	17	164

Table 5: Secondary Price Deviation versus Redemptions

This table shows the results from regressing the lowest daily secondary market price against the daily volume of redemptions for each USDT and USDC. The lowest secondary market price is the lowest hourly price for each coin on each day. The daily volume of redemptions is 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	-0.30**	-0.13**
	(0.13)	(0.06)
Observations	438	892
Adjusted R2	0.14	0.02

Table 6: Distribution of $p(\theta)$

This table shows quantiles of the distributions of the expected recovery value of assets per unit stablecoin. We combine Markit data on CDS spreads for different asset classes from 2008 to 2022, with data on stablecoin issuers' asset class holdings and over-collateralization ratios, using expression (6.2).

coin	date	p10	p25	p50	p75	p90
USDT	2021m9	0.9857	0.9896	0.9929	0.9940	0.9950
USDT	2021m12	0.9873	0.9908	0.9931	0.9941	0.9952
USDT	2022m3	0.9884	0.9915	0.9936	0.9945	0.9956
USDC	2021m8	0.9765	0.9858	0.9907	0.9931	0.9940
USDC	2021m9	0.9769	0.9861	0.9919	0.9944	0.9950
USDC	2021m10	0.9769	0.9861	0.9919	0.9944	0.9950

Table 7: Estimated Run Probabilities

This table shows our estimated run probabilities for different stablecoin issuers at different dates, calculated by combining our estimates of the distribution of $p(\theta)$, expected recovery value of assets per unit stablecoin using CDS data from expression (6.2), with the run cutoffs computed using expression (5.7).

month	coin	runprob
2022m3	USDT	0.0345
2021m10	USDC	0.0014

Table 8: Primary Market Daily Redemption Activity (Tron)

Panels (a) to (f) provide statistics about daily primary market redemption activity on the tron blockchain, including the number of APs, the market share of the top 1 and top 5 APs, and the volume of redemptions. 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
AP Num	1	1	1	2
Top 1 Share	96	100	100	100
Top 5 Share	100	100	100	100
Vol (mil)	450	40	110	460

(b) USDC

	mean	p25	p50	p75
AP Num	33	7	17	28
Top 1 Share	67	45	67	94
Top 5 Share	93	91	98	100
Vol (mil)	2	0	0	2

(c) TUSD

	mean	p25	p50	p75
AP Num	1	1	1	1
Top 1 Share	97	100	100	100
Top 5 Share	100	100	100	100
Vol (mil)	10	0	0	2

Table 9: Primary Market Monthly Redemption Activity (Tron)

Panels (a) to (f) provide statistics about monthly primary market redemption activity on the tron blockchain, including the number of APs, the market share of the top 1 and top 5 APs, and the volume of redemptions. 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
AP Num	5	2	4	6
Top 1 Share	72	53	68	94
Top 5 Share	100	100	100	100
Vol (mil)	4625	651	3575	7515

(b) USDC

	mean	p25	p50	p75
AP Num	446	11	317	391
Top 1 Share	58	33	51	81
Top 5 Share	84	78	85	100
Vol (mil)	41	3	24	70

(c) TUSD

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

Table 10: Primary Market Daily Redemption Activity (Avalanche)

Panels (a) to (f) provide statistics about daily primary market redemption activity on the avalanche blockchain, including the number of APs, the market share of the top 1 and top 5 APs, and the volume of redemptions. 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
AP Num	1	1	1	1
Top 1 Share	100	100	100	100
Top 5 Share	100	100	100	100
Vol (mil)	31	5	30	60

(b) USDC

	mean	p25	p50	p75
AP Num	3	1	2	4
Top 1 Share	88	78	99	100
Top 5 Share	100	100	100	100
Vol (mil)	6	0	0	1

(c) BUSD

	mean	p25	p50	p75
AP Num	2	1	1	2
Top 1 Share	90	86	100	100
Top 5 Share	100	100	100	100
Vol (mil)	0	0	0	0

(d) TUSD

	mean	p25	p50	p75
AP Num	6	3	6	8
Top 1 Share	72	54	73	91
Top 5 Share	99	99	100	100
Vol (mil)	6	1	2	5

Table 11: Primary Market Monthly Redemption Activity (Avalanche)

Panels (a) to (f) provide statistics about monthly primary market redemption activity on the avalanche blockchain, including the number of APs, the market share of the top 1 and top 5 APs, and the volume of redemptions. 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
AP Num	1	1	1	1
Top 1 Share	100	100	100	100
Top 5 Share	100	100	100	100
Vol (mil)	50	1	10	120

(b) USDC

	mean	p25	p50	p75
AP Num	34	18	32	47
Top 1 Share	49	31	42	60
Top 5 Share	94	87	96	99
Vol (mil)	111	3	16	219

(c) BUSD

	mean	p25	p50	p75
AP Num	22	10	18	30
Top 1 Share	37	30	40	42
Top 5 Share	83	73	82	94
Vol (mil)	0	0	0	0

(d) TUSD

	mean	p25	p50	p75
AP Num	66	49	74	85
Top 1 Share	50	36	46	64
Top 5 Share	86	79	91	94
Vol (mil)	154	31	85	260

A Appendix: Additional Institutional Details

A.1 Minting of Stablecoins

Practically, stablecoins are ERC-20 tokens. 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 own control, and then issue stablecoins in primary markets by sending tokens from the "treasury" address to customers' wallets, and allow redemptions when customers send tokens to the treasury address.[14]

A.2 Trading on Crypto Exchanges

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

[14] Treasury address tokens technically count towards the market cap of any given stablecoin, but they are not economically meaningful as part of 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.

the escrow account to the seller's account. In this process, purchased stablecoins are sent directly to the user's self-custodial wallet.