



Supervisory Research and Analysis Unit

Working Paper | SRA 23-02 | Original Publication Date: August 24, 2023 | This Version: March 26, 2024

Runs and Flights to Safety: Are Stablecoins the New Money Market Funds?

Kenechukwu Anadu, Pablo D. Azar, Marco Cipriani, Thomas M. Eisenbach, Catherine Huang, Mattia Landoni, Gabriele La Spada, Marco Macchiavelli, Antoine Malfroy-Camine, and J. Christina Wang





Supervisory Research and Analysis (SRA) Working Papers present economic, financial, and policy-related research conducted by staff in the Federal Reserve Bank of Boston's Supervisory Research and Analysis Unit. SRA Working Papers can be downloaded without charge at: <http://www.bostonfed.org/publications/sra/>

The views expressed in this paper are those of the author(s) and do not necessarily represent those of the Federal Reserve Bank of Boston or the Federal Reserve System.

Runs and Flights to Safety: Are Stablecoins the New Money Market Funds?*

Kenechukwu Anadu¹, Pablo D. Azar², Marco Cipriani², Thomas M. Eisenbach²,
Catherine Huang³, Mattia Landoni¹, Gabriele La Spada², Marco Macchiavelli⁴,
Antoine Malfroy-Camine¹, and J. Christina Wang¹

¹Federal Reserve Bank of Boston

²Federal Reserve Bank of New York

³Harvard Business School

⁴Isenberg School of Management

August 24, 2023

This version: March 26, 2024

Abstract

Similar to the more traditional money market funds (MMFs), stablecoins aim to provide investors with safe, money-like assets. We investigate similarities and differences between these two investment products. Like MMFs, stablecoins suffer from “flight-to-safety” dynamics: we document net flows from riskier to safer stablecoins on days of crypto-market stress and estimate a discrete “break-the-buck” threshold of \$1, below which stablecoin redemptions accelerate. We then focus on two specific stablecoin runs, in 2022 and 2023, showing that the same flight-to-safety dynamics also characterized these episodes. Finally, as flight-to-safety flows occur within MMF families, stablecoin flows tend to happen within blockchains.

JEL classification: G10, G20, G23.

Keywords: Stablecoins, Money Market Mutual Funds, Financial Stability, Crypto Assets, Runs, Liquidity Transformation.

*We thank for helpful comments and suggestions Anders Brownworth, Mark Dickerson, Edward Dumas, Jessica Dwyer, Siobhan Sanders, Larry Wall, and participants to the XII Economics of Payments Conference. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Boston, the Federal Reserve Bank of New York, or the Federal Reserve System. Emails: Ken.Anadu@bos.frb.org, Pablo.Azar@ny.frb.org, Marco.Cipriani@ny.frb.org, Thomas.Eisenbach@ny.frb.org, catherinehuang1@g.harvard.edu, Mattia.Landoni@bos.frb.org, Gabriele.LaSpada@ny.frb.org, mmacchiavelli@isenberg.umass.edu, Antoine.Malfroy-Camine@bos.frb.org, and Christina.Wang@bos.frb.org.

1 Introduction

Stablecoins are crypto assets that seek to maintain a stable price (usually \$1) through various mechanisms, such as backing each token with assets or using a supply–demand matching algorithm.¹ The market capitalization of stablecoins has grown exponentially over the last few years, from \$5 billion in 2019 to about \$180 billion at its peak in 2022.

In several important ways, stablecoins closely resemble more traditional financial vehicles, particularly money-market mutual funds (MMFs). A key similarity between stablecoins and MMFs is that they both provide money-like assets to investors by engaging in liquidity transformation: they issue liabilities with a stable nominal value, but the collateral backing these claims can suddenly become illiquid.² This liquidity transformation makes stablecoins and MMFs vulnerable to runs (Rosengren, 2021; Federal Reserve Board, 2022; Azar et al., 2022).

Indeed, MMFs suffered runs during both the financial crisis in 2008 and the COVID-19 crisis in 2020 (Duygan-Bump et al., 2013; Cipriani and La Spada, 2020; Li et al., 2021).³ Furthermore, like MMFs, several stablecoins invest their reserves in short-term fixed-income instruments, and there is some evidence suggesting that stablecoin issuance exerts notable effects on traditional money market issuance volumes and interest rates (Kim, 2022; Barthélémy et al., 2023).

Another similarity between MMFs and stablecoins is that both differ in the amount of risk they take: some stablecoins are riskier than others. In fact, stablecoins exhibit an even wider range of risk profiles than MMFs. Some stablecoins report that they are backed by safe assets, such as cash and

¹Throughout this paper, we use the term *crypto assets* to refer to assets held on a blockchain. A blockchain is a distributed computer system that uses cryptography to enable multiple parties to agree on who owns which assets at any given time (Nakamoto, 2008).

²Algorithmic stablecoins are a special case in which price stability is not ensured by holding collateral, but rather by an algorithmic mechanism that uses arbitrage to balance demand and supply. The liquidity of these stablecoins, rather than depending on collateral liquidity, depends on investor confidence in the pegging algorithm. See Section 2 and Appendix 2.1.3.

³There are, of course, several differences between MMFs and stablecoins; for example, MMFs are governed by the US Securities and Exchange Commission under Rule 2a-7 of the Investment Company Act of 1940, which sets minimum portfolio liquidity and maturity standards, among other restrictions, whereas stablecoins are not subject to this regulation. Also, while this paper focuses on MMFs and stablecoins, there are other collective investment vehicles with similar structural vulnerabilities, such as private liquidity funds and short-term investment funds (Federal Reserve Board, 2022). On the money-like characteristics of stablecoins, see also (Gorton and Zhang, 2021).

US Treasuries, whose value remains stable or even increases during times of stress. Others, instead, are reportedly backed by riskier collateral, such as corporate debt or even other crypto assets; as the collateral backing some stablecoins loses value, they are likely to lose their peg, potentially triggering a run. A third group of stablecoins maintain their pegs through algorithms aimed at matching supply and demand; if investors’ beliefs about the effectiveness of these algorithms deteriorate, such stablecoins may also suffer runs.

The potential for runs on stablecoins—and the stablecoin runs that did occur in 2022 and 2023—have attracted substantial theoretical interest. Several recent works examine the drivers of stablecoin instability and propose self-governance or regulatory solutions to make stablecoins more stable.⁴ In contrast, with a few exceptions, there is a dearth of empirical evidence on this subject.⁵

In this paper, we contribute to the empirical understanding of the stablecoin market by documenting the similarities between stablecoins and MMFs, focusing on their runnability and flight-to-safety dynamics. First, we document stablecoin flow dynamics during periods of crypto market stress. Specifically, we isolate stress events in crypto markets associated with large declines in the price of Bitcoin and study the response of stablecoin investors to such events. We find that, during these episodes, investors run from riskier stablecoins to safer ones, a dynamic similar to what we observe in MMFs in times of stress.⁶

We also show that the so-called “depegging” dynamics for stablecoins closely resemble the “break-the-buck” dynamics observed for MMFs. Specifically, we estimate that stablecoin redemptions accelerate significantly if the stablecoin’s price drops below \$1: while a stablecoin priced \$1.00 or above experiences flows that are uncorrelated with prices, once its price drops below the

⁴See, for instance, Arner et al. (2020), Bertsch (2023), d’Avernas et al. (2022), Gorton et al. (2022a), and Li and Mayer (2022).

⁵Ma et al. (2023) structurally estimate the run risk of two major stablecoins. Hoang and Baur (2021) quantify the volatility of stablecoin prices. (Oefele et al., 2024) shows suggestive evidence of capital flight from riskier to safer stablecoins around the failure of Silicon Valley Bank (SVB) in 2023. Other papers focus on individual runs (Adams and Ibert, 2022; Uhlig, 2022; Liu et al., 2023) or examine historical evidence from episodes of runs on financial asset that bear some resemblance to modern stablecoins (Frost et al., 2020; Gorton et al., 2022b).

⁶Our characterization of a stablecoin as “safe” or “risky” is based on the composition of its assets (also known as “reserves”) *as reported by the stablecoin issuer*. Reporting quality, however, varies across stablecoin issuers, because there is no standardized or regulatory reporting requirement.

\$1 threshold, its daily outflow rate jumps by 34 basis points. This suggests that investors consider a stablecoin to have depegged when its price drops below \$1, similar to the 95-cent “break-the-buck” threshold for MMFs.

We next show that the flight-to-safety dynamics described above also occur during run episodes that originate in the stablecoin market. The most significant episode so far occurred when Terra, an algorithmic stablecoin and the fourth-largest stablecoin at the time, suffered a run and subsequently collapsed in early May 2022. Following a classic flight-to-safety dynamic, this run had negative spillover effects on other algorithmic stablecoins and stablecoins backed by risky assets, whereas those backed by relatively safer assets experienced net inflows.

Stablecoins suffered another run triggered by the failure of Silicon Valley Bank (SVB) in March 2023. USD Coin (USDC), the second-largest stablecoin at the time, held deposits at SVB. As in the May 2022 run, crypto investors responded quickly to news concerning the financial health of SVB by selling or redeeming their USDC tokens. Unlike the May 2022 Terra run, which occurred among the riskier stablecoins, the March 2023 run had at its epicenter USDC, which had previously been considered among the safest stablecoins due to its portfolio of Treasury securities and bank deposits. The run spilled over into Dai and Frax, both of which were partially collateralized by USDC. We show that investors ran from USDC to other stablecoins that they perceived as safer at the time. In particular, investors fled to stablecoins based outside the U.S. backed by traditional financial assets, especially Tether (USDT), which were considered riskier during the 2022 run.

Overall, these run and flight-to-safety dynamics observed in the stablecoin market are similar to those observed in the MMF industry. During a stress event, investors identify a driver of risk and run away from MMFs that are more exposed to such risk toward relatively safer MMFs or similar vehicles. Indeed, during the Global Financial Crisis in 2008, MMF investors ran on prime funds with larger exposures to Lehman Brothers and asset-backed commercial paper (McCabe, 2010; Kacperczyk and Schnabl, 2013; Duygan-Bump et al., 2013), and toward the safety of government funds. In 2011, during the European debt crisis, investors ran on funds with exposure to European banks (Chernenko and Sunderam, 2014; Ivashina et al., 2015; Gallagher et al., 2020). In 2020, at

the onset of the COVID-19 pandemic, funds closer to regulatory liquidity constraints experienced heavier outflows (Cipriani and La Spada, 2020; Li et al., 2021). In all these episodes, prime MMFs were exposed to the source of risk and thus suffered outflows, while government MMFs were safer and experienced contemporaneous inflows. Importantly, contagion is a feature of both stablecoin and MMF runs; for instance, Cipriani and La Spada (2020) document information contagion from institutional to retail MMFs during the 2020 run.

An additional parallel between MMF and stablecoin investors relates to the pattern of flows. In stablecoins markets, flows from riskier to safer stablecoins tend to occur within the same blockchain, at least within each of the larger and safer blockchains such as Ethereum and Binance Smart Chain, where 90 percent of all stablecoins are traded. This is similar to the pattern of flight-to-safety flows across MMFs, which tend to occur within the same fund complex (Cipriani and La Spada 2020, 2021). In contrast, smaller and riskier blockchains—which only represent less than 10 percent of all stablecoin circulation—exhibit a different pattern. Within these blockchains, outflows from riskier stablecoins are correlated with outflows from safer stablecoins; this suggests that investors deem the blockchains themselves as overall risky and hence depart them for safer (larger) blockchains.

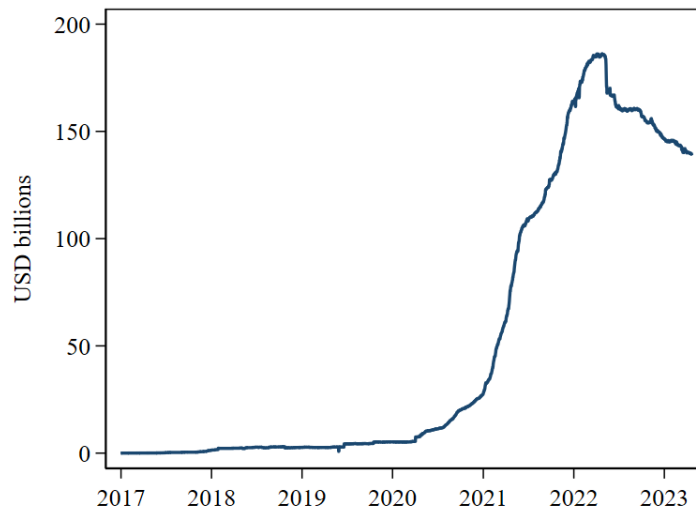
The rest of the paper is organized as follows. Section 2 provides background information on stablecoins and MMFs, and briefly discusses flight-to-safety dynamics. Section 3 describes the data. Section 4 presents our empirical analysis of stablecoin flow dynamics, and Section 5 concludes. Appendix A presents statistics on the largest stablecoins in the industry and describes how they maintain their peg. Appendix B provides tables further supporting our empirical results.

2 Background

2.1 Stablecoins

Over the past several years, stablecoins have experienced extraordinary growth; their market capitalization has increased from \$5 billion in 2019 to about \$125 billion in April 2023 (Figure 1).

Figure 1: Stablecoin industry circulation from January 2017 to April 2023.



Source: CoinGecko

Currently, stablecoins are predominantly used to facilitate the trading of other volatile crypto assets, such as Bitcoin, primarily through trading platforms for digital assets (PWG, 2021).

Stablecoins may be issued by US-based and non-US-based (henceforth, “offshore”) entities. Over the past several years, stablecoins issued by US-based entities typically have been backed by financial assets with little credit or liquidity risk, such as bank deposits and US Treasury securities. By contrast, stablecoins issued by offshore entities tend to be backed by riskier assets, such as commercial paper or other crypto assets, or by algorithmic mechanisms.⁷

Across stablecoins, there are three main types of arrangements: (1) stablecoins backed by traditional financial assets; (2) stablecoins backed by crypto assets; and (3) algorithmic stablecoins.⁸

⁷In 2021, Circle, a US-based entity that co-issues USDC with Coinbase, announced that it was shifting the composition of USDC’s assets to US Treasury obligations and deposits; previously, USDC also had held relatively riskier assets, such as corporate bonds. Also, Tether, which is issued by an offshore entity, has been raising the fraction of its cash and Treasury holdings, although it also holds riskier assets, such as loans and crypto assets.

⁸Stablecoins can also be distinguished by their degree of centralization. Some are issued by a single entity with centralized governance, such as USDC and Tether. By contrast, others, such as DAI and the now-defunct TerraUSD, are issued and administered using smart contracts under decentralized governance structures. Finally, note that there is a fourth type of stablecoin, which is backed by commodities; since these stablecoins typically are not pegged to the US dollar, we exclude them from our analysis.

2.1.1 Financial Asset-Backed Stablecoins

Financial asset-backed stablecoins (henceforth, asset-backed stablecoins) represent the largest type of stablecoins, accounting for about 96 percent of the industry market capitalization in December 2022. These stablecoins, which can be issued by entities based in the United States or offshore, are mainly backed by assets that carry little credit or liquidity risk, such as cash, Treasury securities, certificates of deposit, and commercial paper. Their tokens are minted and burned (redeemed) by a centralized entity. Customers may deposit dollars with the issuer and receive stablecoin tokens issued to their public address on the blockchain. Conversely, customers may redeem their tokens by sending them back to the issuer’s public address on the blockchain and receiving a dollar credit to their bank account.

As Ma et al. (2023) and Lyons and Viswanath-Natraj (2023) discuss, only a restricted set of participants has access to primary market transactions with the issuer, similarly to what happens with exchange-traded funds (ETFs); for general stablecoin investors, transacting on the secondary market—buying and selling stablecoins on crypto exchanges—is often the only option.

US-based asset-backed stablecoins include USD Coin (USDC), Binance USD (BUSD), Pax Dollar (USDP), and Gemini Dollar (GUSD). As mentioned above, non-US-based asset-backed stablecoins, like Tether (USDT) and TrueUSD (TUSD), are usually partially backed by riskier assets, such as precious metals or corporate bonds.

2.1.2 Crypto-Backed Stablecoins

Crypto-backed stablecoins are issued by a smart contract and are backed by crypto assets, such as Bitcoin, Ethereum, and even other stablecoins. Their peg is maintained by complex algorithms (Kozhan and Viswanath-Natraj, 2021; Lyons and Viswanath-Natraj, 2023). When the collateral is volatile, crypto-backed stablecoins are typically overcollateralized. The most popular crypto-backed stablecoin is Dai (DAI); others include Magic Internet Money (MIM) and Liquity USD (LUSD).

2.1.3 Algorithmic Stablecoins

In contrast to asset-backed or crypto-collateralized stablecoins, algorithmic stablecoins are not backed by assets; rather, their peg is maintained by an algorithmic mechanism that uses arbitrage to balance demand and supply. The most popular algorithmic stablecoins work by issuing two tokens: a free-floating cryptocurrency (e.g., Luna) and a stablecoin pegged to the US dollar (e.g., Terra). A smart contract allows users to mint one unit of stablecoin by burning \$1 worth of the free-floating cryptocurrency, and vice versa. For instance, suppose that Luna's price is \$12 and Terra's price drifts down to \$0.9 because of low demand. Then, arbitrageurs can buy 12 Terra tokens for \$10.8 ($= \0.9×12), burn them to obtain one Luna token, and sell it for \$12 for a quick \$1.2 profit. The decreased supply of Terra will bring demand and supply back into balance and hopefully restore the \$1 market price.

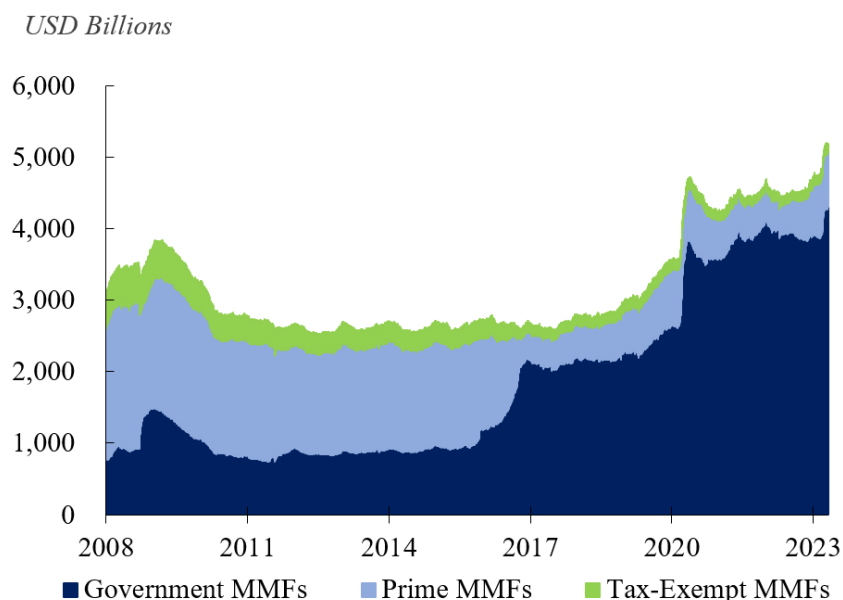
Examples of algorithmic stablecoins include Terra (USTC), Frax (FRAX), and Decentralized USD (USDD). Some algorithmic stablecoins, such as Frax, are hybrids; that is, their peg is partially maintained through outside collateral and partially through an algorithmic mechanism that issues a free-floating cryptocurrency called Frax Shares. In Appendix A, we describe these algorithms in more detail.

2.2 Money Market Funds

MMFs are open-end mutual funds registered with the Securities and Exchange Commission (SEC) that—like stablecoins—seek to maintain a stable price (or minimal price volatility). There are three types of MMFs: (1) government funds, which invest in cash, Treasury and agency securities, and repurchase agreements (repos) collateralized by those assets; (2) prime funds, which, in addition, can invest in non-government money market instruments, such as commercial paper and certificates of deposit; and (3) tax-exempt MMFs, which invest primarily in debt instruments issued by municipalities.⁹

⁹MMFs are restricted by regulation to only hold relatively safe assets with short-term maturities.

Figure 2: Assets under management of US MMFs.



Source: iMoneyNet. Data as of March 15, 2023.

Note: The areas in dark blue, light blue, and green represent the assets under management of government, prime, and tax-exempt MMFs, respectively, in billions of US dollars.

In addition to the characteristics of their portfolio holdings, MMFs may be distinguished by their investor type, with institutional MMFs marketed to corporations and fiduciaries, and retail MMFs limited to “natural persons.”¹⁰

As shown in Figure 2, total net assets (TNA) in publicly offered MMFs have increased over the past few decades, from \$3 trillion in 2008 to \$5.2 trillion in 2023.¹¹

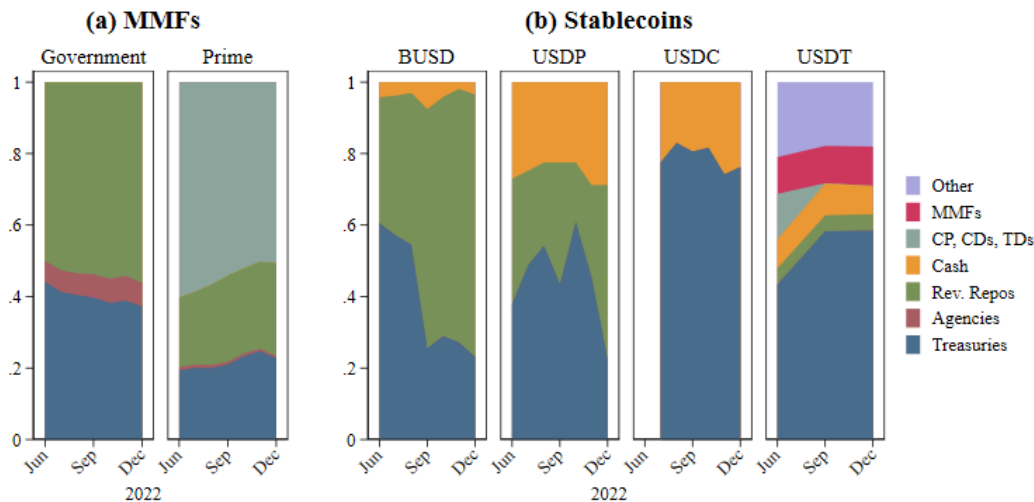
2.3 Stablecoins and MMFs: Differences and Similarities

There are several notable differences between stablecoins and MMFs. First, MMFs are regulated by the SEC, whereas stablecoins are not. Second, MMFs are typically sponsored by large banks or fund families, while stablecoins are sponsored by digital asset issuers. Third, the clienteles of stablecoins and MMFs are different: MMF investors include large, traditional, institutional investors,

¹⁰There are also non-public institutional funds, which are not available to the general public.

¹¹The increase in government fund assets beginning in 2016 was primarily driven by the SEC’s 2014 reforms to MMFs, which had an effective date of October 2016 (Cipriani and La Spada, 2021; Gissler et al., 2021). The increase that began in 2020 was driven by flight-to-safety inflows amid the onset of the COVID-19 pandemic (Cipriani and La Spada, 2020; Anadu and Sanders, 2021).

Figure 3: Composition of money market fund and stablecoin asset portfolios in 2022.



Sources: iMoneyNet and stablecoin issuer attestations.

such as financial and nonfinancial corporations, whereas stablecoin investors are mainly either retail investors or crypto-related companies. Fourth, MMFs are not traded on secondary markets and can be redeemed at par with the fund, while stablecoins are traded in secondary markets (crypto exchanges), and redemption rights and mechanisms vary across stablecoins. Finally, MMFs are backed only by traditional financial assets, while stablecoins may also be backed by digital assets or an algorithm.

Despite these differences, MMFs serve as a useful counterpoint to stablecoins. Among stablecoins backed by traditional financial assets, US-based stablecoins look remarkably similar to government MMFs. Figure 3 shows the composition of MMF and stablecoin portfolios in 2022. Panel (a) shows that about 50 percent of government funds' assets are held in repos, while a little less than 50 percent are held in Treasury securities. Panel (b) of Figure 3 shows that, similar to government funds, US-based stablecoins (BUSD, USDP, and USDC) are primarily backed by Treasury securities and repos. By contrast, the offshore, asset-backed stablecoin Tether (USDT) is backed by a wider mix of assets that includes commercial paper and certificates of deposit, similar to prime MMFs.

Both MMFs and stablecoins issue money-like liabilities that are vulnerable to runs. Cipriani and La Spada (2021) show that MMF investors demand money-like assets and are willing to pay a premium for greater money-likeness. They show this by exploiting the impact of the SEC MMF reform, which in October 2016 introduced redemption gates and fees based on portfolio liquidity for all prime funds and a floating net asset value (NAV) for institutional funds. Both regulatory measures increased the information sensitivity and therefore decreased the money-likeness of prime funds, resulting in greater differentiation between prime and government MMFs. Government funds, which were unaffected by the reform, were afterwards seen as more money-like than prime funds. Indeed, in the months leading up to October 2016, more than \$1 trillion flowed from prime to government MMFs as the total size of the MMF industry stayed roughly level. In summary, investors are willing to accept lower net yields in government funds for access to a more money-like product.

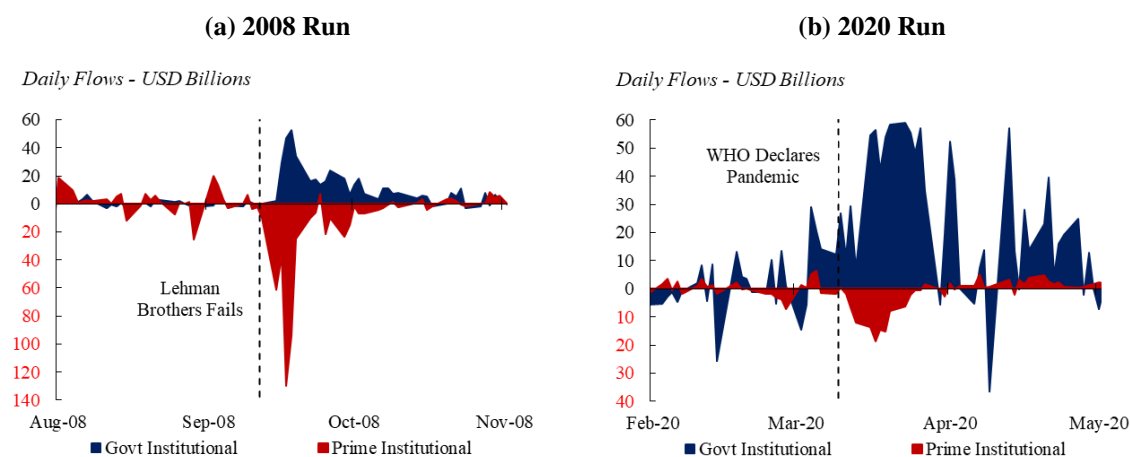
Money-like liabilities are a known source of run risk; indeed, MMFs were subject to runs in 2008 and more recently in 2020. Similarly, for stablecoins, Ma et al. (2023) hypothesize the existence of a tradeoff between run risk and price stability. A higher number of authorized participants results in a greater coordination problem, leading to higher run risk; however, it also results in greater arbitrage activity, helping to keep the stablecoin's secondary market price aligned with the value of the underlying assets.

2.4 Flights to Safety in MMFs

During episodes of stress in the MMF industry, investors run on the funds more exposed to the source of risk; these have usually been prime funds, which hold unsecured private debt such as bank certificates of deposit. The same investors then typically deposit their cash with government funds, which hold either public debt or secured private debt. These flight-to-safety dynamics in MMFs occurred during the 2008 financial crisis, the 2011 European sovereign debt crisis, and the March 2020 COVID-19 crisis.

In particular, the failure of Lehman Brothers on September 15, 2008, triggered a run on the Reserve Primary Fund, an MMF whose holdings of Lehman Brothers commercial paper had

Figure 4: MMFs during periods of stress



Source: iMoneyNet.

Note: Panel (a) shows the total net flows of institutional prime and government MMFs during the 2008 crisis. The dotted black vertical line indicates the beginning of the 2008 run on MMFs when Lehman Brothers failed in September 2008. Likewise, Panel (b) displays the daily net flows of institutional prime and government MMFs during the March 2020 COVID-19 crisis (dashed vertical line).

suddenly collapsed in value. Panic soon spread to other institutional prime funds, especially those exposed to other assets that were losing value, namely asset-backed commercial paper (McCabe, 2010; Kacperczyk and Schnabl, 2013; Duygan-Bump et al., 2013). Panel (a) of Figure 4 shows the total net flows for government and prime share classes offered to institutional investors during the 2008–2010 period.

The COVID-19 shock triggered another run on prime MMFs in early 2020, as shown in Panel (b) of Figure 4; similarly to 2008, MMFs experienced a sizable industry-wide reallocation from prime to government funds. Over the two-week period ending on March 20, 2020, net outflows from institutional prime funds totaled \$90 billion, 27 percent of these funds' assets at the beginning of March, while net outflows from retail prime funds were \$33 billion, 7 percent of their assets (PWG, 2020). During the same period, government funds experienced significant inflows. Cipriani and La Spada (2020) show that, due to low switching costs, investors ran from prime funds to government funds in the same family.

3 Data

We construct a panel of the 12 stablecoins listed in Table 1, using daily data from February 2015 through April 2023 from CoinGecko. For each day-stablecoin observation, we obtain the market capitalization and volume-weighted average price across exchanges in dollars. We calculate each stablecoin’s circulation, defined as the number of tokens outstanding, by dividing market capitalization by the price. We then estimate flows as changes in circulation.¹²

In the earlier part of the sample, our data include only a few stablecoins, mainly Tether. For this reason, we only use the whole sample in our analysis of depegging thresholds (Section 4.2), where we benefit from exploiting the additional time-series variation available. In our analysis of the effects of cross-sectional heterogeneity of stablecoins on investor flows, we use a shorter sample, which we will describe next.

Our main sample consists of daily data on the 12 stablecoins reported in Table 1 from January 2021 to March 15, 2023, though we use a subset of these data for each analysis. We further define two “event” samples: January 1, 2022, through the end of the Terra run on May 16, 2022 (“the 2022 sample”), and November 1, 2022, through the end of the USDC run on March 12, 2023 (“the 2023 sample”). To have the same number of stablecoins in each event sample and to ensure that our results are not driven by the behavior of the smallest stablecoins, we use the 10 largest stablecoins at the start of 2022 and 2023, respectively; we report the stablecoins in the two event samples in the last two columns of Table 1.¹³ The stablecoins in the 2022 and 2023 samples comprise more than 99 percent of market capitalization at the start of 2022 and 2023, respectively.

In each of the 2022 and 2023 samples, roughly the same number of stablecoins fall into each of four major categories of stablecoins: US-based asset-backed, offshore asset-backed, crypto-backed,

¹²CoinGecko reports market capitalization and price. CoinGecko’s glossary defines *market capitalization* as the product of the circulating supply of tokens and its current price, and *circulating supply* as the total supply minus any tokens assigned to the token’s issuer and affiliates. Because of the nature of stablecoins, these adjustments to circulating supply are not meaningful, unlike for other kinds of crypto assets that can be created and assigned without any backing. See CoinGecko’s Methodology and Glossary pages for more information.

¹³The 2022 sample includes Terra, whose stabilization algorithm collapsed in 2022. The 2023 sample includes stablecoins that were not in the top 10 in 2022, including a new algorithmic stablecoin (USDD) and an additional US-based stablecoin (GUSD).

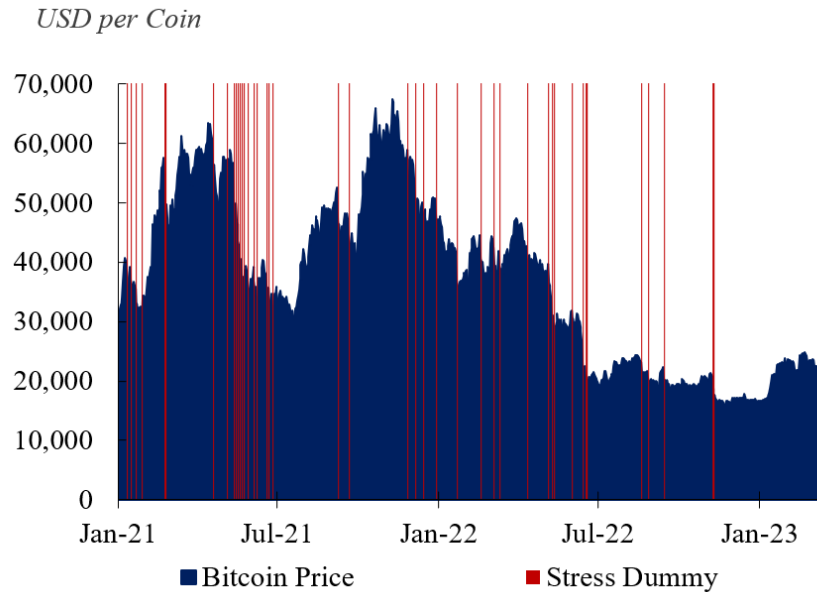
Table 1: Stablecoins in our sample by jurisdiction and stabilization mechanism.

Name	Symbol	US-based	Backing	Market Cap (\$m)	Collateral Type	2022 Sample	2023 Sample
Tether	USDT	N	Assets	73,378	Cash, US Treasuries, Corp. Debt	Y	Y
USD Coin	USDC	Y	Assets	38,426	Cash, US Treasuries, CD	Y	Y
Binance USD	BUSD	Y	Assets	8,381	Cash, US Treasuries	Y	Y
Dai	DAI	N	Crypto assets	6,042		Y	Y
TrueUSD	TUSD	N	Assets	2,032	Cash	Y	Y
Frax	FRAX	N	Algorithm	1,043		Y	Y
Pax Dollar	USDP	Y	Assets	835	Cash and Cash Equiv.	Y	Y
Decentralized USD	USDD	N	Algorithm	721		N	Y
Gemini Dollar	GUSD	Y	Assets	387	Cash, Cash Equiv., MMMF	N	Y
Liquity USD	LUSD	N	Crypto assets	257		Y	Y
TerraUSD	USTC	N	Algorithm	232		Y	N
Magic Internet Money	MIM	N	Crypto assets	83		Y	N

Sources: CoinGecko. The data are as of March 15, 2023.

Note: The table presents summary statistics for the stablecoins used in the analysis. Market Cap is in millions of US dollars. For stablecoins backed by traditional financial assets, we indicate the types of assets in the Collateral Type column.

Figure 5: Crypto Stress Measure



Sources: Authors' computations based on data from CoinGecko. Data as of March 15, 2023.

Note: The blue area represents the price of Bitcoin in US dollars. The vertical red lines indicate the stress days ("shocks") identified using the methodology described in the text. The data are daily, and the sample period is January 2021 to March 2023.

and algorithmic stablecoins. In the 2022 panel, three stablecoins are US-based and backed by safe assets (USDC, BUSD, and USDP), two are offshore and backed by safe assets (USDT and TUSD), three are backed by crypto assets (DAI, MIM, and LUSD), and two are algorithmic (USTC and FRAX).¹⁴ In the 2023 panel, four stablecoins are US-based stablecoins and backed by safe assets (USDC, BUSD, USDP, and GUSD), two are offshore and backed by safe assets (USDT and TUSD), two are backed by crypto assets (DAI and LUSD), and two are backed by algorithms (USDD and FRAX).

Throughout our sample period, we also obtain Bitcoin prices which we use in some of our analyses as an indicator of stress in the crypto ecosystem. In particular, we identify stress days as those in the bottom 5% of the daily return distribution (shown by the vertical red lines in Figure 5.)

Finally, in addition to daily prices and market capitalization from CoinGecko, we also obtain daily data from DefiLlama on the value of stablecoins staked in 10 different blockchains. These

¹⁴As described in Appendix A, TUSD used to be US-based until December 2020 but was later sold to an Asian conglomerate.

data cover the length of the Terra run, from May 8, 2022, through May 16, 2022, and the length of the USDC run, from March 9, 2023, through March 12, 2023. For three blockchains (Terra, Solana, and Tron), data only begins on May 12, 2022 and thus we exclude these blockchains from the 2022 blockchain-level flow analysis in Section 4.4 below.

4 Empirical Analysis

In this section, we provide empirical evidence of flights to safety in stablecoins. The first subsection shows how stress in Bitcoin valuation is followed by stablecoin outflows that vary in intensity for different stablecoin types. The second subsection estimates how changes in price affect stablecoin circulation, showing that many stablecoins “break the buck” and face increased redemptions when their price falls below \$1. The third subsection conducts event studies that quantify the effects of the Terra collapse and the SVB bankruptcy on the stablecoin market. The fourth subsection analyzes stablecoin flows within blockchains. Each of these stablecoin dynamics has a parallel with the behavior of investor flows in the MMF industry.

4.1 Flights to Safety in Stablecoins

In this subsection, we estimate the effect of crypto market shocks on stablecoins. Our shock measure is a dummy variable for stress days in crypto markets, identified as days in the 5th percentile of the distribution of Bitcoin daily price changes over our sample period (see Figure 5). For this analysis, we use our main sample (all twelve stablecoins) over the period starting in January 2021 and ending on March 15, 2023. Prior to the beginning of this period, the market capitalization of all but the major asset-backed stablecoins was negligible.

We use the local projection method (Jordà, 2005) to estimate the response of stablecoin flows to our shock measure. For each horizon h , we estimate the following daily regression on our panel

of stablecoins:

$$\begin{aligned} \text{Flow}_{i,t-1}^{t+h} &= \sum_{n=0}^2 \text{Shock}_{t-n} \left(\beta_{h,n}^1 \text{US}_i + \beta_{h,n}^2 \text{Offshore}_i + \beta_{h,n}^3 \text{Crypto}_i + \beta_{h,n}^4 \text{Algo}_i \right) + \\ &+ \sum_{n=1}^2 \gamma_{h,n} \text{Flow}_{i,t-n-1}^{t-n} + \mu_{h,i} + \varepsilon_{i,t}, \end{aligned} \quad (1)$$

where $\text{Flow}_{i,t-1}^{t+h}$ is the cumulative percentage change in the outstanding number of tokens of coin i at $t+h$ relative to $t-1$, while $\text{Flow}_{i,t-n-1}^{t-n}$ is the daily percentage change in the outstanding number of tokens of coin i at $t-n$ relative to $t-n-1$. Shock_t is an indicator variable that equals 1 if day t is a stress day (as defined in Section 3) and 0 otherwise. The variables US_i , Offshore_i , Crypto_i , and Algo_i are four indicator variables that equal 1 if stablecoin i is respectively US-based, offshore asset-backed, backed by crypto assets, and algorithmic (as described in Table 1).¹⁵ The variable $\mu_{h,i}$ is a coin fixed effect (which is different for each horizon h ; that is, in each regression). We include lagged stablecoin flows to control for serial correlation and possible trends; in Appendix B, we show that information-based selection criteria suggest that the optimal number of lags is two.¹⁶ We run regression (1) for horizons up to eight days after the crypto market shock, that is, $h \in \{0, 1, \dots, 8\}$.

The regression is estimated by weighted least squares using each stablecoin's market capitalization one month prior as the weight so as to give greater weight to the larger stablecoins. We use market capitalization lagged by one month to avoid any mechanical relation between current flows and current weights. Standard errors are clustered at the coin and date levels.¹⁷

The coefficients $\beta_{h,0}^1$, $\beta_{h,0}^2$, $\beta_{h,0}^3$, and $\beta_{h,0}^4$ estimate the effects of a crypto market shock at time t on the cumulative percentage flows into the different stablecoin types between time t and $t+h$. In particular, the $h=0$ coefficient represents the same-day effect of the shock, that is, the effect of a significant drop in the price of Bitcoin in the last 24 hours on the percentage flows in stablecoins during the same period.

¹⁵Since these four categories are mutually exclusive and exhaustive, $\text{US}_i + \text{Offshore}_i + \text{Crypto}_i + \text{Algo}_i = 1$.

¹⁶In Appendix B, we also report the results from robustness checks using different lag lengths; the results are similar.

¹⁷In Appendix B, we repeat the same analysis using different standard errors; the results are similar.

In Figure 6, we plot the response of stablecoins' cumulative percentage flows to our measure of crypto market stress at different horizons and for the four groups of stablecoins described above: US-based ($\beta_{h,0}^1$), offshore asset-backed ($\beta_{h,0}^2$), crypto-backed ($\beta_{h,0}^3$), and algorithmic ($\beta_{h,0}^4$). The corresponding regression estimates are shown in the Appendix (Table B.2). The blue lines represent the point estimates; the solid and dashed red lines represent the 90 percent and 95 percent confidence intervals.

The response of stablecoin flows to crypto market stress depends on collateral type. Following a large decline in the price of Bitcoin, the flows of algorithmic stablecoins have the most negative reaction (Panel (d)), followed by crypto-based stablecoins (Panel (c)), and by a more muted response of offshore asset-backed stablecoins (Panel (b)). Eight days after a shock, the expected excess cumulative flows for these three stablecoin types are -9.2% , -3.0% , and -1.1% . In contrast, the flows of US-based stablecoins (Panel (a)) have a positive reaction, with expected cumulative flows of $+1.9\%$, consistent with a flight-to-safety dynamic. The difference between US-based stablecoins and others is statistically significant (see Table B.2 in Appendix B). These differential outflows suggest that, following negative shocks to non-stablecoin crypto assets, investors shun riskier stablecoins—those backed either by riskier assets or by algorithms relying, at least in part, on investor confidence—and flee to stablecoins backed by safer assets.

Similar flight-to-safety dynamics are documented in the MMF literature (Bouveret et al., 2022; Cipriani and La Spada, 2020, 2021). Following stress events, money fund investors run on the more exposed prime funds and flee to the safer government MMFs.

4.2 Depegging Thresholds

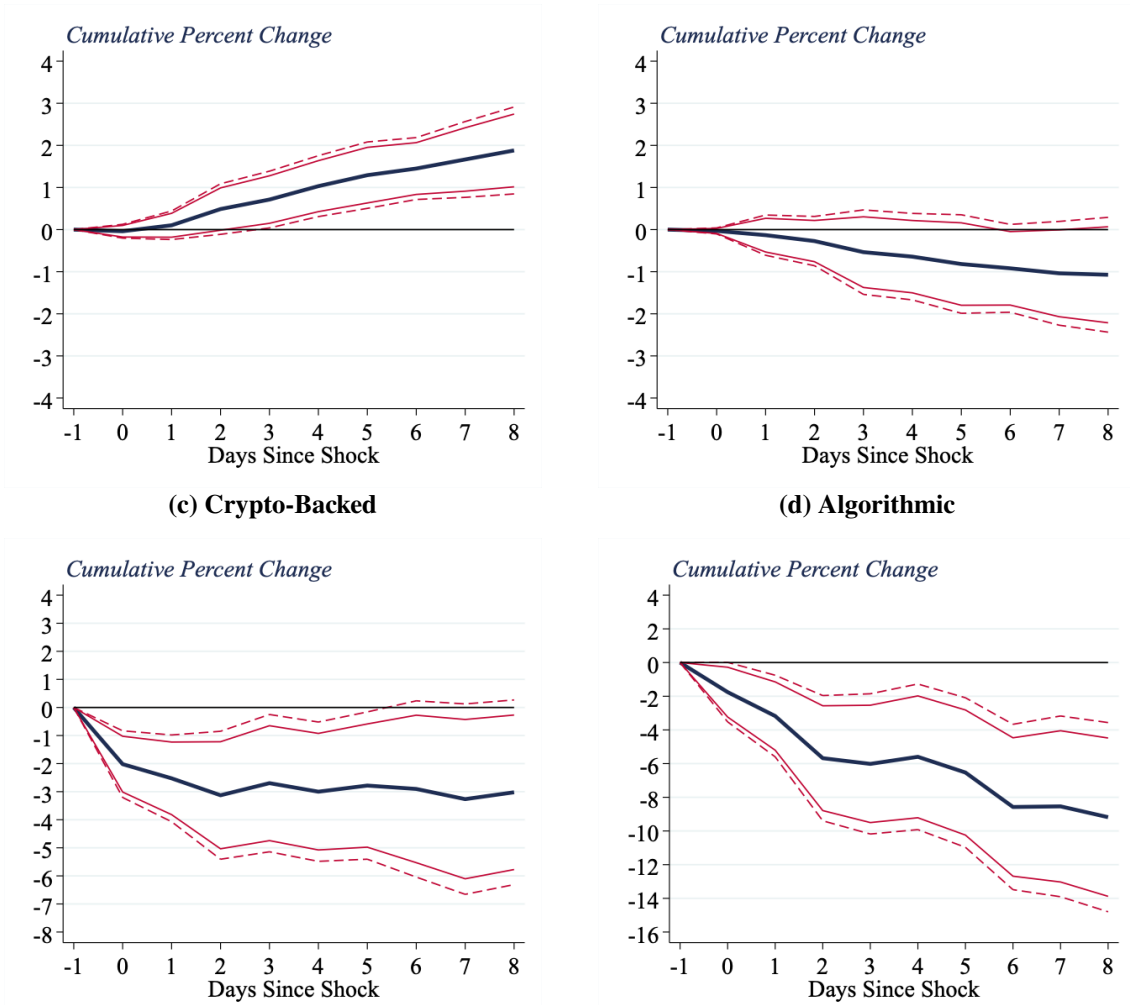
In this section, we provide insight into the threshold at which stablecoins are considered by market participants to have depegged.

There is a well-understood relationship between the level of an MMF's net asset value (NAV) and fund flows. A low NAV due to investment losses results in outflows and vice versa: Outflows

Figure 6: Impulse Response Functions by Stablecoin Type.

(a) US-based

(b) Offshore Asset-Backed



Sources: Authors' computations based on data from CoinGecko.

Note: The blue line represents the estimated impulse response function for investor net flows upon shocks to the price of Bitcoin. The solid and dashed red lines represent the 90 and 95 percent confidence intervals, respectively.

push the NAV down. In particular, a fund is said to “break the buck” if its NAV calculated using market prices drops below \$0.995; this is exactly what happened to Reserve Primary Fund in September 2008 due to its exposure to Lehman Brothers’ debt, which led investors to run on the fund.

To study whether a similar relationship exists between stablecoin prices and flows, we run the following daily regression on our panel of stablecoins:

$$\text{Flow}_{i,t-1}^t = \sum_{m \in M} \beta_m \mathbf{I}(\text{Price} < m)_{i,t-1} + \alpha_i + \gamma_t + \varepsilon_{i,t}, \quad (2)$$

where $\text{Flow}_{i,t-1}^t$ is the daily percentage change in the outstanding number of tokens of coin i at t relative to $t - 1$, $M := \{100, 99.9, \dots, 98.6, 98.5\}$ represent a grid of price levels in cents, $\mathbf{I}(\text{Price} < m)_{i,t-1}$ is an indicator variable equal to 1 if the price for stablecoin i on day $t - 1$ is strictly below m cents, α_i are stablecoin fixed effects, and γ_t are date fixed effects. Standard errors are Driscoll-Kraay with five lags.

Column (5) of Table 2 shows that a stablecoin priced below the 100-cent threshold experiences larger daily outflows than a stablecoin that maintains a \$1 peg or higher by 0.3 percentage points. For all other price thresholds, we do not observe significant effects on flows.¹⁸

In other words, our results show that when a stablecoin price falls below \$1, investors start to redeem, which is reminiscent of investor outflows from MMFs when these funds “break the buck.”

4.3 Event Studies

4.3.1 The 2022 Terra Collapse

In many ways, the May 2022 run on stablecoins mirrored the 2008 and 2020 runs on MMFs. The run on the algorithmic stablecoin Terra (USTC) had strong negative spillovers into the broader

¹⁸Outflows also accelerate below 99.1 cents, by 3.4 percentage points daily, though the effect is only significant at the 10% level.

Table 2: Daily regression of percentage change in circulation against price dummies

	Flows _{it} (%)				
	(1)	(2)	(3)	(4)	(5)
Price _{it-1}	0.01*	0.01	-0.00	-0.00	-0.00
	(1.78)	(1.09)	(-0.33)	(-0.33)	(-0.06)
I(Price < 100) _{it-1}		-0.57***	-0.34***	-0.34***	-0.34***
		(-4.30)	(-2.98)	(-2.98)	(-2.97)
I(Price < 99.9) _{it-1}			0.01	0.01	0.01
			(0.07)	(0.07)	(0.07)
I(Price < 99.8) _{it-1}			-0.36	-0.36	-0.36
			(-1.47)	(-1.47)	(-1.46)
I(Price < 99.7) _{it-1}			0.05	0.06	0.06
			(0.18)	(0.20)	(0.21)
I(Price < 99.6) _{it-1}			-0.08	-0.09	-0.09
			(-0.19)	(-0.19)	(-0.20)
I(Price < 99.5) _{it-1}			-0.81	-0.87	-0.87
			(-1.59)	(-1.09)	(-1.09)
I(Price < 99.4) _{it-1}				0.85	0.86
				(0.85)	(0.86)
I(Price < 99.3) _{it-1}				-0.65	-0.65
				(-0.74)	(-0.74)
I(Price < 99.2) _{it-1}				0.62	0.62
				(0.84)	(0.84)
I(Price < 99.1) _{it-1}				-3.44*	-3.43*
				(-1.95)	(-1.95)
I(Price < 99) _{it-1}				2.65	2.71
				(1.49)	(1.47)
I(Price < 98.9) _{it-1}					-0.08
					(-0.07)
I(Price < 98.8) _{it-1}					-1.66
					(-1.12)
I(Price < 98.7) _{it-1}					1.75
					(1.58)
I(Price < 98.6) _{it-1}					0.35
					(0.24)
I(Price < 98.5) _{it-1}					-0.21
					(-0.14)
Coin FE	Yes	Yes	Yes	Yes	Yes
Date FE	Yes	Yes	Yes	Yes	Yes
R ²	0.18	0.18	0.18	0.18	0.18
Observations	14431	14431	14431	14431	14431

$I(\text{Price} < m)_{it-1}$ is an indicator variable equal to 1 if the average price for stablecoin i on day $t - 1$ is strictly below m cents. t-statistics, based on Driscoll-Kraay standard errors with 5 lags, are in parentheses. ***, **, and * represent 1%, 5%, and 10% statistical significance.

offshore stablecoin industry; the only exceptions were US-based stablecoins, which received inflows, similarly to how government MMFs receive inflows when investors run from prime MMFs.

The run began when Terra's algorithmic stabilization mechanism collapsed. As described in Appendix A.4, arbitrage should peg the price of Terra at \$1 as long as the price of Luna is above zero, but features of the mechanism leave open the possibility for depegging events. In early 2022, asset prices declined, including those of important crypto assets such as Bitcoin, Ethereum, and Luna. At the same time, there was increasing awareness that Anchor, a protocol offered by the developers of Terra, would not be able to keep their promised 20 percent interest rate to Terra depositors.¹⁹ The decreasing value of Luna, combined with the uncertainty about the Anchor protocol rate, led to a loss of confidence in the ability of the Terra algorithm to maintain the US dollar peg and triggered the run on Terra.

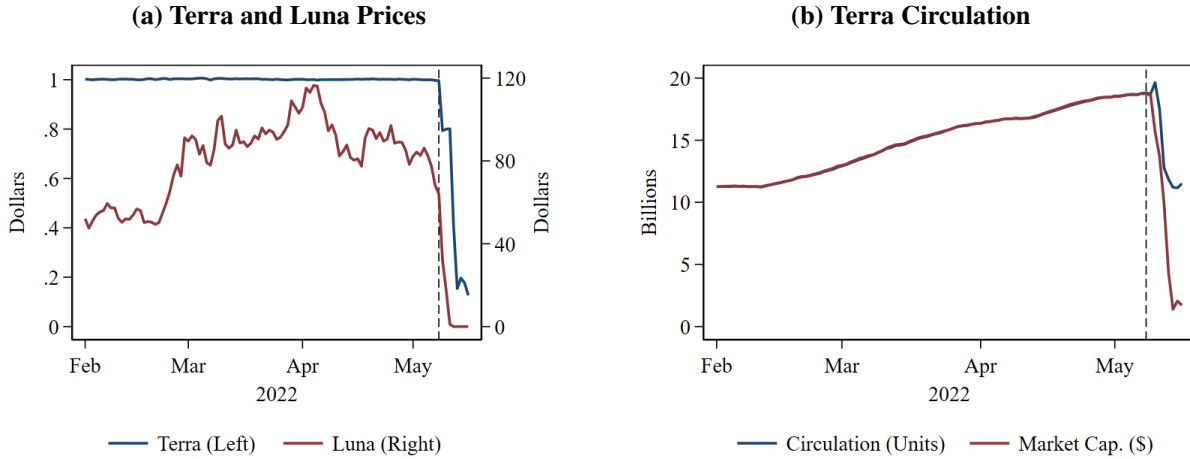
Panel (a) of Figure 7 shows that Luna's price dipped significantly on May 8, 2022, when the token closed at \$64.08. By the next day, Luna's price had halved to \$32.00; by May 12, Luna was trading at \$0.4 cents, or 0.004% of its January 1 value. The same figure shows that Terra's price plummeted following the collapse of Luna: USTC was worth \$0.996 on May 8, \$0.800 on May 9, and only \$0.409 on May 12. After May 17, 2022, the price of Terra never exceeded \$0.10. Importantly, the price of Terra neared zero because investors expected the price of Luna to remain at zero; as a result, the algorithmic stability mechanism was broken.

Panel (b) shows that the run reversed all of Terra's 2022 growth as circulation returned to January 2022 values, and its market capitalization dropped by about 95 percent.

As with MMF runs, the May 2022 stablecoin run was characterized by strong negative spillover into the broader industry. Panel (a) of Figure 8 shows significant outflows from all offshore stablecoins relative to US-based stablecoins in a behavior reminiscent of runs among prime MMFs. Panel (b) breaks down the results by stablecoin type. It shows that non-US asset-backed, algorithmic, and

¹⁹The developers of Terra and the Anchor protocol would lend Terra deposits to borrowers, who used them to buy more volatile crypto assets and invest in decentralized finance protocols. The operator of the Anchor protocol aggregated these higher and more volatile yields and passed them to Terra depositors. The massive growth of Terra in early 2022 was largely due to the interest rates paid by the Anchor protocol; according to decrypt.co, 72 percent of all Terra in circulation was invested in Anchor.

Figure 7: The Terra collapse of 2022.



Source: CoinGecko.

Note: The left panel shows daily closing prices for Luna and Terra. The right panel shows the circulation and market capitalization of Terra in 2022. The dashed vertical lines are drawn at May 8, 2022, the first day of the stablecoin run.

crypto-collateralized stablecoins experienced outflows in the wake of Terra’s depeg event, and the effect was strongest for the algorithmic stablecoin Frax, which suffered a 45 percent loss of its May 1 market capitalization by May 16, 2022.

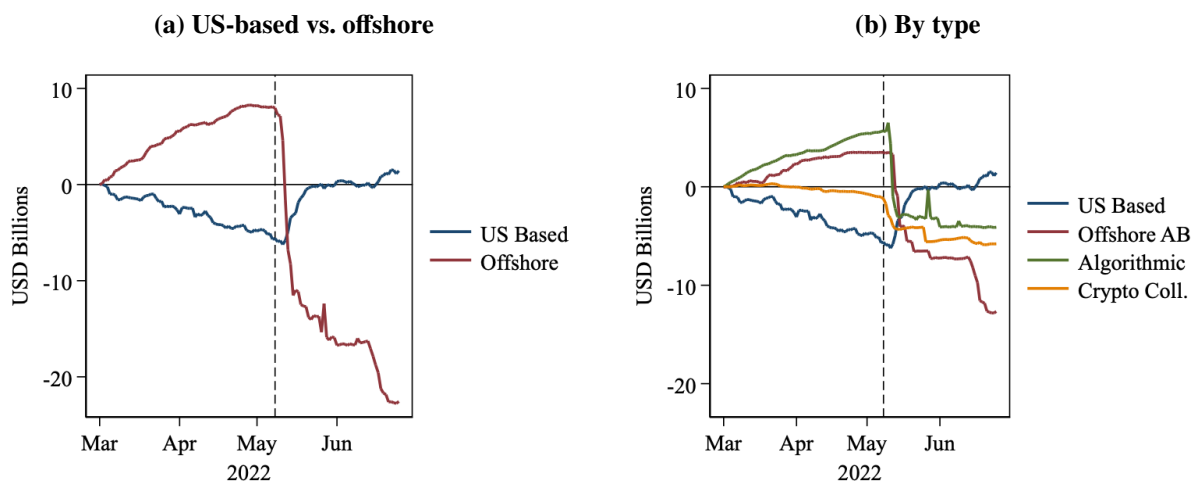
To quantify the outflows suffered by offshore stablecoins relative to their US-based counterparts in May 2022, we run the following daily regression:

$$\text{Flow}_{i,t-1}^t = \beta \text{Run}_t \times \text{Offshore}_i + \alpha_i + \gamma_t + \varepsilon_{i,t}, \quad (3)$$

where $\text{Flow}_{i,t-1}^t$ is the daily percentage change in the outstanding number of tokens of coin i at t relative to $t - 1$, Run_t is an indicator variable for the run period (May 8 through May 16, 2022, inclusive), Offshore_i is an indicator variable for non-US-based stablecoins (asset-backed, crypto-backed, or algorithmic), α_i are stablecoin fixed effects, and γ_t are date fixed effects. Standard errors are Driscoll-Kraay with five lags. The sample is January 1 through May 16, 2022.

Column (1) of Table 3 shows that during the run, the daily percentage change in circulation in offshore asset-backed, crypto-collateralized, and algorithmic stablecoins was 2.3 percentage points

Figure 8: Stablecoin flows around the Terra collapse



Source: CoinGecko.

Note: Cumulative flows by stablecoin type since March 1, 2022. The dashed vertical line is drawn at May 8, 2022, the first day of the stablecoin run. Panel (a) shows all offshore stablecoins together, while Panel (b) breaks them down by type. “Offshore AB” stands for “Offshore Asset-Backed.”

lower than the change in circulation in US-based stablecoins. Column (2) breaks down this analysis by category, showing that algorithmic and crypto-collateralized stablecoins were the most affected by the run, with changes in circulation 5.7 and 3.6 percentage points lower than those in US-based stablecoins; the impact on offshore asset-backed stablecoins was smaller (1.4 percentage points).

4.3.2 Silicon Valley Bank Failure and the 2023 Run on USDC

Since 2020, commercial banks have increased their provision of services to crypto industry clients. As stablecoins grew rapidly from early 2020 through 2022, they increased their deposits with some commercial banks. Figure 9 shows the total deposits of three banks serving stablecoins from 2018 through 2022.²⁰ As shown in the figure, since 2020, these banks’ deposits experienced significant growth, along with the market for stablecoins. If one of these banks were to find itself in trouble, a stablecoin with deposits above the FDIC insurance limit could suddenly be unable to access these deposits and fail to meet its obligations. This would likely trigger a run by its investors.

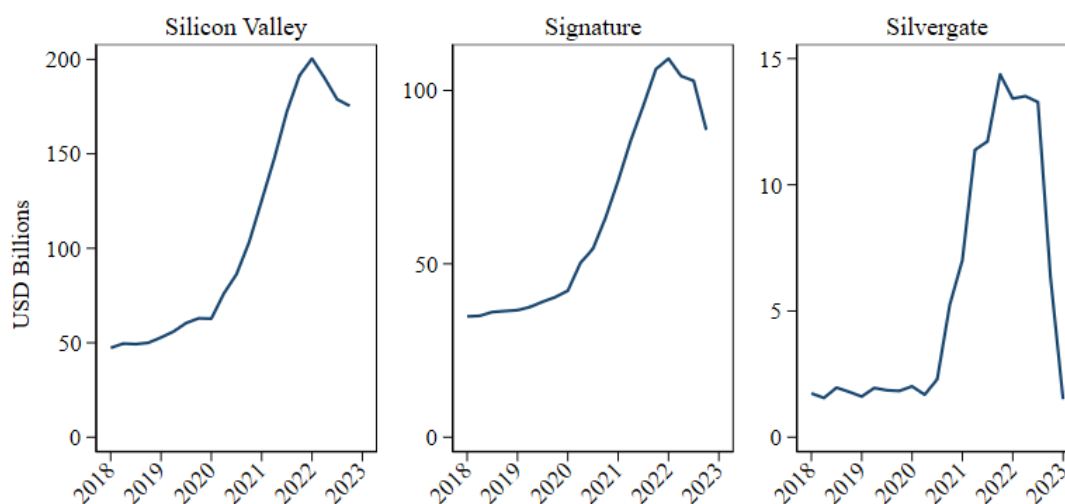
²⁰We identify these banks based on the list of banks disclosed in stablecoin issuers’ attestations.

Table 3: Daily regression of percentage change in circulation against indicator for the May 2022 run period interacted with stablecoin types.

	Flows _{it} (%)	
	(1)	(2)
Run _t × Offshore _i	-2.345** (0.899)	
Run _t × Offshore Asset Backed _i		-1.392* (0.831)
Run _t × Algorithmic _i		-5.719** (2.746)
Run _t × Crypto Collateralized _i		-3.602*** (1.176)
Coin FE	Yes	Yes
Date FE	Yes	Yes
Run Definition	5/8/22-5/16/22	5/8/22-5/16/22
Sample	1/1/22-5/16/22	1/1/22-5/16/22
R ²	0.24	0.29
Observations	1360	1360

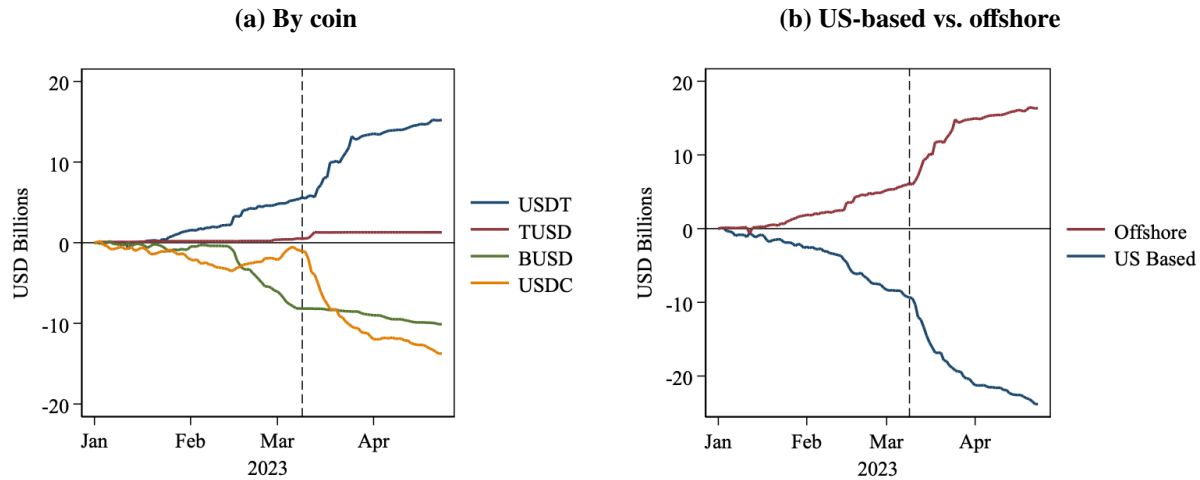
Note: Run_t equals 1 on May 8 through May 16, 2022, inclusive. The US-based stablecoins are USDC, BUSD, and USDP; offshore asset-backed stablecoins are USDT and TUSD; crypto-collateralized stablecoins are DAI, MIM, and LUSD; and algorithmic stablecoins are USTC and FRAX. Driscoll-Kraay standard errors with five lags are in parentheses. ***, **, and * represent 1%, 5%, and 10% statistical significance.

Figure 9: Growth in deposits of banks that serve stablecoin issuers.



Source: Call Reports.

Figure 10: Stablecoin flows around the Silicon Valley Bank failure



Source: CoinGecko.

Note: Cumulative flows since January 1, 2023. Panel (a) shows disaggregated flows by stablecoin, showing only those with significant net flows in March 2023. Panel (b) shows aggregate flows for US based versus Offshore stablecoins. The dashed vertical line is drawn at March 9, 2023, the day SVB was taken over by the FDIC.

The stablecoin USDC found itself precisely in this situation when Silicon Valley Bank (SVB) suffered a run on March 9, 2023, and was taken over by the FDIC on March 10; on March 12, the FDIC, with the Department of the Treasury and the Board of Governors of the Federal Reserve System, announced that SVB resolution would fully protect all depositors, including uninsured ones.²¹ The run on USDC began on March 9 when investors became concerned about USDC's ability to access the cash it had deposited with SVB. From July 2022 through March 2023, USDC listed publicly that it held part of its cash at SVB. When SVB collapsed, USDC's investors quickly ran. Panel (a) of Figure 10 shows that after March 9, the majority of USDC outflows likely went into USDT and TUSD, two offshore asset-backed stablecoins. Note that other stablecoins in our sample did not have sizable and persistent dollar flows, so we exclude them from the plot.

The run on USDC was different from the MMF runs of 2008 and 2020 and from the Terra run of 2022 in two ways. First, the run on USDC did not have negative spillovers into stablecoins similar to USDC. Our 2023 sample contains four US-based asset-backed stablecoins: USDC, BUSD, USDP, and GUSD. In March 2023, GUSD and USDP did not have significant dollar flows, and BUSD

²¹See <https://www.federalreserve.gov/newsevents/pressreleases/monetary20230312b.htm>.

experienced a noticeable slowdown of previously steady outflows that had begun in mid-February.²² Although similar to USDC in terms of jurisdiction and broad portfolio composition, BUSD, USDP, and GUSD promptly reported having no exposure to SVB in March 2023.²³

The second way in which the 2023 stablecoin run distinguished itself from most MMF runs and the 2022 stablecoin run was that assets moved from a stablecoin that was previously considered among the safest ones to stablecoins that were generally seen as riskier. Panel (b) of Figure 10 shows aggregate cumulative flows from January to April 2023, separately for US-based and offshore stablecoins of all kinds. The figure shows clearly that as US-based stablecoins (mostly USDC, with some additional outflows from BUSD’s earlier run) suffered outflows during the March run, offshore stablecoins (asset-backed, crypto-collateralized, and algorithmic) experienced significant inflows.

To quantify the inflows experienced by offshore stablecoins relative to US-based stablecoins during the 2023 run, we use regression (3). Column (1) of Table 4 shows that, compared with US-based stablecoins, the average offshore stablecoin experienced larger daily inflows by 1.3 percentage points from March 9 to March 12 (our definition of the run period). Column (2) separates the analysis by stablecoin type to show that offshore asset-backed, algorithmic, and crypto-collateralized stablecoins all received inflows during the run. The effect on crypto-collateralized stablecoins is especially large in the short term because DAI experienced large daily inflows, beginning on March 11. These inflows, however, reversed by the end of March and were small in dollar amounts relative to inflows to the offshore asset-backed stablecoins Tether and TrueUSD (see again Figure 10).²⁴

²²Outflows from BUSD started after its issuer (Paxos) announced that it would stop minting BUSD at the direction of the New York State Department of Financial Services. BUSD lost market capitalization at a rate of about \$350 million per day from February 14 to March 6, 2023; this loss slowed to about \$38 million per day from March 7 to June 5, 2023.

²³For BUSD and USDP, Paxos issued the following statement on Twitter: “Paxos has no relationship with Silicon Valley Bank. For further certainty, Paxos stablecoins have zero exposure to Silicon Valley Bank failure.” See <https://twitter.com/Paxos/status/1634242623916642304>. Similarly, regarding GUSD, Gemini issued the following statement: “@Gemini does not have any banking relationship with Silicon Valley Bank.” See <https://twitter.com/Gemini/status/1634260777686257664>.

²⁴As a robustness check, we extend the run period to the end of March 2023 and show that the coefficients on the offshore asset-backed and algorithmic categories are comparable to those in Table 4. The coefficient on crypto-collateralized stablecoins is positive but not significant at the 10 percent level, indicating that the effect of the run on crypto-collateralized stablecoins was strong only in the short run.

Table 4: Daily regression of percentage change in circulation against indicator for the March 2023 run period interacted with stablecoin types.

	Flows _{it} (%)	
	(1)	(2)
Run _t × Offshore _i	1.295** (0.575)	
Run _t × Offshore Asset Backed _i		1.162** (0.508)
Run _t × Algorithmic _i		1.163*** (0.410)
Run _t × Crypto-Collateralized _i		3.052** (1.520)
Coin FE	Yes	Yes
Date FE	Yes	Yes
Run Definition	3/9/23-3/12/23	3/9/23-3/12/23
Sample	11/1/22-3/12/23	11/1/22-3/12/23
R ²	0.12	0.12
Observations	1320	1320

Note: Run_t equals 1 on March 9 through March 12, 2023, inclusive. The US-based stablecoins are USDC, BUSD, USDP, and GUSD; offshore asset-backed stablecoins are USDT and TUSD; crypto-collateralized stablecoins are DAI and LUSD; and algorithmic stablecoins are USDD and FRAX. Driscoll-Kraay standard errors with 5 lags are in parentheses. ***, **, and * represent 1%, 5%, and 10% statistical significance.

4.3.3 Investors' Flight to Safety: MMFs vs. Stablecoin Runs

During periods of crypto market stress, stablecoin investors tend to flee to safety. What constitutes “safety,” however, depends on the circumstances. Specifically, whether an asset is considered “safe” depends on the nature of the risk identified by investors as the driver of an adverse event. Following a shock, investors run on investment vehicles that are more exposed to the identified risk while at the same time moving funds to the vehicles with less or no exposure to such risk.

This principle not only explains our finding that stablecoin investors flee to US-based stablecoins following crypto market shocks, but it also explains the dynamics of flows to the two largest stablecoins, Tether and USDC, during the 2022 and 2023 stablecoin runs. In the aftermath of the collapse of TerraUSD in May 2022, USDC—the second-largest stablecoin—experienced large net inflows, as investors at that time deemed it the safest stablecoin because it reportedly held Treasury securities and bank deposits, assets generally regarded as safe. By contrast, Tether—the largest stablecoin—experienced net outflows, likely because it held commercial paper and other securities that are generally considered riskier (Figure 8, Panel (b)).

After the failure of SVB in March 2023, however, USDC experienced large net redemptions (Figure 10), and its price dropped to as low as \$0.88.²⁵ These redemptions occurred because \$3.3 billion (about 8 percent) of USDC’s assets comprised uninsured deposits held at SVB.²⁶ The uncertainty around the fate of those uninsured deposits rendered some of USDC’s assets risky, as they might not have been available to maintain USDC’s \$1.00 peg. The chance of nontrivial losses in the USDC portfolio made it rational for USDC holders to run. Tether did not experience a similar shock to its asset portfolio because it did not have cash deposits at SVB or other banks negatively affected by the SVB run. As a result, Tether was the beneficiary of notable flight-to-safety inflows during this episode, and its price rose to \$1.03.

²⁵This price reflects secondary market transactions. As explained in Section 2.1.1, in addition to secondary market transactions, some investors may redeem or purchase tokens directly with the stablecoin issuer (that is, primary market transactions). We do not know if any USDC tokens were redeemed for less than \$1 in the primary market.

²⁶These numbers are based on USDC’s reported assets of \$42.1 billion as of March 13, 2023. Other assets backing USDC included \$32.4 billion (77 percent of the total) in US Treasury bills and \$5.4 billion (13 percent) and \$1 billion in deposits at BNY Mellon and Customers Bank, respectively (Circle, 2023).

Note, however, that in the MMF industry, flight-to-safety episodes have always seen investors leaving prime MMFs for government MMFs. The reason is that government funds had limited exposure to the risks that triggered the runs on prime funds; in contrast to prime MMFs, government MMFs can hold only government debt and repos backed by it, and they cannot impose redemption gates or liquidity fees.

4.4 Blockchain-Specific Run Dynamics

Cipriani and La Spada (2020) show that during the 2008 and 2020 MMF runs, outflows from prime MMF investors largely went into government MMFs within the same fund family, possibly due to low switching costs. Panel (a) of Figure 11, for example, shows the scatterplot of cumulative government-MMF inflows against cumulative prime-MMF outflows within the same fund family during the March 2020 MMF run, together with a least-squares fit line, indicating a clear positive relationship.

We find a similar behavior during stablecoin runs. There are no “fund families” in the stablecoin world; there are, however, blockchains, within which stablecoins are traded. It is possible to move stablecoins from one blockchain into another, but doing so may incur transaction fees both on the origin and destination blockchain. To study whether stablecoin runs within blockchains display a similar pattern as MMF runs within fund families, Panel (b) of Figure 11 shows the daily outflows from offshore stablecoins in each blockchain versus the daily inflows into US stablecoins within the same blockchain during the 2022 run; that is, each dot represents a blockchain-day pair. Overall, consistent with the pattern observed in MMF families, days of higher outflows from offshore stablecoins are also days of higher inflows into US stablecoins within the same blockchain.

If we examine the dynamics in large and small blockchains separately (Panels (c) and (d) of Figure 11), however, we observe an important difference. For the large blockchains, the pattern is similar to the aggregate one, with outflows from offshore coins corresponding to inflows into US-based coins on the same blockchain. For smaller blockchains, in contrast, we find that days of high outflows from offshore coins are also days of high outflows from US-based coins. This suggests

that in May 2022, investors did not switch from offshore to US-based stablecoins within smaller blockchains, but rather they left these blockchains altogether, consistent with these blockchains being perceived as riskier.²⁷

In Table 5, we formally confirm the visual evidence from Figure 11 by regressing daily inflows into US-based stablecoins in large (small) blockchains on outflows to offshore stablecoins in large (small) blockchains, during the 2022 run.²⁸ The first column shows the results of for all blockchains; the second column shows the results for blockchains with stablecoin circulations above \$10 billion; the third shows results for blockchains with circulation below \$10 billion. These results corroborate the patterns observed above: investors flows from offshore to US stablecoins in larger blockchains, whereas they flee both stablecoin types in smaller blockchains. Namely, during the 2022 stablecoin run, a one-dollar outflow from offshore stablecoins corresponds to a 42-cent inflow in US-based stablecoins within large blockchains; in smaller blockchains, a one-dollar outflow from offshore stablecoins corresponds to a one-dollar outflow from US ones.

In Figure 12 and Table 6, we repeat our blockchain-level analysis for the March 2023 run, when investors left US stablecoins for offshore ones. Panel (a) of Figure 12 shows that, during March 9-12, 2023, days of higher outflows from US stablecoins are also days of higher inflows into offshore stablecoins within large blockchains. Panel (b) shows that, in contrast, higher outflows from US stablecoins are associated with higher outflows from US stablecoins in small blockchains, although the effect is less pronounced than during the 2022 run. Table 6 tests both of these observations formally. Within large blockchains, a one-dollar outflow from US stablecoins (stressed ones) corresponds to a 22-cent inflow in offshore stablecoins during the 2023 run. Within small blockchains, instead, the effect is negative and insignificant.

²⁷In proof-of-stake blockchains, a user is selected as a transaction validator with a probability proportional to the total share of the native token held by that user. Other things equal, smaller proof-of-stake chains are less secure because it would be easier for an adversary to gain control of a large amount of stake and manipulate transactions in their favor via double spending attacks.

²⁸As discussed in Section 3, this analysis is limited to the seven blockchains for which DefiLlama has data throughout the run period.

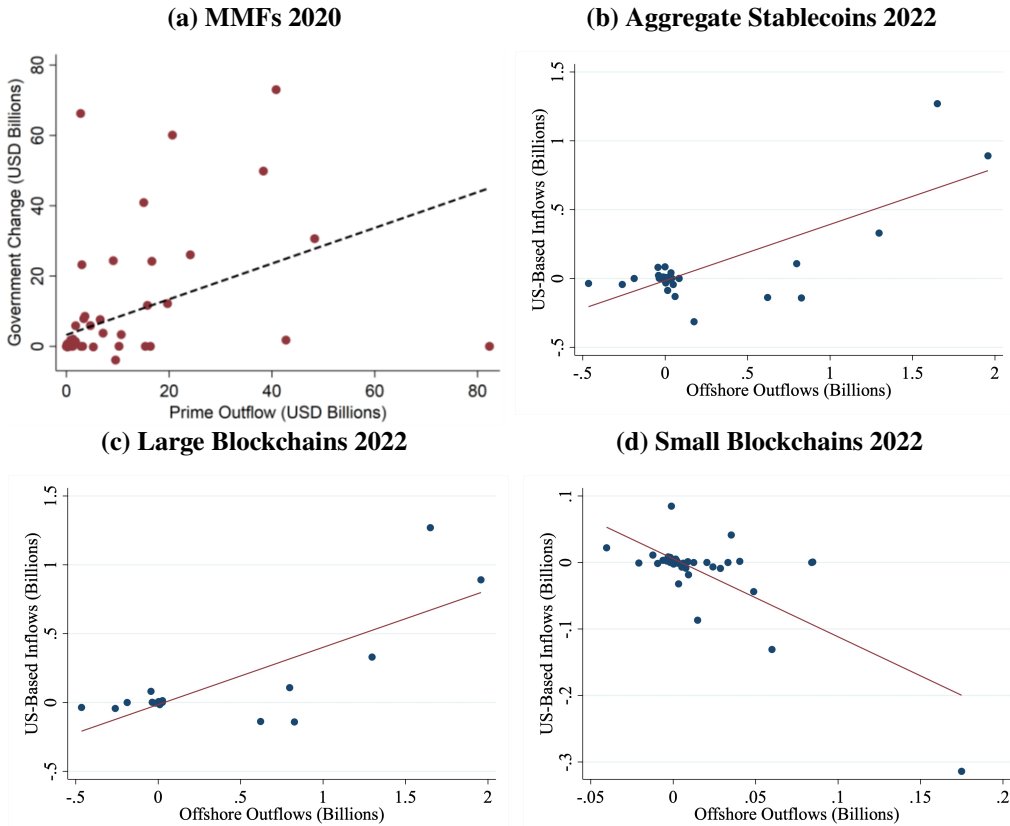
	(1)	(2)	(3)
	US-Based Inflows		
Offshore Outflows	0.408*** (3.58)	0.416** (3.64)	-1.174** (-2.78)
Constant	-0.015 (-1.46)	-0.015 (-0.44)	0.006 (1.23)
Blockchain Size	All	Large	Small
Sample (2022 Run)	5/8/22-5/16/22	5/8/22-5/16/22	5/8/22-5/16/22
N	63	18	45
R^2	0.600	0.616	0.536

Table 5: This table shows the results from regressing aggregate inflows into US-based stablecoins on aggregate outflows from offshore stablecoins during the May 2022 run for the subset of blockchains covered by DefiLlama during all days of the run. This excludes Terra, Solana, and Tron, for which DefiLlama only provides data starting on May 12, 2022. The three columns show the results, respectively, of aggregate flows across all 7 remaining blockchains, all blockchains with a circulation above \$10 billion (Ethereum and Binance Smart Chain), and all blockchains with a circulation below \$10 billion (Arbitrum, Avalanche, Fantom, Optimism, and Polygon). t-statistics, based on standard errors robust to heteroskedasticity, are in parentheses. ***, **, and * represent 1%, 5%, and 10% statistical significance.

	(1)	(2)	(3)
	Offshore Inflows		
US-Based Outflows	0.237*** (15.01)	0.218*** (10.76)	-0.038 (-0.93)
Constant	0.011 (1.02)	0.055 (1.41)	-0.001 (-0.88)
Blockchain Size	All	Large	Small
Sample (2023 Run)	3/9/23-3/12/23	3/9/23-3/12/23	3/9/23-3/12/23
N	40	12	28
R^2	0.702	0.684	0.036

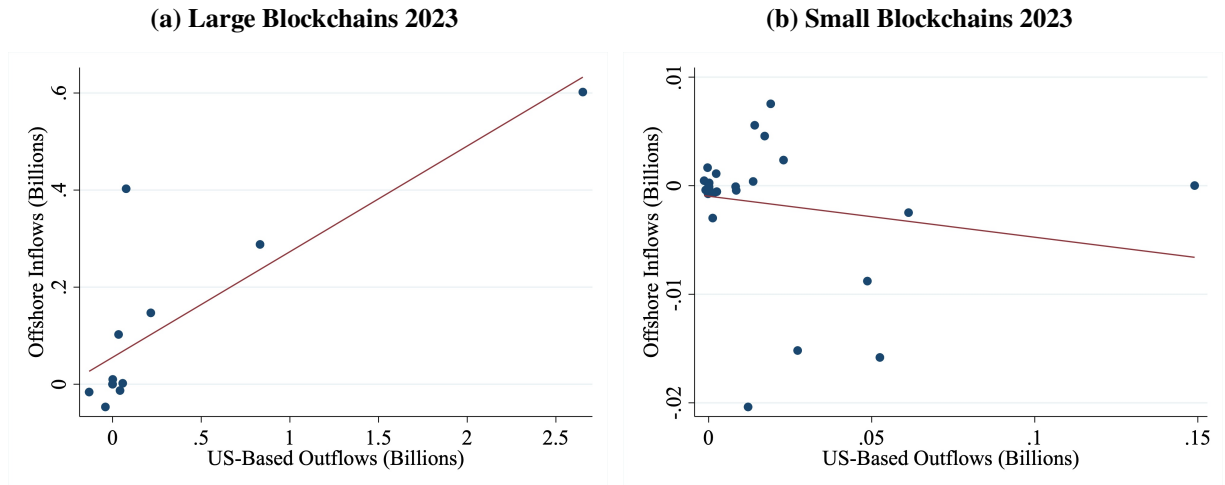
Table 6: This table shows the results from regressing aggregate outflows from US-based stablecoins on aggregate inflows into offshore stablecoins during the March 2023 run using data from DefiLlama. The three columns show the results, respectively, of aggregate flows across all 10 blockchains, large blockchains (Ethereum, Tron, and Binance Smart Chain), and small blockchains (Arbitrum, Avalanche, Fantom, Optimism, Polygon, Solana, and Terra). t-statistics, based on standard errors robust to heteroskedasticity, are in parentheses. ***, **, and * represent 1%, 5%, and 10% statistical significance.

Figure 11: Flows from offshore to US-based stablecoins (prime to government MMFs).



Note: Panel (a) shows intra-family flows in MMFs in March 2020, Panel (b) shows daily aggregate intra-industry stablecoin flows over the May 8–16, 2022, period and Panels (c) and (d) show daily intra-blockchain stablecoin flows in large and small blockchains, respectively, over the May 8–16, 2022, period. Large blockchains are Ethereum and Binance Smart Chain (BSC). Small blockchains are Arbitrum, Avalanche, Fantom, Optimism, and Polygon. Source: DefiLlama.

Figure 12: Flows from offshore to US-based stablecoins.



Note: Daily intra-blockchain stablecoin flows in large and small blockchains, respectively, over the March 9–12, 2023, period. Large blockchains are Ethereum, Binance Smart Chain (BSC), and Tron. Small blockchains are Arbitrum, Avalanche, Fantom, Optimism, Polygon, Solana, Terra. Source: DefiLlama.

5 Conclusion

While flight-to-safety dynamics in money market funds have been extensively documented in the literature—with money flowing from the riskier prime segment of the industry to the safer government segment—little is known about whether such dynamics are also at play among stablecoins. In this paper, we bridge this gap by showing that flight-to-safety dynamics in stablecoins resemble those in the MMF industry. During periods of stress in crypto markets, safer stablecoins experience net inflows, while riskier ones suffer net outflows. This is true not only on average, but also during specific periods of high market stress, such as the 2022 run on Terra, which spilled over into similarly risky stablecoins, and the 2023 collapse of Silicon Valley Bank, which spilled over into stablecoins exposed to the failing bank.

We also show that flight-to-safety flows across stablecoins tend to occur within the same blockchain, in a similar fashion as flight-to-safety flows across MMFs occur within the same fund family; unless the blockchain is small and relatively risky. Moreover, we estimate that when a stablecoin price drops below \$1, investor redemptions accelerate significantly, in a way that is reminiscent of MMFs’ “breaking the buck.”

Our findings show that stablecoins are vulnerable to runs during periods of broad crypto market dislocation as well as idiosyncratic stress events. Should stablecoins continue to grow and become more interconnected with key financial markets, such as short-term funding markets, they could become a source of financial instability for the broader financial system.

References

- Adams, A., Ibert, M., 2022. Runs on algorithmic stablecoins: Evidence from Iron, Titan, and Steel. Finance and Economics Discussion Series Notes.
- Anadu, K., Sanders, S., 2021. Money market mutual funds: Runs, emergency liquidity facilities, and potential reforms. Federal Reserve Bank of Boston Notes.
- Arner, D., Auer, R., Frost, J., 2020. Stablecoins: risks, potential and regulation. BIS Working Paper No. 905, available at <https://www.bis.org/publ/work905.pdf>.
- Azar, P.D., Baughman, G., Carapella, F., Gerszten, J., Lubis, A., Perez-Sangimino, J.P., Scotti, C., Swem, N., Vardoulakis, A., Rappoport W, D.E., 2022. The financial stability implications of digital assets. Federal Reserve Bank of New York Staff Report 2022-058.
- Barthélémy, J., Gardin, P., Nguyen, B., 2023. Stablecoins and the financing of the real economy. SSRN working paper no. 4359660.
- Bertsch, C., 2023. Stablecoins: Adoption and fragility. Sveriges Riksbank Working Paper Series, 2023. Available on SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4466431.
- Bouveret, A., Martin, A., McCabe, P.E., 2022. Money market fund vulnerabilities: A global perspective. FRB of New York Staff Report 1009.
- Chernenko, S., Sunderam, A., 2014. Frictions in shadow banking: Evidence from the lending behavior of money market mutual funds. *Review of Financial Studies* 27, 1717–1750.
- Cipriani, M., La Spada, G., 2020. Sophisticated and unsophisticated runs. FRB of New York Staff Report .
- Cipriani, M., La Spada, G., 2021. Investors' appetite for money-like assets: The MMF industry after the 2014 regulatory reform. *Journal of Financial Economics* 140, 250–269.

Circle, 2023. An update on USDC & Silicon Valley Bank. Circle Blog.

d'Avernas, A., Maurin, V., Vandeweyer, Q., 2022. Can stablecoins be stable? University of Chicago, Becker Friedman Institute for Economics Working Paper .

Duygan-Bump, B., Parkinson, P., Rosengren, E., Suarez, G.A., Willen, P., 2013. How effective were the Federal Reserve emergency liquidity facilities? Evidence from the Asset-Backed Commercial Paper Money Market Mutual Fund Liquidity Facility. *Journal of Finance* 68, 715–737.

Federal Reserve Board, 2022. Financial stability report. Board of Governors of the Federal Reserve System.

Frost, J., Shin, H.S., Wierds, P., 2020. An early stablecoin? the bank of amsterdam and the governance of money. BIS Working Paper No. 902.

Gallagher, E.A., Schmidt, L.D., Timmermann, A., Wermers, R., 2020. Investor information acquisition and money market fund risk rebalancing during the 2011–2012 Eurozone crisis. *Review of Financial Studies* 33, 1445–1483.

Gissler, S., Macchiavelli, M., Narajabad, B., 2021. Providing safety in a rush: How did shadow banks respond to a \$1 trillion shock. Social Science Research Network Working Paper. <https://ssrn.com/abstract=3595417>.

Gorton, G.B., Klee, E.C., Ross, C.P., Ross, S.Y., Vardoulakis, A.P., 2022a. Leverage and stablecoin pegs. National Bureau of Economic Research Working Paper Series No. 30796.

Gorton, G.B., Ross, C.P., Ross, S.Y., 2022b. Making money. National Bureau of Economic Research Working Paper Series.

Gorton, G.B., Zhang, J., 2021. Taming wildcat stablecoins. *University of Chicago Law Review* 90.

Hoang, L.T., Baur, D.G., 2021. How stable are stablecoins? *European Journal of Finance* 234, 1–17.

- Ivashina, V., Scharfstein, D.S., Stein, J.C., 2015. Dollar funding and the lending behavior of global banks. *Quarterly Journal of Economics* 130, 1241–1281.
- Jordà, Ò., 2005. Estimation and inference of impulse responses by local projections. *American Economic Review* 95, 161–182.
- Kacperczyk, M., Schnabl, P., 2013. How safe are money market funds? *Quarterly Journal of Economics* 128, 1073–1122.
- Kim, S.R., 2022. How the cryptocurrency market is connected to the financial market. SSRN working paper no. 4106815.
- Kozhan, R., Viswanath-Natraj, G., 2021. Decentralized stablecoins and collateral risk. WBS Finance Group Research Paper .
- Li, L., Li, Y., Macchiavelli, M., Zhou, X., 2021. Liquidity restrictions, runs, and central bank interventions: Evidence from money market funds. *Review of Financial Studies* 34, 5402–5437.
- Li, Y., Mayer, S., 2022. Money creation in decentralized finance: A dynamic model of stablecoin and crypto shadow banking. Available on SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3757083.
- Liu, J., Makarov, I., Schoar, A., 2023. Anatomy of a run: The Terra Luna crash. National Bureau of Economic Research Working Paper Series No. 31160.
- Lyons, R.K., Viswanath-Natraj, G., 2023. What keeps stablecoins stable? *Journal of International Money and Finance* 131.
- Ma, Y., Zeng, Y., Zhang, A.L., 2023. Stablecoin runs and the centralization of arbitrage. Social Science Research Network Working Paper. <https://ssrn.com/abstract=4398546>.
- McCabe, P.E., 2010. The cross section of money market fund risks and financial crises. FEDS Working Paper.

Montiel Olea, J., Plagborg-Møller, M., 2021. Local projection inference is simpler and more robust than you think. *American Economic Review* 89, 1789–1823.

Nakamoto, S., 2008. Bitcoin: A peer-to-peer electronic cash system. *Decentralized Business Review* , 21260.

Oefele, N., Baur, D.G., Smales, L.A., 2024. Flight-to-quality—money market mutual funds and stablecoins during the march 2023 banking crisis. *Economics Letters* 234.

PWG, 2020. Overview of recent events and potential reform options for money market funds. President’s Working Group on Financial Markets, US Department of the Treasury.

PWG, 2021. Report on stablecoins. President’s Working Group on Financial Markets, US Department of the Treasury.

Rosengren, E.S., 2021. Financial stability. Remarks at Official Monetary and Financial Institutions Forum (Fed Week).

Uhlig, H., 2022. A luna-tic stablecoin crash. NBER working paper No. 30256.

A Appendix: Stability Mechanisms

In this section, we describe the stablecoins that comprise our data and explain the mechanisms through which notable stablecoins maintain their peg. Because of the permissionless nature of blockchain technology, there have been many different attempts to create stablecoins, with varying degrees of success. Our data set contains stablecoin data on 12 stablecoins, which we describe roughly in order of highest to lowest average market capitalization in early 2022, before the Terra crash. Taken together, BUSD, USDC, USDP, USDT, TUSD, FRAX, USTC, DAI, LUSD, and MIM comprised more than 99 percent of the total market capitalization in early 2022; BUSD, USDC, USDP, GUSD, USDT, TUSD, USDD, FRAX, DAI, and LUSD²⁹ comprised more than 99 percent of the total market capitalization in early 2023.

A.1 Tether (USDT)


Tether was launched in 2014 by a Hong Kong–based company called iFinex, which is also the owner of Bitfinex, a large crypto exchange. At launch, Tether promised that each unit of stablecoin issued would be backed by \$1, held in a bank account. Tether has revised the specifics of this promise several times starting in 2019; Figure A.1 shows that Tether claimed it was backed by traditional currency in February 2019 but by April 2019 changed its disclosure to state that Tether is backed by a mixture of assets.

Table A.1 shows Tether’s assets according to its May 2022 attestation.³⁰ In March 2022, Tether had \$82.26 billion in liabilities and \$82.42 billion worth of assets. The assets were \$20.1 billion of commercial paper, \$39.2 billion of Treasury securities, \$6.8 billion of money market funds shares, \$3.1 billion of secured loans, \$3.7 billion worth of corporate bonds, mutual funds and precious metals, and \$4.9 billion worth of other investments, including other crypto assets.

²⁹These are the same 10 stablecoins as in the 2022 analysis, except USTC and MIM were overtaken by GUSD and USDD when ranking stablecoins by market capitalization.

³⁰https://assets.ctfassets.net/vyse88cgwfb1/1np5dpcwuHrWJ4AgUgI3Vn/07fcaeb1cd7ce6df71ce8f5abb09ddb7/Tether_Assurance_Consolidated_Reserves_Report_2022-03-31__1_.pdf

Figure A.1: Comparison of Tether disclosures from February 2019 and April 2019.



Source: Internet Archive snapshots of Tether homepage (February; April).

Note: The disclosure from February 2019 asserts that every Tether is backed by traditional currency, while the disclosure from April 2019 states that Tether is backed by a mixture of assets.

On many occasions, Tether has traded below its one-to-one peg with the dollar, but it has always recovered it. Nevertheless, because it is backed by risky assets and lacks transparency, Tether is perceived as having a large degree of run risk.

A.2 USD Coin (USDC)

USD Coin (USDC) was launched in 2018 by Circle, a company based in Massachusetts. It is operated by a consortium of companies called Centre, which includes Circle and Coinbase. Like Tether, USDC maintains its peg by backing its coin with financial assets; but unlike Tether, it invests mainly in safe government assets. For instance, according to USDC's attestation reports,³¹ in July 2022, there were 54.6 billion USDC units in circulation, backed by \$54.7 billion of assets, 22.4 percent held in cash and 77.6 percent held in short-term Treasury securities. USDC has money transmitter licenses in multiple US states and the District of Columbia, and regularly publishes attestations of their assets.

³¹<https://www.circle.com/en/transparency>

Table A.1: Tether Balance Sheet according to its May 2022 attestation.

Asset Type	Billions of Dollars	Share of Assets (%)
Commercial Paper	20.1	24.4
US Treasuries	39.2	47.6
Money Market Funds	6.8	8.3
Secured Loans	3.1	3.8
Corporate Bonds, Funds and Precious Metals	3.7	4.5
Other Investments, Including Cryptocurrencies	4.9	5.9
Total Assets	82.42	
Total Liabilities	82.26	

Of course, even though USDC is backed by cash and short-term Treasury securities, it is still runnable. In particular, its liabilities are not eligible for FDIC insurance; therefore, a run on USDC does not necessarily have the usual dampeners as an analogous run on a US bank. This proved to be the case in March 2023, when investors ran on USDC after observing that 8 percent of the stablecoin’s assets were at risk in the Silicon Valley Bank collapse. Within hours, USDC’s price crashed to \$0.88, and its market capitalization declined by \$7.9 billion, 18 percent of its market capitalization at the beginning of the day.³² Only after regulators extended the \$250,000 cap on FDIC deposit insurance did market participants regain enough confidence in USDC for the stablecoin to regain its peg.

A.3 Binance USD (BUSD)

The third-largest stablecoin in 2022 was Binance USD (BUSD), issued by a company called Paxos, which is based in New York State and therefore regulated by the New York Department of Financial Services (NYDFS). As of February 21, 2023, Paxos stopped issuing BUSD at NYDFS’s direction. According to the Paxos website,³³ BUSD can still be redeemed for USD or exchanged for USDP,

³²This price reflects secondary market transactions. As explained earlier, in addition to secondary market transactions, some investors may redeem or purchase tokens directly with the SC issuer (that is, primary market transactions.) We do not know if any UDSC tokens were redeemed for less than \$1 in the primary market during this time.

³³<https://paxos.com/2023/02/13/paxos-will-halt-minting-new-busd-tokens/>

another stablecoin issued by Paxos, and Paxos will fully collateralize BUSD with safe assets such as cash, US Treasury bills, and overnight reverse repos indefinitely.

A.4 TerraUSD (USTC)

TerraUSD (USTC) is an algorithmic stablecoin and was the fourth-largest stablecoin up to May 2022, when it suffered a run. In contrast to asset-backed or crypto-collateralized stablecoins, algorithmic stablecoins are not backed by a financial asset. Rather, their peg is maintained through an algorithmic mechanism.

In the Terra stablecoin mechanism, there are two tokens: Terra, which is designed to be stable, and Luna, a crypto asset similar to Bitcoin whose value fluctuates over time but which, in contrast to Bitcoin, can be minted in an unlimited amount. Any investor with a node in the Terra blockchain has access to a smart contract, which allows them to create or redeem one unit of Terra for \$1 worth of Luna; for instance, if the price of Luna is \$10, the smart contract will exchange one Terra for 0.1 unit of Luna. Arbitrage should keep the value of Terra stable at \$1 as long as the price of Luna is greater than \$0.

There are several reasons why such a system may be unstable, including:

1. **Limits to Arbitrage:** There are costs to running a node on the Terra blockchain, and not every investor runs a node.
2. **Positive Feedback Loop:** As the price of Luna decreases, any redemption of Terra will increase the supply of Luna by larger and larger amounts. For instance, if Luna trades at \$0.01, then redeeming one unit of Terra will create 100 units of Luna. As Terra crashed, the supply of LUNA increased from 365 million units on May 9 to more than 6 trillion units by May 13.
3. **Multiple Equilibria:** The mechanism for stabilizing Terra's price has two equilibria: If investors believe Luna will trade for a positive price, then Terra will trade for \$1; but if investors expect the price of Luna to be \$0, then the price of Terra will also be \$0.

Indeed, the limits of the Terra algorithm became apparent in May 2022, when the price of both Terra and Luna crashed close to \$0.

A.5 Dai (DAI)

DAI is issued by MakerDAO, a so-called distributed autonomous organization (DAO) established in 2014. MakerDAO launched two tokens, DAI and Maker (MKR), in 2017 on the Ethereum blockchain. The MKR token is a governance token, with MKR holders able to vote on the parameters that determine the behavior of the DAI stablecoin; investors can buy and sell MKR in an exchange. That is, MakerDAO is similar to a company with no board of directors whose decisions are made by equity holders (that is, the holders of MKR).

The DAI token is a crypto-collateralized stablecoin. New units of DAI are minted and redeemed through smart contracts called vaults. A user can mint new DAI by creating a vault and depositing collateral in it. There are several types of vaults, each specifying the minimum collateral needed to issue a unit of DAI. Each type of vault is established through a vote by MKR owners. The collateral can be Ethereum, Bitcoin, other crypto assets including stablecoins, and increasingly “real world assets” such as tokenized mortgages.

For instance, a user who creates a vault of type ETH-A would deposit at least \$145 worth of Ethereum and generate 100 units of DAI. If the peg holds so that 1 DAI is worth \$1 on secondary markets, then this would generate \$100 worth of DAI. The user can use the DAI they minted to invest in crypto assets or DF protocols on the Ethereum blockchain. Note that although DAI can be created through different vault type, once created, they are all fungible and can be traded freely on secondary markets.

A user who mints DAI has effectively borrowed it from the system through a collateralized loan. To recover their collateral, the user needs to repay the DAI back to the vault smart contract, together with an interest rate, also specified for each vault type by the holders of MKR. As we will see below, the proceeds from this interest rate contribute to a stability buffer.

When DAI is backed by volatile assets, it is usually overcollateralized. If the value of the collateral drops below the minimum collateral specified by the vault type (\$145 worth of Ethereum for \$100 worth of DAI in the previous example), then the collateral can be liquidated via an auction. Any holder of DAI can start the auction and would get a fee to do that. The winner of the auction would receive the collateral in exchange for their DAI holdings. The user who deposited the collateral would need to deposit more Ethereum to avoid the liquidation. As long as the collateral value in each vault is above 100 percent, then the peg of DAI will be maintained because any DAI holder is better off participating in the auctions than selling DAI at a price below \$1.

Dai's smart contract has the following properties:

- **Issuance:** Any user can send collateral (in the form of Bitcoin, Ethereum, or other allowed assets) and receive newly minted units of DAI at the price of \$1. The collateral gets cryptographically locked by the smart contract into a vault, which is associated with the user who minted the DAI.
- **Use:** A user who has minted DAI can proceed to use it to buy other crypto assets or loan the DAI in a decentralized finance protocol and earn higher yield. A popular use case is using DAI to take a leveraged position on ETH. For instance, a user may use \$150 worth of ETH collateral to mint \$100 worth of DAI. Then, the user proceeds to buy \$100 worth of ETH with their newly minted DAI. This gives the user a \$250 position on ETH using an initial capital of \$150, amplifying returns by a factor of $\frac{5}{3}$. Users can repeat this cycle to amplify their total returns by a factor of 3.
- **Redemption:** Conversely, any user who had minted DAI can redeem it by returning their units and receiving back the collateral minus a redemption fee. It is important to note that only the user who minted the DAI associated with a particular vault can redeem the collateral in that vault.
- **Liquidation:** If the value of the collateral in a vault drops below a certain threshold, then any user can trigger an auction. Any DAI holder can bid in the auction for the collateral. All

bids are denominated in DAI; the winning bidder obtains the assets backing DAI, and the equivalent amount of stablecoins is destroyed.

Note that the mechanism outlined above allows an arbitrageur to profit if DAI trades outside the $[\frac{1}{1.45}, 1.45]$ interval. If DAI traded above \$1.45, then a user could buy \$1.45 worth of Ethereum, use it as collateral to generate one unit of DAI, and sell the DAI for more than \$1.45, obtaining a riskless profit. If DAI traded below $\frac{1}{1.45}$, then the arbitrageur could short one unit of DAI, buy $\frac{1}{1.45}$ worth of Ethereum, and use it to generate one unit of DAI to repay their short. While this does not guarantee a very tight peg, there are alternative mechanisms that keep the peg in a much more narrow band.

- **Savings Rate:** DAI provides a DAI savings rate—essentially a deposit rate—which can be raised when the price of DAI is low and lowered when the price of DAI is high.
- **Stability Fees:** Users who mint DAI will continuously pay a stability fee—essentially a borrowing rate—until they redeem their DAI. When the stability fee is high, it incentivizes users to redeem their DAI and take stablecoins out of circulation, increasing the price of DAI. Conversely, lower stability fees encourage the creation of new DAI, reducing the price.
- **Multiple Types of Collateral:** In addition to the earlier Ethereum example, there are other types of collateral with different collateralization ratios. Many of these other types of collateral are asset-backed stablecoins, with collateralization ratios around 100 percent, allowing for a tighter peg. We observe that this means that DAI's peg is linked to the stability of asset-backed stablecoins.
- Even with all these mechanisms, DAI also has a peg stability module, which allows users to exchange DAI for USD Coin at a 1:1 ratio—without locking any collateral. This is in addition to using other asset-backed stablecoins like Paxos as collateral. Thus, the value of DAI is intrinsically linked to the value of USD Coin and other US-based stablecoins.

A.6 Magic Internet Money (MIM)

Magic Internet Money (MIM) is a crypto-collateralized stablecoin that uses interest-bearing crypto-asset derivatives as collateral. It was launched in 2021 and operates on the Ethereum blockchain. Like DAI, MIM is crypto-collateralized: Any user can mint new MIM by depositing collateral in a smart contract called a cauldron and can redeem MIM by returning it to the smart contract and reclaiming the collateral. The parameters governing the behavior of the cauldrons are determined using a governance token called SPELL, which is analogous to MakerDAO's MKR.

In contrast to DAI, where most of the crypto-collateral used is a standard crypto asset like ETH, most of the crypto-collateral in the MIM protocol is in the form of interest-yielding tokens issued by other decentralized finance protocols. In this way, MIM adds another layer of complexity and potential instability to the decentralized finance ecosystem.

A.7 TrueUSD (TUSD)

TrueUSD is an asset-backed stablecoin that was originally issued by TrustToken, a company based in San Francisco, and whose reserves were stored in the United States.³⁴ In 2020, the TUSD brand was sold to Techteryx, an Asian conglomerate based in China³⁵. After this transfer, the collateral was held at a variety of banks, including in the United States, the Bahamas, and Hong Kong, making TUSD transition from US-based to Offshore.³⁶

A.8 Frax (FRAX)

Frax was launched in 2020 and follows a hybrid design. In its original incarnation—Version 1—Frax was partially collateralized by USDC and partially collateralized by Frax Shares (FXS), a

³⁴In an attestation from 2018 (<https://trusttokenteam.medium.com/nov30-b4261325c468>), the issuing company attests that the collateral funds are held by two escrow agents in the United States.

³⁵<https://trueusd.medium.com/scaling-trust-announcing-tusds-next-stage-of-growth-f1fb58d62b>

³⁶<https://trueusd.medium.com/trueusd-attestations-49092b7cb500>

free-floating token like LUNA. Version 2 is collateralized by a wider variety of assets, including tokens representing ownership shares of decentralized exchange contracts where FRAX is traded.

Frax Version 1. The key state variable in FRAX Version 1 is the collateral ratio $\rho \in [0, 1]$. Given a collateral ratio of ρ , a unit of FRAX can be minted by depositing ρ units of USDC and $(1 - \rho)$ dollars worth of FXS into a smart contract. In the other direction, a user can redeem a unit of FRAX via this smart contract and receive ρ units of USDC and $(1 - \rho)$ units of FXS. The collateral ratio can be increased or decreased if certain conditions are met. For example, if FRAX is trading above \$1, then the collateral ratio can be decreased so that less USDC collateral is needed to mint one unit of FRAX. If FRAX is trading below \$1, then the collateral ratio can be increased. This makes it more difficult to mint new units of FRAX and increases the incentive to redeem existing units (by providing more USDC when FRAX units are redeemed). These collateral changes can be triggered by any user who calls the corresponding functions in the FRAX smart contract. However, these functions can only be called if the corresponding price conditions (FRAX above \$1 or FRAX below \$1) are met.

FRAX Version 2. The main distinction between FRAX Version 1 and Version 2 is that Version 2 relies on a wide array of crypto assets as the backing collateral. However, there are other important distinctions that increase its interconnections with the rest of the crypto ecosystem:

1. Since multiple tokens can be used as collateral, it is possible that the value of the collateral is above or below 100 percent. If the value of collateral is above 100 percent, there is a function in the smart contract—called FXS 1559—allowing some FXS units to be redeemed for collateral.
2. The protocol rehypothecates some of the USDC collateral by investing it in decentralized finance protocols such as Aave, Compound, and Yearn.
3. The protocol also rehypothecates USDC collateral by placing it in the Curve or Uniswap decentralized exchanges.

4. New FRAX can be minted by borrowing it using collateral, in a similar way that DAI is minted.

A.9 Pax Dollar (USDP)

Paxos issues another stablecoin, Pax Dollar (USDP), which was the ninth-largest stablecoin in April 2022. Like BUSD, USDP is backed by cash, US Treasury bills, and overnight reverse repos.

A.10 Liquity USD (LUSD)

Liquity USD (LUSD) is another crypto-collateralized stablecoin that operates on the Ethereum platform. In contrast with DAI and MIM, LUSD smart contracts only accept Ethereum as collateral and do not charge an interest rate. Instead, there is a one-time fee at the time of borrowing. The collateral ratio needed to generate LUSD is 110 percent.

A.11 Gemini Dollar (GUSD)

Like BUSD, Gemini Dollar (GUSD) is an exchange-branded token that is issued and custodied by Paxos. Like BUSD and USDP, the funds used to back GUSD are custodied in US financial institutions.

A.12 Decentralized USD (USDD)

USDD is an algorithmic stablecoin backed by several tokens, including Tron (TRX), USDT, and USDC. USDD trades on the TRON blockchain and was introduced in May 2022. Like FRAX, USDD is a hybrid stablecoin. It has an algorithmic mechanism that allows users to exchange one unit of USDD for \$1 worth of TRX at any time. In addition, it is backed by a peg stability module (PSM) holding reserves of USDT and USDC, allowing users to exchange one unit of USDD for one unit of USDT or USDC.

USDD has broken its peg multiple times, including in May and June 2022 and March 2023, but has recovered its peg since then.

B Appendix: Regression Tables

B.1 Local projections: regression results and robustness checks

Table B.2 reports the coefficient estimates corresponding to the plots in Figure 6. The bottom six rows of Table B.2 report the p -values from F tests for the difference between the different coefficients. For instance, in the “ $h=1$ ” column, the p -value for “US \neq Algo” is 0.032, indicating that the difference in cumulated flows between US-based stablecoins and algorithmic stablecoins one day after a Bitcoin price shock is significant at the 5 percent level. These test results, taken together with the divergent sign of the point estimates (positive flows for US-based stablecoins and negative for the rest), indicate that the flows to U.S.-based stablecoins are of a different nature.

In addition to our main specification of Table B.2, we run a number of robustness checks. The first one concerns the number of lags, which are included to control for serial correlation in both the dependent and independent variables. To choose the optimal number of lags, we follow a procedure based on the discussion in Jordà (2005) and Montiel Olea and Plagborg-Møller (2021). We run a vector autoregression (VAR) for the $[\text{Flow}_{i,t}, \text{Shock}_{i,t}]$ vector for each stablecoin i . We estimate the VAR with with $p = \{1, 2, 3, \dots, 31\}$ lags (to allow for monthly seasonality) and choose the optimal number of lags based on an information criterion (IC). We use the Akaike IC, or AIC, and the Bayesian IC, or BIC, which in practice agree for every stablecoin. For nine out of twelve stablecoins, the IC indicates two lags as the optimal number. We choose this as our main specification under the parsimonious assumption that the optimal lag length for all stablecoins should be the same.

For the other three stablecoins, the IC points to much longer lags (14, 30, and 30). Furthermore, Montiel Olea and Plagborg-Møller (2021) suggest that the VAR lag length p should be chosen “conservatively ... there is no asymptotic efficiency cost of controlling for more than p_0 lags if the

Table B.2: Regression results: local projections of cumulated flow on Bitcoin shock (value-weighted)

		Dependent Variable: Cumulative Percent Outflows h Days After Shock								
		$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
US-Based		-0.04 (-0.46)	0.10 (0.60)	0.49 (1.59)	0.71* (2.07)	1.03** (2.80)	1.29*** (3.21)	1.45*** (3.87)	1.67*** (3.63)	1.88*** (3.57)
Offshore		-0.03 (-0.85)	-0.13 (-0.54)	-0.27 (-0.91)	-0.54 (-1.05)	-0.64 (-1.23)	-0.82 (-1.37)	-0.92 (-1.73)	-1.04 (-1.65)	-1.07 (-1.54)
Algorithmic		-1.76* (-1.95)	-3.18** (-2.57)	-5.68** (-2.99)	-6.02** (-2.83)	-5.60** (-2.54)	-6.53** (-2.88)	-8.58*** (-3.42)	-8.54*** (-3.12)	-9.18*** (-3.20)
Crypto-Collateralized		-2.02*** (-3.33)	-2.52*** (-3.20)	-3.13** (-2.69)	-2.70* (-2.16)	-3.00** (-2.37)	-2.78* (-2.08)	-2.90* (-1.81)	-3.27* (-1.89)	-3.02* (-1.80)
Coin FE		Y	Y	Y	Y	Y	Y	Y	Y	Y
N. Obs.		8,676	8,676	8,676	8,676	8,676	8,676	8,676	8,676	8,676
Adj. R^2		0.04	0.06	0.08	0.09	0.09	0.09	0.09	0.09	0.09
p : US \neq Offshore		0.940	0.524	0.162	0.116	0.067	0.050	0.018	0.022	0.029
p : US \neq Algo		0.098	0.032	0.013	0.016	0.024	0.012	0.004	0.008	0.007
p : US \neq Crypto		0.007	0.009	0.019	0.042	0.027	0.032	0.042	0.039	0.043
p : Offshore \neq Algo		0.080	0.013	0.009	0.008	0.017	0.006	0.003	0.005	0.004
p : Offshore \neq Crypto		0.007	0.005	0.024	0.051	0.039	0.040	0.112	0.086	0.094
p : Algo \neq Crypto		0.763	0.403	0.019	0.007	0.043	0.004	0.000	0.000	0.000

Note: This table reports selected estimated coefficients from Equation (1) representing the cumulated flow impulse response function for US-based (β_0^1), offshore asset-backed (β_0^2), crypto-backed (β_0^3), and algorithmic (β_0^4) stablecoins. Observations are weighted by pre-event market capitalization. All regressions have coin fixed effects. Standard errors are clustered at the coin and date level. T-statistics are in parentheses. Stars indicate statistical significance at the 10% (*), 5% (**), and 1% (***) levels.

true model is a VAR(p_0).” Montiel Olea and Plagborg-Møller (2021) also suggest using simple robust standard errors instead of clustered standard errors. Thus, in Table B.3, we report the results of a specification analogous to Equation (1), but including 31 daily lags and using robust standard errors. The point estimates are essentially identical and more statistically significant.

In unreported results, we estimate every combination of robust and clustered standard errors with 1, 2, 3, 7, 14, and 31 daily lags. None of the point estimates are materially different from the reported ones, and the statistical significance remains for almost all coefficients in almost all specifications.

Table B.3: Regression results: local projections of cumulated flow on Bitcoin shock (value-weighted)

		Dependent Variable: Cumulative Percent Outflows h Days After Shock								
		$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
US-Based		-0.05 (-0.25)	0.06 (0.22)	0.39 (1.36)	0.58* (1.72)	0.86** (2.08)	1.02** (2.07)	1.19** (2.10)	1.36** (2.23)	1.52** (2.28)
Offshore		-0.05 (-0.77)	-0.16 (-0.85)	-0.34 (-1.54)	-0.71** (-2.04)	-0.85** (-2.07)	-1.09** (-2.25)	-1.24** (-2.42)	-1.42** (-2.57)	-1.50** (-2.45)
Algorithmic		-2.64* (-1.84)	-4.12** (-2.43)	-6.75** (-2.53)	-6.91** (-2.33)	-6.71** (-2.15)	-7.56** (-2.37)	-9.85*** (-2.77)	-9.82** (-2.52)	-10.45*** (-2.59)
Crypto-Collateralized		-2.06*** (-4.60)	-2.56*** (-4.38)	-3.31*** (-4.39)	-3.01*** (-3.13)	-3.46*** (-3.33)	-3.40*** (-3.16)	-3.75*** (-3.30)	-4.24*** (-3.65)	-4.08*** (-3.46)
Coin FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
N. Obs.	8,426	8,426	8,426	8,426	8,426	8,426	8,426	8,426	8,426	8,426
Adj. R^2	0.12	0.14	0.16	0.17	0.17	0.17	0.18	0.19	0.19	0.19
p : US \neq Offshore	0.990	0.486	0.044	0.008	0.003	0.003	0.002	0.001	0.001	0.001
p : US \neq Algo	0.073	0.015	0.008	0.012	0.016	0.016	0.008	0.002	0.005	0.003
p : US \neq Crypto	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
p : Offshore \neq Algo	0.070	0.020	0.017	0.038	0.063	0.063	0.045	0.017	0.033	0.028
p : Offshore \neq Crypto	0.000	0.000	0.000	0.024	0.019	0.019	0.050	0.045	0.028	0.052
p : Algo \neq Crypto	0.697	0.385	0.215	0.210	0.322	0.322	0.217	0.103	0.170	0.129

Note: This table reports selected estimated coefficients from Equation (1) representing the cumulated flow impulse response function for US-based (β_0^1), offshore asset-backed (β_0^2), crypto-backed (β_0^3), and algorithmic (β_0^4) stablecoins. The specification is modified to include 31 lags for the dependent and independent variables. Observations are weighted by pre-event market capitalization. All regressions have coin fixed effects. Standard errors are robust. T-statistics are in parentheses. Stars indicate statistical significance at the 10% (*), 5% (**), and 1% (***) levels.