Public information and stablecoin runs

Rashad Ahmed
OCC

Iñaki Aldasoro
BIS

Chanelle Duley
University of Auckland

August 27, 2025

Abstract

We show that the informational quality of stablecoin (SC) issuers' reserve assets affects the type of coordination game being played among SC holders and its equilibria. When the volatility of reserve assets is unknown, par convertibility is resilient to small shocks but fails with large negative public shocks to reserve asset values, even if they are initially high. Public information disclosure increases (reduces) run risk for sufficiently low (high) holders' priors about reserve quality. Transparency and quality of reserve assets have distinct effects on issuer failure risk. Our results point to a trade-off between peg stability and issuer fragility.

JEL Codes: C70, D83, D84, E42, G01, G20.

Keywords: Stablecoins, crypto, global games, bank runs.

We thank Toni Ahnert, Jeff Allen, Kartik Anand, Christoph Bertsch, Lin William Cong, Giulio Cornelli, Sebastien Derivaux, Sebastian Doerr, Prasanna Gai, Priscilla Koo Wilkens, Mattia Landoni, Kebin Ma, William Mann, Cyril Monnet, Nico Oefele, Alexandros Vardoulakis, and conference/seminar participants at the third Bonn-Mannheim Workshop on Digital Finance, the Economics of Payments XII Conference, Bank of Canada/Riksbank 2nd Conference on the Economics of CBDC, IBEFA Summer Meeting 2024, CEAR-Atlanta Fed 2024 Financial Services Industry Conference, Federal Reserve Board, Office of the Comptroller of the Currency, Bank for International Settlements, Norges Bank Workshop "The role of central banks in stablecoin regulation", HIAS workshop "Money and payments: evolution or revolution?", Deutsche Bundesbank, Reserve Bank of New Zealand, De Nederlandsche Bank, University of Auckland and University of Cape Town for helpful comments and suggestions, and Jose Maria Vidal Pastor for help with data. Rashad Ahmed (rashad.ahmed@occ.treas.gov) is at the U.S. Department of the Treasury, Office of the Comptroller of the Currency, 400 7th St SW, Washington, DC. Iñaki Aldasoro (inaki.aldasoro@bis.org) is at the Bank for International Settlements, Monetary and Economic Department, Centralbahnplatz 2, CH-4002 Basel, Switzerland (*corresponding author*). Chanelle Duley (chanelle.duley@auckland.ac.nz) is at the University of Auckland, Economics Department, 12 Grafton Road, 1010 Auckland, New Zealand. The views expressed here are those of the authors only, and not necessarily those of the Bank for International Settlements, the Office of the Comptroller of the Currency, or the U.S. Department of the Treasury.

1 Introduction

Stablecoins (SCs) are crypto tokens that live on distributed ledgers and promise to be always worth a dollar.[1] Variously likened to banks (Gorton and Zhang, 2023), exchange rate pegs (Levy Yeyati and Katz, 2022) and a combination of money market and exchange traded funds (Anadu et al., 2024; Ma et al., 2023; Oefele et al., 2023), the key defining feature of stablecoins is their promise to deliver par convertibility to the sovereign unit of account (Aldasoro et al., 2023). To make that promise credible, stablecoin issuers hold a variety of reserve (collateral) assets, including fiat-denominated money market instruments, Treasuries, bank deposits and other cryptoassets (including other stablecoins).

Public information and perceptions regarding the quality, transparency and volatility of reserves are thus key for stablecoin peg stability. This was evident during the March 2023 US banking turmoil. As stress in Silicon Valley Bank (SVB) mounted, the stablecoin USD Coin (USDC, known to have a yet-undisclosed amount of deposits at SVB) broke par (see Figure 1, red line). The situation deteriorated when Circle, the issuer of USDC, disclosed it held $3.3 billion of its cash reserve at SVB (first vertical dashed line). This episode illustrates how a negative shock arising from increased *transparency* of reserves led to a shift in the aggregate behavior of coin holders. USDC had served as a stable reserve asset for other stablecoins such as Dai, which shortly thereafter lost its peg as well (solid black line): a negative shock to the perceived *quality* of reserves that did not affect other (dollar-backed) stablecoins.[2] The situation only improved when the US government announced a backstop for SVB (third vertical dashed line).

This paper analyzes how stablecoins' peg stability is affected by the perceived quality and volatility of reserve assets, as well as reserve disclosure and transparency. In an environment where the volatility of the portfolio of reserve assets is unknown, the dollar value of these assets is characterized by a fat-tailed distribution. Par convertibility is maintained in the

[1] Stablecoins arose from the need for a safe and stable unit of account within the volatile crypto ecosystem. They saw a meteoric rise from $5 billion market capitalization in early January 2020 to $190 billion in early May 2022, followed by a steady decline over a couple of years. More recently they resumed growth, standing at a record $203 billion as of December 2024. They established themselves as the medium of exchange of, and gateway into, crypto. We focus on stablecoins pegged to the US dollar, as they make the lion's share of the market.

[2] Crypto-backed stablecoins such as Dai or Frax provide constant public visibility of the composition of reserves through on-chain mechanisms. Others like Circle do not continually disclose reserve assets and instead publish regular reports and attestations. We refer to *transparency* as relating to the frequency and credibility of this kind of public information, which contrasts with reserve *volatility* – large-scale fluctuations in the dollar value of reserves that may or may not be made transparent to the public. We use Dai for illustrative purposes in Figure 1 due to better quality of data at the minute frequency. For empirical analysis we will focus instead on another USDC-backed stablecoin (Frax), given several characteristics of Dai make it unfit for purpose.

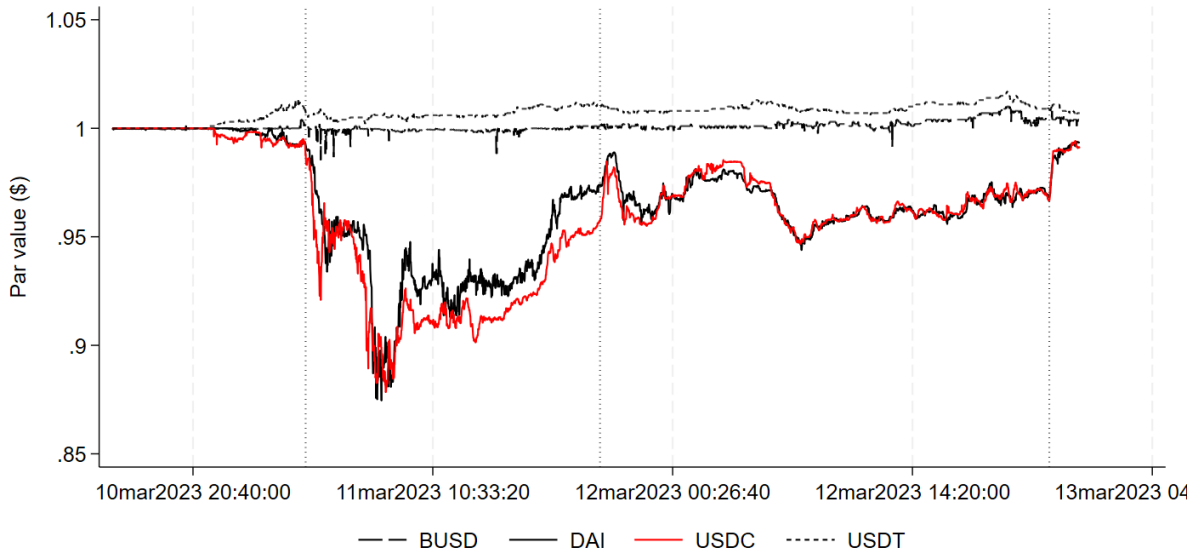

Figure 1: Stablecoin pegs around the run on Silicon Valley Bank.

Notes: Based on minute-by-minute data. The first vertical dashed line denotes Circle's disclosure that $3.3 billion of its cash reserve was held at SVB; the second line denotes the announcement by Circle that normal liquidity operations would resume by Monday 13 March; and the third line denotes the announcement of a backstop by the US government. Figure 15 in Appendix C zooms in on the event to highlight more clearly the sequencing (USDC first, then Dai). Figure 16 shows this was a common feature of USDC-backed stablecoins, whereas those backed by other cryptoassets remained stable. Source: Cryptocompare.com.

face of small shocks (a *resilience* effect) but collapses under large negative public shocks, even for high initial collateral values (a *change* effect). In turn, increased transparency has a nuanced effect on run risk. When stablecoin holders' priors about reserve quality are sufficiently weak (which arise if the peg was recently broken), greater transparency amplifies run risk. Conversely, when priors about reserve quality are strong (which arises in the absence of severe pressure on the peg in previous periods), transparency can reduce run risk. An implication is that a long history of stability with frequent disclosures is overall conducive to peg stability. Drawing on several case studies and using a synthetic control approach to address endogeneity concerns, we find empirical support for the implications of the model.

We model a run on a stablecoin issuer as a global game of regime change (Morris and Shin, 2003).[3] This class of models is well suited to the study of stablecoins because, by design,

[3]Global games have been extensively studied in applications ranging from bank runs (Goldstein and Pauzner, 2005; Rochet and Vives, 2004) and currency crises (Morris and Shin, 1998; Angeletos et al., 2006), to political protests (Edmond, 2013; Little et al., 2015; Chen and Suen, 2016), the emergence of tax havens (Konrad and

they operate as a unilateral exchange rate peg to a reference asset (usually fiat money). The breaking of par convertibility can be avoided when a sufficiently small proportion of stablecoin holders demand conversion. We extend the global games approach to study an environment characterized by model uncertainty, reflected in unknown structural parameters governing the distribution of reserve assets. This transforms the coordination game into a generalized, *second-generation* global game that features multiple, locally-unique equilibria (Morris and Yildiz, 2019). Our approach allows us to disentangle the distinct effects that transparency (i.e., public information about reserve assets) and reserve quality (i.e., variance of reserve assets around their reference value) have on run risk.

The model extends across two dates. At the initial date, the dollar value of a stablecoin issuer's collateral is realized and stablecoin holders observe noisy signals about it. At the final date, stablecoin holders decide whether to convert their SC back to fiat money, or to maintain their holding. The issuer observes the aggregate size of redemption requests and is not able to defend par whenever the liquidation value of reserve assets is smaller than the value demanded by stablecoin holders. We assume that owing to positive but small transaction costs, stablecoin holders prefer to maintain a holding of coins (e.g. because holding yields benefits such as the option value of remaining in the crypto universe using stablecoins as collateral for investment/speculation purposes). But since the issuer may be unable to defend par and collapse altogether, each stablecoin holder may prefer to dash for cash to avoid losing the value of their claims. The strategic interactions among stablecoin holders depend on their *rank beliefs* – the probability each one assigns to the event that others have more pessimistic signals about the issuer's ability to meet conversion requests. As a result, the degree of flight risk to which the issuer is exposed depends on the nature of stablecoin holder benefits and the type of collateral held by the issuer.

Stablecoin issuers vary in the type of reserves they hold: fiat-denominated instruments, commodities and other crypto assets, including other stablecoins.[4] We present an environment that can accommodate this variety, where the volatility of reserve assets is not known

Stolper, 2016), and financial crises (Bebchuk and Goldstein, 2011). Recent work on stablecoins, discussed in more detail below, also employs global game techniques (e.g. d'Avernas et al., 2023; Bertsch, 2023; Ma et al., 2023; Gorton et al., 2022; Li and Mayer, 2021).

[4]In the extreme, algorithmic stablecoins employ policies that increase or decrease the supply of the stablecoin, often using other highly volatile crypto-assets – effectively attempting to guarantee stability by shifting volatility to a paired token. To be sure, the backing of algorithmic coins cannot be likened to that of, say, a stablecoin backed by U.S. Treasury securities. However, the market capitalization of paired tokens (think of Luna in the case of TerraUSD) operates effectively as a proxy for reserve assets from the perspective of stablecoin holders (Liu et al., 2023), and as such it can act as a catalyzer of shifts in stablecoin holders' beliefs.

and agents perceive it to be fat-tailed, i.e. there is model uncertainty (Morris and Yildiz, 2019). Fat tails capture the relatively high probability that stablecoin holders assign to extreme realizations of reserve asset values, generating uncertainty about the issuer's ability to honor redemption requests. We show how, as the perceived volatility of reserve assets becomes small (i.e. variance uncertainty declines), the model environment converges to global game models of regime change in the tradition of Morris and Shin (1998) where the perceived distribution of reserve assets is thin-tailed.

Heightened uncertainty about fundamentals undoes the uniqueness that characterizes traditional global games, leading to local multiplicity when the value of reserve assets enters an intermediate range. We embed the model in a dynamic global game and appeal to hysteresis equilibrium (Morris and Yildiz, 2019; Romero, 2015) to form predictions about the effect of such tail risk on the probability of a stablecoin run. Importantly, whereas the thin-tailed global game predicts runs whenever collateral value drops below a threshold *level*, the fat-tailed global game requires both a fall in collateral value below a threshold and a sufficiently large *change* relative to previous periods to make running a uniquely rationalizable response by stablecoin holders.

We then examine how run risk is related to the degree of transparency of the issuer's collateral portfolio. Increasing the degree of transparency of reserves – through, for example, issuing public broadcasts or portfolio audits on a regular basis – provides stablecoin holders with more *common knowledge* about the variance of reserve assets. We find that the effect of public disclosure on run risk is ambiguous. Greater transparency can lead to increased run risk whenever stablecoin holders have sufficiently low priors about reserve assets to begin with. Conversely, heightened transparency lowers run risk when priors are strong. Our results highlight the nuanced but important role that public information has on stablecoin runs.

Our model generates a set of testable implications. Whenever priors about reserve assets are not weak, pressure on the peg should decrease following public disclosures that provide information about their quality. Conversely, the broadcasting of public information on reserves when there are doubts about the issuer's reserve adequacy should lead to greater peg instability. More generally, when stablecoin reserves' value enters a "ripe for run" region, holders pay close attention to past events and peg stability depends on shock size: the par promise is resilient to small negative shocks to collateral value (a *resilience* effect), but it

breaks down in the face of large negative public shocks, even when initial collateral values are relatively high (a *change* effect).

For most analyses, we estimate the causal effect of information disclosure events on stablecoin peg stability. In particular, we assess the effects of changes in collateral quality (real or perceived) based on public information disclosures on stablecoins' average daily absolute price deviations from $1 (Lyons and Viswanath-Natraj, 2023). To address endogeneity concerns we employ regression methods inspired by the synthetic control literature to construct counterfactual time series of stablecoin peg deviations that we compare against realized peg deviations before and after each event (Abadie and Gardeazabal, 2003; Doudchenko and Imbens, 2016). Using the pre-event sample, we first estimate dynamic regressions of the treated stablecoin's peg deviations on a set of control units: measures of cryptoasset and (when feasible) financial market volatility along with the peg deviations of other non-treated stablecoins.[5] The sequence of post-event counterfactual peg deviations are generated by applying the parameter estimates to the post-event observations of the control units.

The first case study analyzes the sudden de-peg suffered by USDC during the turmoil at Silicon Valley Bank in March 2023 (Figure 1). We interpret the disclosure of Circle's $3.3 billion of reserves deposited with SVB as an exogenous public information shock regarding the stablecoin's reserve adequacy, since priors were weak in light of the knowledge that Circle banked with SVB. Our model predicts a deterioration in USDC's peg stability following the disclosure. And indeed we find strong evidence of significantly larger absolute peg deviations compared to counterfactual peg deviations using minute-by-minute data.

Our second set of analyses also look at transparency and disclosures, but focuses on Tether, the largest stablecoin. In particular, we conduct two exercises. First, we exploit a series of publicized concerns about Tether's reserve adequacy in 2018, signaling a reduction in the perceived quality of its reserves. Tether temporarily lost its peg due to rising concerns that it was not 100% backed by US dollar assets (contrary to repeated, yet unsubstantiated, claims by the issuer) and the alleged co-mingling of customer funds. Second, we zoom in on Tether's peg stability before and after the February 28 and March 30 2021 attestation report releases, the first such reports after two and a half years without any disclosure on reserves and the only time where two such reports were released in two consecutive months. We focus on these attestation reports for two reasons. First, given they were the first after a

[5] We carefully select these control units such that they are likely to impact the price stability of the treated stablecoin but unlikely to be themselves affected by the policies or conditions of the treated stablecoin.

long period without any news on reserves, they carry high informational value. Second, the multiplicity and frequency of attestations (i.e. "treatments") thereafter makes a synthetic control analysis difficult.

The results from the second case study are consistent with model predictions. For one, our model predicts a deterioration of Tether's peg stability following the publicized concerns over its reserve adequacy in 2018. Indeed, we do find significant deviations in the actual series which are statistically significantly different from the synthetic counterfactual. Moreover, when zooming in on the two consecutive attestation reports in the sample, we find a statistically significant reduction in actual peg deviations relative to counterfactuals following the attestation releases.

Our third case study revisits the SVB event but examines Frax instead of USDC. Frax is a crypto-collateralized stablecoin, which among other cryptocurrencies held USDC as collateral going into March 2023. As a result, the sharp de-pegging of USDC served as a reserve volatility shock for Frax investors. Consistent with the *change* effect predicted by the model arising from a negative shock to Frax collateral values, we find that Frax peg deviations increased significantly following Circle's disclosure and the associated de-peg of USDC.

Finally, the fourth case study examines the TerraUSD stablecoin. TerraUSD was the largest algorithmic stablecoin and third largest stablecoin overall at its peak before imploding in May 2022.[6] TerraUSD's peg was underpinned by Luna, a self-issued and highly volatile cryptoasset. Luna effectively operated as a reserve asset for Terra, though (to put it mildly) it is of course not of the same quality as those that back the stablecoins in the previous event studies (i.e. actual assets). Consistent with the predictions of the model, we find that large negative public shocks – measured from the time-varying volatility of Luna's daily returns – were associated with larger peg deviations in TerraUSD's price even before its collapse. This link becomes stronger whenever the "equity value" of Terra, defined as the market capitalization of Luna minus that of TerraUSD, is low and hence the stablecoin is more vulnerable to a run (i.e. it is in a "ripe for run" region). These findings help rationalize the May 2021 UST depegging event (consistent with the *resilience* effect) as well as the final collapse in May 2022 (consistent with the *change* effect).

Our paper contributes to the understanding of stablecoins along multiple dimensions.

[6] See Uhlig (2022) for a model of TerraUSD and Liu et al. (2023) for a detailed analysis of the Terra ecosystem and its demise, consistent with our model predictions.

We consider how (i) the quality of issuers' reserve assets and (ii) public information about these assets change both stablecoin exposure to run risk and the type of global game being played among stablecoin holders. We provide a stylized setting that allows us to speak to various stablecoin arrangements by providing a general model of global games that nests the most standard type used in much of the literature on stablecoins. And we develop and empirically test theoretical predictions using event studies and a synthetic control approach across stablecoins. While previous literature situates run risk in the broader stablecoin ecosystem, we study this run risk in detail, providing a richer classification of different types of stablecoins and their distinct run dynamics.

Our results underscore the similarities and differences between increased precision in public information (via heightened disclosure and transparency) and reserve asset quality on run risk. One implication of our general model and our results is that, over time, stability can breed more stability as it supports path dependence on strong priors about reserve assets. Coupled with a policy of systematic disclosure, one can expect the main stablecoins to converge from the generalized model to one where the distribution of reserve assets is perceived as thin-tailed. In other words, the choice of reserve assets (coupled with disclosure) dictates the type of coordination game being played among stablecoin holders and issuer.

A second implication is that *transparency* has a distinct effect on run risk from the *asset quality* affecting the composition of reserves. Whereas the former dictates the propensity of SC holders to run on the issuer (i.e., a proxy for market risk which, in turn, affects issuer peg *stability*), the latter shapes the density over the region where the issuer fails for a given level of flight risk (thus affecting the issuer's *fragility*). Together, these results suggest a trade-off between high issuer fragility but a stable peg on the one hand, and low issuer fragility but more flighty investors on the other. This conceptual distinction arises only in the generalized model, and suggests that regulation of issuer disclosure should not happen without consideration of collateralization. Given that stablecoin issuers are currently largely outside the scope of traditional financial regulation, our work contributes to the policy debate over whether stablecoins should hold loss-absorbing resources to buffer against run risk. More broadly, insights from our model could potentially be applied to other settings with run risk, including money market funds and fractional reserve banking.

Roadmap. Section 2 reviews the related literature in detail and highlights our contribution. Section 3 introduces the model and discusses: how large shocks trigger runs (Section

3.4), the effect of transparency through public disclosure on stablecoin run risk (Section 3.5) and how reserve asset quality affects run risk (Section 3.6). Section 4 combines market data on stablecoins with event studies to empirically test the predictions of the model. Section 5 concludes. All proofs are contained in Appendix A.

2 Related literature

Our paper contributes to a growing literature on stablecoins.[7] Arner et al. (2020) provide an early non-technical overview that combines economic and legal perspectives, Mell and Yaga (2022) provide a comprehensive assessment of stablecoins from a technical standpoint, Barthélémy et al. (2023) show how stablecoin issuers' reserve management can spill over to the real economy, and Makarov and Schoar (2022) provide a broader overview of cryptoassets.[8] Lyons and Viswanath-Natraj (2023) highlight arbitrage design as a source of stability for Tether and argue that decentralization of issuance and access to arbitrage trades with the issuing treasury are key for peg stability. While they focus on arbitrage to assess what makes a stablecoin stable, our question is rather what makes them unstable, with a focus on the role of reserve transparency and quality. Li and Mayer (2021) derive stablecoin management strategies by issuers and agent demand for stablecoins in a dynamic model that focuses on the interplay between the endogenously determined stablecoin price (the exchange rate), reserve management and the issuer's equity shares or governance tokens. Whereas Li and Mayer (2021) shed light on issuer incentives around governance token issuance and debasement to mitigate stablecoin price volatility, we focus on stablecoin holder strategic incentives to maintain holdings for a given level of reserves and issuer equity. d'Avernas et al. (2023) in turn theoretically study stablecoin issuance as a commitment problem, potentially solvable through smart contracts. They consider various protocols for issuance and redemption and zoom in on *ex ante* design. Instead, our focus is on how these factors may affect run dynamics for a given design, explicitly featuring coordination, strategic uncertainty and equilibrium selection. Moreover, their model does not focus on the role of reserve transparency and the volatility of collateral, an important component of our analysis.

[7] More broadly, it also relates to, and builds on, an extensive literature on bank runs and currency attacks mentioned above.

[8] See also Caramichael and Liao (2022), Eichengreen et al. (2023) and Levy Yeyati and Katz (2022), among others. Klages-Mundt and Minca (2021, 2022) provide early studies of deleveraging spirals and stablecoin runs within the computer science literature.

Closest to our paper are the recent contributions by Bertsch (2023), Ma et al. (2023) and Gorton et al. (2022). Bertsch (2023) uses a first generation regime change global game to capture stablecoin fragility stemming from concerns about the quality of the issuer's assets, giving rise to run risk; this is the same fragility we are concerned with. Bertsch (2023) endogenizes the liability side of the stablecoin issuer's balance sheet *ex ante*, before the issuer is exposed to run risk. The issuer faces demand for stablecoins that arises from heterogeneity in preferences by groups of consumers over holding different monies (Agur et al., 2022), producing heterogeneous payoffs in the global game. The adoption stage links stablecoin demand to outside options such as bank deposits, and allows Bertsch (2023) to study the relationship between factors that increase adoption and shape stablecoin issuer fragility. We abstract from tensions in the adoption stage, focusing instead on how differences on the asset side of the issuer affect higher-order beliefs and run incentives among a fixed mass of stablecoin holders. Our contribution is to unpack theoretically and test empirically the effect of changes in issuers' asset quality, transparency and volatility on incentives to run.

Ma et al. (2023) also feature a first generation global game in the spirit of Goldstein and Pauzner (2005), with Diamond and Dybvig (1983) style liquidity shocks. In addition, they include a layer of arbitrageurs as firewalls, with the efficiency of the secondary market being an indicator of the fragility of the stablecoin's price. Their key finding is that higher price volatility comes with lower run risk, and *vice versa*, emphasizing the role that the concentration of arbitrageurs plays in staving off run risk under normal market conditions.[9] Our setup, by contrast, focuses on shocks to fundamentals that cast doubt over the solvency of the issuer, leading the entire market, arbitrageurs included, to collapse when issuers can no longer honor the promise of par convertibility. We demonstrate how transparency and volatility of reserve assets affect run risk, rather than how the two-layer market infrastructure of stablecoins affects stability of the stablecoin price in terms of market design.[10] Finally, our model allows for a broader class of both private and fundamental distributions, can

[9]Concentration of arbitrageurs leads to a trade off between price stability and run risk in their model. When there are fewer authorized arbitrageurs, arbitrage is less efficient and therefore price deviations are more frequent. But opening the door to more arbitrageurs increases run risk because it increases the first mover advantage for arbitrageurs. We find a similar trade-off between price stability and run risk, but facilitated by the interaction between reserve asset quality and the precision of information about reserve assets (transparency), for a given number of arbitrageurs.

[10]In Appendix B, we present an extension to include a secondary market in which stablecoins can be traded, and show that the existence of a secondary market can make arbitrageurs relatively more likely to demand redemption. More importantly, our core results continue to hold in this expanded setting.

accommodate non-fiat dollar backed stablecoins, and features a "large shock" result that is absent from first-generation global games.

Gorton et al. (2022) argue that stablecoins are able to fulfill their par promise despite being exposed to run risk and not paying any interest because levered traders (a third agent in their model) provide compensation to stablecoin holders for lending their coins.[11] As in Bertsch (2023), Gorton et al. (2022) model a classic first generation global game with uniform distributions. Therefore, they do not consider how variation in the information structure influences run risk, which is our focus. Moreover, investors in their model may learn about the proportion of safe assets only from issuer disclosures but do not study the implications of this theoretically (rather, it is posed as an empirical challenge). We contribute by formalizing how opaque portfolios held by issuers influence the beliefs of stablecoin holders (formed by the convolution of various assets) and, thus, contribute to run risk. We discuss when disclosure lowers/heightens run risk and how the magnitude of public shocks (informational or fundamental) are critical in triggering runs.

The interaction between the precision of public information and run risk is well established in the global games literature (Ahnert and Kakhbod, 2017; Metz, 2003; Prati and Sbracia, 2010; Szkup and Trevino, 2015). Our contribution over and above these results is twofold. Firstly, we disentangle the effects of transparency about reserve assets from the degree of tail risk in their quality. Secondly, we use a model that features a much less restrictive informational environment than in standard global games (in the spirit of Morris and Yildiz (2019)), allowing for a clearer distinction between the effects of transparency and quality of fundamentals. We extend this model into a regime-change setting and study the effect of changes in structural parameters on the equilibria. Our application to stablecoins offers a rich landscape from which to test these distinct channels.

3 Model

In this section, we present our model of stablecoins to analyze how information and the distribution of reserves influence runs. The model features uncertainty about the variance of

[11] Our model abstracts from the role of crypto speculators. However, we indirectly capture the interconnectedness that arises with such speculators by modeling the benefits to stablecoin holders from rolling over their holdings as positively dependent on the dollar value of assets in the crypto space. With this payoff structure, we argue that there is some correlation between the collateral value of the issuer and the expected payoffs to stablecoin holders from transacting in the crypto ecosystem – whether it be via direct speculation or lending to speculators.

reserve assets, which introduces fat tails into the distribution of the value of reserve assets. Fat-tailed distributions could capture, for example, sudden changes in reserve volatility, or fundamental uncertainty about reserve assets. In the extreme, as variance uncertainty declines and eventually converges to zero, the model we present nests more traditional regime change global game models in the spirit of Morris and Shin (1998, 2003) and Goldstein and Pauzner (2005). Some of the results we present (e.g. on transparency and disclosure) apply as well in such more simplified settings, whereas others are unique to our model environment. We draw on Morris and Yildiz (2019) to model the resultant global game and use hysteresis equilibrium to analyze the role of large shocks in triggering runs.

3.1 Model elements

The model features a unit continuum of risk-neutral stablecoin (SC) holders, $i \in [0, 1]$, and a single issuer of stablecoins. One stablecoin is issued to each SC holder.[12] For simplicity we assume that SC holders have redemption rights with the issuer, and therefore abstract from the role of the secondary market. In Appendix B, we show that our results are robust to the inclusion of a two-layer market structure, similar to Ma et al. (2023).[13] Stablecoins are backed by a vector of reserve assets with a combined dollar value θ. SC holders perceive θ as a random variable drawn from the portfolio-weighted convolution of all reserve asset distributions, with realizations of θ equal to the portfolio-weighted realized values of individual reserve assets (see Appendix C.1 for details).[14]

Cash is converted to stablecoins and redeemed at a one-to-one conversion rate, subject to the issuer continuing to operate. The transaction costs collected by the issuer for conversion are denoted by τ, with $0 < \tau < 1$.[15]

[12] While beyond the scope of our paper, the model can be extended to study the effect of large players ("whales") on run risk. See Corsetti et al. (2000) for an example in a currency crisis setting.

[13] In practice, for the biggest stablecoins only a selected group of authorized participants (arbitrageurs) has direct redemption rights with the issuer, and provide liquidity in the secondary market. Ma et al. (2023) show how more efficient arbitrage can exacerbate run risk. Since our focus is squarely on runs, we zoom in on the dynamics between stablecoin holders as a function of the transparency and volatility of SC reserves, in circumstances where pressure on the peg is strong, and abstract away from the intermediate arbitrageur layer.

[14] Our model can easily accommodate an analysis of the effect of increasing the weight of low-risk assets in the reserve portfolio. Tether's transition away from commercial paper and towards short-term safe assets is an example. Formally, an increase in the weight on low-risk assets lowers the variance in reserve asset values, which, as we show in Proposition 3, lowers the risk of issuer failure due to a run.

[15] We assume that transaction costs are fixed at the time SC holders decide whether or not to demand conversion. In practice, transaction costs are likely endogenous (e.g. some issuers may charge so-called "gas fees" on redemption, potentially increasing with transaction volume or the economic environment). Endogenizing the value of transaction fees would complicate the analysis and make issuers relatively more resilient to small shocks to fundamentals, but our core insights would be unchanged. We provide comparative statics on the value of τ in

The game extends over two stages. In the first stage, Nature selects θ and the issuer observes the dollar value of reserve assets. In the second stage, SC holders decide whether to demand conversion to cash, $a_i = 1$, or to maintain their stablecoin holding, $a_i = 0$. During this stage, the issuer processes aggregate conversion requests, $A \in [0, 1]$, and becomes unable to defend par whenever $A > \theta$, i.e., whenever the conversion value of reserve assets, θ, is smaller than the value of fiat currency demanded. The inability to meet redemptions by relying on the liquidation of reserve assets effectively renders the issuer insolvent. But note that the key problem facing the issuer is one of *liquidity*: what matters is not the value of reserve assets relative to liabilities (i.e. a solvency concern), but rather whether assets can command enough resources to liquidate liabilities.

The interaction between SC holder redemptions and the issuer's balance sheet produces a tripartite classification of fundamentals (Morris and Shin, 1998). There is a lower threshold, $\theta_L = 0$, such that for all $\theta \in (-\infty, 0)$, the issuer is fundamentally insolvent even if no SC holders demand conversion (i.e., $A = 0$). Similarly, there is an upper threshold, $\theta_U = 1$, such that for all $\theta \in [\theta_U, \infty)$, the issuer is solvent even if all SC holders demand conversion (i.e., $A = 1$). We consider the stablecoin to be in a "ripe for run" region whenever $\theta \in [\theta_L, \theta_U)$.

The payoffs received by SC holders based on their actions for each possible state of the issuer are summarized in Table 1. The payoff per unit accruing to SC holders is captured by $\pi : \mathbb{R} \to \mathbb{R}$ with $\pi(\theta) \geq 1 - 2\tau$ for all θ.[16] Such a payoff could reflect returns from using stablecoins as collateral in decentralized finance applications, from realizing gains in the crypto space without costly conversion back to fiat currency, or revenue from illicit activities such as tax evasion. The returns may depend on fundamentals in the crypto environment, such that $\pi'(\theta) \geq 0$. The link between the value of the assets used in the reserve portfolio and the crypto returns to SC holders is not necessary for our results (indeed, we allow for $\pi'(\theta) = 0$), but rather illustrates interdependency between issuer and holder, particularly when crypto assets are used as collateral.

In states where the issuer is solvent, SC holders prefer to maintain their holdings, and receive a strictly higher expected payoff from demanding conversion in states where the

the Appendix.

[16] This mapping restricts the image of π to real numbers and ensures that two-sided limit dominance is satisfied for all $\theta \in \mathbb{R}$ and for any arbitrarily small transaction cost τ. An example of a function satisfying this condition is $\pi(\theta) = e^{\theta} + 1$, which suggests that SC holders receive at least their claim of \$1 plus additional benefits from transacting in the crypto environment. A well-defined dominance region is crucial for the uniqueness of a global game equilibrium, since it relies on iterative deletion of strictly dominated strategies (Milgrom and Roberts, 1990). See Goldstein and Pauzner (2005) for a discussion of one-sided strategic complementarity.

issuer is rendered insolvent. If a SC holder demands conversion and the issuer is rendered insolvent, the SC holder faces a sequential service constraint, and is able to obtain her fiat money converted at the pegged value, net of transaction costs, with a probability, $f : [0,1] \to [0,1]$, that is inversely proportional to aggregate withdrawals as she has to wait in line. In the event of issuer failure, she receives zero if abstaining from the run. If the issuer remains solvent, SC holders who demand conversion to cash re-enter the market by buying the stablecoin again, maintaining their dollar claim on the issuer but incurring fees on their exit and re-entry, whereas they receive $\pi(\theta)$ if they abstained from demanding conversion.[17]

	Issuer solvent	Issuer insolvent
$a_i = 1$	$1 - 2\tau$	$f\left(\frac{\theta}{A}\right)(1-\tau)$
$a_i = 0$	$\pi(\theta)$	0

Table 1: Stablecoin holder payoffs. Action $a_i = 1$ denotes demanding conversion; action $a_i = 0$ denotes maintaining a holding.

3.2 Information structure

Under complete information, there are multiple equilibria for all $\theta \in [\theta_L, \theta_U)$. If SC holders anticipate $A > \theta$, it is a best response to demand conversion (since $1 - \tau > 0$). If they anticipate $A \leq \theta$, it is a best response to maintain their holding (since $\pi(\theta) > 1 - 2\tau$).

SC holders, however, do not observe the quality of reserve assets θ directly. Instead, they receive noisy private signals about θ. We assume that θ is normally distributed, and we also assume this distribution and its mean, $\underline{y} < y < \overline{y}$, are common knowledge to all SC holders, while the true variance, σ_θ^2, is unknown.[18] The assumption of common knowledge of all structural parameters is standard in most global games models, as it allows to focus on the higher-order uncertainty inherent in these settings (see Morris and Shin (2003)). We take a Bayesian approach to emphasize the uncertainty prevalent in the stablecoin universe.

We assume a collective belief by stablecoin holders that the variance is drawn from an inverse chi-squared distribution. This reflects a degree of model uncertainty in the environment, whereby SC holders do not fully understand the properties of reserve assets sustaining

[17] We normalize the benefit from holding coins to SC holders to zero with a defunct issuer, as there are currently no resolution procedures in place for stablecoins (see Bains et al. (2022)).

[18] The restriction on y ensures that the global game is not dominance-solvable as in Morris and Yildiz (2019). See the Appendix for derivations of $\underline{y}$ and $\overline{y}$. The assumption of normality greatly simplifies notation without loss of the key insights, and allows us to place more probability mass on certain values, as opposed to, say, using the uniform distribution. One way to interpret the θ is as deviations from average values, to allow for negative numbers.

the issuer's promise.[19] Effectively, stablecoin holders perceive that the reserve portfolio is drawn from a distribution with fat tails (i.e., with positive excess kurtosis). In addition to being able to shed light on the role of transparency and public information disclosures, our model can also help explain why markets sometimes remain relatively optimistic even in the face of gradually deteriorating reserve assets, or why events seemingly unrelated to fundamentals can trigger runs.

3.3 Equilibrium analysis

In the second stage of the game, SC holders observe noisy private signals about the quality of reserve assets:

$$x_i = y + z_i, \tag{1}$$

where $z_i = \sigma_x \varepsilon_i + \eta$, with independently and identically distributed components $\varepsilon \sim \mathcal{N}(0,1)$ and $\eta \sim t(\nu) = G(\eta;\nu)$, with degrees of freedom, $\nu > 2$.[20] The degrees of freedom are proportional to the number of past observations that SC holders have for θ. The second term in equation (1) is an aggregated noise component that SC holders disentangle to form beliefs about fundamentals and the mass of aggregate withdrawals.

Let

$$R(z) = Pr[z_j \leq z | z_i = z] = \frac{\int \Phi(\varepsilon)\, \phi(\varepsilon) g(z - \sigma_x \varepsilon) d\varepsilon}{\int \phi(\varepsilon) g(z - \sigma_x \varepsilon) d\varepsilon}, \tag{2}$$

denote the *rank belief* of SC holder i. Agent i's rank belief represents the proportion of other SC holders who observe a lower signal than i, conditional on her own signal. Rank beliefs form a mapping between an individual's private signal and the withdrawal mass she expects.[21] With strategic complementarity among SC holders, and the fact that a sufficiently large mass of withdrawals causes the issuer to fail, rank beliefs form a crucial part of the equilibrium in switching strategies. Given the variance uncertainty SC holders have over θ, the rank belief function is *non-monotonic* in z and approaches $\frac{1}{2}$ as z grows large in either direction (see Figure 2).

Our first result, summarized in Proposition 1, comprises equilibrium thresholds for both SC holders and the SC issuer that are *locally unique*. For all SC holders, there is a common

[19] Representing model uncertainty in this way has been widely used in the finance literature (Weitzman (2007)).

[20] The inverse χ^2 distribution is a conjugate prior for the variance of the normal distribution which forms a posterior student's t-distribution over fundamentals.

[21] For a comprehensive exposition of rank beliefs and "(q,p)-evident events" that determine rationalizability, see Morris et al. (2016).

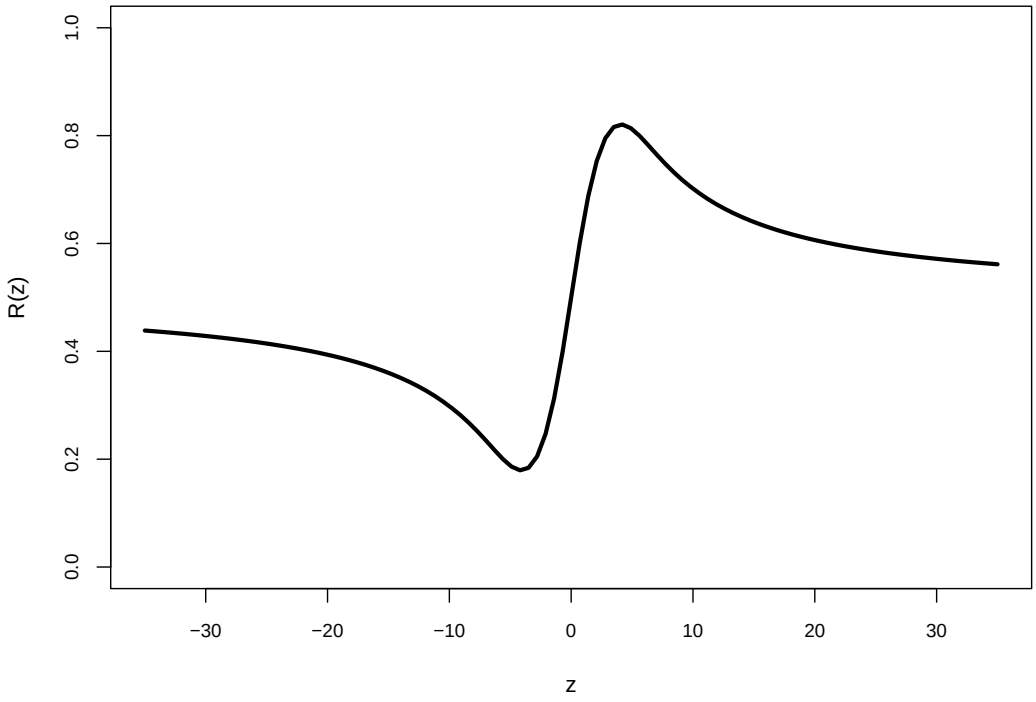

Figure 2: Rank belief function with fat-tailed distribution of reserve assets.

Notes: The rank belief function for SC holders when they perceive tail risk in the reserve portfolio. Rank beliefs represent conditional beliefs about the proportion of other SC holders who observe lower signals. They are non-monotonic and approach uniformity in the limit as $|z|$ (the absolute value of aggregate noise) grows large.

signal threshold, $x_c \in \{\hat{\hat{x}}, \hat{x}\}$, that drives SC behavior and produces consistent higher-order beliefs about the strategies of other SC holders. For the SC issuer there is a fundamental threshold, $\theta_c \in \{\hat{\hat{\theta}}, \hat{\theta}\}$, that dictates the point at which the mass of redemption requests becomes so large that the issuer is rendered insolvent. The SC holder of type x_c holds a posterior belief over the probability of issuer failure greater than or equal to $\rho(\theta_c) \equiv \frac{\pi(\theta_c)+2\tau-1}{\pi(\theta_c)+2\tau-1+(1-\tau)f(\theta_c/R(z_c))}$, where $\rho(\theta_c)$ is the belief over the event of issuer failure that renders her indifferent between demanding conversion and maintaining the holding. The threshold noise component $z_c = x_c - y$ produces a rank belief that is equal to the anticipated withdrawal mass, since all SC holders with lower signals are expected to run on the issuer.

Proposition 1. *Fix $\tau < \hat{\tau}$ and $\sigma_x^2 > \hat{\sigma}_x^2$. There are two locally unique thresholds, $\hat{\hat{\theta}} = y + \hat{\hat{\eta}} > \theta_L$ and $\hat{\theta} = y + \hat{\eta} < \theta_U$, such that the issuer is rendered insolvent if $\theta < \hat{\hat{\theta}}$ and remains solvent if $\theta \geq \hat{\theta}$. There are multiple equilibria for all $\theta \in [\hat{\hat{\theta}}, \hat{\theta})$.*

For each threshold on the value of reserve assets, $\hat{\hat{\theta}}$ and $\hat{\theta}$, there is a corresponding

unique individual signal, $\hat{\hat{x}}$ and $\hat{x}$, such that each SC holder runs on the issuer if $x_i < \hat{\hat{x}}$, and maintains a holding if $x_i \geq \hat{x}$. Both actions are rationalizable for all $x_i \in [\hat{\hat{x}}, \hat{x})$.

Proposition 1 pins down equilibrium thresholds in the global game among SC holders and the issuer. SC holders use Bayesian inference to determine the probability that others have equally or more pessimistic signals than their own, allowing them to weigh the expected benefits and costs of early withdrawal. For $a_i = 1$ (demand conversion) to be uniquely rationalizable, and given a critical portfolio return θ_c, SC holder i must believe with a sufficiently high probability that at least θ_c other SC holders will demand conversion.

Conditional on a switching strategy around $\hat{\hat{\theta}}$, threshold $\hat{\hat{x}}$ represents the signal that renders a SC holder just indifferent between redeeming and rolling over. Two conditions characterize this indifference. First, for all $x_i < \hat{\hat{x}}$, the payoff from withdrawing is at least as large as from maintaining a holding under belief $F(\hat{\hat{\theta}}|x_i)$ which, in turn, reinforces other SC holders' decisions to demand conversion. Second, for all signals below $\hat{\hat{x}}$, the expected withdrawal mass, A, is at least as large as the value of reserve assets, $\mathbb{E}[\theta|x_i]$. For the SC holder of type $\hat{\hat{z}} = \hat{\hat{x}} - y$, this mass is $A = R(\hat{\hat{z}})$. The decision rule among SC holders aggregates into a threshold for the dollar value of collateral such that the issuer experiences a large-scale run that leads to failure whenever $\theta < \hat{\hat{\theta}}$. Similarly, threshold $\hat{\theta}$ is the lowest value of reserve assets that ensures issuer survival when SC holders form a switching strategy around threshold $\hat{x}$.

Extreme beliefs generate contagion. In particular, even if held by a small proportion of SC holders with very pessimistic or very optimistic signals, extreme beliefs catalyze a decision to withdraw or hold that extends towards SC holders with more moderate signals (i.e. signals that are closer to the true value of fundamentals, θ). Given that all SC holders who observe sufficiently low signals will demand conversion, those whose signals are slightly higher will also find it optimal to demand conversion, since their own beliefs about the likelihood that the issuer survives are low and they anticipate that everyone with even lower signals will definitely run.[22] This, in turn, causes SC holders with even larger signals to follow suit, and so on. On the other extreme, SC holders who have a strictly dominant

[22] Formally, there exists a lower threshold, $\underline{x}$ such that

$$\mathbb{P}[\theta \leq 0 | x_i = \underline{x}] = \frac{\pi(0) + 2\tau - 1}{\pi(0) + 2\tau - 1 + (1-\tau) f(0/R(\underline{z}))},$$

where $\underline{z} = \underline{x} - y$. When $x_i \leq \underline{x}$, SC holders have an optimal strategy to withdraw even if the anticipated withdrawal mass is zero.

strategy to maintain a holding justify the decisions of others with more moderate signals to refrain from the run.[23] This justifies those with beliefs that reserve assets are of lower value to refrain from running, until a SC holder is just indifferent. The thresholds $\hat{\hat{x}}$ and $\hat{x}$ bound two regions of uniquely rationalizable actions, resulting in determinate outcomes when reserve assets are subject to large positive or negative shocks.

The technical restrictions in Proposition 1 pin down local uniqueness. Transaction costs must be sufficiently low, $\tau < \hat{\tau}$, to produce strategic complementarity among SC holders so that they form consistent beliefs about the actions of others. In a restricted case where $\sigma_x^2 \leq \hat{\sigma}_x^2$, we obtain the well-established unique equilibrium found in the global games literature (Morris and Shin, 1998, 2003).[24] If we take $\nu \to \infty$, the $t-$distribution approaches a standard normal distribution, rank beliefs become monotone in z, and as long as private noise is small enough, the thresholds $\hat{\theta}$ and $\hat{\hat{\theta}}$ converge to the unique equilibrium θ^*.[25]

3.4 Large shocks and equilibrium shifts

We now focus on equilibrium shifts (i.e., when do we expect a run on stablecoin issuers and under what conditions can the peg be re-established?) and the role that sudden, large shocks to reserves play in such shifts. The result of the static global game can be easily transformed towards a dynamic application.[26] At the beginning of each period $t \geq 0$, there is an expected dollar value of reserve assets, y_t. Common shocks (η_t) and idiosyncratic shocks (ε_{it}) are independently drawn across time and SC holders, who decide $a_{it} = \{0, 1\}$ (i.e., roll over or withdraw) after observing a signal x_{it} as in the static game above, aggregating to a withdrawal mass A_t. The commonly held prior over the dollar value of reserve assets

[23] The analogous upper threshold $\overline{x}$ is given by

$$\mathbb{P}[\theta \leq 1 | x_i = \overline{x}] = \frac{\pi(1) + 2\tau - 1}{\pi(1) + 2\tau - 1 + (1-\tau) f(1/R(\overline{z}))},$$

where $\overline{z} = \overline{x} - y$. When $x_i > \overline{x}$, SC holders always find it optimal to maintain their holding, even if the anticipated withdrawal mass is one.

[24] The intuition for the equilibrium switching point is identical in both the unique-equilibrium and generalized settings. However, certain properties of the threshold break down when we allow for fat tails. We contrast the comparative statics in the unique equilibrium case with the results in Proposition 1 in Appendix A.1.

[25] It is worth noting that the locally unique thresholds $\hat{\theta}$ and $\hat{\hat{\theta}}$ are qualitatively different from θ^*, because in the fat-tailed environment the individual indifference condition is decoupled from the critical mass condition, with different comparative static results reported in Appendix A.1. See also Duley and Gai (2023) for a formal treatment. Our generalized environment permits a larger degree of idiosyncratic noise than the traditional environment.

[26] To be clear, we do not extend the model into a rich dynamic arrangement of the kind studied in d'Avernas et al. (2023) or Li and Mayer (2021). Instead, we simply re-frame the model as a sequence of plays of the static model, linking a global games framework across two dates $t-1$ and t, along the lines of Morris and Yildiz (2019).

in the current period, $y_t = Y(\theta_{t-1})$, is determined by a known process, $Y : \mathbb{R} \to \mathbb{R}$, with $\theta_t = y_t(\theta_{t-1}) + \eta_t$. We denote the sample size of past observations by $|\mathbf{\Theta}_t|$ where $\mathbf{\Theta}_t \equiv \{\theta_{t-1}, \theta_{t-2}, \dots \theta_0\}$, that determine the degrees of freedom $\nu_t \propto |\mathbf{\Theta}_t|$ in distribution $G(\cdot)$.[27]

To address the indeterminacy of outcomes in the presence of multiple equilibria in Proposition 1, we focus on *hysteresis equilibrium*. This equilibrium corresponds to games that exhibit path dependence, so that agents' actions are systematically shaped by the aggregate outcomes or payoff parameters in previous periods.[28] With this equilibrium selection method, we resolve the indeterminacy of outcomes for moderate values of θ_t by arguing that agents will play the less aggressive decision rule (i.e., run if $x_{it} < \hat{\hat{x}}_t$) if their choice was to maintain a holding in the previous period, and will play the more aggressive decision rule (i.e., run if $x_{it} < \hat{x}_t$) otherwise.[29] Corollary 1 summarizes this.

Corollary 1. *Under hysteresis equilibrium, each SC holder maintains a holding if and only if* $x_{it} \geq x_c$ *where*

$$x_c = \begin{cases} \hat{\hat{x}}_t(y_t) & \text{if } t = 0 \text{ or } A_{t-1} < \frac{1}{2} \\ \hat{x}_t(y_t) & \text{otherwise.} \end{cases} \tag{3}$$

This equilibrium formulation suggests *inertia* in majority behavior by SC holders. If the issuer was solvent in the previous period and a SC holder observes a modest deviation from her prior, maintaining her holding is uniquely rationalizable. Aggregating this behavior, whenever the value of reserve assets experiences small shocks, there is a *resilience* effect by SC holders and the peg is maintained. However, a large negative shock (i.e., $\theta_t < \hat{\hat{x}}_t(y_t)$) induces a *change effect*, making withdrawing uniquely rationalizable for the median SC holder who observes $x_t = \theta_t$.

Figure 3 illustrates withdrawal dynamics under hysteresis equilibrium. Time is on the horizontal axis, whereas the vertical axis captures a sample path of fundamentals (value

[27] That is, common past observations of θ provide public information that SC holders use to learn about the population parameters governing the distribution of reserve assets. The shape of the t distribution is influenced by this sample of past observations. Our approach differs from the dynamic global game application in Angeletos et al. (2007) in which fundamentals do not change over time, but instead players observe a sequence of past plays which helps them learn about the value of fundamentals.

[28] See Bebchuk and Goldstein (2011), Rajan (1994) and Romero (2015) for similar characteristics of persistence in coordination games.

[29] The nature of the global game is unchanged. We do not alter the timing of the game nor the payoff parameters. Instead, we assume the history of past play (as summarized by the observed mass of conversion requests and commonly held beliefs over reserve asset quality in the previous period) is instructive in resolving multiplicity.

of reserve assets) and aggregate behavior of SC holders (determining the stablecoin price). With small variance in idiosyncratic noise, aggregate behavior is always close to one or zero. There are two periods in which the peg is maintained: at the beginning, the majority rolls over their holdings even though there are downward drifts in the value of reserve assets. However, a large enough negative shock in the value of reserve assets, coupled with withdrawing becoming a $p-$dominant response for the median SC holder, causes a majority of SC holders to run on the stablecoin and its price drops significantly.[30] Subsequent increases in the value of reserve assets do not recover the peg as long as withdrawing remains $p-$dominant for the median SC holder. A large positive shock that exceeds $\hat{x}(y_t)$ near $t = 80$ triggers an optimistic shift in higher-order beliefs for the median SC holder, causing a majority reversion to an equilibrium SC holding close to one.

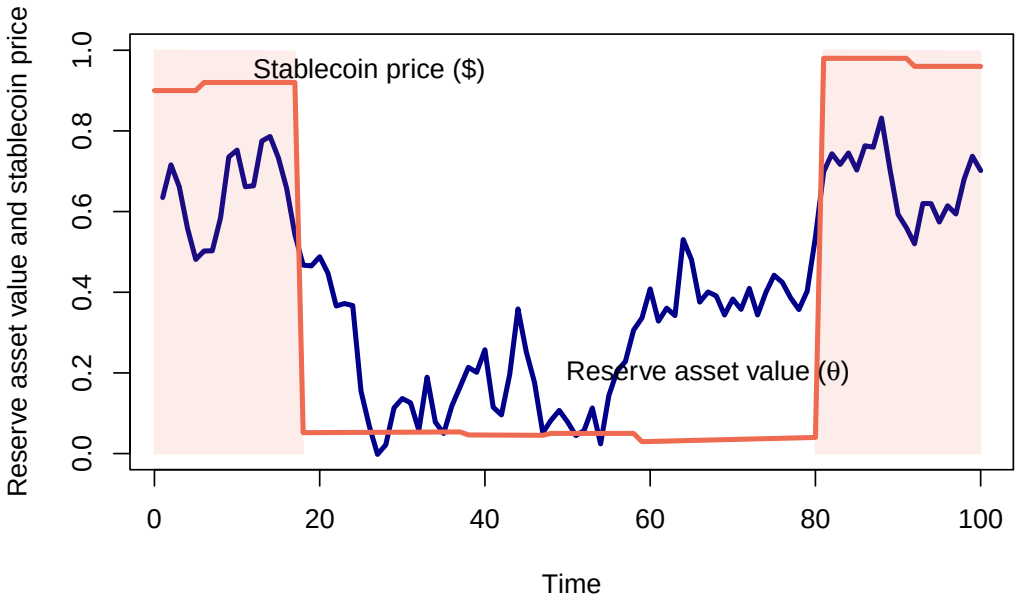

Figure 3: Equilibrium shifts under hysteresis equilibrium.

Notes: The issuer is solvent only during the early and final stages (shaded time intervals) for a given path of the dollar value of reserve assets (blue solid line).

The hysteresis equilibrium in Corollary 1 gives rise to two predictions.

[30] Demanding conversion is said to be $p-$dominant whenever it is a best response to the conjecture placing probability p on the event that the issuer is rendered insolvent (Morris et al., 2016). In our formulation, SC holders must believe with probability at least $p = \rho(\theta_t) \equiv \frac{\pi(\theta_t)+2\tau-1}{\pi(\theta_t)+2\tau-1+(1-\tau)f(\theta_t/A)}$ that the issuer will fail to make withdrawing a rationalizable response.

Prediction 1. Resilience effect. *When the dollar value of reserve assets enters a region where it is "ripe for a run" (i.e., the dollar value of reserves backing each coin, θ, is between 0 and 1), the SC peg is resilient to small shocks in reserve asset values (i.e., $\theta_t(y_t) \geq \hat{\hat{\theta}}_t(y_t)$).*

Prediction 2. Change effect. *When the dollar value of reserve assets enters a region where it is "ripe for a run" (i.e., the dollar value of reserves backing each coin, θ, is between 0 and 1), a large negative shock in reserve assets triggers a run on the SC, even if the value was initially close to* \$1 *(i.e. $\theta_t(y_t) < \hat{\hat{\theta}}_t(y_t)$). Analogously, after a destabilizing run on the SC, the peg may recover following a large enough positive shock in reserve asset values, even if the value of reserve assets is not* \$1 *(i.e., when $\theta_t(y_t) \geq \hat{\theta}_t(y_t)$).*

To test the change effect described above, we analyze how USDC's exposure to SVB served as a large reserve volatility shock to Frax investors in section 4.5, and we compare the difference between the effects of small and large shocks to fundamentals by testing peg deviations of TerraUSD in Section 4.6.

3.5 Reserve transparency and run risk

Does a commitment to transparency by stablecoin issuers in the form of broadcasting reserve portfolios mitigate or aggravate the probability of a run (i.e., $\mathbb{P}[\theta \leq \theta_c]$)? An increase in transparency helps to resolve uncertainty among SC holders over the parameters governing the distribution of reserve assets. Both the disclosure of public information about reserve assets in period t and past realizations of the dollar value of reserve assets bring the perceived distribution closer to the true distribution (determined by the issuer's reserve asset portfolio). Disclosure by the issuer is assumed to be commonly observed by all SC holders, and thus acts as a *public signal* that increases transparency and affects perceived fundamental parameters. This is conceptually distinct from changes to the dispersion in private information, σ_x, which comes from the likes of social media, word of mouth, and so on.[31] To compare the effects of transparency on run risk with that of issuers holding varying-quality reserve assets, we provide a comparative static result on reserve asset value thresholds $\hat{\hat{\theta}}_t$ and $\hat{\theta}_t$ that correspond with the signal thresholds in Proposition 2 below.

Proposition 2. *Fix $\tau < \hat{\tau}$ and $\sigma_x^2 > \hat{\sigma}_x^2(\nu)$. The release of public information or a longer time since the issuer's launch, that increase degrees of freedom (ν_t), lowers $\hat{\hat{x}}_t$ and $\hat{\hat{\theta}}_t$, and*

[31] Indeed, one of the benefits of our generalized setting over the traditional global game environment is that we allow for substantial variance in idiosyncratic signals.

increases $\hat{x}_t$ *and* $\hat{\theta}_t$.

Proposition 2 presents the role of transparency and issuer age under hysteresis equilibrium. Figure 4 illustrates the effect on thresholds of the underlying global game that characterize the hysteresis equilibrium. Increased transparency towards SC holders who believe that reserves are subject to high volatility has a negative influence on the probability of a run if no run occurred in the previous period, and increases the probability of a run continuing if one took place last period. In other words, runs on the issuer become less likely as SC holders receive more precise public information, but recoveries from a run (i.e. SC holders finding it optimal to maintain a holding) also become more difficult once a run is triggered. To the extent that SC holders use past events to learn about the issuer, this result gives theoretical support to evidence that issuers that have been around for longer are subject to less price volatility than their younger peers (Kosse et al., 2023). Descriptive evidence, for e.g. Tether, is in line with this, as over time and in the wake of systematic disclosures the peg becomes more stable, as evident in Figure 19 in the Appendix. We explore this more systematically in the next section.

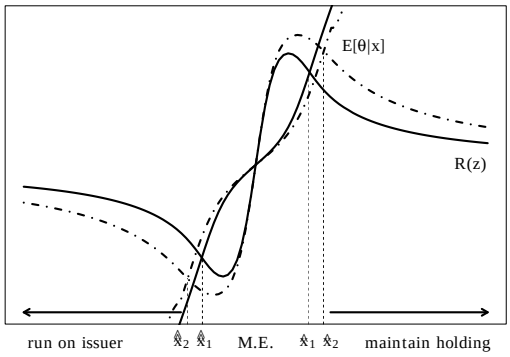

(a) An increase in transparency widens individual indifference conditions that define switching point signals for stablecoin holders.

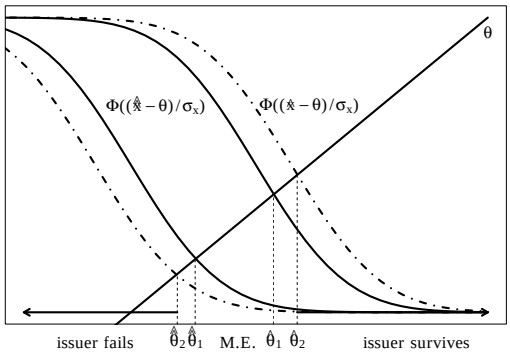

(b) A widening of switching points also shifts the mass of withdrawals, $\Phi(\cdot)$, at every given value of reserve assets, θ, widening failure points, $\hat{\hat{\theta}}$ and $\hat{\theta}$.

Figure 4: The greatest and least Bayes-Nash equilibrium signal thresholds, $(\hat{x}, \hat{\theta})$ and $(\hat{\hat{x}}, \hat{\hat{\theta}})$, are sensitive to disclosure or increases in transparency by the stablecoin issuer. An increase in transparency raises the degrees of freedom, ν, that form a structural parameter in the unknown value of reserve assets, lowering the lower switching point, $\hat{\hat{x}}_1$ to $\hat{\hat{x}}_2$, and increasing the upper point, $\hat{x}_1$ to $\hat{x}_2$.

What does this mean for the probability of a run? Intuitively, when market expectations are poor, a commitment to higher transparency of information about reserves makes

relatively smaller pessimistic signals carry greater weight in SC holders' strategic reasoning. Every SC holder anticipates that every other SC holder will be relatively flighty when a run has already occurred in the previous period. Under these conditions, greater precision of public information amplifies the contagion process that determines the switching point $\hat{x}$. As a result, the critical dollar value $\hat{\theta}$ increases, widening the event space where a run takes place, $(-\infty, \hat{\theta})$. Such a dynamic was very clearly at play during the breaking of USDC's peg as stress in SVB mounted (Figure 1): the large shock brought on by the bank distress induced a run, and increased transparency about the value of dollar reserves (i.e. cash held at SVB), acted as a focal point for SC holders to coordinate on sell-offs.

Proposition 2 thus delivers clear testable predictions. In particular, it suggests that, for a given level of fundamentals and a depegging in the prior period, $A_{t-1} \geq \frac{1}{2}$, pressure on the peg should grow following an increase in transparency over reserve assets in the issuer's portfolio. To test this prediction, we analyze the stability of USDC's stablecoin peg following public disclosure about the size of reserves held in SVB deposits in Section 4.3, and analyze market responses to the increased public scrutiny over Tether's reserve portfolio following the release of audit reports in Section 4.4.

From this proposition a corollary also emerges that links results in our setting to more traditional global games settings:

Corollary 2. *Fix $\tau < \hat{\tau}$. An increase in ν_t also increases $\hat{\sigma}_x^2$. For $\nu_t > \hat{\nu}_t$, where $\hat{\nu}_t$ solves $\sigma_x^2 = \hat{\sigma}_x^2(\hat{\nu}_t)$, thresholds converge to a unique equilibrium characterized by (x_t^*, θ_t^*).*

Corollary 2 shows how, all else equal, increases in transparency beyond a critical level cause the thresholds $\hat{\hat{x}}$ and $\hat{x}$ to converge to the well known unique equilibrium in the global games literature. It suggests, therefore, that in a generalized setting, increases in transparency change not only the tendency of SC holders to run on a given signal, but also the nature of the coordination game being played. The convergence to uniqueness follows by virtue of rank beliefs becoming uniform and the greatest and least switching points, $\hat{x} \geq x^*$, and $\hat{\hat{x}} \leq x^*$ respectively, converging to x^*.

3.6 Reserve quality and run risk

Results so far have emphasized issues related to the transparency of reserves. There is an important distinction to be made, however, between reserve transparency and quality. In this section, we show that changes in the distribution of the underlying assets that back an

issuer's stablecoins are qualitatively distinct from transparency in their effect on self-fulfilling run dynamics among SC holders.

Corollary 1 shows that, conditional on issuer survival in the previous period, the probability that an issuer fails due to a self-fulfilling run can be measured by $\mathbb{P}[\theta \leq \hat{\hat{\theta}}_t] = \Phi\left(\frac{\hat{\hat{\theta}}_t - y_t}{\sigma_\theta}\right)$. Proposition 3 presents a comparative static result on the variance of reserve assets and the probability of issuer failure:

Proposition 3. *Suppose that $t = 0$ or $A_{t-1} < \frac{1}{2}$. A reduction in the variance of the distribution of reserve assets lowers the probability of issuer failure:*

$$\frac{\partial \Phi(\cdot)}{\partial \sigma_\theta} = \frac{-(\hat{\hat{\theta}}_t - y_t)}{\sigma_\theta^2} \phi\left(\frac{\hat{\hat{\theta}}_t - y_t}{\sigma_\theta}\right) \geq 0\,. \tag{4}$$

Insofar as the tightness of the distribution of θ captures the quality of reserve assets (since negative deviations from the reference value, \$1, mean that the issuer cannot honor the promise of par conversion), an improvement in asset quality (i.e., a reduction in σ_θ) lowers the probability of failure, irrespective of the transparency of the issuer. This result arises from the distinction between *collective beliefs* over the issuer's balance sheet, which determine the failure *thresholds*, and the *probability of failure*, which increases with the density of the probability distribution over the failure range, $(-\infty, \hat{\hat{\theta}}]$, for a given threshold.

Summary of model results and implications. Before moving to the empirical applications, we take stock of the various results obtained from the model, and contrast them with those from a more traditional version where variance is known.

Figure 5 summarizes our results by contrasting both the threshold(s) that characterize the equilibria of the game and the type of coordination game being played. When an issuer holds reserve assets whose value is drawn from a fat-tailed distribution and a majority of SC holders did not demand conversion in the previous period, Corollary 1 demonstrates that the probability of issuer failure is proportional to the size of the interval $(-\infty, \hat{\hat{\theta}}_t)$. As model uncertainty decreases and the structural parameters become common knowledge, Corollary 2 establishes convergence to the well-known unique equilibrium with a failure threshold between the two fat-tailed switching points where issuer failure is proportional to the size of the interval $(-\infty, \theta_t^*)$. The distance of the second interval is (weakly) larger than the first. But note that this does not necessarily imply that the probability of issuer failure, defined

by $\mathbb{P}[\theta \leq \theta^*]$, is also larger than that of an issuer with more opaque reserve assets, $\mathbb{P}[\theta < \hat{\hat{\theta}}]$.[32] While transparency influences the size of these intervals by affecting the perceived quality of reserve assets, the true quality of these assets is what matters for the *probability* that an issuer fails. This is what distinguishes the effects of the *quality* of reserve assets on run risk from *transparency* over the composition of reserve assets. Figure 6 illustrates this using Frax, USDC and Tether as examples.

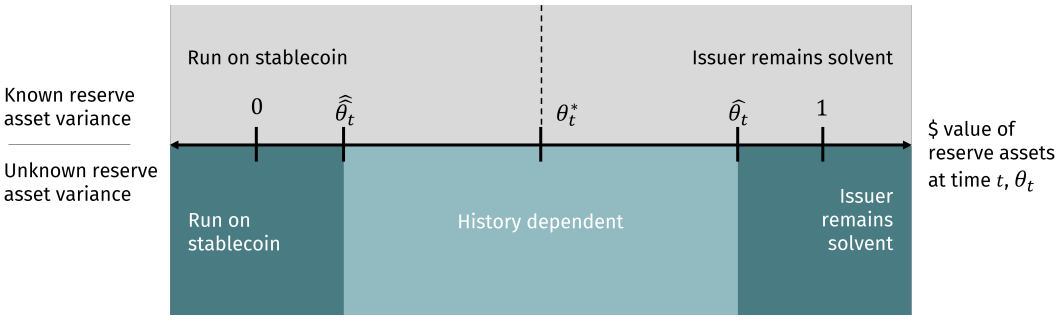

Figure 5: Model summary – switching thresholds in ripe for run region.

Notes: Switching thresholds in the ripe for run region $[0, 1)$ that determine run risk on issuers whose reserve asset variance is known (gray area, traditional global game) and those whose variance is unknown (blue area, generalized global game).

The summary in Figure 5 highlights a notable difference between the unique equilibrium in traditional global games and our model environment. While results from a model with known reserve variance suggests issuer solvency should teeter precariously around the fundamental threshold θ_t^*, a stablecoin whose characteristics are not well understood among SC holders features runs and recovery of confidence only following *large public shocks*.

This difference may seem surprising, as it implies that issuers with more opaque reserve assets are more resilient to small shocks, but the inertia has a simple intuition. A SC holder of slightly pessimistic type x_i does not trust that her signal is representative of weak reserve asset values, and believes that others harbor similar doubts about the information recovered from their own signals. This prevents her from running on signals that would induce a run in the unique equilibrium case. By contrast, a signal that is large in absolute terms, such that $x_i < \hat{\hat{x}}_t$, causes the SC holder to interpret her signal as a fundamental shock rather than idiosyncratic noise. With these beliefs, she no longer knows where she ranks in

[32] Notice, too, that the same inertia that prevents SC holders from running in the first place also prevents a recovery in the event of a run once the dollar value of reserves picks up. The interval $[\theta_t^*, \infty)$ over which an issuer recovers the peg, is larger for issuers with tightly distributed reserve assets than those with high volatility assets, for which the interval is $[\hat{\theta}_t, \infty)$. In this way, the heightened strategic uncertainty introduced by fat tails in the perceived distribution of reserve assets acts as a double-edged sword for stablecoin issuers.

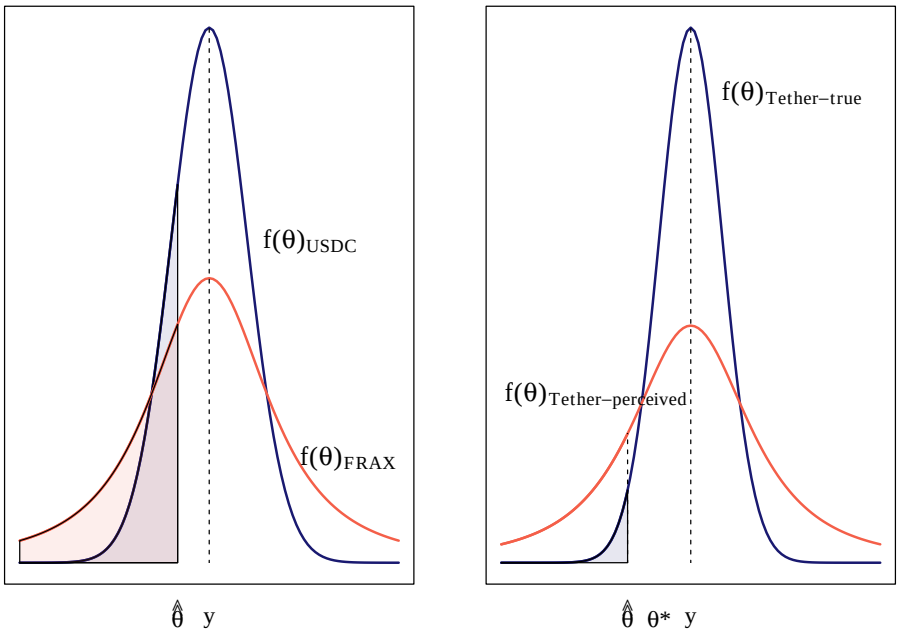

Figure 6: An illustration of reserve quality and transparency.

Notes: The left panel illustrates the reserve quality effect. It plots the probability density functions (PDFs) for a stablecoin with high quality reserves (USDC) and more volatile reserves (Frax). While the run threshold ($\hat{\theta}$) is determined by collective beliefs shaped by transparency, the probability of issuer failure is proportional to not only the threshold but also the (true) variance of the reserve asset value. High-quality reserves (blue shaded area) have a lower probability of failure than low-quality reserves (orange shaded area) for a given level of transparency. The right panel illustrates the reserve transparency effect. It plots the PDFs for a high-quality reserve stablecoin (say, Tether in its current version) with transparent (true) and opaque (perceived) reserve quality respectively. The run threshold ($\hat{\theta}$) is determined by the perceived distribution, but the probability of issuer failure is dependent on the true distribution (blue shaded area).

the population (i.e., her rank beliefs are diffuse). Because running produces higher payoffs than refraining under such beliefs, she is prompted to demand conversion. In this way, both opacity about the reserve portfolio and volatility in the realized dollar value of reserve assets serve to anchor runs around large, public shocks.

Together, Proposition 2 (how increased transparency affects run risk) and Proposition 3 (the true dispersion in reserve asset value) suggest a possible trade-off between peg stability and issuer fragility. We show that on the one hand, conditional on model uncertainty about the reserve portfolio and that a run is not already occurring, lower transparency by the issuer lowers the tendency of SC holders to run on the issuer. In this sense, opacity can be "good" for the issuer as the peg is more stable. But it also creates moral hazard. When the disciplining effect of run risk is diminished, issuers are likely to seek higher returns in more volatile assets, leading to greater fragility (left panel of Figure 6) since $\mathbb{P}[\theta \leq \hat{\hat{\theta}}; \sigma_\theta]$ is larger.

On the other hand, greater transparency leads SC holders to run on more moderate signals, but also disciplines the issuer to hold assets that are more tightly distributed around the reference value, which lowers fragility (right panel of Figure 6).

4 Empirical evidence

In this section we take the various predictions from the model to the data by analyzing peg stability and the role of transparency, volatility and large shocks with four case studies. We start by looking at two case studies that speak to implications also found in models with known variance. As discussed in the previous section, this case is nested within our baseline model. The first uses minute-level data on the sudden de-peg suffered by USDC during the turmoil at Silicon Valley Bank (SVB) in March 2023 (Figure 1). For the second event we look into daily data for Tether. In particular, we consider two different analyses: (i) the period of heavy scrutiny over Tether's collateral adequacy that began in October 2018 and its effect on peg stability, and (ii) the effect of two early attestation report releases that carried high informational value. We then move to two case studies that speak to implications only found in our baseline model environment. The third case study goes back to the SVB collapse in March 2023 but looks at hourly data on Frax, a decentralized stablecoin that is backed in part by USDC and lost its peg in the wake of USDC's de-peg event due to its USDC collateral exposure. Finally, we also examine the stability of the now defunct TerraUSD stablecoin and how it related to the volatility of its collateral asset Luna, both in the presence of small and large shocks.

4.1 Identification strategy

A simple pre-post analysis of changes in stablecoins' peg stability over time faces an important identification challenge. In particular, due to the endogenous nature of stablecoin collateral policy shifts, such an approach may not estimate a causal effect of the event of interest if additional factors that may also drive stablecoin peg stability were changing over the same period.

To overcome this problem, for most of our analyses we construct counterfactual values of peg stability in the respective post-event samples through a synthetic control inspired approach (Abadie and Gardeazabal, 2003; Doudchenko and Imbens, 2016). Like in the

synthetic control literature, we estimate pre-post changes in "synthetic" counterfactual peg stability measures that can then be compared to the actual pre-post change in peg stability in the spirit of a difference-in-difference exercise. But unlike difference-in-differences, which requires multiple treated and control units, synthetic control techniques are designed for settings with just one treated unit such as ours.

Our approach, however, differs from traditional synthetic controls in two important respects. First, given the limited number of liquid stablecoins in the crypto ecosystem, the number of control units (i.e. other stablecoins) that can be used to estimate the counterfactual is very small – depending on the period considered, even zero. We therefore rely on data beyond stablecoins that are likely to be important determinants of stablecoin variability to estimate our counterfactuals (see discussion below).[33] Second, one advantage of our approach is that the relatively small number of covariates used to construct counterfactual predictions implies that the risk of overfitting the counterfactual estimate is much lower compared to a traditional synthetic control exercise, where the number of control units is large and often exceeds the number of observations.

The outcome measure of interest across our empirical analyses is the stability of a stablecoin's dollar peg, obtained from secondary market data. Peg deviations capture the strength of the pressure on a stablecoin's peg, and act as a broad proxy for x_c – the mass of investors that have a best response to demand conversion for a given dollar value of reserve assets, θ.[34] The positive association we posit between secondary market prices and run risk is further substantiated in Appendix B where we introduce a two-layered market into our theoretical model. As in Lyons and Viswanath-Natraj (2023), daily stablecoin peg stability is measured in absolute price deviations from \$1:

$$d_t = |1 - p_t^s|, \tag{5}$$

where d_t is the minute, hourly, or daily price deviation and p_t^s is the closing price of the stablecoin corresponding to the same frequency.

[33] In this sense our approach can be viewed as a hybrid between traditional synthetic control event studies and factor models, as it takes pre-treatment peg stability outcomes as benchmarks when choosing weights for control units and uses correlations between treated and control units to predict treated counterfactuals (Xu, 2017; Chen, 2023).

[34] Since x_c and θ_c are affiliated, an increase in θ_c corresponds with an increase in x_c and vice versa. For previous work that uses measures of investor flightiness to test global game predictions, see Prati and Sbracia (2002) and Metz (2003).

To construct the synthetic control we first estimate a dynamic regression of peg stability as a function of its lagged value and additional variables that shape peg stability over the pre-event sample period:

$$d_t = \alpha + \phi d_{t-1} + \boldsymbol{\beta} \mathbf{X}_t + e_t, \tag{6}$$

where d_t is the absolute peg deviation defined in (5), α captures differential averages between d_t and the estimated counterfactual $\hat{d}_t$, and $\mathbf{X}_t$ includes control units or covariates that could impact stablecoins.

We consider covariates that capture both cryptoasset market conditions and broader financial conditions. To proxy for crypto market conditions we use intraday range-based volatilities of the two largest cryptoassets (Bitcoin and Ether).[35] To proxy for broader financial market conditions we rely on indicators from more conventional asset classes.[36] These include the VIX index (as a proxy for broad financial market risk appetite), the MOVE index for interest rate volatility (which shapes switching costs between money-like instruments potentially serving as stablecoin substitutes), option-implied volatility of gold (which serves as a gauge of uncertainty as well as of the riskiness of gold-backed stablecoins), and option-implied volatility of the US dollar-euro exchange rate (capturing international US dollar market conditions and also serving as a global risk barometer).

A threat to identification in our setting arises from the possibility that control units in $\mathbf{X}_t$ are affected by the treated outcome or the policy change specific to the outcome variable. This is also known as a violation of the Stable Unit Treatment Value Assumption (SUTVA). We deal with this threat by carefully selecting variables in $\mathbf{X}_t$ that are unlikely to be impacted by policy changes in the stablecoins studied. Take for instance the case studies examining USDC and Frax. In light of their relatively small market capitalization relative to the much larger BTC and ETH markets, it is unlikely that changing dynamics in the former can affect the latter (let alone conventional financial markets). In the case of Tether, while it features a considerably larger market capitalization ($80 billion in 2022)[37], it is still substantially smaller than ETH and BTC (around $500 billion and $1 trillion over

[35] These range volatility estimates are defined as: $rv_t^k = \ln p_{t,high}^k - \ln p_{t,low}^k$, where rv_t^k is the time t range volatility of cryptoasset $k \in \{BTC, ETH\}$ measured as the log-difference of the day's high and low prices given by $p_{t,high}^k$ and $p_{t,low}^k$, respectively.

[36] The main results are largely unchanged when excluding these indicators, suggesting that the risk of overfitting is relatively low.

[37] This excludes very recent records, where it surpassed $100 billion and which fall outside our sample.

the same period) and dwarfed by traditional financial markets. Finally, given the large size difference between Tether and Frax/USDC and the outsized role played in stablecoin market developments by Tether, we assume that Tether's peg stability can impact that of Frax or USDC *but not* vice versa. We thus include Tether's absolute peg deviations in $\mathbf{X}_t$ when studying other stablecoins' peg stability, but not the reverse.

The pre-event sample counterfactual estimate is given simply by the fitted value recovered from (6), $\hat{d}_{t,pre}$. The post-event counterfactual estimate, $\hat{d}_{t,post}$, is in turn given by:

$$\hat{d}_{t,post} = \hat{\alpha} + \hat{\phi} E[d_{t-1,post} | d_{T,pre}, \mathbf{X}_{t-1,post}] + \widehat{\boldsymbol{\beta}} \mathbf{X}_{t,post} \tag{7}$$

where $\hat{\alpha}$, $\hat{\phi}$, and $\widehat{\boldsymbol{\beta}}$ are estimates of α, ϕ, and $\boldsymbol{\beta}$ from the pre-event period, respectively. We cannot directly incorporate post-event lagged peg deviations (d_{t-1}) to construct the counterfactual because those values are "treated". Instead, we recursively estimate the expected value of $d_{t-1,post}$ conditional on the last value of absolute peg deviations from the pre-event sample ($d_{T,pre}$) and lagged (post-event) values of the covariates ($\mathbf{X}_{t-1,post}$). The full counterfactual series of absolute peg deviations that we use to compare against realized peg deviations is then constructed by joining the pre and post event estimates, $[\hat{d}_{t,pre}, \hat{d}_{t,post}]$. This counterfactual path allows us to estimate an average treatment effect in the post sample period using a t-test for the difference in means between $\hat{d}_{t,post}$ and $d_{t,post}$, since the pre-event mean-differences average out to zero by design.

4.2 Data

The empirical analyses rely on minute, hourly and daily frequency exchange-level cryptoasset data. Minute and hourly price data for various stablecoins and other cryptoassets are sourced from cryptocompare.com. Daily prices are taken from a variety of sources. Ether and Bitcoin data are from the Bitfinex exchange and Coingecko, whereas daily Tether prices are from the Kraken exchange and Coingecko. These exchange-level price data are also sourced from cryptocompare.com and the choice of exchange is based on coin-level exchange data quality and transparency.[38] We also make use of daily market capitalization statistics for Tether, TerraUSD and Luna from Coingecko. Data on conventional financial market variables are from FRED.[39]

[38] Figures 17 and 18 in Appendix C present price series for the main cryptoassets considered.

[39] The interest rate, gold, and US dollar-Euro implied volatility indices are derived using the traditional VIX formulation on US Treasury bond, gold ETF, and Euro ETF options, respectively. They are identified by the

4.3 USDC loses peg during the SVB crisis

Our first case study uses minute-by-minute data around the large de-pegging event suffered by USDC in March 2023. Stress in the regional bank SVB began to build up in early March. The situation took a turn for the worse on Thursday March 9th when the share price of SVB fell sharply after the company announced it planned to raise additional capital by issuing stock. Some venture capital firms advised startups (key customers of SVB) to withdraw their money. A run ensued and the FDIC announced one day later it was putting SVB into receivership. Against this background, and after an attempt to withdraw funds did not prosper, Circle (issuer of USDC) disclosed it held $3.3 billion of its cash reserve at SVB.[40] Almost immediately, USDC fell dramatically (Figure 1, first vertical dashed line). This episode is a clear illustration of a negative shock due to increased *transparency*, which led to a shift in the aggregate behavior of coin holders.

Our model predicts peg instability against such increased transparency with weak priors. To assess this, we fit the dynamic regression specification described in equation (6), where USDC's absolute peg deviations are modeled as a function of their own lag and the set of controls discussed above for which minute-frequency data are available: Tether's absolute peg deviations, Bitcoin volatility and Ether volatility. We set a window of roughly 72 hours before and after the disclosure. Specifically, the pre-treatment estimation period uses data sampled at the minute frequency from March 7 to March 10, 2023 just prior to the disclosure. The post-treatment sample period then spans through March 13th.

Results are in line with the predictions of the model. Figure 7 plots realized peg deviations (dark blue) against counterfactual peg deviations (red), before and after Circle's disclosure. Before the disclosure, the counterfactual series is indistinguishable from the actual series, both sitting at $1. USDC's peg deviations increased markedly following Circle's disclosure, whereas the counterfactual series remained flat.

Table 2 reports t-statistics on average differences in absolute peg deviations over pre- and post-event periods. Following the disclosure by Circle, average daily peg deviations grew to $0.03 compared to a pre-event average of effectively zero and a counterfactual post-event average of zero. The difference in averages in the post-treatment period between actual and

ticker codes MOVE, GVZ and EVZ.

[40] USDC had already slightly de-pegged some hours before, as the fact that they had a relationship with SVB was public information, even if not widely known. However, the extent of the exposure was not known until Circle's disclosure. Through the lens of our model, this captures an environment in which market priors, y, were likely weak, leading an increase in public information precision to aggravate investor flight and run risk.

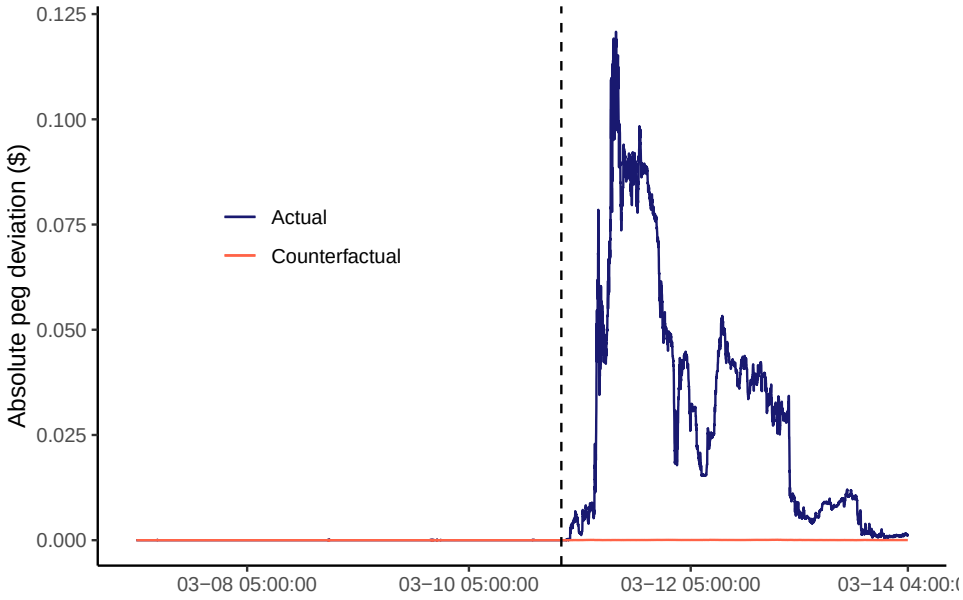

Figure 7: USDC loses peg around SVB crisis.

Notes: The figure plots actual absolute peg deviations of USDC (dark blue) and the counterfactual path (red) estimated from Equations 6 and 7 using minute-frequency data on USDC absolute peg deviations, USDT absolute peg deviations, BTC range volatility and ETH range volatility. The dashed vertical line represents the disclosure by Circle (issuer of USDC) that it held $3.3 billion of its reserve at SVB.

counterfactual peg deviations is highly statistically significant.

Mean absolute peg deviations	Pre event 03-07-2023 to 03-10-2023	Post event 03-11-2023 to 03-13-2023
Actual	0.000	0.030
Counterfactual	0.000	0.000
Difference	0.000	0.030***
t-statistic	0.000	70.032
Observations	5520	4499

Table 2: Average effect (in $) on USDC peg deviations after Circle discloses deposits held with Silicon Valley Bank, minute-by-minute before and after March 10, 2023. Significance at the 10%, 5%, and 1% level given by '*', '**' and '***', respectively.

4.4 Tether's reserve adequacy and attestations

Our second set of event studies examines the largest stablecoin, Tether. In this case, we present two distinct pieces of evidence to support the implications of the model. First, we look at peg deviations during a time, in the early life of Tether, when there was no public

information and doubts mounted on its reserve adequacy. Second, we move forward in time and zoom in on the only instance when two reserve attestation reports were released in two consecutive months (early 2021), which also happened after a long spell (two and half years) without any information on Tether's reserves. Arguably, the information content of these two consecutive releases was among the largest for all attestation reports.

Doubts about Tether's reserve adequacy. We look at the period around October 2018, when Tether experienced temporary peg instability connected to worries over its backing. This occurred in the aftermath of a June 2018 audit report that was completed by a law firm, rather than a certified accounting firm.[41] Concerns over whether Tether was fully backed by US dollars (as claimed by the issuer, without any supporting evidence) continued to fester and worries also rose regarding Tether's relationship with Bitfinex and the possible co-mingling of customer funds. These concerns eventually led to a sharp de-pegging on October 15, 2018, followed by another one in April 2019 when the NY Attorney General announced its lawsuit against iFinex, the parent company of Tether and Bitfinex.[42]

In this event the reserve quality of the stablecoin was called into question, hence *a priori* we expect to observe greater peg instability. We consider September 30, 2018 as our event date as it stands days before Bitfinex publicly responded to rumors of insolvency, followed by an October 11 announcement it was temporarily shutting down fiat deposits in the face of payment processing complications, triggering severe concerns over Tether's reserve adequacy.[43][44] We interpret such increased worries over Tether's backing during this period as weakening priors about Tether's reserves. The pre-event sample spans February 15, 2018 to September 30, 2018 and the post-event sample spans October 1, 2018 to June 30, 2019. Tether's market capitalization surpassed $3 billion before the event and by October 2018

[41] The law firm, Free, Sporkin and Sullivan, added a disclaimer that, "FSS is not an accounting firm and did not perform the above review and confirmations using Generally Accepted Accounting Principles," and, "The above confirmation of bank and Tether balances should not be construed as the results of an audit and were not conducted in accordance with Generally Accepted Auditing Standards."

[42] During the case, Tether's lawyers stated that as of April 30, just 74% of Tether was backed by US dollar assets, contrary to repeated previous statements by the issuer. The case reached a settlement in 2021.

[43] See "A Response to Recent Online Rumors" published by Bitfinex on October 7, "Bitcoin Jumps after Credit Scare; Fidelity Enters Crypto Sphere", published by MarketWatch on October 15, 2018, and "Crypto Markets Roiled as Traders Question Tether's Dollar Peg", published by Bloomberg on October 15, 2018. In February 2021, The New York State Attorney General reported in their investigation of Tether that the October 2018 shutdown caused Bitfinex to suffer a "massive and undisclosed loss of funds".

[44] The choice of treatment date is not clear-cut in this exercise because of the ongoing scrutiny over Tether's reserve adequacy over this period. However, the results discussed below are robust to selecting alternative treatment dates. For example, the results also obtain if the treatment date is chosen to be that of the first audit report released in June 2018.

Mean absolute peg deviations	Pre event 02-15-2018 to 09-30-2018	Post event 10-01-2018 to 06-30-2019
Actual	0.011	0.009
Counterfactual	0.011	0.001
Difference	0.000	0.008***
t-statistic	0.101	15.042
Observations	228	274

Table 3: Average effect (in $) on Tether peg deviations after rising concerns over reserve adequacy, before and after September 30, 2018. Significance at the 10%, 5%, and 1% level given by '*', '**' and '***', respectively.

it had fallen by $1 billion, suggesting that these concerns were highly consequential for the stablecoin (left-hand panel of Figure 8).

We employ a dynamic regression specification to estimate the synthetic counterfactual path of Tether's absolute peg deviations before and after the event date. Specifically, we regress absolute peg deviations of Tether on its lagged values, range volatilities of BTC and ETH, and the conventional financial market variables discussed above. Because Tether is the largest stablecoin, we do not include variables based on other smaller stablecoins to mitigate potential endogeneity issues.

The effects of the concerns regarding reserve quality are very clearly visible in terms of (lack of) peg stability in the right-hand panel of Figure 8. This panel plots Tether's actual absolute peg deviations over the pre and post event periods in dark blue and its corresponding counterfactual path in red (as in the left-hand panel, the vertical line marks September 30, 2018, the beginning of rising concerns about Tether's reserve adequacy). Actual peg deviations spike in the wake of the event and remain considerably large thereafter, whereas the counterfactual remains stable. We test these differences more formally in Table 3, which reports mean absolute peg deviations over the pre and post event samples along with those of the counterfactual estimate. Tether's realized peg deviations of $0.009 were significantly larger in the post-event sample compared to the counterfactual mean deviation of $0.001 (in other words, realized peg deviations were about 9 times, or 800% larger than in the counterfactual).

Tether's peg around attestation release dates. The adequacy and quality of Tether's reserves have been historically clouded by mystery, in no small part because the stablecoin operated outside the purview of regulators. One way the operating firm aimed

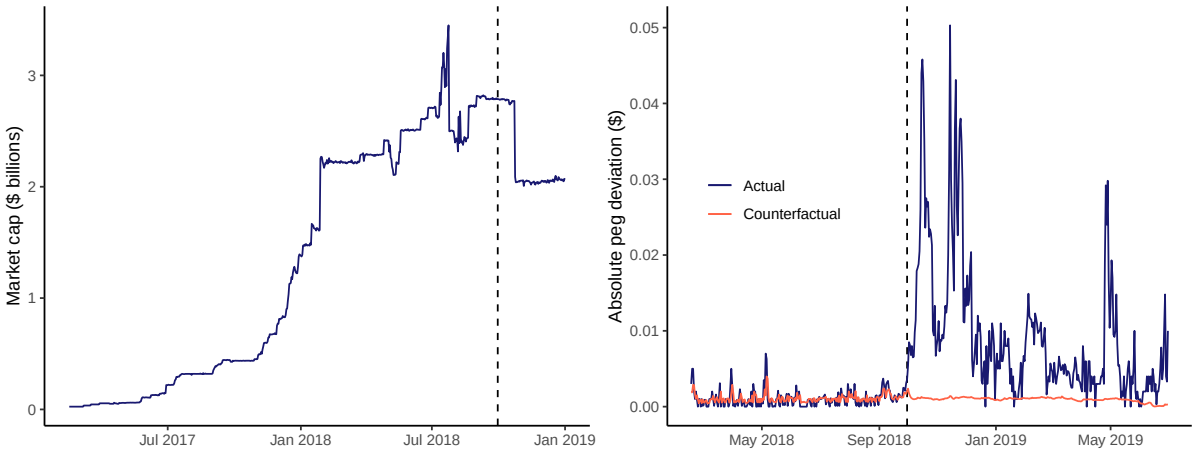

Figure 8: Doubts about Tether's reserve adequacy.

Notes: The left panel traces the market capitalization of Tether. The right panel plots actual absolute peg deviations of Tether (dark blue) against the counterfactual path (red) estimated from Equations 6 and 7. The dashed vertical line is September 30, 2018, the beginning of concerns over Tether's reserve adequacy.

to gain credibility was through regular reserve attestations, or periodic audits. The earliest available attestation report of Tether's reserves is dated June 1, 2018 by firm Free, Sporkin, and Sullivan LLP (see footnote 41 and preceding discussion). However, as the case study just discussed shows, FSS made it publicly clear that they are not a certified accounting firm.

Tether resumed the publication of attestation reports in early 2021, after two and half years without any information. Specifically, there was a gap in audit releases between October 2018 and February 2021. Tether's peg remained relatively stable over this period, which we interpret as relatively strong priors held by its holders despite the lack of reserve transparency, likely due to buoyant market conditions in the latter part of the period. This dry spell was broken when Moore (then Tether's auditor) published attestations of Tether's reserves on February 28, 2021 and then again on March 30, 2021 – the only period when Tether released audits in two consecutive months.[45] We take advantage of this period to examine whether Tether's peg stability was impacted by the first audit release after several years. We estimate the same synthetic control specification as in the preceding exercise, although now the pre-treatment period is from October 30, 2020 to February 28, 2021 and the post-treatment period covers March 1, 2021 through April 30, 2021.

[45]Tether continued to periodically publish reports thereafter, changing auditors multiple times over the early years of reporting. As a result, questions over reserve adequacy persist, even to this day.

Figure 9 and Table 4 present the results. The first and second vertical dashed lines in Figure 9 respectively denote February 28 and March 31 attestation report releases. Counterfactual peg deviations were larger than actual peg deviations following the February 28 release, and the difference between them grew even larger following the March 31 release. Information disclosure about reserve quality following a long period without any such information contributed to reducing the run risk perceived by investors, as manifested in the volatility of Tether's peg. Table 4 reports the average reduction in peg deviations observed in Figure 9 following the February 28, 2021 audit release. The difference between counterfactual and actual series is statistically significant at the 1% level.

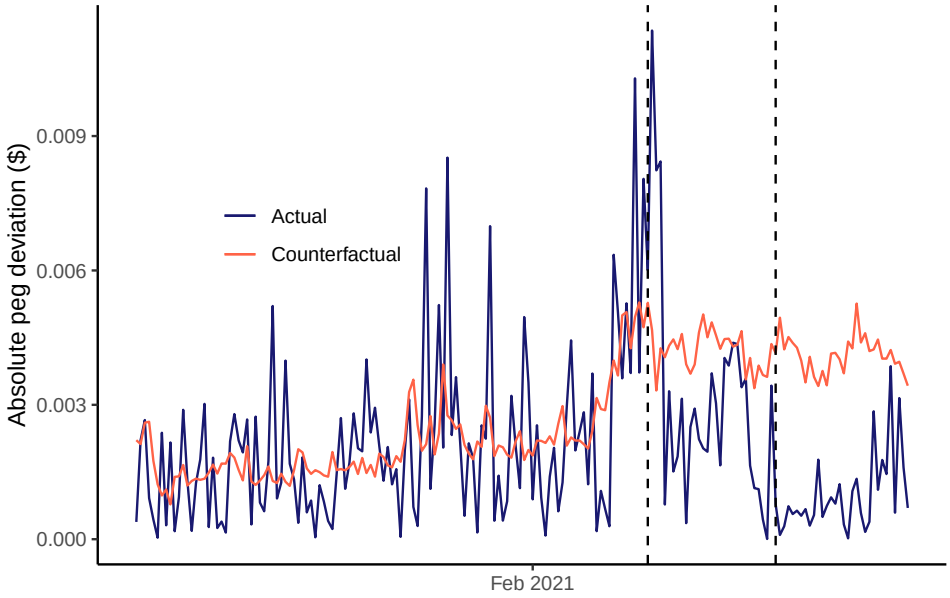

Figure 9: Tether absolute peg deviations before and after the February 28, 2021 audit release.
Notes: The figure plots the actual absolute peg deviations of Tether (dark blue) against the counterfactual path (red) estimated from Equations 6 and 7. The first dashed vertical line is February 28, 2021, the date of the first audit release of Tether's reserves since 2018. The second dashed vertical line is March 30, 2021, the date of the subsequent audit release. These were the only reports released in consecutive months.

4.5 Frax's peg around the SVB crisis

The third case study goes back to events around the collapse of SVB, but focuses on the behavior of Frax. Unlike USDC or Tether, Frax is a smaller crypto-collateralized stablecoin that was originally designed to be partially backed by cryptocurrency collateral and partially

Mean absolute peg deviations	Pre event 10-30-2020 to 02-28-2021	Post event 03-01-2021 to 04-30-2021
Actual	0.002	0.002
Counterfactual	0.002	0.004
Difference	0.000	-0.002***
t-statistic	0.000	-7.974
Observations	121	61

Table 4: Average effect (in $) on Tether peg deviations after Tether's February 28, 2021 audit attestation release. Significance at the 10%, 5%, and 1% level given by '*', '**' and '***', respectively.

algorithmic, with the collateral ratio varying over time based on demand for the stablecoin. However, in February 2023 the Frax community voted to launch an update to the stablecoin protocol with the goal of eliminating the algorithmic component and hence becoming a fully crypto-collateralized stablecoin.

Frax makes for a particularly unique case study around the SVB event because it is a stablecoin that is partially backed by USDC. Therefore the depegging of USDC around SVB can be seen as an unanticipated collateral volatility shock for Frax, potentially impacting its own peg stability. Moreover, Frax does not feature confounding mechanisms such as those present for Dai, which is also backed by USDC and which also de-pegged around SVB (Figure 1).[46]

To assess the impact of USDC collateral volatility on Frax peg stability, we follow a similar approach as above. In particular, we fit the dynamic regression specification described in equation (6), where Frax's hourly absolute peg deviations are modeled as a function of their own lag and the set of controls discussed above, for which hourly data are available: Tether's absolute peg deviations, Bitcoin volatility and Ether volatility. Like in the USDC case studied above, the treatment is the disclosure by Circle of the amount of reserves held at SVB. The pre-treatment estimation period uses data sampled at the hourly frequency from March 8 to March 10, 2023, just prior to the disclosure. The post-treatment sample period spans through March 13th.

Figure 10 shows actual peg deviations (dark blue) following USDC's disclosure versus

[46] As shown in Figure 16, de-pegging was common to stablecoins backed by USDC, but not to those backed by other (even otherwise more volatile) crypto collateral. Other crypto-collateralized stablecoins such as Dai also hold USDC as collateral. But Dai also has a price stability mechanism (PSM) in effect with USDC, allowing for 1-for-1 convertibility. As a result, it is difficult to attribute Dai's peg instability around SVB directly to collateral volatility, as the presence of a PSM ties the price of USDC and Dai together.

Mean absolute peg deviations	Pre event 03-08-2023 to 03-10-2023	Post event 03-10-2023 to 03-13-2023
Actual	0.004	0.019
Counterfactual	0.010	0.004
Difference	0.000	0.010***
t-statistic	0.000	4.916
Observations	71	60

Table 5: Average effect (in $) on Frax peg deviations after Circle discloses deposits held with Silicon Valley Bank, hour-by-hour before and after March 10, 2023. Significance at the 10%, 5%, and 1% level given by '*', '**' and '***', respectively.

counterfactual peg deviations (red). It is quite visible how Frax's actual peg deviations rose significantly following the disclosure, while counterfactual peg deviations did not rise by nearly as much. Table 10 shows that Frax's average absolute peg deviations grew significantly following Circle's disclosure. Moreover, the difference between actual and counterfactual average absolute peg deviations is statistically significant at the 1% level.

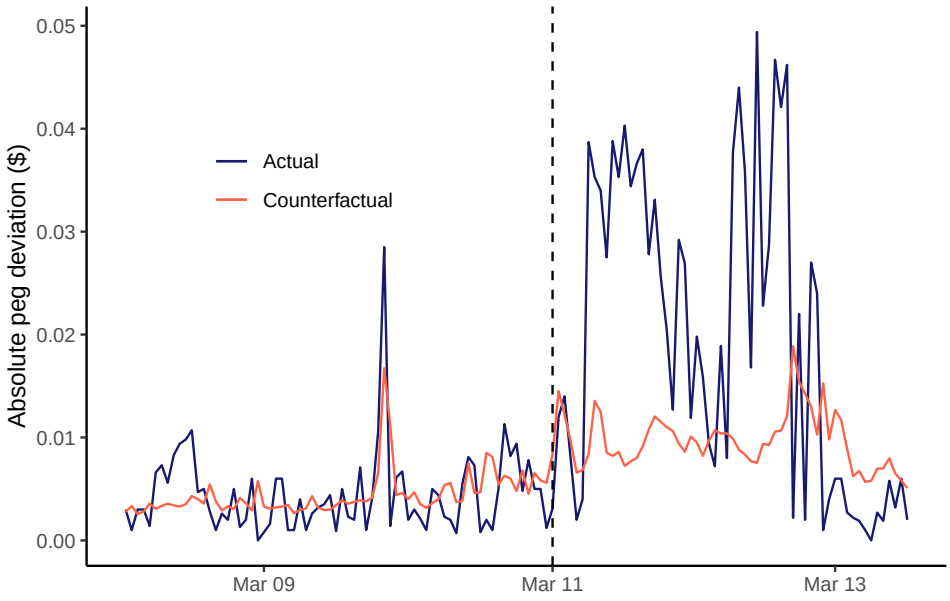

Figure 10: Frax loses peg around SVB crisis.

Notes: The figure plots the actual absolute peg deviations of Frax (dark blue) against the counterfactual path (red) estimated from Equations 6 and 7 using hourly data on Frax absolute peg deviations, USDT absolute peg deviations, BTC range volatility and ETH range volatility. The dashed vertical line represents the disclosure by Circle (issuer of USDC) that it held $3.3 billion of its reserve at SVB.

4.6 Reserve volatility and peg stability: TerraUSD and Luna

Lastly, we also test our model predictions using data on TerraUSD, the now defunct algorithmic stablecoin. When the variance of the reserve asset returns is unknown, beliefs about it change over time. In the context of the model, large reserve asset shocks are destabilizing when the stablecoin lies within a "ripe for run" region but not outside this region, whereas small shocks are not destabilizing.

Luna's market capitalization was indeed volatile. The left panel in Figure 11 presents evidence of this volatility, while highlighting market capitalization still generally remained well above the value of Terra's outstanding stablecoin liabilities, except for two periods: May 2021 and May 2022 (when TerraUSD permanently de-pegged). During both periods (denoted by grey vertical areas), the market capitalization of Luna approached that of TerraUSD. We define this difference as Terra's "equity value".[47] For example, the 2021 episode when Terra equity values quickly approached zero was caused by the bursting of the cryptoasset bubble, with the price of Bitcoin falling 40% in a matter of days, bringing Luna's price down with it. The notion of Terra's equity value is compelling because if Luna's market capitalization falls below that of TerraUSD, then there is insufficient Luna available to be sold to cover all TerraUSD liabilities, and therefore Terra as a stablecoin issuer is more likely to become insolvent in the eyes of TerraUSD holders (Liu et al., 2023).

As the variance of the reserve asset is unknown, we assume that beliefs over it are formed using observable data on changes in reserve asset prices. As a result, TerraUSD holders' best guess of the variance of the reserve asset Luna is likely a function of Luna's price return history, and this estimate can change over time with the arrival of new information. Validating the assumption that variance is not known and non-constant, the right-panel of Figure 11 shows strong evidence of time-varying variance for Luna. Specifically, we use a GARCH(1,1) estimate of the conditional standard deviation of daily returns (i.e. volatility) of Luna through May 9, 2022, right before the permanent de-peg when volatility exploded to over 400%, resulting in a time-series plot that is quite literally off the chart. The non-constant conditional volatility of Luna generates fat tails in the unconditional distribution of the reserve asset returns, consistent with the statistical properties of Luna's price returns.

Our model also implies that stablecoins are in a "ripe for run" region they are susceptible to large shocks toppling them. Here the concept of Terra equity defined above is useful, as

[47] Concretely, we define it as follows: *Terra Equity = Luna Market Cap - TerraUSD Market Cap.*

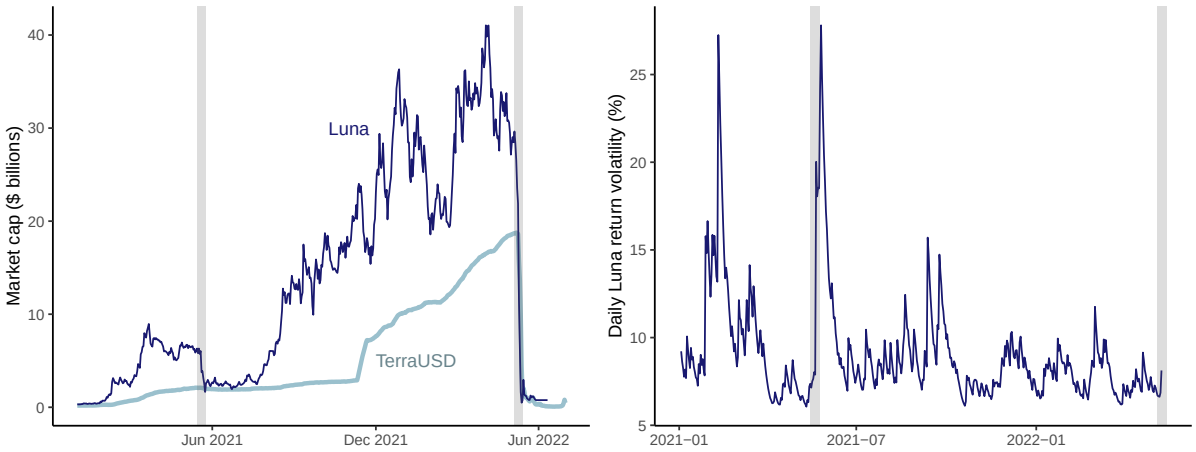

Figure 11: TerraUSD market capitalization, Luna market capitalization and volatility.

Notes: The left-panel plots the market capitalization of TerraUSD (light blue, thick) and Luna (dark blue, thin), including the period of the terminal depegging event. We define the difference between Luna and TerraUSD market capitalization as Terra equity. The right-hand plots the daily return volatility of Luna, estimated under a GARCH(1,1) specification through May 9, 2022 (pre de-pegging event). By May 11 Luna estimated daily volatility reached 50% and exceeded 400% by May 14 (not shown for better visibility). The shaded regions indicate May 15-25, 2021 and May 5-15, 2022, respectively.

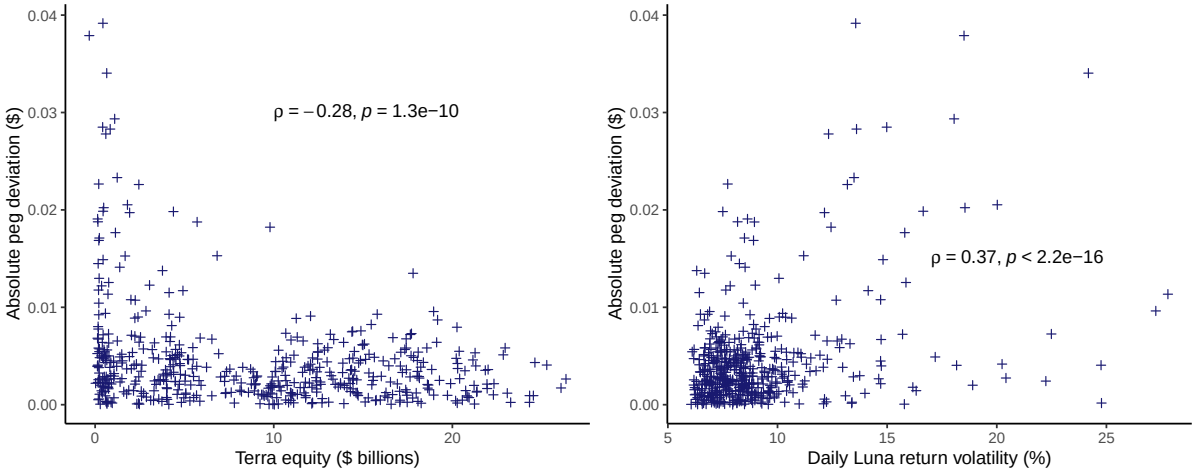

Figure 12: Terra equity, volatility and peg deviations.

Notes: The left panel plots Terra equity (defined as market capitalization of Luna minus that of TerraUSD) versus absolute peg deviations of TerraUSD. The right panel plots the daily return volatility of Luna estimated under a GARCH(1,1) specification versus absolute peg deviations of TerraUSD. Daily data from January 1, 2021 to May 9, 2022.

it can be seen as a measure that determines vulnerability to a run. In other words, the run region is more likely to be approached as Terra's equity value approaches zero. Indeed, as Terra's equity value falls, absolute deviations of TerraUSD's dollar peg increases and are largest when Terra's equity value approaches zero (left panel of Figure 12). Without

explicitly taking a stance on defining a "ripe for run" threshold, it is possible to simplify to a continuous setting where one can think of Terra's run risk, quantified using absolute peg deviations, as decreasing in Terra's equity value. Luna volatility, in turn, is positively associated with TerraUSD's absolute peg deviations, as shown in the right panel of Figure 12. In other words, when the range of possible Luna reserve asset returns increases, so do absolute peg deviations of TerraUSD. We interpret large (small) values of Luna return volatility as realizations of large (small) reserve asset shocks.

Taking these stylized facts together helps motivate a simple regression analysis to test model implications. We test the following two predictions: (i) reserve asset shocks are positively correlated with TerraUSD peg deviations, i.e. large shocks are more destabilizing than small shocks; and (ii) the effect of reserve asset shocks on TerraUSD peg stability are stronger within the "ripe for run" region, i.e. when Terra's equity value is smaller. A parsimonious regression model to test these predictions can be set up as follows:

$$d_t = \alpha + \phi d_{t-1} + \beta_1 vol_{t-1} + \beta_2 [vol_{t-1} \times equity_{t-1}] + e_t, \tag{8}$$

where variables on the right-hand side are lagged to help reduce the risk of endogeneity arising from simultaneity between TerraUSD and Luna. TerraUSD's absolute peg deviations are given by d_t, Luna's GARCH(1,1) conditional volatility is given by vol_t and Terra's equity value is given by $equity_t$. Daily data from January 1, 2021 to May 9, 2022 are used, so our estimates exclude the final de-peg event of May 2022.[48]

The interpretation of coefficients is as follows. A positive estimate of β_1 suggests that TerraUSD's peg deviations are larger when reserve asset volatility is higher (larger shocks). A negative estimate of β_2 in turn indicates that for any given level of Luna volatility, its impact on TerraUSD's peg stability is weaker when Terra's equity value is larger. Our specification is motivated by the interpolated surface plot shown in Figure 13 which uses local linear smoothing to fit a surface relating Terra's equity and Luna's volatility to TerraUSD's peg deviations. Empirically, it can be seen that even prior to the permanent de-peg that occurred in May 2022, peg deviations were largest when both Luna's volatility was high and Terra's equity value was low.

Regression results support the two predictions (Table 6). Indeed as suggested by Figures

[48] Results are not sensitive to the choice of using logged or non-logged volatility, nor the choice to include or exclude a lagged dependent variable.

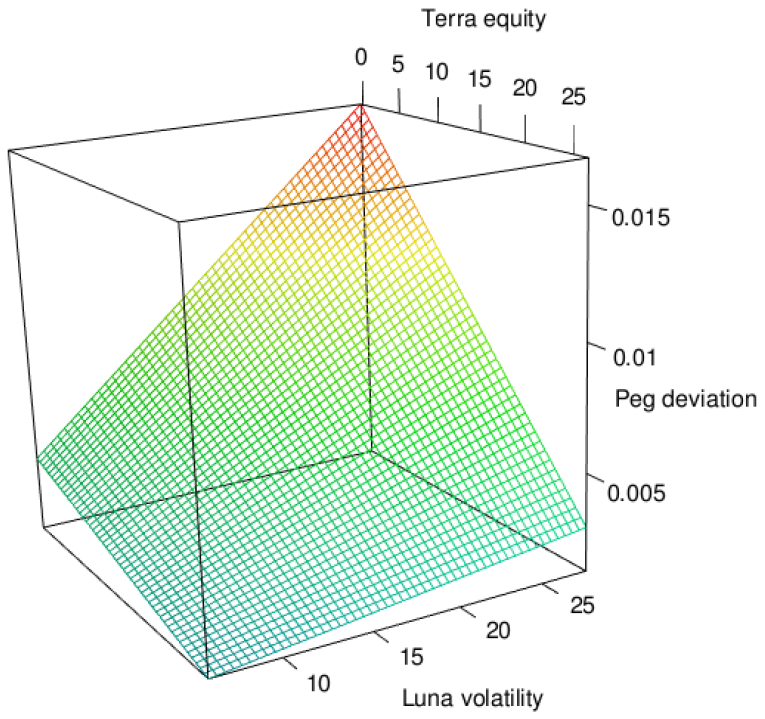

Figure 13: Terra equity, Luna volatility and TerraUSD absolute peg deviations.

Notes: Terra equity is defined as the market capitalization of Luna minus the market capitalization of TerraUSD. Luna volatility corresponds to the daily conditional return volatility estimated from a GARCH(1,1) model. The surface is interpolated via a locally estimated scatter plot smoothing regression with degree of 1 and span of 0.95.

12 and 13, as reserve asset volatility rises, so do absolute peg deviations. However, the impact of higher reserve asset volatility on peg deviations depends significantly on the equity value of Terra. When equity is low (i.e. when approaching the "ripe for run" region), the same level of Luna volatility has a substantially larger impact on TerraUSD's peg stability.

Covariate	Estimate	SE
Intercept	0.0013	(0.0007)
d_{t-1}	0.4978***	(0.0403)
vol_{t-1}	0.0186*	(0.0067)
$vol_{t-1} \times equity_{t-1}$	-0.0011***	(0.0003)

Table 6: Regression estimates from Equation 8 where the dependent variable is d_t, TerraUSD absolute peg deviations. Significance at the 10%, 5%, and 1% level given by '*', '**' and '***', respectively. Standard errors adjusted for heteroskedasticity and autocorrelation. Estimates and standard errors (SE) on vol_{t-1} and $[vol_{t-1} \times equity_{t-1}]$ are multiplied by 100. The sample runs from from January 1, 2021 to May 9, 2022, and includes 492 observations.

To put results in perspective, let us consider a scenario where Luna daily volatility rises

to from 5% to 15%. Under an equity value of $15 billion, TerraUSD's peg deviations would increase from $0.00013 to $0.0004.[49] However, under an equity value of just $1 billion, TerraUSD peg deviations would rise from $0.0009 to $0.0026, a 29-fold increase. On May 12, 2022, TerraUSD broke its peg and fell from $1 to roughly $0.78 and Terra's equity value was wiped out. By May 14, TerraUSD crashed to roughly $0.12 and Luna volatility rose sharply, exceeding 400%. Our simple linear model estimated on data before the final de-peg qualitatively captures the peg stability risk of Terra that eventually became realized, although unsurprisingly non-linearities would likely need to be considered to quantitatively match the nature of Terra's final de-pegging event.

5 Conclusion

Stablecoins were designed to provide a stable unit of account within the crypto ecosystem. Yet despite the various strategies used to defend their promise of par convertibility to the sovereign unit of account, that promise was broken on multiple occasions, irrespective of the type of reserves held. Besides the notorious failure of the algorithmic stablecoin TerraUSD, the March 2023 banking crisis simultaneously highlighted the key role of reserve transparency and volatility.

In this paper, we analyze the various ways in which information (in the form of public broadcasts and learning from past observations) shapes the risk of coordination failure by stablecoin holders through their beliefs about peg stability. Using global games to model the strategic interactions among stablecoin holders and issuers, we argue that the characteristics of the reserves held determine the type of game being played. We show that the effect of public disclosure and large shocks on run risk is ambiguous. Greater transparency can lead to greater run risk when market expectations are pessimistic; conversely, transparency strengthens a stablecoin peg when priors are strong. We also predict inertia in aggregate behavior by stablecoin holders. Individuals pay close attention to past events, and small public shocks to fundamentals induce a reversion to historical outcomes, while large negative (positive) shocks trigger wide-spread runs (recoveries). Our assessment of the effect of recent publicized changes to collateral holdings by prominent stablecoin issuers on stablecoin price stability provides strong support to the model's predictions.

[49] TerraUSD absolute peg deviations, d_t, have a standard deviation of $0.0052 from January 1, 2021 to May 9, 2022.

Our work has implications for policy. The model and empirical findings indicate that stablecoins are subject to run risk and can (and do) de-peg, failing on their promise of par convertibility. These stability risks are connected to reserve adequacy, and perceptions thereof – inadequate or illiquid for some, non-existent for others such as algorithmic coins. As a result, our work speaks to ongoing debates on whether stablecoins should hold loss absorbing resources to buffer against these risks. Banks and money market funds also issue dollar-like liabilities, but unlike them, stablecoin issuers do not have access to public liquidity backstops – likely for good reason (Aldasoro et al., 2023). The demise of SVB demonstrates how even stablecoin issuers with *ex-ante* fully liquid reserves are not really risk-free (as they are exposed to credit and liquidity risk) and are subject to the type of financial stability risks that are ultimately borne by the public in the absence of loss absorbing resources. To be sure, such loss absorbing resources will not in and of themselves solve stablecoin run risk, but will likely mitigate losses for users, especially in the absence of explicit backstops. Finally, our work highlights that reserve asset transparency and quality play distinct roles in shaping peg stability and the probability of issuer failure, pointing to an important distinction between reserve *disclosure* and *quality* requirement that can inform policy design.

Our paper points to interesting avenues for future work. For example, future research could consider the interplay between conflicting public messages about the quality of stablecoin issuer collateral, and how public information interacts with shocks to idiosyncratic noise that, together, shape run risk. The first issue pertains to instances where, for example, public audits reveal information that contradicts disclosures by the issuer. The second issue involves the impact of informational sources such as social media on the idiosyncratic beliefs of market participants. Such work may be facilitated by the continuous development of stablecoins with an increasingly diverse range of reserves and collateral policies. Our model could also be usefully extended into a bank setting to understand whether, and to what extent, banks are well placed to issue stablecoins.

References

Abadie, Alberto and Javier Gardeazabal (2003) "The economic costs of conflict: A case study of the Basque Country," *American Economic Review*, 93 (1), 113–132.

Agur, Itai, Anil Ari, and Giovanni Dell'Ariccia (2022) "Designing central bank digital currencies," *Journal of Monetary Economics*, 125, 62–79, https://doi.org/10.1016/j.jmoneco.2021.05.002.

Ahnert, Toni and Ali Kakhbod (2017) "Information choice and amplification of financial crises," *The Review of Financial Studies*, 30 (6), 2130–2178.

Aldasoro, Iñaki, Perry Mehrling, and Daniel H. Neilson (2023) "On par: A money view of stablecoins," BIS Working Paper 1146.

Anadu, Kenechukwu, Pablo Azar, Catherine Huang et al. (2024) "Runs and flights to safety: Are stablecoins the new money market funds?"Technical report, Federal Reserve Bank of New York Staff Report No. 1073, April.

Angeletos, George-Marios, Christian Hellwig, and Alessandro Pavan (2006) "Signaling in a global game: Coordination and policy traps," *Journal of Political Economy*, 114 (3), 452–484.

——— (2007) "Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks," *Econometrica*, 75 (3), 711–756.

Angeletos, George-Marios and Iván Werning (2006) "Crises and prices: Information aggregation, multiplicity, and volatility," *American Economic Review*, 96 (5), 1720–1736.

Arner, Douglas, Raphael Auer, and Jon Frost (2020) "Stablecoins: Potential, risks and regulation," BIS Working Papers 905, Bank for International Settlements, https://ideas.repec.org/p/bis/biswps/905.html.

Bains, Parma, Arif Ismail, Fabiana Melo, and Nobuyasu Sugimoto (2022) "Regulating the crypto ecosystem: The case of stablecoins and arrangements," International Monetary Fund FinTech Notes, 2022(008), A001.

Barthélémy, Jean, Paul Gardin, and Benoit Nguyen (2023) "Stablecoins and the Financing of the Real Economy," Working papers 908, Banque de France, https://ideas.repec.org/p/bfr/banfra/908.html.

Bebchuk, Lucian A and Itay Goldstein (2011) "Self-fulfilling credit market freezes," *The Review of Financial Studies*, 24 (11), 3519–3555.

Bertsch, Christoph (2023) "Stablecoins: Adoption and Fragility," Working Paper Series 423, Sveriges Riksbank (Central Bank of Sweden), https://ideas.repec.org/p/hhs/rbnkwp/0423.html.

Caramichael, John and Gordon Y. Liao (2022) "Stablecoins: Growth Potential and Impact on Banking," International Finance Discussion Papers 1334, Board of Governors of the Federal Reserve System (U.S.).

Chen, Heng and Wing Suen (2016) "Falling dominoes: A theory of rare events and crisis contagion," *American Economic Journal: Microeconomics*, 8 (1), 228–255.

Chen, Jiafeng (2023) "Synthetic control as online linear regression," *Econometrica*, 91 (2), 465–491.

Corsetti, Giancarlo, Amil Dasgupta, Stephen Morris, and Hyun Song Shin (2000) "Does one Soros make a difference? The role of a large trader in currency crises," *Review of Economic Studies*, 71 (1), 87–113.

d'Avernas, Adrien, Vincent Maurin, and Quentin Vandeweyer (2023) "Can Stablecoins Be Stable?," working paper, (available at SSRN), https://ssrn.com/abstract=4226027.

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

Doudchenko, Nikolay and Guido W. Imbens (2016) "Balancing, Regression, Difference-In-Differences and Synthetic Control Methods: A Synthesis," NBER Working Papers 22791, National Bureau of Economic Research, Inc, https://ideas.repec.org/p/nbr/nberwo/22791.html.

Duley, Chanelle and Prasanna Gai (2023) "Macroeconomic tail risk, currency crises and the inter-war gold standard," *Canadian Journal of Economics/Revue canadienne d'économique*, 56 (4), 1551–1582.

Edmond, Chris (2013) "Information manipulation, coordination, and regime change," *Review of Economic studies*, 80 (4), 1422–1458.

Eichengreen, Barry, My T. Nguyen, and Ganesh Viswanath-Natraj (2023) "Stablecoin Devaluation Risk," WBS Finance Group Research Paper Available at SSRN: https://ssrn.com/abstract=4460515, Warwick Business School.

Goldstein, Itay and Ady Pauzner (2005) "Demand–Deposit Contracts and the Probability of Bank Runs," *Journal of Finance*, 60 (3), 1293–1327.

Gorton, Gary B, Elizabeth C Klee, Chase P Ross, Sharon Y Ross, and Alexandros P Vardoulakis (2022) "Leverage and Stablecoin Pegs,"Technical report, National Bureau of Economic Research.

Gorton, Gary B and Jeffery Zhang (2023) "Taming wildcat stablecoins," *University of Chicago Law Review*, 90.

Klages-Mundt, Ariah and Andreea Minca (2021) "(In)Stability for the Blockchain: Deleveraging Spirals and Stablecoin Attacks," *Cryptoeconomic Systems*, 1 (2).

——— (2022) "While Stability Lasts: A Stochastic Model of Non-Custodial Stablecoins," arxiv:2004.01304v3.

Konrad, Kai A and Tim BM Stolper (2016) "Coordination and the fight against tax havens," *Journal of International Economics*, 103, 96–107.

Kosse, Anneke, Marc Glowka, Ilaria Mattei, and Tara Rice (2023) "Will the real stablecoin please stand up?," BIS Papers No 141, November.

Levy Yeyati, Eduardo and Sebastian Katz (2022) "The stablecoin paradox," voxeu.org, https://cepr.org/voxeu/columns/stablecoin-paradox.

Li, Ye and Simon Mayer (2021) "Money Creation in Decentralized Finance: A Dynamic Model of Stablecoin and Crypto Shadow Banking," CESifo Working Paper Series 9260, CESifo, https://ideas.repec.org/p/ces/ceswps/_9260.html.

Little, Andrew T, Joshua A Tucker, and Tom LaGatta (2015) "Elections, protest, and alternation of power," *The Journal of Politics*, 77 (4), 1142–1156.

Liu, Jiageng, Igor Makarov, and Antoinette Schoar (2023) "Anatomy of a Run: The Terra Luna Crash," NBER Working Papers 31160, National Bureau of Economic Research, Inc, https://ideas.repec.org/p/nbr/nberwo/31160.html.

Lyons, Richard K and Ganesh Viswanath-Natraj (2023) "What keeps stablecoins stable?" *Journal of International Money and Finance*, 131, 102777.

Ma, Yiming, Zhao Yeng, and Anthony Lee Zhang (2023) "Stablecoin runs and the centralization of arbitrage," working paper, (available at SSRN), https://ssrn.com/abstract=4398546.

Makarov, Igor and Antoinette Schoar (2022) "Cryptocurrencies and Decentralized Finance (DeFi)," NBER Working Papers 30006, National Bureau of Economic Research, Inc, https://ideas.repec.org/p/nbr/nberwo/30006.html.

Mell, Peter and Dylan Yaga (2022) "Understanding Stablecoin Technology and Related Security Considerations," NIST Internal Report 8408, National Institute of Standards and Technology (NIST).

Metz, Christina E (2003) *Information Dissemination in Currency Crises*: Springer, New York.

Milgrom, Paul and John Roberts (1990) "Rationalizability, learning, and equilibrium in games with strategic complementarities," *Econometrica: Journal of the Econometric Society*, 1255–1277.

Morris, Stephen and Hyun Song Shin (1998) "Unique equilibrium in a model of self-fulfilling currency attacks," *American Economic Review*, 587–597.

——— (2003) "Global Games: Theory and Applications," in Dewatripont, M., L. Hansen, and S. Turnovsky eds. *Advances in Economics and Econometrics (8th World Congress of the Econometric Society)*: Cambridge, UK: Cambridge University Press.

Morris, Stephen, Hyun Song Shin, and Muhamet Yildiz (2016) "Common belief foundations of global games," *Journal of Economic Theory*, 163, 826–848.

Morris, Stephen and Muhamet Yildiz (2019) "Crises: Equilibrium shifts and large shocks," *American Economic Review*, 109 (8), 2823–54.

Oefele, Nico, Dirk G Baur, and Lee A Smales (2023) "Are Stablecoins the Money Market Mutual Funds of the Future?," Available at SSRN 4550177.

Prati, Alessandro and Massimo Sbracia (2002) "Currency crises and uncertainty about fundamentals,"Technical report, IMF Working paper No. 2002/003.

——— (2010) "Uncertainty and currency crises: Evidence from survey data," *Journal of Monetary Economics*, 57 (6), 668–681.

Rajan, Raghuram G (1994) "Why bank credit policies fluctuate: A theory and some evidence," *the Quarterly Journal of Economics*, 109 (2), 399–441.

Rochet, Jean-Charles and Xavier Vives (2004) "Coordination failures and the lender of last resort: Was Bagehot right after all?," *Journal of the European Economic Association*, 2 (6), 1116–1147.

Romero, Julian (2015) "The effect of hysteresis on equilibrium selection in coordination games," *Journal of Economic Behavior & Organization*, 111, 88–105.

Szkup, Michal and Isabel Trevino (2015) "Information acquisition in global games of regime change," *Journal of Economic Theory*, 160, 387–428.

Uhlig, Harald (2022) "A Luna-tic Stablecoin Crash," NBER Working Papers 30256, National Bureau of Economic Research, Inc, https://ideas.repec.org/p/nbr/nberwo/30256.html.

Van Zandt, Timothy and Xavier Vives (2007) "Monotone equilibria in Bayesian games of strategic complementarities," *Journal of Economic Theory*, 134 (1), 339–360.

Weitzman, Martin L (2007) "Subjective expectations and asset-return puzzles," *American Economic Review*, 97 (4), 1102–1130.

Xu, Yiqing (2017) "Generalized synthetic control method: Causal inference with interactive fixed effects models," *Political Analysis*, 25 (1), 57–76.

A Derivations and proofs

A.1 Proof of Proposition 1

First, we establish that the game is supermodular with each Bayes-Nash equilibrium in strategies that are monotone in type. The joint density, $p(\eta, z_i) = g(\eta)\phi(\frac{z_i-\eta}{\sigma_x})$, is log supermodular by the concavity of $\log\phi$. This establishes that η and z_i are affiliated, so posterior beliefs over θ are increasing in shock component z_i in the sense of first-order stochastic dominance.

By affiliation of z_i and η, SC holders' posterior beliefs over θ are also monotonically decreasing in signal x_i by their proportionality to the shock components. Conditional on observing x_i, the posterior belief over the event that θ is less than or equal to some critical point, $\theta_x = y + \eta_c$, is given by

$$\mathbb{P}[\theta \leq \theta_c | x_i] = G(\eta_c | z_i) = \frac{\int_{-\infty}^{\eta_c} g(\eta)\phi(z_i - \eta)d\eta}{\int_{-\infty}^{\infty} g(z_i - \sigma_x s)\phi(s)ds}\,. \tag{9}$$

The payoff differential between running on the issuer and maintaining a holding is given by

$$\psi(z) = \int_{-\infty}^{\eta_c} \left[(1-\tau) f\left(\frac{y + \mathbb{E}[\eta|z]}{y + \mathbb{E}[\eta|z]}\right)\right] g(\eta|z)d\eta + \int_{\eta_c}^{\infty} (1 - 2\tau - \pi(y+\eta))g(\eta|z)d\eta\,, \tag{10}$$

where we have used the equilibrium condition that at signal x_c, the expected withdrawal mass is equal to the expected value of reserve assets, i.e., $R(x_c - y) = \mathbb{E}[\theta | x = x_c] = y + \mathbb{E}[\eta|z_c] = \theta_c$. The expected value of reserve assets given signal x, is

$$\mathbb{E}[\theta|x] = y + \frac{\int \eta g(\eta)\phi\left(\frac{x-y-\eta}{\sigma_x}\right)d\eta}{\int g(s)\phi\left(\frac{x-y-\eta}{\sigma_x}\right)ds}\,. \tag{11}$$

A sufficient condition for equation (10) to have a unique solution at $\psi(z) = 0$ is if $\frac{\partial\psi}{\partial z} \leq 0$ which establishes supermodularity of the payoff function. This is satisfied whenever:

$$\tau < \hat{\tau} \equiv \tau : \int_{-\infty}^{\eta_c} \frac{\partial g}{\partial z}(1-\tau) f(1) + \int_{\eta_c}^{\infty} \frac{\partial g}{\partial z}(1 - 2\tau - \pi(y+\eta))d\eta = 0\,. \tag{12}$$

Next we show consistency of higher-order beliefs. We work backwards, first analyzing the behavior of the issuer for any mass of early withdrawals before solving the equilibrium

strategies of SC holders.

For a given mass of redemption requests, A, the issuer is able to process withdrawals by selling down reserve assets at a unit value θ per stablecoin. Whenever $A > \theta$, the issuer is rendered insolvent. Therefore, conditional on $A = \theta_c$, where θ_c is some critical level at which withdrawal requests are just equal to the value of reserve assets, it is in every SC holder's best interest to demand conversion in the hope of reclaiming funds before the issuer becomes insolvent (i.e., to run on the stablecoin).

Given some critical reserve asset value $\theta_c = y+\eta_c$, let $\mathcal{P}$ be the set of all signals associated with a non-negative expected payoff from demanding conversion, so that

$$\mathcal{P} \equiv \left\{ x \left| \int_{-\infty}^{\theta_c} \left((1-\tau) f\left(\frac{\mathbb{E}[\theta|x]}{R(x)} \right) \right) g(\theta|x) d\theta + \int_{\theta_c}^{\infty} (1 - 2\tau - \pi(\theta)) dG(\theta_c|x) > 0 \right. \right\}. \quad (13)$$

We say that demanding conversion is $p-$dominant for SC holder i whenever $x_i \in \mathcal{P}$, and maintaining a holding is $p-$dominant whenever $x_i \in \mathcal{P}'$ (the complement of $\mathcal{P}$). Since the left-hand side of the expression in (13) is strictly decreasing in x, by the intermediate value theorem, for any θ_c, there is a corresponding indifference condition characterized by a unique signal x such that the expected payoff differential is exactly zero.

To derive the critical threshold, θ_c, SC holders must also appeal to higher order reasoning to justify the decision to withdraw their holdings. Let $\mathcal{Q}$ be the set of all signals associated with rank beliefs that exceed the expected value of reserve assets, so that

$$\mathcal{Q} \equiv \left\{ x \left| R(x_j - y) > y + \mathbb{E}[\eta|x_j - y] \ \forall x_j \leq x \right. \right\}. \quad (14)$$

Whenever $x_i \in \mathcal{Q}$, SC holder i believes that at least $q = y + \mathbb{E}[\eta|x - y]$ others will withdraw early, where $\mathbb{E}[\eta|x_j - y]$ is another SC holder's expected common shock component conditional on observing x_j.

This mass will be sufficient to render the issuer insolvent and, crucially, anyone with a more pessimistic signal has a belief at least as strong as i's that the issuer will become insolvent. This is because the rank belief function provides i's expected mass of withdrawals conditional on i's signal being the switching point x_c. Denoting the complement of $\mathcal{Q}$ by $\mathcal{Q}'$, SC holder i believes that too few others will withdraw to cause the issuer to fail whenever $x_i \in \mathcal{Q}'$.

Together, conditions $x \in \mathcal{P}$ and $x \in \mathcal{Q}$ are necessary and sufficient to make withdrawing

uniquely rationalizable at signal x.[50] Owing to supermodularity of the game among SC holders, there exists a greatest and least Nash equilibrium that bound all rationalizable strategies (Van Zandt and Vives, 2007). The least equilibrium switching point, $\hat{\hat{x}}$, is the unique solution to

$$\hat{\hat{x}} = \sup_x \{x \in \mathcal{P} \cap \mathcal{Q}\}. \tag{15}$$

The greatest equilibrium switching point, $\hat{x}$, is the unique solution to

$$\hat{x} = \inf_x \{x \in \mathcal{P}' \cap \mathcal{Q}'\}. \tag{16}$$

Since x and θ are stochastically affiliated, we have $\hat{\hat{x}} \leq \hat{x}$, and a sufficient condition for a unique equilibrium, x^*, at which the greatest and least Bayes-Nash equilibria converge, is $\sigma_x < \hat{\sigma}_x(\nu)$, where

$$\hat{\sigma}_x(\nu) \equiv \sigma_x : \frac{\partial \mathbb{E}[\eta|z;\sigma_x,\nu]}{\partial z}\Big|_{z=0} - \frac{\partial R(z;\sigma_x,\nu)}{\partial z}\Big|_{z=0} = 0\,, \tag{17}$$

since, for any σ_x and ν, $\frac{\partial \mathbb{E}[\eta|z;\sigma_x,\nu]}{\partial z}$ is minimized at $z = 0$, while $\frac{\partial R(z;\sigma_x,\nu)}{\partial z}$ is maximized at $z = 0$, and by the intermediate value theorem which ensures at least one solution. In what follows, we set $\sigma_x \geq \hat{\sigma_x}(\nu)$ to ensure multiple equilibria.

For each equilibrium switching point, $\hat{x}$ and $\hat{\hat{x}}$, there is a corresponding critical dollar value of reserve assets, $\hat{\theta}$ and $\hat{\hat{\theta}}$ respectively, given by the unique solution to

$$\theta = \Phi\left(\frac{x-\theta}{\sigma_x}\right). \tag{18}$$

The issuer always becomes insolvent when $\theta < \hat{\hat{\theta}}$, always survives when $\theta \geq \hat{\theta}$, and faces an indeterminate outcome whenever $\hat{\hat{\theta}} \leq \theta < \hat{\theta}$.

Comparative statics

For a given critical value of fundamentals, θ_c, the continuity of the payoff function guarantees that there is a unique $p-$dominance threshold x^*, at which a stablecoin holder is indifferent between demanding conversion and maintaining a holding when $\mathbb{E}[\theta^*|x^*]$ others are expected to withdraw. Fix model parameters such that $\overline{x}_{\mathcal{Q}} > x^* > \underline{x}_{\mathcal{Q}}$, where $\underline{x}_{\mathcal{Q}}$ is the lowest, and

[50] Analogously, $x \in \mathcal{P}'$ and $x \in \mathcal{Q}'$ is necessary and sufficient to make maintaining a holding uniquely rationalizable at x.

$\overline{x}_{\mathcal{Q}}$ is the greatest of the solutions to

$$R(z) = y + \mathbb{E}[\eta|x_j]. \tag{19}$$

In this case, it is *beliefs about aggregate behavior* rather than individual payoff parameters that are crucial in determining the probability of issuer insolvency.[51] As such, using the implicit function theorem, changes in market expectations about reserve quality, y, produce the following effects on switching strategies:

$$\begin{aligned} \frac{\partial \hat{x}(z)}{\partial y} &= \frac{1}{R'(z) - \mathbb{E}'[\eta|x]} < 0 \\ \frac{\partial \hat{\hat{x}}(z)}{\partial y} &= \frac{1}{R'(z) - \mathbb{E}'[\eta|x]} < 0, \end{aligned} \tag{20}$$

where we have used the definition for switching point $\hat{x}$ from (16). The sign of (20) is negative since $R'(z) < 0$ and $\mathbb{E}'(\cdot) > 0$ at both $\hat{x}$ and $\hat{\hat{x}}$.[52] By affiliation of x and θ in (18), the fundamental thresholds are also decreasing in y.

Note that the solutions to (19) are independent of the transaction cost τ. Therefore, as long as $\overline{x}_{\mathcal{Q}} > x^* > \underline{x}_{\mathcal{Q}}$ holds, increases in transaction costs *do not affect switching strategies* defined by $\hat{x}$ and $\hat{\hat{x}}$, and so the fundamental thresholds $\hat{\theta}$ and $\hat{\hat{\theta}}$ are also invariant to small increases in transaction costs.

However, since the condition $\overline{x}_{\mathcal{Q}} > x^* > \underline{x}_{\mathcal{Q}}$ is sensitive to τ, and with the well-known result that both x^* and θ^* are decreasing in τ, it follows that a sufficiently large increase in transaction costs could lead to a case where $x^* < \underline{x}_{\mathcal{Q}} < \overline{x}_{\mathcal{Q}}$. When this happens, we have $\partial \hat{\hat{x}}/\partial \tau < 0$ and, hence, $\partial \hat{\hat{\theta}}/\partial \tau < 0$. When the condition $\overline{x}_{\mathcal{Q}} > x^* > \underline{x}_{\mathcal{Q}}$ no longer holds, the lower switching strategy, $\hat{\hat{x}}$ is defined by the $p-$dominance threshold x^*, rather than $\underline{x}_{\mathcal{Q}}$. In this case, by condition (13), an increase in τ has the following effect on the lower threshold:

$$\left.\frac{\partial \hat{\hat{x}}}{\partial \tau}\right|_{\tau < \hat{\tau}} = \frac{\int_{-\infty}^{\theta_c} f\left(\frac{\mathbb{E}[\theta|x]}{R(x)}\right) g(\theta|x) d\theta + \int_{\theta_c}^{\infty} 2g(\theta|x) d\theta}{\frac{\partial G(\theta_c|x)}{\partial x}(1-\tau) - \int_{\theta_c}^{\infty} g'(\cdot)[\pi(\theta) - (1-2\tau)] d\theta} < 0, \tag{21}$$

[51] The condition $\overline{x}_{\mathcal{Q}} > x^* > \underline{x}_{\mathcal{Q}}$ is not necessary for our results. The effect of changes in y on switching point x^* is well established in the literature (Prati and Sbracia, 2002, 2010). We focus on this characterization for the independence of the equilibria on marginal changes in transaction costs, as derived below.

[52] To see this consider: $\frac{\partial R(z)}{\partial z} = \frac{\phi * g(z) \cdot \Phi \phi * \frac{\partial g}{\partial z}(z) - \Phi \phi * g(z) \cdot \phi * \frac{\partial g}{\partial z}(z)}{(\phi * g(z))^2}$, where we use $\phi * g$ to represent convolution of ϕ and g. First, $\hat{\hat{z}} \leq \overline{z}(\underline{R}) < 0$ as per the proof of Proposition 4 by Morris and Yildiz (2019), where $\underline{R}$ is the minimum rank belief, given distributions g and ϕ. Further, by the definition of $\underline{R}$ and the uniform limit rank beliefs property, $\partial R(z)/\partial z \leq 0$ for all $z \leq \overline{z}(\underline{R})$. By symmetry, the same properties hold at $\hat{z}$.

where $G(\theta_c|x)$ is decreasing in x by stochastic affiliation of θ_c and x, and $g'(\cdot) \geq 0$ at the lower threshold since $\hat{\hat{x}} \leq \hat{\hat{\theta}}$.

Therefore, while the common prior has the same dampening effect on the flightiness of SC holders as in unique-equilibrium global games, transaction costs are only effective at staving off a run when they are raised by a sufficiently large degree, or when the switching point is defined by individual payoff parameters, rather than aggregate beliefs (i.e., when $x^* \leq \underline{x}_{\mathcal{Q}} < \overline{x}_{\mathcal{Q}}$).

A.2 Proof of Corollary 1

Equilibrium shifts to majority redemption if and only if withdrawing early is uniquely rationalizable for the median SC holder. Since signals are symmetric around the mean, this type is $x_{it} = \theta_t$. Using the results in Proposition 1, maintaining a holding is uniquely rationalizable for this type whenever $\theta_t \geq \hat{\hat{x}}_t(y_t)$ given a new issuance or that there was not a run in the previous period. Otherwise, maintaining a holding continues to be uniquely rationalizable provided $\theta_t \geq \hat{x}_t(y_t)$.

A.3 Proof of Proposition 2

From the proof of Proposition 1 and Corollary 1, we have that $\hat{\hat{\theta}}_t$ is defined implicitly by

$$\hat{\hat{\theta}}_t = \Phi\left(\frac{\hat{\hat{x}}_t - \hat{\hat{\theta}}_t}{\sigma_x}\right), \tag{22}$$

and that $\hat{\theta}_t$ is defined implicitly by

$$\hat{\theta}_t = \Phi\left(\frac{\hat{x}_t - \hat{\theta}_t}{\sigma_x}\right), \tag{23}$$

given thresholds $\hat{\hat{x}}_t$ and $\hat{x}_t$ respectively. Equations (22) and (23) show that $\hat{\hat{\theta}}_t$ ($\hat{\theta}_t$) is increasing in $\hat{\hat{x}}_t$ ($\hat{x}_t$).

By the implicit function theorem, it suffices to examine the effect of an increase in transparency on the equilibrium condition.[53] The partial derivative of (19) with respect to

[53] This is because $R'(z) < 0$ and $\mathbb{E}'(\cdot) > 0$ at both $\hat{\hat{x}}_t$ and $\hat{x}_t$, and so the partial derivative of the equilibrium condition (19) with respect to x is negative.

ν_t is given by:

$$\frac{\partial R(z_t)}{\partial \nu_t} - \frac{\partial \mathbb{E}[\eta_t | x_{jt}]}{\partial \nu_t}\,. \tag{24}$$

The term $\partial R(z_t)/\partial \nu_t$ causes a 'widening' of the rank belief function as follows:

$$\frac{\partial R}{\partial \nu_t} = \frac{\Phi\phi * \frac{\partial g}{\partial \nu_t}(z_t)\phi * g(z_t) - \Phi\phi * g(z_t)\phi * \frac{\partial g}{\partial \nu_t}(z_t)}{[\phi * g(z_t)]^2}\,. \tag{25}$$

By symmetry of $G(\eta_t; \nu_t)$ around 0, and since an increase in ν_t causes a reduction in the density of the tails of the distribution, the expression in (25) is negative at $\hat{\hat{x}}_t$ and is positive at $\hat{x}_t$.

The effect on the second term in (24) is given by:

$$\begin{aligned}\frac{\partial \mathbb{E}[\eta_t | x_{jt}]}{\partial \nu_t} = \frac{1}{(\phi * g(z_t))^2}\Bigg\{ &\int g(s)\phi(z_t - s)ds \cdot \int \eta_t \frac{\partial g}{\partial \nu_t}(\eta_t)\phi(z_t - \eta_t)d\eta_t \\ &- \int \eta_t g(\eta_t)\phi(z_t - \eta_t)d\eta_t \cdot \int \frac{\partial g}{\partial \nu_t}(s)\phi(z_t - s)ds\Bigg\}.\end{aligned} \tag{26}$$

Define $\hat{\hat{z}}_t = \hat{\hat{x}}_t - y_t$ and $\hat{z}_t = \hat{x}_t - y_t$. By affiliation of η_t and z_t, and since $\hat{\hat{z}}_t < 0 < \hat{z}_t$, $\partial \mathbb{E}(\cdot)/\partial \nu_t > 0$ at $\hat{\hat{z}}_t$, while $\partial \mathbb{E}(\cdot)/\partial \nu_t < 0$ at $\hat{z}_t$. Together, the signs in (25) and (26) make the expression in (24) negative at $\hat{\hat{z}}_t$ and positive at $\hat{z}_t$. We thus have that $\hat{\hat{x}}_t$ is decreasing in ν_t and $\hat{x}_t$ is increasing in ν_t. By equations (22) and (23), $\hat{\hat{\theta}}_t$ is decreasing in ν_t and $\hat{\theta}_t$ is increasing in ν_t.

A.4 Proof of Corollary 2

In the proof of Proposition 1, we define the restriction on idiosyncratic noise, $\hat{\sigma}_x^2$, that determines multiplicity of equilibria. By its definition in equation (17), the threshold is sensitive to changes in degrees of freedom, ν, that parameterize the t-distribution $G(\eta)$. Taking the partial derivative of $\hat{\sigma}_x$ with respect to ν using the implicit function theorem, we have

$$\frac{\partial \hat{\sigma}_x}{\partial \nu} = \frac{-1}{\left[\frac{\partial^2 \mathbb{E}[\eta|z]}{\partial z \partial \sigma_x}\Big|_{z=0} - \frac{\partial^2 R(z)}{\partial z \partial \sigma_x}\Big|_{z=0}\right]} \left(\frac{\partial^2 \mathbb{E}[\eta|z]}{\partial z \partial \nu}\Big|_{z=0} - \frac{\partial^2 R(z)}{\partial z \partial \nu}\Big|_{z=0}\right) > 0\,. \tag{27}$$

The term in square brackets is negative and the term in round brackets is positive by virtue of the fact that the marginal change in the conditional expectation function given a change

in z decreases (increases) as σ_x (ν) increases, while the marginal change in the rank belief function given a change in z increases as σ_x (ν) increases.

Define $\hat{\nu}$ as the degree of freedom such that $\hat{\sigma}_x(\hat{\nu}) = \sigma_x$. Then whenever $\nu > \hat{\nu}$, $\sigma_x < \hat{\sigma}_x$ and so $\hat{\hat{x}}$ and $\hat{x}$ converge at x^*. At this – and every – point, $R(z)$ is approximately 1/2, ensuring that there is approximate common knowledge of approximately uniform rank beliefs (Morris et al., 2016), which, by the proof of Proposition 1, makes the condition x^* a unique switching point with corresponding reserve asset value, θ^*, that causes an issuer to fail.

In the game where a unique equilibrium is obtained, the effects of increases in precision of common shocks ($\alpha \equiv \frac{1}{\sigma_\theta}$) and private shocks ($\beta \equiv \frac{1}{\sigma_x}$), are then determined by the well known results of Metz (2003). We can write the SC holder indifference condition at the critical level of fundamentals, θ^*, as follows

$$x^* = \frac{\alpha+\beta}{\beta}\theta^* - \frac{\alpha}{\beta}y - \frac{\sqrt{\alpha+\beta}}{\beta}\Phi^{-1}\left(\rho(\theta^*)\right). \tag{28}$$

The effect of a marginal increase in the precision of public information, α, on switching point x^* is given by

$$\frac{\partial x^*}{\partial \alpha} = \frac{\theta^*}{\beta} - \frac{y}{\beta} - \frac{1}{2\beta\sqrt{\alpha+\beta}}\phi^{-1}(\rho(\theta^*)). \tag{29}$$

The sign of (29) is ambiguous. In particular, the equilibrium switching point, which indicates the propensity of SC holders to demand conversion for a given level of fundamentals, is increasing in the precision of public information if

$$\theta^* > y + \frac{1}{2\sqrt{\alpha+\beta}}\Phi^{-1}(\rho(\theta^*)). \tag{30}$$

Using the aggregate equilibrium condition,

$$\theta^* = \Phi\left(\sqrt{\beta}(x^* - \theta^*)\right),$$

and by the implicit function theorem, the effect of a change in x^* on θ^* is given by

$$\frac{\partial \theta^*}{\partial x^*} = \frac{-1}{1+\sqrt{\beta}\phi(\cdot)}(-\phi(\cdot)) > 0, \tag{31}$$

since $\phi(\cdot)$ is weakly positive for all reals. This means that if the propensity of SC holders to demand conversion increases in response to an increase in the precision of public information,

then so too does the overall probability of a run since $\mathbb{P}[\theta \leq \theta^*]$ is increasing in θ^*.

A.5 Proof of Proposition 3

We consider an environment characterized by no run having taken place in the prior period under hysteresis equilibrium by setting $t = 0$ or $A_{t-1} < \frac{1}{2}$. This implies that the equilibrium played by SC holders and the issuer has switching points $(\hat{\hat{\theta}}, \hat{\hat{x}})$. Equation (4) is weakly positive since $\phi(\cdot)$ is weakly positive over all reals, and since $\hat{\hat{\theta}}_t \leq y_t$.

To see that $\hat{\hat{\theta}}_t \leq y_t$, consider first that if the prior, $y_t = \theta_{t-1}$, lies in the unit interval, then at the start of the period, the issuer is "ripe for a run". The lowest value that y_t can take without rendering the game dominance solvable is $\min_y\{y : y \geq \underline{y}\}$, where $\underline{y}$ is defined as the lowest y for which there exists a shock $z < 0$ such that $R(z) \leq y + \mathbb{E}[\eta|z]$.[54] By definition of threshold $\hat{\hat{\theta}}$ in the proof of Proposition 1, we thus have that $\hat{\hat{\theta}}_t \leq y_t$ and so $\frac{\partial \hat{\hat{\theta}}}{\partial \sigma_\theta} \geq 0$.

B Run dynamics in a two-layered market structure

The impact of reserve volatility and transparency on run risk is robust to the inclusion of a two-layered market structure. In this section, we introduce a secondary market that determines the stablecoin's price on an exchange, and contrast equilibrium strategies with those where there are direct redemption rights with the issuer. For the sake of brevity and simplicity, we focus on a special case of our model where $\nu > \hat{\nu}$ so that we have a unique equilibrium, but our results extend to the general setting with fat tails and multiple equilibria.

As in the main text, suppose there is a continuum of SC holders, $i \in [0, 1]$. A share $\lambda \in [0, 1]$ are *arbitrageurs* who can redeem coins directly with the issuer at a promised price of \$1. The remaining $(1 - \lambda)$ SC holders are *secondary market participants* who can only buy and sell coins on an exchange at price $p(\lambda, \theta) = p \in \mathbb{R}_+$. While coins are bought and sold freely on the secondary market, the issuer is bound to service redemption requests only from the λ arbitrageurs. As before, the issuer draws on reserve assets that have a combined

[54] Analogously, $\overline{y}$ is the smallest prior that supports $R(z) \geq y + \mathbb{E}[\eta|z]$ for some $z > 0$.

dollar value θ per coin issued. Therefore, the issuer fails if and only if

$$\theta < \lambda A(\theta) = \lambda \int_i a_i di \,. \tag{32}$$

The game extends over three stages. First, Nature selects $\theta \in \mathbb{R}$. Second, both arbitrageurs and secondary market participants observe private signals about the value of reserve assets. Secondary market participants determine whether to sell their coin ($b_j = 1$) for an immediate payoff or hold ($b_j = 0$) for future flow payoffs, based on beliefs about the issuer's liquidity. Third, arbitrageurs observe prices and sell on the secondary market ($a_i = 0$) or redeem directly with the issuer on the primary market ($a_i = 1$).[55] The issuer processes redemption requests in the primary market, and state-contingent payoffs accrue.

Arbitrageurs are large in the primary market but small in the secondary market, so they take the mapping from fundamentals to prices as given when making their decisions.[56] Prices in the secondary market are determined by the joint selling pressure from secondary market participants and arbitrageurs. Prices (expressed in terms of the currency to which the stablecoin is pegged) take the following form at the time arbitrageurs make redemption decisions

$$p(\lambda, \theta) \equiv 1 - (1 - \lambda)B(\theta) = p \,, \tag{33}$$

where $B(\theta) = \int b_j dj$ is the mass of sell-offs from secondary market participants without redemption rights.[57] State-dependent payoffs to arbitrageurs are summarized in Table 7. To ensure that redemption of coins with the issuer is not strictly dominated, we make the following additional assumption:

Assumption 1 (**transaction costs**). $\tau < \frac{1}{2}$.

In the event the issuer is solvent, the benefit to arbitrageur i from redeeming directly with the issuer is increasing in the difference between the \$1 pledged by the issuer and

[55] Like Ma et al. (2023), we assume that arbitrageurs face high inventory costs from holding coins, and therefore the option to hold is strictly dominated in all states of the world.

[56] In particular, issuers mint a given number of coins which we normalize to one, and sell these to λ primary-market investors who have redemption rights with the issuer. Over time, some of these investors sell their coins in to the secondary market, generating more concentration in the primary market (i.e., lowering λ) and increasing the size of the secondary market (i.e., increasing $1 - \lambda$). This is how each agent in our model is endowed with a stablecoin at the beginning of the first stage.

[57] In general, prices can serve as public signals that can re-introduce multiplicity into global games (Angeletos and Werning, 2006). We abstract from this by taking the limit of vanishing private noise and fixing arbitrageurs' signals to be sufficiently precise relative to secondary market investors, so that prices do not reveal the value of fundamentals. In practice, this suggests that there are enough noise traders that the signal-noise ratio of the secondary market is relatively low.

	Issuer solvent	Issuer insolvent
$a_i = 1$	$1 - 2\tau$	$1 - 2\tau$
$a_i = 0$	$p - \tau$	0

Table 7: Arbitrageur payoffs. Action $a_i = 1$ denotes demanding conversion; action $a_i = 0$ denotes selling in the secondary market.

the price of the stablecoin traded on the secondary market (both net of transaction costs, $2\tau > 0$).[58] When the issuer is solvent, arbitrageurs prefer to sell into the secondary market whenever their arbitrage profits are sufficiently low relative to the transactions costs from redeeming. By contrast, when the issuer is insolvent, arbitrageurs derive a strictly higher benefit from redeeming over selling their coins whenever Assumption 1 is satisfied.

Secondary market participants derive the same benefits, $\pi(\theta)$ from holding and on-lending stablecoins as in the main text. This payoff is conditional on the stablecoin maintaining a store of value until the end of the game, which happens only in the state where the issuer survives. Alternatively, stablecoins can be sold on the secondary market at price $p(\lambda, \theta)$, which incurs transaction cost τ. Payoffs to secondary market participants are summarised in Table 8.

	Issuer solvent	Issuer insolvent
$b_i = 1$	$p(\lambda, \theta) - \tau$	$p(\lambda, \theta) - \tau$
$b_i = 0$	$\pi(\theta)$	0

Table 8: Non-arbitrageur stablecoin holder payoffs. Action $b_i = 1$ denotes selling on the secondary market; action $b_i = 0$ denotes maintaining a holding.

Since equation (33) shows that prices are decreasing in selling pressure, this implies strategic substitutability among secondary market participants in states where the issuer is solvent and the market functions as intended. In the event the issuer remains solvent, secondary market participants prefer to hold their stablecoins whenever the yield from on-lending exceeds the market price they would receive from selling the coin, net of transaction cost τ. But in the event large-scale redemptions cause the issuer to fail, secondary market participants prefer to sell their coins whenever $p(\lambda, \theta) > \tau$.

We focus on low volatility reserves, where $\theta \sim \mathcal{N}(y, \frac{1}{\alpha})$, but our analysis extends to the

[58] In practice, fees incurred on a secondary market differ substantially from those incurred in redemption on a primary market (Lyons and Viswanath-Natraj, 2023). For notational simplicity, we assume that fees in the primary market are simply double those in the secondary market. Our results would still go through under a different fee structure, $\tilde{\tau}$, provided that fees are not exorbitant, $\tau < \tilde{\tau} < 1$.

general setting with $\theta = y + \eta\,,\ \eta \sim t(\nu)$. The signal structure of the market can now be decomposed among secondary market investors and arbitrageurs. Let each secondary market investor observe signal $x_j \sim \mathcal{N}(\theta, \frac{1}{\beta_2})$ where β_2 is the precision of secondary market signals. Arbitrageurs receive independently and identically distributed signals $x_i \sim \mathcal{N}(\theta, \frac{1}{\beta_1})$, where β_1 denotes the precision of arbitrageur signals.

A Bayesian Nash Equilibrium for the two-layered market comprises a fundamental threshold, $\tilde{\theta}^*$, switching signal for arbitrageurs, x_1^*, and switching signal for secondary-market investors, x_2^*, such that:

1. In the final stage, the redemption decision, x_1^*, by arbitrageurs is optimal;
2. In the second stage, purchasing decision, x_2^*, by secondary market investors is optimal;
3. In the first stage, given thresholds x_1^*, x_2^* and price $p(\lambda, \tilde{\theta}^*)$, the draw of θ by Nature causes the issuer to fail whenever $\theta < \tilde{\theta}^*$.

Proposition 4. *In the limit of vanishing private noise, there is a unique switching point, x_1^*, for arbitrageurs such that each arbitrageur redeems if and only if $x_i \leq x_1^*$ and sells into the secondary market otherwise. There is also a unique secondary-market switching point, x_2^*, such that each non-redeeming investor sells her coin if and only if $x_j \leq x_2^*$ and holds otherwise. There is a unique fundamental threshold such that the issuer fails whenever $\theta \leq \tilde{\theta}^*$.*

Proof. We start by deriving the failure condition for the issuer. Since only arbitrageurs can redeem with the issuer, given a switching point x_1^*, the issuer fails whenever the realized mass of redemptions is sufficiently large, $\theta < \tilde{\theta}^*$, where

$$\tilde{\theta}^* = \lambda \Phi \left(\sqrt{\beta_1}(x_1^* - \tilde{\theta}^* \right) . \tag{34}$$

We next turn to the equilibrium strategy of arbitrageurs. Arbitrageur i prefers to redeem if and only if

$$1 - 2\tau > \int_{\tilde{\theta}^*}^{\infty} [p(\lambda, \theta) - \tau] f(\theta | x_i) d\theta \,. \tag{35}$$

Using equation (33) and substituting in the realized mass of selling secondary market investors (which we derive below), the indifference condition characterized by type x_1^* is given by

$$1 - 2\tau = \int_{\tilde{\theta}^*}^{\infty} [1 - (1 - \lambda)\Phi(\sqrt{\beta_2}(x_2^*(\theta) - \theta)) - \tau] f(\theta | x_1^*) d\theta \,. \tag{36}$$

We focus on vanishing private noise $\beta_1, \beta_2 \to \infty$, and $\beta_1/\beta_2 = c$ where c is a constant.[59] Under the resulting Laplacian beliefs (Morris and Shin, 2003), the left-hand side of condition (36) is constant, while the right-hand side is monotone decreasing in arbitrageur signal x_i, by affiliation of θ and x_i, so that condition (36) defines, implicitly, a unique switching point x_1^*.

Turning to non-redeeming secondary-market investors, the indifference condition that uniquely determines switching point x_2^* is given by:

$$\begin{aligned} \int_{-\infty}^{\infty} [p(\lambda, \theta) - \tau] f(\theta | x_2^*) d\theta &= \int_{\tilde{\theta}^*}^{\infty} \pi(\theta) f(\theta | x_2^*) d\theta \\ 1 - (1-\lambda) R_{22}(x_2^*) - \tau &= \int_{\tilde{\theta}^*}^{\infty} \pi(\theta) f(\theta | x_2^*) d\theta \,, \end{aligned} \tag{37}$$

where

$$R_{22}(x) = \Phi\left(\frac{\alpha\sqrt{\beta_2}}{\sqrt{(\alpha+\beta_2)(\alpha+2\beta_2)}} (x - y) \right) , \tag{38}$$

is the rank belief of investor j in the population of secondary market investors. As $\beta_2 \to \infty$, $R_{22} \to \frac{1}{2}$ for all x, and so the left-hand side of (37) simplifies to $\frac{1}{2}(1+\lambda) - \tau$, while the right-hand side is strictly decreasing in x, leading to a unique indifference condition at x_2^*.

□

Proposition 4 provides two insights. First, it illustrates that the self-fulfilling beliefs leading to a run extend to a two-layered market. To the extent that beliefs in the issuer's survival can be sustained, the strategic substitutability induced by the price mechanism ensures a well-functioning market in which upwards price pressure from arbitrage stabilizes the peg. However, when fundamentals fall below a critical threshold, $\tilde{\theta}^*$, there is strategic complementarity in the secondary market. The risk-dominant action – sell – becomes uniquely rationalizable for secondary market investors. This, in turn, increases the profit opportunity for arbitrageurs, leading to a sufficiently large mass of redemptions to cause the issuer to fail.

Second, it allows us to explore how the existence of a secondary market affects run risk. In particular, we can analyze whether the secondary market makes a run more or less likely. To do this, we consider how the probability of failure, as measured by $\mathbb{P}[\theta \leq \theta^*]$, is affected

[59] A study focusing on differential information that arbitrageurs and secondary market participants hold is beyond the scope of this paper. In such a case, a unique equilibrium may still be obtained provided $\beta_1, \beta_2 > \frac{\alpha^2}{2\pi}$, where π refers to the mathematical constant. Future work could explore whether the markets are differentially informed and how this might shape run risk.

by the size of the secondary market, $(1-\lambda)$, relative to those with redemption rights, λ, for a given supply of coins in circulation.

Proposition 5. *Let $\beta_1, \beta_2 \to \infty$ and $\beta_1/\beta_2 = c$. The presence of a secondary market makes arbitrageurs more aggressive. Switching point x_1^* is decreasing in λ, the proportion of arbitrageurs relative to secondary market participants. The presence of a secondary market also increases the probability of failure, $\frac{\partial \tilde{\theta}^*}{\partial \lambda} < 0$.*

Proof. Let

$$h(x, \lambda) \equiv 1 - 2\tau - \int_{\tilde{\theta}^*}^{\infty} \left[\frac{1}{2}(1+\lambda) - \tau \right] dF(\theta | x_1^*) = 0$$

define the equilibrium switching point x_1^*. By the implicit function theorem, the effect of an increase in λ on switching point x_1^* is given by

$$\frac{\partial x_1^*}{\partial \lambda} = \frac{-1}{\partial h / \partial x_1^*} \left(\frac{\partial h}{\partial \lambda} \right) < 0 \,. \tag{39}$$

The denominator is negative, as $F(\tilde{\theta}^* | x)$ is decreasing in x and $\partial h / \partial \lambda < 0$.

Next, consider the effect on the probability of failure. Let

$$g(\theta, \lambda) \equiv \tilde{\theta}^* - \lambda \Phi \left(\sqrt{\beta_1} (x_1^*(\lambda) - \tilde{\theta}^*) \right) = 0$$

define the equilibrium fundamental threshold $\tilde{\theta}^*$. An increase in λ has the following effect on $\tilde{\theta}^*$:

$$\frac{\partial \tilde{\theta}^*}{\partial \lambda} = \frac{-1}{\partial g / \partial \tilde{\theta}^*} \left(\frac{\partial g}{\partial \lambda} + \frac{\partial g}{\partial x_1^*} \frac{\partial x_1^*}{\partial \lambda} \right) . \tag{40}$$

The denominator of the factor, $\partial g / \partial \tilde{\theta}^*$, is positive and so $\tilde{\theta}^*$ is decreasing in λ whenever the expression inside brackets is positive. This is the case if and only if

$$\frac{1}{1 + \lambda - 2\tau} > \frac{\Phi \left(\sqrt{\beta_1} (x_1^* - \tilde{\theta}^*) \right)}{\lambda \sqrt{\beta_1} \phi \left(\sqrt{\beta_1} (x_1^* - \tilde{\theta}^*) \right)} , \tag{41}$$

which holds in the limit $\beta_1 \to \infty$ and by Assumption 1. Therefore $\frac{\partial \tilde{\theta}^*}{\partial \lambda} < 0$. □

The direct effect, $\frac{\partial g}{\partial \lambda} < 0$, increases $\tilde{\theta}^*$, and hence the probability of failure, due to increased liabilities on the balance sheet of the issuer (i.e., a larger number of investors

with redemption rights). So the presence of a secondary market attenuates the risk of failure by introducing non-redeeming SC holders that dilute the market. The indirect effect, $\frac{\partial g}{\partial x_1^*}\frac{\partial x_1^*}{\partial \lambda} > 0$, reduces $\tilde{\theta}^*$ and the probability of failure. So the larger the secondary market, the higher the risk of failure, because arbitrageurs become more aggressive due to the arbitrage opportunity arising from selling pressure on the secondary market. On balance, the indirect effect dominates in the limit of vanishing private noise, making the probability of failure increasing in the size of the secondary market.

Proposition 4 shows that the secondary market introduces an additional profit incentive for arbitrageurs to demand redemption with the issuer beyond the panic effect arising from beliefs that the issuer will fail. The effect on arbitrageurs is not reliant on the limit of vanishing private noise.[60] In general, as long as the posterior that the issuer cannot honor redemptions is sufficiently large, a secondary market does not alleviate the risk of self-fulfilling runs on stablecoin issuers.

C Additional material

C.1 Reserve assets as a portfolio-weighted convolution

In the context of our model, stablecoins are backed by a vector of reserve assets that have a combined dollar value captured by θ. Stablecoin holders perceive θ to be a random variable drawn from the portfolio-weighted convolution of all component reserve asset distributions, with realizations of θ equal to the portfolio-weighted realized values of individual reserve assets used in the reserve portfolio, $\theta = \omega_1\theta_1 + \omega_2\theta_2 + \ldots + \omega_n\theta_n$, where ω_j denotes the share of the portfolio allocated to asset j.

[60] Away from the limit, a necessary condition for uniqueness of x_2^* is

$$\left[\frac{\pi'(\theta)}{\pi(\theta)} - \frac{\partial p(\lambda,\theta)/\partial\theta}{p(\lambda,\theta)}\right]\frac{\pi(\theta)}{\pi'(\theta)}p(\lambda,\theta) > \tau\,,$$

i.e., the returns to secondary market participants from a change in fundamentals feature a sufficiently large semi-elasticity at each level of θ relative to the price elasticity of the coin at θ.The condition is derived by formulating the indifference condition as follows,

$$F(\tilde{\theta}^*|x_2^*) = 1 - \frac{p(\lambda;\tilde{\theta}^*) - \tau}{\pi(\tilde{\theta}^*)},$$

and restricting the right-hand side to be increasing in $\tilde{\theta}^*$, ensuring a unique switching point x_2^*. In other words, returns from on-lending stablecoins are relatively more volatile than the price variation in the secondary market, which is, in general, consistent with empirical evidence. Formally, the condition ensures supermodularity of the payoff differential between selling and holding for secondary market investors. However, the overall effect on the probability of failure would then depend on the size of β_1 relative to β_2.

To fix ideas, think of a portfolio with three reserve assets. In this case, the joint probability density function (PDF) $f(\theta)$ is given by

$$f(\theta) = (f_1 * (f_2 * f_3))(\theta)\,,$$

where $f_1(\cdot)$, $f_2(\cdot)$ and $f_3(\cdot)$ are the PDFs of collateral assets θ_1, θ_2 and θ_3 respectively. Figure 14 illustrates this example for a given calibration of individual PDFs and equal portfolio weights. We define the convolution $f_1 * (f_2 * f_3)$ of f_1, f_2, and f_3 as follows:

$$\begin{aligned} f(\theta) &= (f_1 * (f_2 * f_3))(\theta) \\ &= \int_{\mathbb{T}} f_1(\omega_1\theta_1) \left[\int_{\mathbb{T}} f_2(\omega_2\theta_2) f_3(\omega_3[\theta - \omega_1\theta_1 - \omega_2\theta_2]) d\theta_2 \right] d\theta_1\,. \end{aligned}$$

The PDF $f(\theta)$ thus summarizes the portfolio of reserve assets. The model can easily be used to assess the effect of increasing the weight of any asset in the reserve portfolio. For example, Tether's transition away from commercial paper and towards short-term safe assets is an example of an increase of low-risk assets in the portfolio. The effects are presented in a result discussed in the paper: greater weighting on low-risk assets decreases run risk.

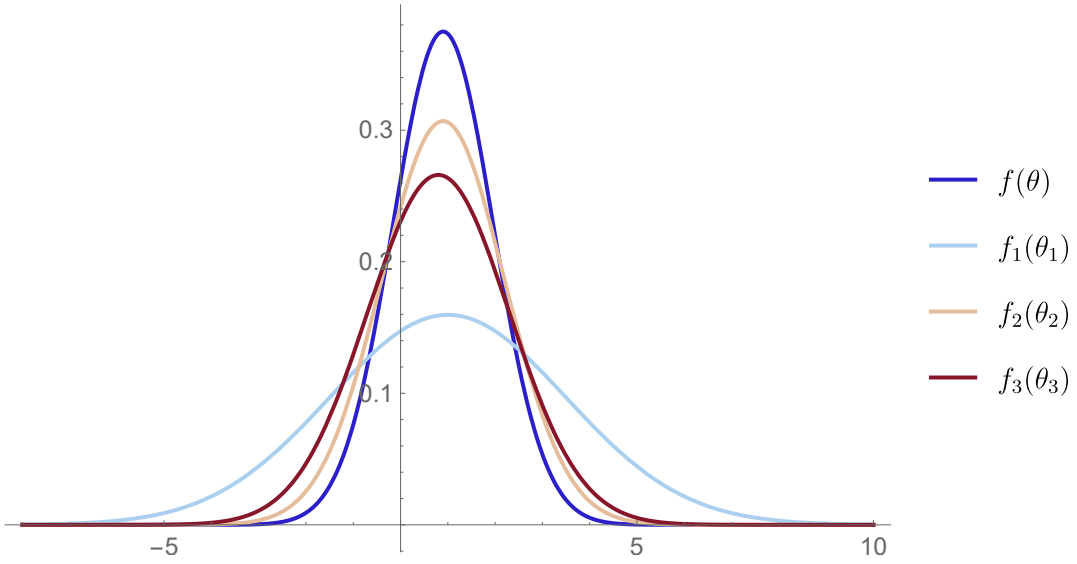

Figure 14: Probability density function of the value of reserve assets (θ) as a convolution of the distribution of assets in the portfolio $(\theta_1, \theta_2, \theta_3)$. In this example, $\theta_1 \sim \mathcal{N}(1, 2.5); \theta_2 \sim \mathcal{N}(0.9, 1); \theta_3 \sim \mathcal{N}(0.8, 1.5)$, with equal portfolio weights for each asset.

C.2 Additional graphs

Figure 15 takes a closer look at March 2023 depegging event by zooming in the early stages. This helps to more clearly highlight the sequencing, with USDC moving first and Dai moving later. The co-movement USDC between Dai and USDC may also partly reflect the price stability mechanism set up by Dai that enables par exchange between the two stablecoins. However, the price stability mechanism is unlikely to be the sole driver of the co-movement between the two stablecoins in the wake of turmoil at SVB, as Figure 16 shows that other stablecoins backed by USDC but without such mechanisms in place also depegged, whereas stablecoins backed by other cryptoassets remained stable during this period.

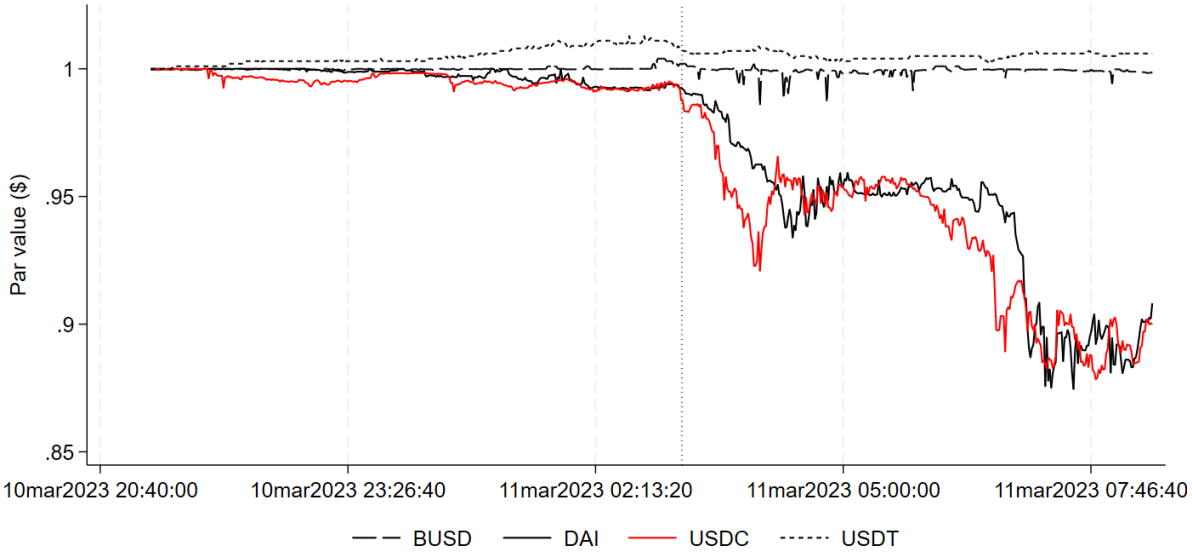

Figure 15: Stablecoin pegs around the run on Silicon Valley Bank under the microscope.

Notes: Based on minute-by-minute data. The vertical dashed line denotes the disclosure by Circle that $3.3 billion of its cash reserve was held at SVB. Source: Cryptocompare.com.

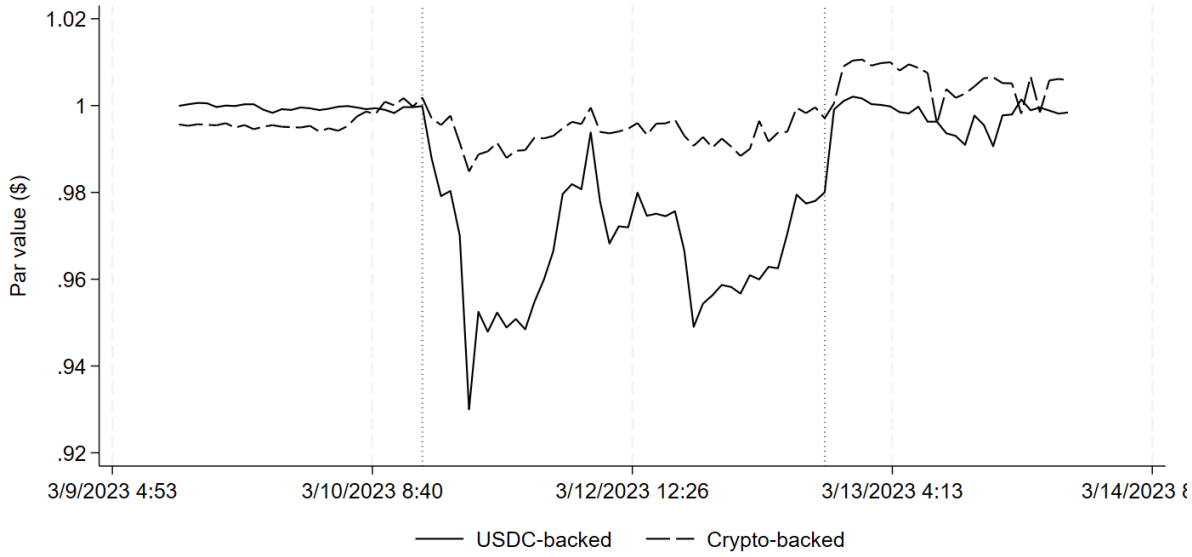

Figure 16: Stablecoin pegs around the run on SVB: USDC-backed versus crypto-backed.

Notes: Based on hourly data. USDC-backed is a simple average of the value of Dai, Frax and Origin Dollar (OUSD). Crypto-backed is a simple average of the value of CeloUSD, Liquity (LUSD), Tron's USDD and sUSD. The first vertical dashed line denotes the disclosure by Circle that $3.3 billion of its cash reserve was held at SVB; the second vertical line denotes the announcement of a backstop by the U.S. government. Source: Cryptocompare.com.

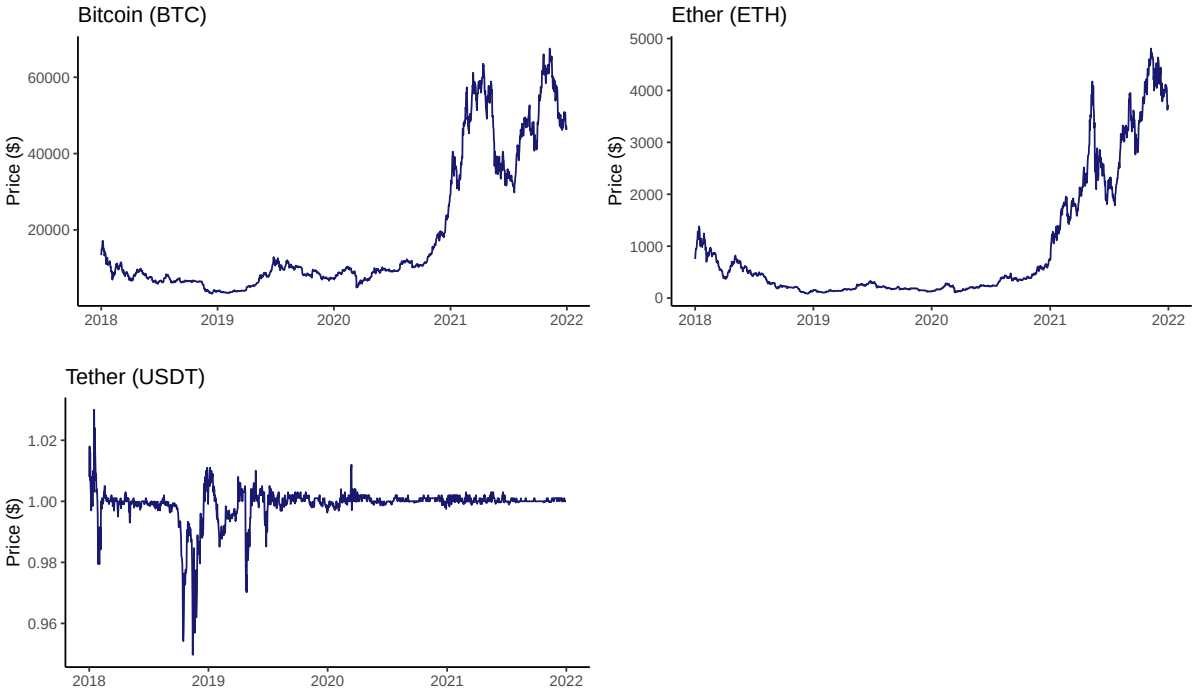

Figure 17: Cryptoasset prices.

Notes: Daily time series of cryptoasset prices from January 1, 2018 to December 31, 2021. Dai series does not begin until May 9, 2018. Source: Cryptocompare.com.

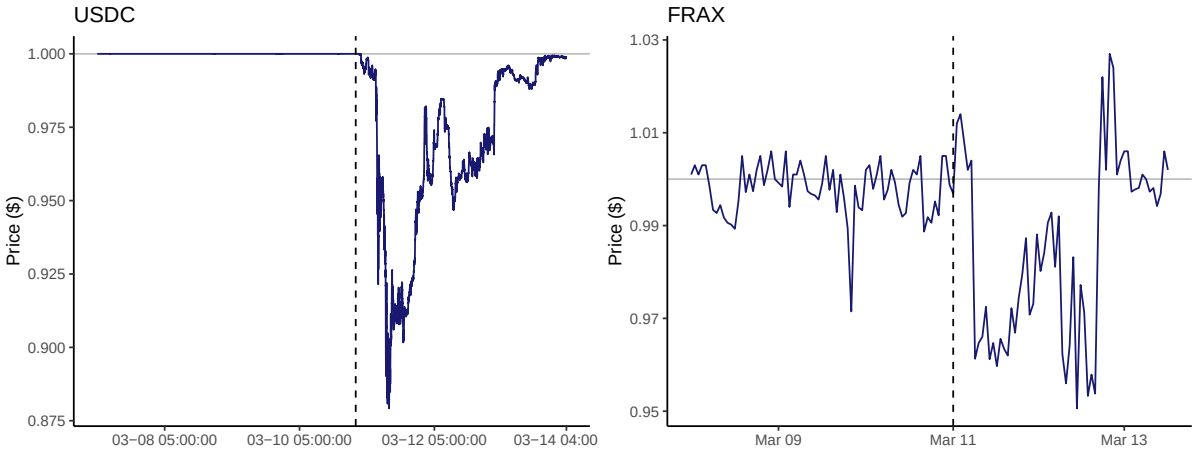

Figure 18: Stablecoin prices around the run on SVB: USDC and FRAX.

Notes: Based on minute-frequency data for USDC (left-panel) and hourly-frequency data for FRAX (right-panel). The vertical dashed line denotes the disclosure by Circle that $3.3 billion of its cash reserve was held at SVB. Source: Cryptocompare.com.

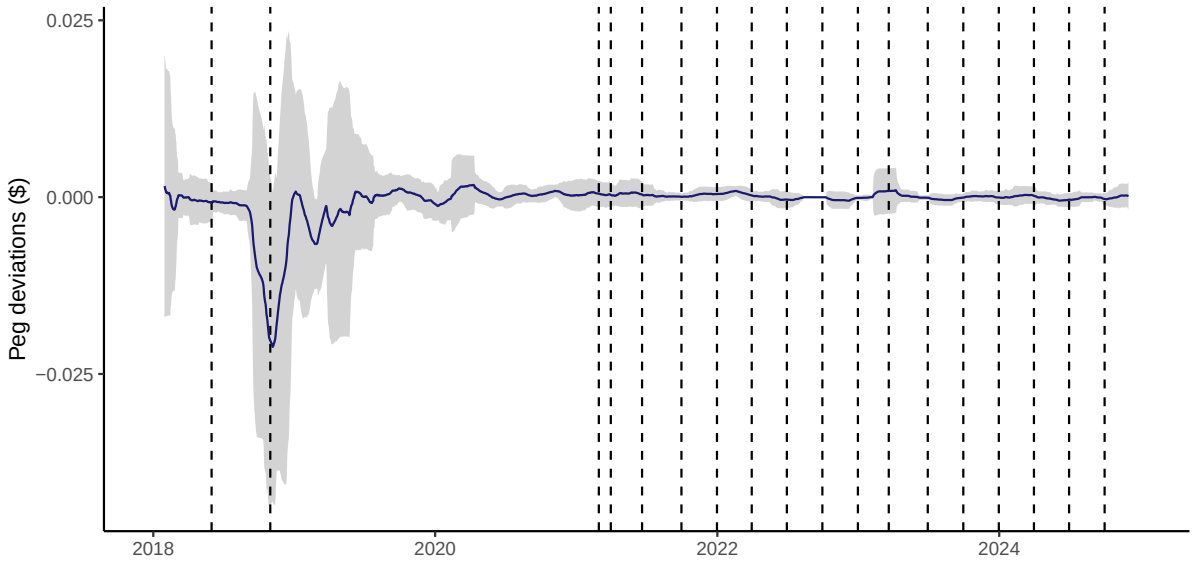

Figure 19: USDT price deviations from $1.

Notes: Based on daily data. Tether price deviations from $1 are computed using a 60-day centered moving average with shaded regions covering ±2 standard deviations. Vertical dashed lines are days with Tether reserves attestation releases. Source: Cryptocompare.com.