

Leverage and Stablecoin Pegs*

Gary B. Gorton^{†1}, Elizabeth C. Klee^{‡2},
Chase P. Ross^{§2}, Sharon Y. Ross^{¶2}, and Alexandros P. Vardoulakis^{||2}

¹Yale & NBER

²Federal Reserve Board

This draft: January 13, 2025

First draft: December 28, 2022

Abstract

Stablecoins are a new form of private money. They are fragile but largely trade at par. How? We present a model and empirical work to examine a novel source of demand for stablecoins. Stablecoin owners are indirectly compensated for run risk by lending their coins to crypto speculators. The stablecoin can then support its \$1 peg, but this arrangement links crypto speculation to traditional financial markets where stablecoins invest their reserves.

JEL Codes: E40, E51, G12, N21

Keywords: money, leverage, stablecoins, cryptocurrencies, global games, speculation

*We thank Oliver Hyman-Metzger and Arazi Lubis for excellent research assistance. For comments and suggestions, thanks to Joseph Abadi, David Arseneau, Christoph Bertsch, Eduardo Davila, John Geanakoplos, Pedro Gomis Porqueras, Daniel Graves, Todd Keister, Michael Palumbo, Christine Parlour, Greg Phelan, David Rappoport, Thomas Rivera, Raluca Roman, Scott Schuh, Donghwa Shin, Alp Simsek, Antoinette Schoar, K. Sudhir, Dimitrios Tsomocos, Quentin Vandeweyer, Chris Waller, Sean Wilkoff, Russell Tsz-Nga Wong, and seminar and conference participants at the Office of Financial Research, the Fed Board 2022 Summer Workshop on Money, Banking, Payments, and Finance, the Fed System Committee on Financial Institutions, Regulation, and Markets, the Philadelphia Fed, the Fed Board FS workshop, the 2023 MFA, the 2023 Yale Cowles Conference on General Equilibrium, the OCC, the 2023 Edinburgh Economics of Financial Technology conference, the 2023 NASMES, the 2023 CEBRA, the 2023 Oxford Saïd-Risk Center at ETH Zürich Macro-finance Conference, the 2023 MoFiR Workshop on Banking, the 2023 EFA, the 2023 Wharton Conference on Liquidity and Financial Fragility, the Economics of Payments XII Conference, Manhattan College, Banque de France/TSE/Panthéon ASSA, and 2024 SED. The views expressed in this paper are those of the authors and do not necessarily represent those of Federal Reserve Board of Governors, or anyone in the Federal Reserve System.

[†]gary.gorton@yale.edu

[‡]elizabeth.c.klee@frb.gov

[§]chase.p.ross@frb.gov

[¶]sharon.y.ross@frb.gov

^{||}alexandros.vardoulakis@frb.gov

I'm trying to buy the dip, but this dip is so big that now my stablecoins are dipping too.

Trifusi0n, Reddit, May 10, 2022

1 Introduction

Stablecoins are a new form of private money. They are fragile, usually pay no interest, and are rarely used for payments. Past incarnations of fragile private money nearly always traded at a discount to par. But stablecoins largely trade at par. We examine a novel source of demand for stablecoins, and we show theoretically and empirically how it helps stablecoins maintain their peg most of the time while being exposed to *priced* run-risk and depegging in times of stress.

Historically, privately-produced money took two forms. First, the money could be backed by agents' personal wealth, like Scottish bank notes backed by the wealth of the bank's partners with unlimited liability. Similarly, English inland bills of exchange were backed by the bill's endorsers, again with unlimited liability (Gorton, 2024). This money had limited geographical circulation because the receiving party needed to recognize the endorsers. Second, the private money could be issued by banks, like pre-Civil War U.S. banknotes. Banks backed the monies with state bonds and loan portfolios. These private bank notes circulated great distances but with discounts from par.

Stablecoins, like private banknotes, promise redemption at par on demand (Gorton and Zhang 2021, Gorton et al. 2022). Unlike unbacked digital assets, such as Bitcoin, stablecoins are usually backed by reserves, potentially risky and illiquid, and are denominated in fiat currency.¹ These two features—redemption at par on demand and illiquidity of reserves—render stablecoins susceptible to runs similar to fragile banks, mutual funds, and money market funds. Stablecoin holders should demand compensation for run risk. But stablecoins have not typically paid any interest, unlike bank deposits and money market funds. Tether, in particular, which has been the biggest collateralized stablecoin by market capitalization, has not paid any interest to its holders despite its non-trivial run risk. How are investors compensated for holding risky stablecoins?

We posit that a novel source of demand for stablecoins stems from their role in facilitating speculation in other digital assets. Crypto speculators can borrow stablecoins to make levered

¹We focus on so-called collateralized stablecoins that hold reserves instead of algorithmic stablecoins. The reserves could be traditional financial assets (commercial paper, reverse repurchase agreements, Treasuries) or crypto-related assets. Collateralized stablecoins constitute the majority of stablecoins by market capitalization, even before the failure of the largest algorithmic stablecoin, TerraUSD, in May 2022. See Azar et al. (2024) for details about the different types of stablecoins. Uhlig (2022) and Liu et al. (2023) provide additional details on algorithmic stablecoins with an emphasis on the collapse of TerraUSD.

bets on other cryptocurrencies. Lenders of stablecoins receive high lending rates, often above 20 percent at an annual rate and about 10 percent on average, even in a period of low interest rates. Stablecoin holders earn indirect compensation for run risk because they can lend the stablecoin to traders. In other words, the primary market issuer does not pay interest, but secondary lending markets compensate for the issuer’s run risk.² It should be noted that the US dollar (or liabilities that settle in US dollars with “no-questions-asked”) has traditionally facilitated leverage in the traditional financial system. Traditional financial institutions, however, are unable or reluctant to lend US dollars to finance crypto leverage, so this role has been assumed by stablecoins.

We develop a model that computes the premium that stablecoin holders require to maintain a traded price of \$1 and connect it to the rate speculators offer to borrow the stablecoin. Higher speculative demand for cryptocurrencies increases the lending rate on stablecoins and, hence, their price. A less liquid portfolio of reserves increases run risk, reducing their price. The two forces interact by pushing the peg in opposite directions. Given a demand for crypto, the stablecoin issuer chooses the liquidity of its reserves and the supply of tokens to maximize its profits while achieving a price of \$1 for the issued tokens. We characterize these mechanisms by proposing a model of stablecoins that nests bank-run models akin to Goldstein and Pauzner (2005) and Kashyap et al. (2024) and models of leveraged collateralized trading similar to Gromb and Vayanos (2002) and Fostel and Geanakoplos (2008).

Then, we turn to empirical work. First, we connect speculative demand for cryptocurrencies to stablecoins’ lending rate. We approximate speculative demand using the funding rate from perpetual futures, liquid crypto derivatives that bet on cryptocurrencies. We show that increases in speculative demand result in statistically and economically significant increases in stablecoin lending rates: a one percentage point increase in the perpetual future funding rate translates to roughly a 20 basis point rise in the lending rate. To account for unobserved endogeneity, we instrument speculative demand using the viewership of Major League Baseball (MLB) games, exploiting the sponsorship deal between MLB and FTX, a major cryptocurrency exchange. We find similar results.

Second, we empirically establish two stabilization mechanisms to maintain the peg: adjusting the liquidity of reserves and letting the lending rate re-adjust via token redemptions or issuance. Data for redemptions and issuance are available at high frequency, enabling us to map out how lower speculative demand leads first to more stablecoin redemptions and

²Other reported use cases of stablecoins could introduce non-pecuniary benefits and drive part of the demand for them. For example, stablecoins act as a store of value between crypto trades; they may generate payment services and offer a store of value in countries with volatile domestic currencies, or they may facilitate illicit finance and even terrorism. See Ostroff and Malsin (2022). We show that our theoretical and empirical results on the importance of speculation for stablecoin demand are robust to such additional use cases.

second to higher lending rates and peg stabilization while controlling for other factors driving the demand for stablecoins such as cryptocurrency trading activity.

Third, we apply the model to the May 2022 crypto turmoil following the collapse of TerraUSD and the near run on Tether. We document how Tether experienced large redemptions and traded considerably below its peg, but once lending rates increased and stabilized at a higher level, Tether’s peg stabilized as well, consistent with our theoretical predictions.

Our analysis highlights the link between crypto and the real economy. Stablecoin issuers invest their reserves to earn profits but must adjust their reserves—possibly quickly—to maintain their peg in response to crypto shocks. Such reallocations can disrupt the money markets in which stablecoins invest their reserves, such as Treasuries, commercial paper, and repos.³ The reverse link from traditional to crypto markets is also incorporated in the model: a bad shock for the traditional risky financial asset held in the stablecoin’s reserves leads to an increase in the stablecoin’s run risk. This was the case in March 2023 during the collapse of Silicon Valley Bank, where Circle, the issuer of the major stablecoin USDC, held part of its reserves as uninsured deposits.⁴

Related Literature. Our paper relates to the two strands of literature on leveraged collateralized trading and bank runs. We contribute by bringing these two strands together and show how their interaction can generate the leverage-money nexus described earlier.

Brunnermeier and Pedersen (2009), Fostel and Geanakoplos (2008), Geanakoplos (2010), and Gromb and Vayanos (2002) investigated the importance of private wealth to meet haircut requirements in collateralized trading. Still, the literature has mostly focused on shocks to the value of collateral that can lead to unfolding leverage and a drop in asset prices. We show that the fragility of the “assets” in which agents hold their private wealth and use to meet haircut requirements—stablecoin tokens in our case—can also be an independent source of instability, interacting with speculative motives.⁵

In contrast, the bank-run literature has studied fragile liabilities issued by financial institutions and has used state-of-the-art global game techniques to link institutions’ balance

³JP Morgan estimated that Tether, the largest stablecoin, was one of the largest investors in the U.S. commercial paper market in June 2021. See <https://www.ft.com/content/342966af-98dc-4b48-b997-38c00804270a>. A FOIA request from the New York Attorney General’s office showed that Tether’s commercial paper portfolio in late 2021 was principally invested in Chinese financial companies (\$5.6b). Still, they also held \$1.4b of non-financial paper (mostly from European companies) and paper issued by three U.S. companies.

⁴See “Circle’s stablecoin banked at SVB and guess what happened next” (March 11, 2023). <https://www.ft.com/content/7c9b2234-c298-4508-b59a-fce49f6bc40a>.

⁵To keep our model tractable, we do not incorporate a channel through which leverage affects asset prices. However, such a channel could be incorporated in our analysis and could be an interesting avenue for future work.

sheet and profitability to the probability of a run (Goldstein and Pauzner 2005, Kashyap et al. 2024, Schilling 2023; and Morris and Shin 2003, Carlsson and van Damme 1993). However, this literature has not connected the run risk in institutions’ liabilities to the premium they can earn in secondary markets, which is a key focus and novelty of our paper. There is a good reason for this. The traditional liabilities of financial institutions are not bearer instruments, unlike private banknotes during the Free Banking Era. Tokenization renders such liabilities fungible so investors can use them for external activities to earn additional premia. We show how speculative motives in crypto markets matter for stablecoin peg stability using a global game model and empirical evidence. Bolt et al. (2023) show that a fixed exchange rate or peg can be stabilized using global game techniques in the context of a central bank (a fiat money issuer) failing.

More narrowly, our paper contributes to the emerging literature on stablecoin stability. A paper close and complementary to ours is Ma et al. (2023). They study the market mechanism through which stablecoins can be redeemed for fiat currency, as most investors cannot directly redeem with the issuer. They conclude that the number of authorized arbitragers allowed to interact with the stablecoin issuer and request redemptions is crucial for stablecoin stability. As we elaborate later, their mechanism could easily be incorporated into our framework without altering our key results. In addition, the focus of our paper is different as we are interested in the implications of crypto speculation, rather than of the redemption mechanism, for stablecoin peg-stability. Bertsch (2023) also considers a global game to model stablecoin run risk but focuses on the use of stablecoins in payments and their desirability over deposits. In addition to the differences in focus, our paper empirically examines the use of stablecoins in crypto speculation and tests the model predictions. d’Avernas et al. (2022) also study the stability of stablecoin pegs focusing on the role of commitment by the issuer, but abstract from run risk considerations, and their analysis is theoretical. Kozhan and Viswanath-Natraj (2021) show that using safer collateral has led to greater peg stability for the decentralized stablecoin DAI. Our paper features run risk, which is endogenized using a global-game method.

Our measurement of speculative demand for cryptocurrencies, including our instrumental identification, may be of independent interest to the literature, along with the connection between speculation, stablecoin size, and stability. We also contribute to the literature on asset pricing and the market structure of decentralized protocols and blockchains (Makarov and Schoar 2022, Lehar and Parlour 2021).

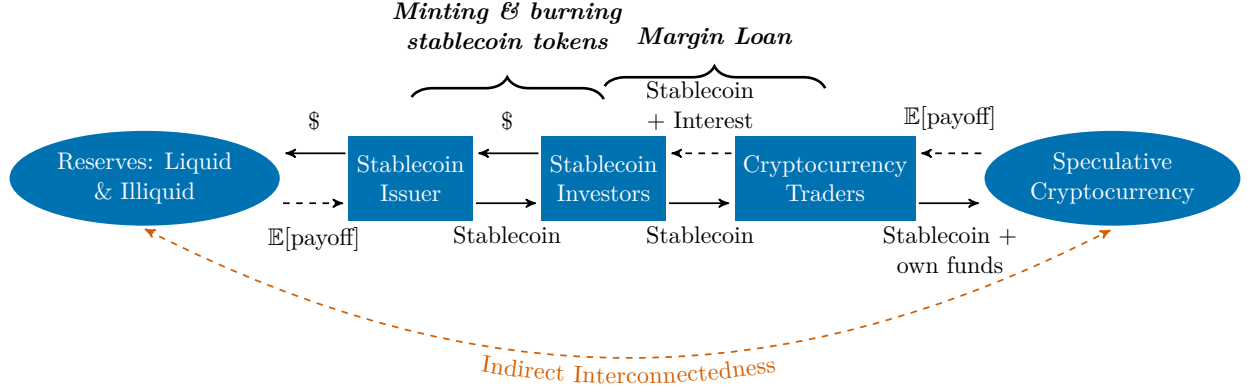


Figure 1: Model Sketch.

2 Model

Model Overview This section presents the model to study the leverage-money nexus summarized in Figure 1. Our model features as separate agents a profit-maximizing stablecoin issuer (Tether) and stablecoin investors who can re-use and lend the token in secondary markets. The stablecoin issuer raises funds from investors and invests them in a portfolio of liquid and illiquid assets. The liquid asset always trades at par, while the illiquid asset may trade at a discount. We use global game techniques to pin down a unique probability of a run that depends on the issuer's balance sheet. With this probability, we can compute the premium that stablecoin holders require to maintain a price of \$1 and connect it to the rate speculators offer to borrow the stablecoin. Higher speculative demand for cryptocurrencies increases the lending rate on stablecoins and, hence, their price. A less liquid portfolio of reserves increases run risk, reducing their price. The two forces interact by pushing the peg in opposite directions. Given a demand for crypto, the stablecoin issuer chooses the liquidity of its reserves and caters to investors' demand for stablecoins by issuing or redeeming tokens to maximize its profits while guaranteeing that the tokens will be traded at a price of \$1.

Model Setup There are three periods ($t = 0, 1, 2$), four assets, and three agents. Two of the assets are traditional, a liquid asset and an illiquid asset, and the other two are digital assets, a stablecoin and a cryptocurrency. All assets are perfectly divisible. The first type of agent is the stablecoin issuer and manager. The second type consists of a continuum of investors, which are identical ex ante, but heterogeneous ex post as described below. The third type consists of a continuum of traders that want to take a leveraged long position in the cryptocurrency.

The stablecoin raises funds in fiat currency from investors at $t = 0$ in exchange for tokens,

which are the liabilities of the stablecoin issuer and the first digital asset in the model. Denote by s the number of tokens in circulation. Given that the issuer will offer tokens at a unitary price, s also captures the total funds invested in the stablecoin. In turn, the stablecoin issuer invests the funds in a portfolio of traditional assets.⁶ Both assets are in perfectly elastic supply; their returns are denominated in fiat currency, and they can be bought for one unit of fiat currency at $t = 0$. The liquid asset yields a gross return of one at $t = 2$ and can be sold at any time before $t = 2$ for the price of one. The payoffs and the liquidation value of the illiquid asset depend on the realization of a fundamental state $\theta \sim U[0, 1]$ with its true value realized at $t = 1$, but not publicly revealed. If $\theta \geq \bar{\theta}$, with $\bar{\theta}$ exogenously set, the illiquid asset yields $X > 1$ at $t = 2$ with certainty and can be liquidated at $t = 1$ also for X . If $\theta < \bar{\theta}$, the illiquid asset yields $X > 1$ at $t = 2$ only with probability θ and zero otherwise, while its liquidation value drops to $\xi < 1$. The assumption about the liquidation value and payoffs of the illiquid asset follows Goldstein and Pauzner (2005). Without loss of generality and to simplify the notation when we derive the stablecoin price and the issuer’s optimization problem later on, we will consider $\bar{\theta} \rightarrow 1$. Finally, denote by ℓ the portion of the portfolio invested in the liquid asset. We will also assume that both liquid and illiquid assets can be resold between $t = 0$ and $t = 1$ —that is, before the shock on the illiquid asset—for the price of one to allow the issuer to rebalance their reserves portfolio, i.e., adjust ℓ if warranted.

Investors are identical ex ante and have deep pockets. At $t = 0$, each investor can hold tokens issued by the stablecoin in exchange for fiat currency. Tokens are initially issued in exchange for fiat currency and are redeemable on demand for fiat currency at a fixed exchange rate of one.⁷ Following Diamond and Dybvig (1983), an individual investor receives with probability δ an idiosyncratic preference shock urging them to redeem their tokens to use their funds for purposes outside the digital asset ecosystem. We will call these investors *impatient*. By the law of large numbers, the total expected redemption at $t = 1$ from impatient investors is equal to δs . The remaining $(1 - \delta)s$ investors do not have a pressing need to redeem but can decide to do so based on private noisy signals $x_i = \theta + \epsilon_i$, with $\epsilon_i \stackrel{iid}{\sim} U[-\epsilon, \epsilon]$, about the realization of θ at $t = 1$. We will call these investors *patient*.

The benefit to patient investors from not redeeming is that they can lend the tokens

⁶The issuer may also have chosen to invest in crypto assets, extend loans, or even invest in stablecoins of other issuers and lend them out for a profit. These alternatives may be important in practice and would be captured under the illiquid asset in our model. For our analysis, it only matters that the issuer may invest in illiquid assets and that the stablecoin tokens issued against these reserves can be re-used by investors.

⁷The assumption that investors can redeem their tokens directly with the stablecoin issues and at their face value is made for simplicity. One could introduce alternative redemption frictions in our model. For example, we could easily introduce a fee for redemptions or authorized arbitrageurs who stand between the stablecoin issuer and investors facilitating redemptions as in Ma et al. (2023). Both frictions would result in a lower, potentially state-contingent, payoff for redemptions. As such, the incentive to redeem would decrease, but all the mechanisms in our model and the implications of our analysis would remain unaltered.

to traders that want to take exposure to the cryptocurrency, expecting a gross return $y = \int \tilde{y} dH(\tilde{y})$, where \tilde{y} is the cryptocurrency return realization. Traders borrow the stablecoin from patient investors, and, combining it with their funds, they buy a cryptocurrency on margin, as described in detail below. For simplicity and without loss of generality, we make four assumptions, which can be relaxed. First, we assume that patient investors do not want to hold the cryptocurrency directly.⁸ Second, we assume the lending of tokens takes place after patient investors have decided whether to redeem and have learned whether the issuer is insolvent or not.⁹ Denote by R the expected return per unit of lending the token, which will be endogenously determined. Third, we assume traders have access to an outside option with gross return ρ . Fourth, the distribution of cryptocurrency returns is independent of the distribution of θ .

The funding structure of stablecoin issuers is fragile because the liquidation value of its reserves may not be enough to fully cover potential redemptions by all token holders. As such, the stablecoin issuer is exposed to run risk from self-fulfilling beliefs. This gives rise to multiple equilibria described in the bank run literature. To resolve this indeterminacy, we model a global game where each individual token holder receives a private noisy signal x_i and decides to redeem or not based on their posterior about θ and their beliefs about the actions of others. We will solve for a threshold equilibrium such that token holders decide to redeem if their signal x_i is below a threshold. Using this threshold, we can compute the ex ante probability at $t = 0$ that the stablecoin may experience a run at $t = 1$ and the price at which stablecoins will trade. A higher run probability pushes the price of tokens down, while a higher lending rate pushes the price up, other things equal.

We derive the expected return from lending the stablecoin token, and given this return, we compute the probability of a run on the stablecoin issuer and the stablecoin price at which investors would be willing to trade tokens. Thereafter, we examine and evaluate the peg-stabilization mechanisms in response to shocks.

Although our paper focuses on this specific use-case of stablecoins—the leverage-money nexus—our model can easily be adjusted to include additional services, such as payment services, as we show in an Online Appendix. Such services may become more important

⁸One way to microfound this assumption is to have investors that are less optimistic than traders about cryptocurrency returns, so they would rather lend to traders who would like to take leverage as in Fostel and Geanakoplos (2008), Geanakoplos (2010), and Simsek (2013). Alternatively, investors could be made sufficiently risk averse at $t = 2$ or incur costs when holding the cryptocurrency directly. We present an extension along the latter dimension in the Online Appendix in Section A.2.

⁹In practice, lenders can recall their loans within small intervals of time, for example, with one-hour notice in FTX or instantly, if the loan terms allow it, in Aave, justifying this assumption. Our analysis would be unaltered if lending also happened before $t = 1$, but investors quickly called back their stablecoins to redeem them. See the Online Appendix Section A.4 for an extension where lending takes place before $t = 1$, but the stablecoins are locked in until $t = 2$ corresponding to longer-term loans that cannot be quickly recalled.

as stablecoins mature and develop a convenience yield, which should be expected to be lower than the aforementioned lending rates.¹⁰ Similarly, it is straightforward to also allow the stablecoin issuer to pay interest directly to investors despite the fact this has not been the practice, as we mentioned earlier.¹¹ As expected, these additional non-pecuniary and pecuniary benefits would further boost the demand for stablecoin, allowing it to grow but, as we show, would not affect the novel mechanism we highlight in our paper, which is how speculative crypto demand interacts with stablecoin peg-stability. Hence, we have kept the baseline model exposition simple and relegated the modeling of additional stablecoin demand sources to the Online Appendix in section A.3.

2.1 Expected return from lending the stablecoin

In this section, we study how the expected return from lending stablecoin tokens is determined in equilibrium. We proceed backward, assuming that the stablecoin has not suffered a run at $t = 1$ and that the realization of θ has become common knowledge. Patient token holders can lend their tokens to traders who want to take a leveraged position in the cryptocurrency. For each dollar of cryptocurrency they buy, traders need to post m percent of margin. Assume the cryptocurrency exchange exogenously sets m . This is not critical for our results. We show in the Online Appendix Section A.2 how the derivation of m can be endogenized.¹²

The expected payoff to the trader from a leveraged position in one dollar of the cryptocurrency is $y - (1 - m)R$. Note that R is the expected return to lending—not the contractual lending rate—that incorporates the case the trader defaults. To see this, denote by R_c the contractual lending rate and suppose $\tilde{y} \sim H(\tilde{y})$ with $y = \int \tilde{y} dH(\tilde{y})$. The trader defaults if $\tilde{y} < y' \equiv (1 - m)R_c$. Conditional on issuer’s solvency, the expected return to the trader is $\int_{\tilde{y} \geq y'} (\tilde{y} - (1 - m)R_c) dH(\tilde{y}) = y - \int_{\tilde{y} < y'} \tilde{y} dH(\tilde{y}) - (1 - m)R_c \int_{\tilde{y} \geq y'} dH(\tilde{y}) = y - (1 - m)R$, because investors receive the collateral for $\tilde{y} < y'$. Given risk-neutrality, we focus on R but control for cryptocurrency volatility in our empirical analysis, where we use contractual borrowing rates to account for the difference between R and R_c .

¹⁰Van den Heuvel (2022) estimates the convenience yield of deposits over same maturity Treasuries, which should partially capture the convenience yield from payments, to be about 80 basis points in recent years. Our model implies that reliance on smaller convenience yields should push the issuer toward a safer and more liquid reserves portfolio, which has been the case for stablecoins that attempt to specialize in payments.

¹¹Traditionally, stablecoins have not paid interest, but more recently, crypto firm Figure is seeking SEC approval to issue an interest-bearing stablecoin. In Europe, the Markets in Crypto Assets (MiCA) legislation prohibits paying interest on stablecoin tokens.

¹²An additional advantage of buying the cryptocurrency with a stablecoin token is that there are small or no haircuts on the pledgeable dollar value. Thus, the dollar value that traders must post to meet margin m is equal to m under a zero haircut, and they only need to borrow $1 - m$ stablecoin tokens per dollar of exposure to the cryptocurrency. For other cryptocurrencies, the haircuts are higher, and traders must post a higher dollar value than m to meet the margin requirement (see Table A.1 in the Online Appendix).

Traders compete and, in equilibrium, offer an R that makes them break even with their outside option ρ , which depends on the aggregate funds invested in the alternative technology, but traders take it as given. If they break even, traders will not invest in the outside option. The equilibrium R is then given by equating levered profits per unit of investment, $(y - (1 - m)R)/m$, to the unlevered outside option profit per unit of investment, ρ , or

$$R = \frac{y - m\rho}{1 - m}. \quad (1)$$

We assume that the outside option consists of a technology, F , common to all traders, with decreasing marginal returns depending on the aggregate amount of funds invested ($F' > 0, F'' < 0$). Denote by e the total funds of traders and by $m(1 - \lambda)s/(1 - m)$ the total funds invested in leveraged cryptocurrency trades, where λ is the number of tokens redeemed at $t = 1$ and not available for lending. Then, ρ in equilibrium is given by

$$\rho = F' \left(e - \frac{m}{1 - m}(1 - \lambda)s \right). \quad (2)$$

We assume that the lower bound for the outside option, when $s = 0$, satisfies $F'(e) > y$, such that traders have an incentive to use leverage, resulting in $R < y$.¹³ Moreover, there might be a \bar{s} such that $y - m\rho < 0$ for $s > \bar{s}$. The number of tokens in circulation may be so high that there is no benefit to lending them. Thus, we will also assume that \bar{s} is high enough such that there is room for the number of tokens to adjust while still paying positive interest when lent out.

Combining 1 and 2, we get that the lending rate as a function of outstanding tokens $(1 - \lambda)s$ at $t = 2$, where λ is the percentage of early redemptions:

$$R(\lambda, s) = \frac{y - mF' \left(e - \frac{m}{1 - m}(1 - \lambda)s \right)}{1 - m}. \quad (3)$$

Before determining how runs on the stablecoin issuer occur, we present the following comparative statics that will be important for what will follow in Sections 2.2 and 2.4 (see Section A.1.1 for detailed derivations):

$$\frac{dR(\lambda, s)}{dy} > 0 \ \& \ \frac{dR(\lambda, s)}{dm} < 0 \ \& \ \frac{dR(\lambda, s)}{ds} < 0 \ \& \ \frac{dR(\lambda, s)}{d\lambda} > 0. \quad (4)$$

Changes in y and m could be interpreted as higher demand for cryptocurrencies and

¹³Otherwise, the relevant outside option may be $\rho = y$ since traders would prefer to invest directly in the cryptocurrency. In that case, $R = y$, which is independent of the stablecoin supply s , and the redemption channel for peg-stability would be absent.

higher cryptocurrency volatility or risk, respectively. Both changes affect the lending rate R , destabilizing the stablecoin peg. Changing the number of tokens, s , is one way to undo the change in the lending rate and re-stabilize the peg; for example, a decrease in s following a decrease in y could bring the lending rate back to its original value. We elaborate on this stabilization mechanism in detail in section 2.4 after we show how the liquidity of the stablecoin's reserves matters for run risk and, thus, for its peg stability.

2.2 Probability of a stablecoin run

Redemptions and Issuer's Portfolio. The issuer collects funds at $t = 0$ equal to s , invests ℓ percentage of them in the liquid asset, and invests the rest in the illiquid asset. If the issuer defaults, the stablecoin holders cannot lend their tokens and earn a lending rate but are distributed pro-rata the remaining assets of the issuer. The expected payoff to an individual patient investor if only impatient investors redeem, i.e., $\lambda = \delta$, is equal to $\theta R(\delta, s) + (1 - \theta) \max(\ell - \delta/1 - \delta, 0)$, since the issuer defaults with probability $1 - \theta$.

Yet, the issuer may fail even when the illiquid asset pays out X , which occurs with probability θ , if enough patient investors decide to redeem their tokens. If $\theta \geq \bar{\theta}$, the issuer has enough liquidity to serve all possible redemptions. If $\theta < \bar{\theta}$, every individual patient investor will redeem independent of what other investors choose to do.¹⁴ For intermediate realizations of $\theta \in [\bar{\theta}, \bar{\theta})$, the issuer does not have enough liquidity to serve all possible redemptions because the liquidation value of the illiquid asset drops to $\xi < 1$. The liquidity position of the issuer at $t = 1$ is $L(\lambda) = [\ell + (1 - \ell)\xi - \lambda]s$, where λs are the total redemptions, with $\lambda \in [\delta, 1]$, depending on how many patient investors decide to redeem. Hence for $\lambda > \bar{\lambda}(\xi)$ the stablecoin issuer does not have enough liquidity to serve all redemptions, where $\bar{\lambda}$ is the solution to $L(\bar{\lambda}) = 0$, i.e.,

$$\bar{\lambda} = \ell + (1 - \ell)\xi. \quad (5)$$

Conditional on having enough liquid resources ($\lambda \leq \bar{\lambda}$) and with probability θ , the profits of the stablecoin issuer at $t = 2$ as a function of λ are given by

$$\Pi(\lambda) = \left[X(1 - \ell) \left(1 - \frac{\max(\lambda - \ell, 0)}{\xi(1 - \ell)} \right) + \max(\ell - \lambda, 0) - (1 - \lambda) \right] s. \quad (6)$$

That is, the issuer extracts all seigniorage after repaying remaining stablecoins at par. The issuer first uses the liquid asset for redemptions and then starts liquidating the illiquid asset. This is optimal since the liquid asset has a higher liquidation value, while the risky one has a

¹⁴ $\bar{\theta}$ is given by $\frac{[1 - \max(\ell - \delta/1 - \delta, 0)]}{[R(\delta, s) - \max(\ell - \delta/1 - \delta, 0)]}$.

higher expected payoff. For $\lambda \leq \ell$, $\Pi(\lambda) > 0$. For higher λ and because $d\Pi(\lambda)/d\lambda < 0$, the stablecoin issuer becomes insolvent at $t = 2$ for $\lambda > \hat{\lambda}(\xi)$, given by $\Pi(\hat{\lambda}(\xi)) = 0$, i.e.,

$$\hat{\lambda} = \frac{X(\ell + \xi(1 - \ell)) - \xi}{X - \xi}. \quad (7)$$

Moreover, we can re-write $\hat{\lambda}(\xi) = (X\bar{\lambda}(\xi) - \xi)/(X - \xi) < \bar{\lambda}(\xi)$, i.e., the issuer becomes insolvent before running out of liquidity.

Finally, note that conditional on having enough liquid resources ($\lambda \leq \bar{\lambda}$) and with probability $1 - \theta$, the issuer always defaults, and any unused liquid assets $\max(\ell - \lambda, 0)$ are distributed pro-rata to the remaining $1 - \lambda$ investors.

Redemptions and Investor Payoffs. A patient investor needs to decide at $t = 1$ whether to redeem their token. The payoff differential between not redeeming and redeeming depends on the beliefs about θ and λ :

$$\nu(\theta, \lambda) = \begin{cases} \theta R(\lambda, s) + (1 - \theta) \max\left(\frac{\ell - \lambda}{1 - \lambda}, 0\right) - 1 & \text{if } \delta \leq \lambda \leq \hat{\lambda} \\ \theta \frac{X(1 - \ell)(1 - \frac{\lambda - \ell}{\xi(1 - \ell)})}{1 - \lambda} - 1 & \text{if } \hat{\lambda} < \lambda \leq \bar{\lambda} \\ -\frac{\ell + (1 - \ell)\xi}{\lambda} & \text{if } \bar{\lambda} < \lambda \leq 1 \end{cases} \quad (8)$$

If the belief about λ is below $\hat{\lambda}$, an investor that does not redeem receives the payoff from lending out the token, $R(\lambda, s)$, with probability θ . If the issuer defaults, with probability $1 - \theta$, the residual assets are distributed pro-rata, yielding $\max((\ell - \lambda)/(1 - \lambda), 0)$ per investor. Redeeming yields one dollar, since the issuer has enough liquidity to meet redemptions. For $\lambda \in (\hat{\lambda}, \bar{\lambda}]$, investors that do not redeem receive pro rata the remaining assets in insolvency; investors that redeem receive one dollar since the issuer has enough liquidity to serve all redemptions. The issuer depleted all liquid assets for this level of λ , so investors are only distributed the proceeds from the remaining illiquid assets at insolvency. Finally, for $\lambda > \bar{\lambda}$, the benefit from not redeeming is zero, as the stablecoin issuer will be fully liquidated; the benefit from redeeming is joining the line in the run and being able to redeem at par with probability $(\ell + (1 - \ell)\xi)/\lambda$, according to sequential servicing. Figure 2 plots for a certain parametrization and some value of θ , the payoff differential $\nu(\theta, \lambda)$ as beliefs about λ vary.

Defending the Peg. Figure 2 is useful to understand the (in)ability of the issuer to defend the peg. For $\lambda \leq \hat{\lambda}$, the issuer will first liquidate the liquid asset and then the illiquid asset to meet redemptions. Thus, investors can earn the expected lending rate $R(\lambda, s)$, which is




Figure 2: Payoff differential $\nu(\theta, \lambda)$ as beliefs about λ vary

increasing in the number of redemptions. A higher lending rate is a stabilizing force that helps the issuer defend the peg. But for $\lambda > \hat{\lambda}$, the issuer becomes first insolvent and then illiquid, and cannot defend the peg, because the liquidation value of the illiquid asset drops to $\xi < 1$. The decision to redeem, which we derive next, will depend on beliefs about total redemptions and, hence, the ability of the issuer to defend the peg.

Redemption Decision and Determinants of Run Risk. Following standard steps from the global games literature we can derive the following condition that determines a unique run threshold θ^* (see Section A.1.2 for detailed derivations and the proof of existence and uniqueness of θ^*):

$$\begin{aligned} \bar{\Delta}^* = & \int_{\delta}^{\hat{\lambda}} \left[\theta^* R(\lambda, s) + (1 - \theta^*) \max \left(\frac{\ell - \lambda}{1 - \lambda}, 0 \right) - 1 \right] \frac{d\lambda}{1 - \delta} \\ & + \int_{\hat{\lambda}}^{\bar{\lambda}} \left[\theta^* \frac{X(1 - \ell) \left[1 - \frac{\lambda - \ell}{\xi(1 - \ell)} \right]}{1 - \lambda} - 1 \right] \frac{d\lambda}{1 - \delta} - \int_{\bar{\lambda}}^1 \frac{\ell + (1 - \ell)\xi}{\lambda} \frac{d\lambda}{1 - \delta} = 0. \end{aligned} \quad (9)$$

The run threshold θ^* implies a run probability θ^* , which is a function of the lending rate R given in (1) and the ratio of liquid assets in total assets, ℓ . θ^* changes with y , m , s , and ℓ as follows (see Section A.1.3 in the Online Appendix for detailed derivations):¹⁵

$$\frac{d\theta^*}{dy} < 0 \ \& \ \frac{d\theta^*}{dm} > 0 \ \& \ \frac{d\theta^*}{ds} > 0 \ \& \ \frac{d\theta^*}{d\ell} < 0. \quad (10)$$

¹⁵ $\frac{d\theta^*}{d\ell} < 0$ is true under a weak sufficient condition derived in the Online Appendix, supported by the data.

2.3 Stablecoin Price

In this section, we compute the price for one stablecoin token given the lending rate R derived in Section 2.1 and the run threshold θ^* derived in Section 2.2.

Given that stablecoin tokens are traded in secondary markets, we compute the price at which investors are willing to trade their stablecoin tokens. Note that this price is different from the cost of getting a token from the issuer or the payoff from (successfully) redeeming a token at the issuer, both of which are set to one. Moreover, we derive the stablecoin price before the realization of θ and the resolution of uncertainty about the possibility of a run.

Denote by P the market price of traded stablecoin tokens, which is given by

$$P = \int_{\theta^*}^1 \left\{ (1 - \delta) \left[\theta R(\delta, s) + (1 - \theta) \max \left(\frac{\ell - \delta}{1 - \delta}, 0 \right) \right] + \delta \right\} d\theta + \int_0^{\theta^*} (\ell + (1 - \ell)\xi) d\theta, \quad (11)$$

as $\bar{\theta} \rightarrow 1$. The market capitalization of the stablecoin is equal to $P \cdot s$.

The first term in (11) is the expected payoff conditional on no run on the issuer. For patient investors, this is equal to the expected value of being able to lend out the token, and the expected repayment should the issuer default, while for impatient ones, it is equal to the par value of the token. Note that the lending rate is equal to $R(\delta, s)$ because δ impatient investors have redeemed their tokens at $t = 1$. The second term is the payoff conditional on a run, which is equal to the liquidation value of the issuer's asset portfolio for a dollar of tokens held and is the same for patient and impatient investors as all withdraw in a run.

Hence, P reflects investors' valuation for one token and is equal to the secondary market price that investors would be willing to trade stablecoins among themselves *before* the realization of uncertainty at $t = 1$. Note that during a run, our model predicts that the secondary market price should drop to $\ell + (1 - \ell)\xi$, i.e., the expected payoff from redeeming a token. As already mentioned, trading frictions may exist for redeeming tokens, for example, due to the presence of authorized arbitragers (Ma et al. 2023), under which the payoff for $\theta < \theta^*$ in (11) would need to be adjusted. As mentioned, we abstract from such considerations to keep the analysis simple and leave these interesting extensions for future work.

The effect of the demand and riskiness of cryptocurrencies, as well as the size and liquidity of the stablecoin on P , are summarized as follows (see Section A.1.4 for detailed derivations):

$$\frac{dP}{dy} > 0 \ \& \ \frac{dP}{dm} < 0 \ \& \ \frac{dP}{ds} < 0 \ \& \ \frac{dP}{d\ell}. \quad (12)$$

Intuitively, a higher y , and lower m or s , increases the payoff conditional on a run not occurring and decreases the probability of a run, both of which push P up. Similarly, a higher

ℓ decreases the probability of a run but also increases the probability of being paid conditional on a run occurring, both of which push P up. In sum, the stablecoin price may fluctuate not only in response to shocks in the expected return and riskiness of the cryptocurrency but also due to adjustments in the size and liquidity of the stablecoin. In section 2.4, we discuss the mechanisms through which the size and liquidity of the stablecoin can stabilize the peg in response to crypto shocks.

2.4 Peg stability

We show how the issuer can maintain the peg in response to shocks. We make a distinction between stabilizing the peg, discussed in this section, and defending the peg, discussed in section 2.2. We view stabilizing the peg as the actions the issuer can take to maintain the peg before the realization of fundamentals' uncertainty, that is, between $t = 0$ and $t = 1$. By contrast, defending the peg corresponds to the (in)ability of the issuer to survive a run at $t = 1$ when uncertainty about θ and ξ is realized. This distinction is important because stablecoins are traded continuously, and they may trade above or below their peg even outside run episodes, which are only characterized by peg devaluations. At the center of the distinction is whether the issuer incurs portfolio rebalancing costs when changing the share of liquid assets held in reserves. Such costs are captured in the model by a potentially lower liquidation value for the illiquid asset after, but not before, $t = 1$.

As such, we focus here on how crypto shocks that may arrive between $t = 0$ and $t = 1$ can destabilize the peg, that is, move P above or below 1. Given that the liquid and illiquid assets can be sold for one before $t = 1$ and there are no other portfolio rebalancing costs, we do not need to track the portfolio allocation of the stablecoin issuer before the shock; we only need to compute the new portfolio allocation after the shock, resembling a comparative statics exercise. We could easily complicate the analysis by introducing an intermediate period between $t = 0$ and $t = 1$ when crypto-related shocks materialize and track the portfolio allocations, yet the results we derive below for peg stability would remain unaltered.

There are two mechanisms to maintain the peg. The first mechanism relates to how ℓ should vary and proxies for how close to money stablecoins are. The second mechanism relates to how s should adjust and is driven by the usefulness of stablecoins in leveraged crypto trades as captured by the lending rate R . The issuer directly controls ℓ but passively issues or redeems tokens based on investors' demand for stablecoins, which affects s . If the price of stablecoins in secondary markets is higher than 1, then investors will purchase more tokens from the issuer at a price of 1 and sell them at the market to make a profit, which increases s . If the price in secondary markets is lower than 1, then investors will redeem with the issuer to receive 1 and make a profit, which decreases s . Hence, the issuer sets the price

for issuing/redeeming tokens and caters to investors' demand. However, when choosing ℓ , the issuer understands how investors will react—demanding new tokens or redeeming their existing ones,—which continues until the secondary price is equal to the price that the issuer issues/redeems tokens, i.e., equal to 1. In other words, the issuer can infer the level of ℓ , given a level of s , that is needed to stabilize the secondary price to 1. This problem is equivalent to one under which the issuer jointly chooses ℓ and s subject to a constraint that the secondary price is equal to 1. As such, we will be referring below to the issuer choosing jointly ℓ and s . It should be clear that the issuer does not set directly s , but can choose it indirectly.

Before deriving the optimal choices, we examine how changes in y and m between $t = 0$ and $t = 1$ matter for how ℓ and s should be adjusted to maintain $P = 1$. A change in y , with m constant, could be interpreted as higher demand for cryptocurrencies. A change in m , with y constant, could be interpreted as higher volatility or risk for cryptocurrencies.

Proposition 1. *A decrease in demand for the cryptocurrency or an increase in its riskiness pushes the stablecoin price below its peg. To stabilize the peg, the stablecoin issuer can either increase ℓ keeping s the same, keep the same ℓ and allow lower demand for tokens to manifest in more redemptions and lower s , or jointly increase ℓ and decrease s . The opposite is true for higher cryptocurrency demand or lower riskiness.*

The proof of Proposition 1 is straightforward. From equation (A.17), we know that a decrease in y or an increase in m results in a lower price of the tokens P . Therefore, ℓ needs to increase to maintain the peg for a certain s (equation A.18). Alternatively, the issuer may keep the same ℓ . This results in lower demand for the stablecoin tokens, and more investors will want to redeem their tokens. A lower number of tokens, s , results in a higher P , and the stablecoin price can return to its peg.

At the heart of the stabilization mechanism to crypto shocks is the interplay between prices and quantities: the issuer caters to the demand of investors by issuing more tokens at one dollar when the secondary market price P is greater than one and redeeming tokens at a dollar when P is lower than one. This hard-wired adjustment in the quantity of tokens outstanding offers the issuer an additional degree of freedom to choose the level of ℓ that maximizes their seigniorage but at the same time introduces volatility in the quantity of tokens outstanding which we exploit in the empirical analysis in Section 4.

As an aside, recall that the issuer cannot freely adjust ℓ to defend the peg at $t = 1$ in the cases that the liquidation value of the illiquid asset drops to $\xi < 1$, introducing portfolio rebalancing costs (see Section 2.2). The stabilization mechanism via ℓ would go in the opposite direction as the issuer would first sell the liquid asset to meet redemptions, effectively increasing the share of illiquid assets in reserves. The stabilization mechanism via s would

still be operational, but its effectiveness would vanish after the level of withdrawals that pushes the issuer to insolvency; we show in Figure 2 that the payoff from not redeeming increases as s decreases, as long as the issuer remains solvent. Defending the peg at $t = 1$ should not be confused with stabilizing the peg in normal times studied below.

Token Supply and Stablecoin Liquidity under Observability. We derive the equilibrium ℓ and s when both ℓ and s are observable. Observability implies that the issuer can write a complete contract, such that the issuer internalizes how ℓ and s affect the stablecoin price P . Absent other frictions that may be important in practice but from which we have abstracted, observability delivers the optimal solution under commitment and, thus, the issuer would want to make ℓ observable (and verifiable). Doing so would result in higher profits compared to the solution under non-observability and lack of commitment. We elaborate below on developments in the digital asset ecosystem that can push toward the observability of reserves.¹⁶ The case that ℓ is unobservable is derived in the Online Appendix Section A.5.¹⁷ Our key points continue to hold under unobservable ℓ , but, contrary to the case of observability, the stablecoin is not viable for low enough levels of speculative demand.

The issuer maximizes their profits from seigniorage (recall that $\bar{\theta} \rightarrow 1$),

$$\max_{\ell, s} \int_{\theta^*}^1 \theta \Pi(\delta) d\theta, \quad (13)$$

subject to $P = 1$ as discussed above. $\Pi(\delta)$ are the profits when only impatient investors withdraw, and the issuer does not default—with probability θ —given by (6) for $\lambda = \delta$.

The issuer internalizes how ℓ and s affect the run threshold and the lending rate, which can be expressed as functions of ℓ and s using (9) and (3) and substituted in the issuer’s optimization problem. Combining the optimality conditions for ℓ and s yields

$$\frac{1 - (\theta^*)^2}{2} \left(\frac{d\Pi(\delta)}{d\ell} - \Pi(\delta) \frac{dP/d\ell}{dP/ds} \right) + \theta^* \Pi(\delta) \left(\frac{d\theta^*}{ds} \frac{dP/d\ell}{dP/ds} - \frac{d\theta^*}{d\ell} \right) = 0, \quad (14)$$

which together with the *peg stability condition* $P = 1$ yields the optimal (ℓ, s) .

As mentioned, P can deviate from 1 in response to shocks until the issuer resets ℓ and s to restore the peg. The optimal ℓ and s derived above require observability. In principle, it is feasible for stablecoin arrangements to be backed by on-chain assets—either other cryptocurrencies or tokenized traditional financial assets—such that ℓ is observable in real-time along with s . Yet, the biggest stablecoins are currently backed by off-chain financial

¹⁶See Kashyap et al. (2024) for similar modeling of complete and incomplete deposit contracts with run risk.

¹⁷ s is always observable, given that the number of tokens in circulation is reported on the blockchain in real-time.

assets and disclose their reserves only infrequently, at best.¹⁸ Thus, the peg stabilization studied above would work when ℓ is observable at disclosure dates. Between these dates, peg stability after a shock could again be achieved, but the issuer will have an incentive to deviate from the choice of ℓ studied herein, discussed in the Online Appendix as mentioned.

Finally, using (14) we can show that the issuer will optimally choose $\ell < 1$ and thus expose the stablecoin to runs, as long as speculative demand to crypto is high enough and, thus, the return from lending the stablecoin is higher than 1. First, consider that $y > 1 + m[F'(e) - 1]$ such that $R(\delta, s) > 1$ for some $s > 0$. Now, suppose that $\ell = 1$ such that $\theta^* \rightarrow 0$. Then, (14) is negative, because $d\Pi(\delta)/d\ell < 0|_{\ell=1}$. Intuitively, the issuer makes zero profits for $\ell = 1$, so accepting some run risk by decreasing ℓ is optimal; and investors' participation constraint is not violated because $R(\delta, s) > 1$ can support some level of run risk. Next, consider $y \leq 1 + m[F'(e) - 1]$, which means that $R(\delta, s) \leq 1$ for any s . Investors would never hold the stablecoin for the purpose of lending it out and will only keep it if the run probability is zero and the token is always worth 1. The issuer can guarantee that by setting $\ell = 1$; otherwise investors would immediately redeem their tokens if there were a shock pushing y below the aforementioned threshold.¹⁹ Proposition 2 summarizes these results.

Proposition 2. *Consider that both ℓ and s are observable. For $y > 1 + m[F'(e) - 1]$, $\ell < 1$ is optimal. Otherwise, the issuer sets $\ell = 1$.*

The aforementioned result hinges upon two modeling choices. First, the liquid asset pays zero interest in our baseline analysis. It is easy to see that the issuer's profits would be positive for $\ell = 1$ if the liquid asset paid interest. The reason is that the issuer pays no interest to stablecoin investors and, thus, extracts all seigniorage. Still, by continuity, there are positive (sufficiently close to zero) interest rates such that the issuer will optimally choose $\ell < 1$. However, for high enough interest rates, the issuer may prefer to avoid taking any run risk and set $\ell = 1$. We have included in the Online Appendix section A.6 analysis that confirms these additional results. Our empirical analysis later on considers the period of December 2020 to November 2022, during which interest rates were (near) zero until March 2022 and started increasing thereafter. This configuration of interest rates is consistent with our model prediction that $\ell < 1$ as well as the stylized fact that stablecoin issuers did not invest solely in liquid assets during this period (see section 4). However, as the interest rate on liquid assets kept increasing and stabilized at higher levels in 2024, stablecoin issuers started

¹⁸There might be pressure from the industry to start disclosing the composition of reserves more frequently. Following the collapse of UST and the run on USDT in May 2022, USDC has started reporting their reserves weekly, potentially putting pressure on other stablecoin issuers to follow suit in the future.

¹⁹Although not explicitly modeled in order to keep the analysis simple, one could think of reasons the issuer may prefer setting $\ell = 1$ rather than having all investors redeem, presumably because the negative shock on y is transitory and scaling up the stablecoin from scratch may entail considerable fixed costs.

holding predominantly safer liquid assets such as Treasuries and Treasury repo, consistent with the analysis in the Online Appendix. It is an open question whether issuers will start investing again in less liquid assets should interest rates normalize to lower levels and crypto speculation pick up.

Second, we have assumed a monopolistic issuer with full market power. By choosing ℓ the issuer indirectly controls the demand for tokens s and, thus, the level of compensation to investors via the lending rate. Competition from other issuers would result in safer stablecoins and a higher supply of tokens. To see this, consider the following experiment. Start with a monopolistic issuer that chooses an observable $\ell < 1$ as outlined in Proposition 2. Suppose now that a second issuer enters the market and issues tokens that are perfect substitutes for the ones of the incumbent issuer. Then, the entrant can choose an observable $\ell' > \ell$ attracting all the investors from the incumbent given the lower run risk. The incumbent can respond by also choosing ℓ'' and so on until both issuers choose $\ell = 1$. In equilibrium, the total supply of tokens will also increase as issuers choose higher ℓ , reaching a saturated level. It should be noted that the aforementioned argument relies on the assumption that the tokens of competing issuers are perfect substitutes. Under imperfect substitutability, competition may result in higher ℓ than the one in Proposition 2 but still lower than one. Indeed, the current stablecoin market dynamics may suggest a bifurcation between the two biggest stablecoins, Tether and USDC, that try to cater to different types of clientele and use cases. Section A.3 in the Online Appendix shows how our model can be extended to encompass additional use cases that can differ by stablecoin. As such, individual issuers can maintain some market power, and our mechanism can remain operational even under stablecoin competition.

A corollary of Proposition 2 under additional use cases is that imperfectly substitutable stablecoins would choose different levels of ℓ and, thus, be exposed to different levels of run risk parameterized by θ^* . Then, it is conceivable that stablecoin investors would prefer to fly from riskier stablecoins to safer ones during a market turmoil instead of flying back to fiat. The reason is that a realization of θ that makes the riskier stablecoins fragile may not be severe enough to generate concerns for the safer ones. These flight-to-quality dynamics were apparent during the TerraUSD 2022 debacle when investors flew from Tether to USDC (Liao, 2022) or from USDC to Tether during the Silicon Valley Bank (SVB) crisis in March 2023 out of concerns about the uninsured deposits Circle held at SVB.²⁰

²⁰“Crypto’s brush with disaster after SVB collapse”, March 17, 2023, <https://www.ft.com/content/f48999ce-6237-48e9-aaab-3926d0c80797>.

3 Institutional Details and Data

3.1 Leverage

Leverage is a critical feature of crypto markets. Crypto traders often speculate with leverage, and centralized exchanges provide leverage as a key service. Crypto traders can get leverage in several ways. We focus on two products offered by centralized exchanges: margin trading and futures derivatives. There are several other mechanisms to get leverage at centralized exchanges and on the blockchain: levered tokens and options, to name a few. We focus on futures and margin trading as they are two long-standing and large sources of leverage. While data on the levered trading volumes are scarce, open interest across all crypto derivatives was \$250 billion in June 2022, the vast majority of which likely comes from perpetual futures derivatives.²¹ Data on margin lending are even more incomplete but likely exceed tens of billions of dollars.

Margin Trading. Margin in crypto is like margin trading in traditional finance: levered traders borrow the coin for a specified time at a given interest rate and use it with their own funds to take a position on a cryptocurrency. Margin trading can be used to take long or short positions. The main difference from traditional margin is that offshore crypto exchanges generally do not comply with Regulation T or other similar requirements.

Futures Derivatives. Traders can also get leverage using futures—not unlike traditional finance futures—many of which are *perpetual* futures that do not have an explicit expiration date, as indicated by their name. Perpetual futures are likely the largest and most liquid type of offshore cryptocurrency derivatives and, at times, can offer more than 100 times leverage.

For a traditional vanilla future, the future and spot prices converge as the expiration date approaches. Such phenomena do not happen with perpetual futures. Instead, perpetual futures use a *funding premium* to keep the spot and future price linked. If the future trades at a premium to the spot price, the investors that are long the future must pay a funding premium to investors that are short. On the rare occasion that the future price trades at a discount to the spot price, the investors short the future must pay a funding premium to investors that are long. For simplicity, we will say the funding premium is positive when investors that are long the future pay a fee to investors that are short. The details vary across exchanges, but the funding payments are paid daily or more frequently. Perpetual futures are typically stablecoin-settled, meaning that the perpetual future is quoted and settled in a

²¹<https://coinmarketcap.com/rankings/exchanges/derivatives/>

stablecoin, and funding payments are paid in the stablecoin. A BTC/USDT perpetual future, for example, is settled in USDT.

Investors may prefer to obtain leverage from either futures or margin trading. Margin trading has two advantages. First, the borrowed coins are fungible and can be used to settle spot transactions. Second, margin trading allows investors to take levered trading positions at spot market prices. Alternatively, futures allow investors to take a levered position in a coin, but they do not obtain the underlying coin until the future’s expiration unless it is a perpetual future, in which case the investor never receives the underlying coin. Limits to arbitrage cause persistent dislocations, often preventing the future price from equaling the spot price. Futures, however, are generally larger markets and allow levered exposure for extended periods.

3.2 Data

Our analysis will focus on the largest stablecoin, Tether, and on centralized exchanges (CEX), such as Binance and FTX (our analysis ends just before FTX’s collapse). There are several reasons for this. First, centralized exchanges remain the most popular way to trade crypto-assets on secondary markets, even though decentralized (DeFi) exchanges have grown in popularity over the past few years (see Watsky et al. 2024). Second, Tether has been by far the largest stablecoin, and it is used considerably more for crypto trading and speculation than the next largest collateralized stablecoin, USDC. In addition, the correlation between Tether coins in circulation and crypto asset returns is non-trivial and much more pronounced than the correlation between USDC coins in circulation and crypto asset returns.²² Third, Tether has not paid interest even though it has been scrutinized and questioned about its non-trivial run risk. This fact supports our choice of Tether for our empirical analysis. Finally, Tether is the predominant stablecoin used to quote and settle perpetual futures, which enable speculative levered bets on cryptocurrency. For example, 82 percent of derivative trading on Binance—the biggest centralized exchange—was settled in Tether, with daily volume of \$39 billion in April 2022. By comparison, 7.6 percent was settled in BUSD and 0 percent in USDC.

We collect prices, volume, and market capitalizations of cryptocurrencies from CoinGecko.²³ We collect margin lending rates from FTX. We focus mainly on Tether (USDT), but we also show results for Dai (DAI), which is the third biggest stablecoin, though much smaller. We

²²See Liao et al. (2023) for a comparison between Tether and USDC relating to crypto trading and speculation as well as König (2023).

²³For details on how Coingecko aggregates information across several exchanges to calculate prices, see <https://www.coingecko.com/en/methodology>.

do not include USDC in our analysis because it is not widely used in centralized exchanges and less so for crypto speculation.²⁴

How big is stablecoin margin lending? It is difficult to estimate precisely, as data from centralized exchanges is unavailable for many of the largest exchanges, implying that estimates are subject to uncertainty. Still, our limited information suggests that USDT’s total margin lending has been, on average, about \$10 billion, amounting to between 10 and 20 percent of the total market capitalization of USDT.²⁵ However, note that a large portion of USDT outstanding is held in inactive or “cold” wallets, which may account for 50% of total wallets. If we account for USDT actively used, it, effectively doubles the percentage of USDT used for lending to between 17 and 34 percent.²⁶ This percentage increases further if we focus on USDT held at centralized exchanges as calculated by GlassNode, which peaked at \$15 billion in June 2022. This last figure suggests that more than half of USDT held at CEX—and thus for the purpose of trading—were used for lending.

We collect perpetual future funding premia from Binance, which is likely their largest market.²⁷ We collapse higher-frequency data on lending rates, funding premia, and prices to a daily frequency using daily averages after converting them to U.S. Eastern time.²⁸ We use implied volatility for BTC and ETH calculated by T3 using option prices. Our sample runs from December 1, 2020, to November 5, 2022. The sample does not include the period of FTX’s collapse, which began on November 6, 2022.²⁹

Table 1 presents the summary statistics for the main variables, including the stablecoin prices and lending rates and the funding premium for BTC/USDT and ETH/USDT. The

²⁴In the Online Appendix Section A.7, we show that FTX’s lending rates are highly correlated with other lending rates in the digital asset ecosystem and so are not idiosyncratic to FTX. Moreover, Table A.2 in the Online Appendix regresses FTX’s Tether lending rates on decentralized finance (defi) platforms’ lending rates. In general, FTX’s lending rates are more correlated with the largest lending platforms, as measured by total value lock. It should be noted, however, that a direct comparison between margin loan contracts in centralized exchanges and decentralized lending platforms is complicated because of different degrees of collateralization as well as underlying liquidity; as explained earlier, Tether is mostly used for taking leverage in centralized exchanges.

²⁵As mentioned, CEX trading dominates DeFi trading during our sample period. Our calculations are based on FTX data and assume that FTX’s share of total centralized exchange trading volume (using data from Cryptocompare) equals its share of total margin lending across centralized exchanges.

²⁶This definition of cold USDT-wallets corresponds to wallets that have not moved their USDT for 3 months. Extending the period of inaction to 6 months reduces the percentage of cold wallets to 35 percent. Source: Glassnode.

²⁷The funding premium at Binance is the sum of two components: a fixed interest rate and a premium. The premium is a function of the spread between the future and spot price.

²⁸We remove one hour of outlier data from USDT’s lending rate on August 10, 2021, at 6 am because it is implausibly high and likely a data entry error. Lending rates are available in hourly snapshots until September 2022, when FTX’s API switched to providing daily instead of hourly data.

²⁹The run began after a tweet by the CEO of Binance that Binance would sell its FTT tokens: https://twitter.com/cz_binance/status/1589283421704290306.

	Days (N)	Mean	Std. Dev.	Min	Max
<i>Stablecoin Prices</i> (\$)					
USDT (Tether)	705	1.0010	0.0022	0.9919	1.0114
DAI (Dai)	705	1.0013	0.0024	0.9912	1.0109
<i>Margin Lending Rates</i> (annualized percent)					
USDT	705	7.95	9.99	1.00	66.03
DAI	650	7.22	10.28	0.88	73.99
<i>Perpetual Futures Funding Rate</i> (annualized percent)					
BTC/USDT	705	18.98	30.48	−32.02	172.30
ETH/USDT	705	21.23	40.32	−239.65	215.45

Table 1: Summary Statistics. Table gives summary statistics for stablecoin prices from Coingecko, stablecoin margin lending rates from FTX, and perpetual futures funding rates from Binance. Sample runs from December 1, 2020 to November 5, 2022.

average price of both stablecoins is close to \$1, as expected. The average lending rate is 8 percent for USDT and 7 percent for DAI. Relative to prices, the lending rates are more volatile. The average funding premium is 19 percent for BTC/USDT and 21 percent for ETH/USDT, indicating that the future price typically exceeds the spot price for both contracts.

As an aside, we show in the Online Appendix in Table A.3 that the lending rates for USDT and DAI are not related to short-term interest rates in the traditional financial system, proxied using the effective fed funds rates. Changes in the stablecoin lending rates have no relationship with changes in the fed funds rate on a daily basis.

4 Empirical Results

We have three sets of results. First, we show that stablecoin lending rates are tightly linked to speculative demand for cryptocurrencies. Second, we test Proposition 1, which shows how stablecoins maintain their peg by linking cryptocurrency demand and risk to the stablecoin issuer’s liquid asset share and token issuance or redemptions. Third, we apply the model to the May 2022 turmoil in crypto markets following the collapse of TerraUSD.

Some clarifications are in order when taking the model predictions to the data. In particular, we made a distinction between stabilizing the peg and defending the peg, based on the ability of the issuer to re-balance their portfolio without incurring costs that could trigger run dynamics. This distinction is meaningful since the “liquid asset portfolio share channel” for peg stabilization matters only in the former case, while the “redemption channel” is operational both when stabilizing and defending the peg. However, it is not straightforward to

empirically distinguish these two cases for the redemption channel, especially without detailed data about the type of assets that the issuer sells to meet redemptions and the potential price discount. At the same time, it is more likely that the more severe depeggings during narrative self-fulfilling run episodes, like the run on Tether in May 2022, are associated with what we describe in the theoretical analysis as “defending the peg”. Thus, in our empirical analysis below, we will try to isolate such narrative run episodes to study separately how the redemption channel and lending rates may help stablecoin weather a run.

Finally, note that we have focused our analysis on the effect of speculative demand on the ability of stablecoins to manage their peg in response to crypto shocks. However, the role of stablecoins in facilitating speculative levered positions should also matter for the adoption of stablecoins over the crypto cycle along with the other reasons why crypto market participants choose to hold stablecoins. Figure A.1 in the Online Appendix plots our measure of speculative demand (described below) along with the daily change in Tether’s market cap, pointing to a positive contribution of speculative demand to Tether’s adoption. However, we should note that studying the drivers of stablecoin adoption is outside the scope of our paper. See Gorton et al. (2022) for a relative evaluation of several factors that contribute to the adoption of stablecoins as a means of payment.

4.1 Lending Rates and Expected Speculative Returns

We show that when the expected return for the speculative asset— y in the model—increases, the stablecoin lending rate grows, as depicted in equation (A.1).³⁰ Speculators’ expected returns are challenging to measure. Because cryptocurrency expected returns are not directly observable, we use perpetual futures funding rates to infer the speculative demand. Futures funding rates reflect the cost investors face to take leverage. We argue that the magnitude of the annualized funding rates is directly related to speculative cryptocurrency demand because no other liquid products provide similar levels of leverage as the perpetual futures.

We proxy for expected returns using the BTC/USDT perpetual future on Binance, likely the largest perpetual future contract in the world. Figure 3 shows the time series of the annualized funding rate of Binance’s BTC and ETH USDT-settled perpetual futures. The funding rate is typically small but positive, indicating that investors who want to take levered long positions must pay a fee. In the Online Appendix Section A.7, we check that Binance’s BTC/USDT perpetual futures funding rate is a robust proxy for expected returns.³¹ One concern is that using the BTC/USDT perpetual futures as a proxy of y overweighs

³⁰In the Online Appendix Table A.4, we also confirm the model’s prediction about the relationship between the outside option ρ and stablecoin volumes. We find that net stablecoin issuance coincides with higher ρ , described in Table A.4 and consistent with the prediction that ρ is increasing in s as shown in (2).

³¹See Table A.5, which presents the correlation of our main measure of y with several other potential measures.




Figure 3: Perpetual Futures Funding Rate. Figure plots the annualized funding rate of USDT-settled Bitcoin perpetual futures for Bitcoin and Ether on Binance. A positive funding rate indicates that long-future investors make payments to short-future investors. Series are seven-day trailing averages.

idiosyncrasies specific to Bitcoin. But the BTC/USDT and ETH/USDT perpetual futures funding rates are tightly linked with a correlation coefficient of 0.89. Binance also has perpetual futures that settle in Binance USD, another stablecoin. The funding rates across perpetual futures are highly correlated regardless of which stablecoin is used for settlement.

Yet another concern is that Bitcoin and Ether are special because they are relatively liquid and mature markets compared to other digital assets. We include a column showing the measures are highly correlated with the perpetual futures funding rate for Dogecoin. Dogecoin is the largest memecoin, a digital asset that ostensibly started as a joke, and exhibits comparatively higher volatility than Bitcoin and Ether.

Finally, we rule out the possibility that Binance’s futures funding rates mainly reflect idiosyncrasies specific to Binance, rather than aggregate expected returns for cryptocurrency beyond Binance. We compare Binance’s perpetual future funding rates with analogous rates from FTX and find that funding rates are similar and highly correlated across the exchanges, confirming that the funding rates are not principally capturing exchange-specific factors. In addition, we show that perpetual futures funding rates are closely linked to expected returns embedded in crypto futures traded on the CME.

Figure 4 shows a binscatter of the perpetual future funding rate and the USDT stablecoin lending rate: the two are strongly positively related. More formally, we test the model’s prediction that lending rates are increasing in y , equation (A.1), by regressing Tether’s lending




Figure 4: Stablecoin Lending Rates and Cost of Leverage. Figure plots a binscatter of daily observations of the annualized funding rate of BTC/USDT perpetual futures on Binance relative to the annualized USDT lending rate on FTX.

rate on FTX on the perpetual futures funding rates using

$$\text{USDT Lending Rate}_t = \alpha + \beta \text{ Futures Funding Rate}_t + \gamma X_t + \varepsilon_t,$$

where X_t is a vector of controls. Table 2 shows the regression results. The first row shows that a 1pp increase in the futures funding rate is associated with an increase in stablecoin lending rates between 0.12 and 0.26pp, depending on the control variables. A one-standard-deviation increase in the future funding rate (30pp) corresponds to lending rates increasing by roughly 3.6pp, using the estimates in column 3. Across all specifications, there is a positive and significant relationship between lending rates and our proxy for expected returns. We include a measure of risk, Bitcoin’s implied volatility, as a control since the lending rates may not be risk-free and may be higher when the underlying collateral, Bitcoin, is more volatile, increasing the risk of counterparty default.³² Moreover, we control for the contemporaneous Bitcoin return to account for other factors driving spot returns at the daily level.

The regression results suggest a relationship akin to the interest rate parity in international finance, governing the relationship between interest rates and spot and future exchange rates. This is a feature, not a bug, of our analysis and strengthens our thesis that lending rates

³²In the Online Appendix, Table A.6 includes robustness tests by regressing stablecoin lending rates on measures of expected returns inferred from CME cryptocurrency futures instead of futures funding rates.

	USDT			USDT and DAI		
	(1)	(2)	(3)	(4)	(5)	(6)
Futures Funding Rate _t	0.26*** (14.41)	0.19*** (8.10)	0.12*** (5.17)	0.23*** (12.96)	0.17*** (8.92)	0.12*** (7.15)
Stablecoin Lending Rate _{t-1}			0.43*** (6.01)			0.30*** (7.50)
BTC Implied Volatility _t			0.01 (0.20)			0.03 (1.18)
R_t^{BTC}			0.05 (0.86)			0.09 (1.49)
N	705	705	704	1,355	1,355	1,353
R^2	0.61	0.71	0.77	0.45	0.59	0.64
Month FE	No	Yes	Yes	No	Yes	Yes
Coin FE	No	No	No	No	Yes	Yes

Table 2: Stablecoin Interest Rates and Expected Returns. Table presents regression $R_{i,t} = \alpha + \beta y_t + \gamma' X + a_i + b_t + \varepsilon_{i,t}$ where $R_{i,t}$ is stablecoin i 's margin lending rate from the FTX exchange, y_t is Binance's BTC/USDT perpetual future funding rate, and X is a vector of controls including the lag of the stablecoin i 's lending rate, the risk of Bitcoin measured by Bitcoin futures' implied volatility, and Bitcoin's contemporaneous price return, a_i is a stablecoin fixed effect, and b_t is a time fixed effect. Observations are daily. t -statistics are reported in parentheses using robust standard errors and clustered by week, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

are driven by leveraged speculation. As mentioned in section 3.1, one way to take a levered bet in cryptocurrencies is to go long on a perpetual future. Another way is margin lending, whereby the speculator borrows a stablecoin to take a leveraged position in a cryptocurrency. By arbitrage, the two alternative speculative strategies should be closely linked to each other. We should note, however, that such arbitrage is imperfect due to frictions and because margin loans and perpetual futures do not have matched maturities.

4.2 Peg Stability

Proposition 1 predicts that stablecoin issuers can maintain their peg following shocks to speculative demand with two tools: either adjusting their liquid asset share ℓ or issuing/redeeming tokens s . We empirically verify both mechanisms.

Liquid Asset Portfolio Share Channel The model shows that stablecoin issuers can offset negative shocks to cryptocurrency demand by increasing their portfolio share of liquid assets, all else equal. We assumed that the issuer's liquid asset holdings, ℓ , are public knowledge, at least at disclosure dates. In practice, it is rarely the case that a stablecoin issuer gives disclosures with enough granularity to verify its liquid asset share. Disclosures are infrequently published, and there are some doubts about their accuracy. Tether, for example, started providing regular disclosures only after an investigation by the New York Attorney




Figure 5: Tether Liquid Share vs. Perpetual Futures Funding Rate and USDT Lending Rate. Left panel plots the perpetual futures funding rate against USDT’s liquid portfolio share in the same quarter. Right panel plots the average annualized lending rate for Tether on FTX by quarter against USDT’s liquid portfolio share in the same quarter. Liquid portfolio share is calculated using public disclosures and is the share of reserves held in cash, bank deposits (including fiduciary deposits), reverse repurchase agreements, and Treasury bills.

General.³³

Despite these limitations, we show that there is a negative relationship between y (and R) and ℓ using public disclosure data from Tether, which has released seven quarterly disclosures with enough granularity to estimate Tether’s ℓ . Figure 5 is a scatterplot comparing the liquid asset share against the perpetual futures funding rate (y) and the USDT lending rates (R). We define Tether’s liquid asset portfolio share ℓ as its share of reserves held in cash, bank deposits (including fiduciary deposits), reverse repurchase agreements, and Treasury bills. While the data are limited to seven quarterly data points, there is a clear negative relationship that ℓ is higher when expected returns and Tether’s lending rate are lower. When crypto demand or the stablecoin lending rates are low, stablecoin issuers hold more liquid assets to maintain the stablecoin’s peg. Over time, alongside greater attention toward Tether’s underlying reserves, Tether’s quarter-end liquid asset share increased from around 26% in early 2021 to roughly 70% by late 2022. Tether’s liquid asset share at quarter ends has been in the 70-76% range in 2023.

One concern is that during our sample period, 2020 to 2022, the Federal Reserve raised

³³In the absence of disclosures, investors may infer that $\ell = 0$. In that case, the stablecoin can keep its peg so long as y is sufficiently high.

interest rates, and Tether may have simply substituted its portfolio toward Treasuries as their yields increased. However, the spread between 3-month AA financial commercial paper and 3-month T-bills increased from 8bps in December 2020 to 31bps in November 2022, mitigating this concern. As a point of comparison, we show that prime money market funds, which are very often compared to stablecoins (Anadu et al., 2023), also did not shift toward Treasuries during this same time period. Figure A.2 in the Online Appendix shows that the share of prime money market funds’ portfolio held in Treasuries fell from 25 percent to less than 10 percent; including overnight reverse repurchases with the Federal Reserve (not available to stablecoins), the share of liquid assets remained flat. This comparison further mitigates concerns that Tether shifted towards safer assets due to higher safe interest rates.

It is true that Tether may have also decreased its risky reserves because of pressure from the industry and the public to become more transparent. This is consistent with our model, which says that observable safe reserves reduce run risk. We document that the consistent decrease of risky reserves was accompanied by a decrease in lending rates, which are an important part of the compensation for holding the prone-to-runs stablecoin.

Redemption Channel Stablecoin issuers do not provide continuous information on their liquid assets, and quick adjustments in their liquid asset share may be difficult and costly when portfolio rebalancing costs are high. Proposition 1 shows that stablecoin issuers can maintain their peg by adjusting the supply of the tokens while holding ℓ fixed. Information on the token’s supply is public, and the supply often fluctuates in the short term. We calculate a stablecoin i ’s net issuance on date t as

$$\Delta s_{i,t} = \left(\frac{\text{Market Cap}_{i,t}}{P_{i,t}} - \frac{\text{Market Cap}_{i,t-1}}{P_{i,t-1}} \right).$$

Net redemptions equal $-1 \times \Delta s_{i,t}$. We divide the market capitalization by the stablecoin’s price because we are interested in the face value of the stablecoin’s liabilities, which the issuer can directly affect. If we did not divide by prices, it would appear that the stablecoin had issued more coins when its price increased, even if the stablecoin issuer took no action.

Table 3 shows summary statistics for redemptions for the largest stablecoins and orders the stablecoins in descending order based on their average 2021 market capitalization. The largest three stablecoins have net redemptions between 25 percent and 39 percent of days, even though stablecoins have grown rapidly over the period. The average redemption for the three ranges between 0.3 percent (USDT) and 1.4 percent (BUSD). TerraUSD (USTC) had the largest one-day net redemption of \$4.7 billion, about 27 percent of its market cap, during its collapse in May 2022. In the post-2019 period, each stablecoin has faced large

single-day redemptions: 4.1 percent for Tether, 8.2 percent for USDC, and 11.9 percent for BUSD amounting to \$3.4 billion, \$3.8 billion, and \$460 million. Figure A.3 in the Online Appendix plots the daily redemptions and issuance of Tether—that will be the focus of our analysis—as a percent of its face value. These flows are volatile and economically meaningful.

Coin	Days		Average Redemption		95%ile Redemption		Largest After 2019	
	Total	% with redemptions	\$ mln	% of face	\$ mln	% of face	\$ mln	% of face
USDT	2,788	26	62.0	0.3	230.4	1.0	3,432.4	4.1
USDC	1,493	39	79.9	0.8	360.9	2.3	3,809.4	8.2
BUSD	1,142	34	65.5	1.4	229.4	5.4	460.4	11.9
DAI	1,083	37	34.4	1.1	145.5	4.0	718.2	13.6
USTC	765	27	99.8	0.9	323.5	3.1	4,748.3	27.1
MIM	493	38	28.8	1.5	53.0	3.0	1,473.6	78.8
TUSD	1,689	43	7.1	1.3	38.0	5.1	235.1	19.0
PAX	1,502	45	7.6	1.5	34.8	5.6	178.7	15.3
LUSD	580	42	11.9	1.9	44.4	6.3	585.0	40.4
HUSD	1,143	38	9.0	2.5	42.2	8.9	171.8	35.7
USDN	1,009	40	4.2	0.8	9.8	2.0	547.2	80.5
FRAX	685	40	10.3	1.1	43.7	4.6	594.8	22.9
ALUSD	586	50	2.2	0.8	7.1	2.8	93.7	27.3
GUSD	1,508	42	4.9	3.4	23.3	17.8	169.3	41.4
USDP	589	48	1.8	3.0	9.5	18.9	64.9	52.4
MUSD	851	54	1.2	2.4	4.8	10.2	16.0	27.7
USDK	1,227	49	0.1	0.3	0.3	0.9	1.6	6.8
RSV	941	51	0.0	0.5	0.2	1.8	3.5	16.6

Table 3: Redemption Summary Statistics. Table presents summary statistics about daily net redemptions for several stablecoins. Rows ordered by average market capitalization in 2021, beginning with the largest (USDT). Sample runs from the date Coingecko has data for the coin until November 5, 2022.

Indeed, the magnitudes of stablecoins’ redemptions are economically large compared to the traditional banking system. Gorton and Zhang (2021) and Gorton et al. (2022) argue that Free Banking era banks and stablecoin issuers are similar because they both created private money—private bank notes and stablecoins—and both did so without a lender of last resort or deposit insurance. The average liquid asset share of New York banks from 1818 to 1861 was 5.7 percent, using data from Weber (2018), defined as cash, cash items, and U.S government bonds. Single-day redemptions on the scale of those faced by stablecoins would have plausibly exhausted the Free Banking system’s liquid assets.

We empirically establish the redemption channel in three steps. First, we regress changes in the token’s log face value supply at t on cryptocurrency return and riskiness at $t - 1$,

which are factors that should affect the stablecoin issuance and redemptions following our theoretical analysis in Section 2.4. This step captures how a change in speculative demand at $t - 1$ affects stablecoin supply at t . Second, we regress the stablecoin lending rates t on the predicted supply changes at t from the first step, while controlling for other contemporaneous factors. This step captures the effect that a change in stablecoin supply induced by a change in speculative demand has on lending rates, which is at the core of the redemption channel. Third, we investigate how the change in lending rates helps stabilize the peg conditional on a de-pegging episode.

Recall that the redemptions channel is operational both in normal times when the issuer is trying to stabilize the peg and in times of stress, during which self-fulfilling runs may unfold when the issuer is trying to defend the peg. We present results for step 1 and step 2 for the full sample, mixing these two cases, and separately for normal times by excluding May 2022, which should include observations corresponding to the narrative run episodes on Tether. The results are qualitatively and quantitatively similar.³⁴

The first-step regression is

$$\Delta \ln(s_{i,t}) = \alpha + \beta_1 y_{t-1} + \beta_2 \sigma_{t-1} + \gamma' X + a_i + b_t + \varepsilon_{i,t}. \quad (15)$$

y is the perpetual futures funding premium, σ is Bitcoin’s risk measured by its futures’ implied volatility, $s_{i,t}$ is the face value of stablecoin i , X is a vector of controls including lags of stablecoin’s face value and issuance as well as BTC’s trading volume, a_i and b_t are stablecoin and time fixed effects. Our theoretical analysis predicts the stablecoin’s supply will increase in y , $\beta_1 > 0$, and decrease in σ , $\beta_2 < 0$. We lag the independent variables by a day to ensure they are in the stablecoin issuer’s information set. Moreover, we include monthly fixed effects to capture possibly slower-moving changes in ℓ , as the proposition’s redemption channel holds ℓ fixed, as well as other factors driving the demand of stablecoins on a monthly frequency.

Table 4 shows the results. The first four columns focus on USDT, and the last four include USDT and DAI. Implied volatility has a negative coefficient, and the funding rate coefficient is consistently positive, so stablecoin redemptions are larger when implied volatility is higher and when funding premia, our proxy for the demand of cryptocurrency speculation, is lower. The results are similar across all specifications including month and coin fixed effects and including lags of redemptions and the token’s face value. $\Delta \ln(s_{i,t})$ is in basis points, so a 10pp increase in the funding premium, all else equal, corresponds to subsequent stablecoin

³⁴We do not present separate results for step 3 given the small sample size for this empirical test as explained below. But we separately study the May 2022 event in Section 4.3 and show the de-pegging and re-pegging dynamics along with Tether’s lending rates.

	USDT				USDT and DAI			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Funding Premium _{<i>t-1</i>}	0.70*** (4.47)	0.68*** (4.15)	0.63*** (3.90)	0.58*** (3.80)	1.05*** (5.07)	1.02*** (5.09)	0.92*** (4.60)	0.92*** (4.45)
Bitcoin Implied Volatility _{<i>t-1</i>}		-0.86** (-2.01)	-0.82* (-1.81)	-0.50 (-1.29)		-1.44*** (-3.36)	-1.31*** (-2.93)	-1.36*** (-2.84)
$\Delta \ln(s_{i,t-1})$			-0.02 (-0.47)	-0.03 (-0.60)			0.07 (0.80)	0.00 (0.03)
$\ln(s_{i,t-1})$			-110.88 (-1.48)	-130.92* (-1.73)			-83.74*** (-3.76)	-70.01*** (-3.13)
$\ln(\text{BTC Volume}_{t-1})$			-3.58 (-0.32)	11.45 (1.31)			-6.82 (-0.54)	12.62 (1.36)
$\Delta \ln(\text{BTC Volume}_{t-1})$			9.00 (1.15)	6.66 (0.82)			15.57 (1.31)	9.96 (0.94)
<i>N</i>	704	704	704	673	1,408	1,408	1,408	1,346
<i>R</i> ²	0.34	0.35	0.35	0.36	0.17	0.18	0.19	0.18
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Coin FE	n/a	n/a	n/a	n/a	No	Yes	Yes	Yes
Sample	Full	Full	Full	Excl. 5/22	Full	Full	Full	Excl. 5/22

Table 4: Peg Stability Priors. Table presents regression $\Delta \ln(s_{i,t}) = \alpha + \beta_1 y_{t-1} + \beta_2 \sigma_{t-1} + \gamma' X + a_i + b_t + \varepsilon_{i,t}$ where y is the perpetual futures funding premium, σ is the risk of Bitcoin measured by Bitcoin futures' implied volatility, $s_{i,t}$ is the face value of stablecoin i , X is a vector of controls including lags of the stablecoin's face value and issuance, the log level of BTC trading volume $\ln(\text{BTC Volume}_{t-1})$, and changes in the log level of BTC trading $\Delta \ln(\text{BTC Volume}_{t-1})$. a_i is a stablecoin fixed effect, and b_t is a time fixed effect. $\Delta \ln(s_{i,t})$ is in basis points. Observation at the daily level by coin. Sample in first four columns is only USDT and in last four columns is both USDT and DAI. Columns 4 and 8 exclude the month of Terra's collapse, May 2022. t -statistics are reported in parentheses using robust standard errors and clustered by week, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

issuance between 6 and 11 basis points. One might be concerned that our results accrue from alternative uses of stablecoin in crypto trading that are correlated with the expected BTC return and volatility but not fully captured by the month fixed effects we have included. In specification (3) we include the trading volume of BTC, and its change, as controls to account for auxiliary stablecoin demand, on a daily level, accruing from the use of stablecoins as a vehicle facilitating storing funds between trades. Although the month fixed effects we include may already absorb much of the variation from alternative stablecoin uses, controlling for BTC daily trading volume should capture residual variation. Columns (4) and (8) study the dynamics in normal times by excluding the period around TerraUSD's collapse, May 2022, and find similar results.

Moving to step two, the model predicts that lending rates will increase after the token's supply falls. Thus, we regress stablecoin lending rates on predicted changes in the token's supply $\widehat{\Delta \ln(s_{i,t})}$, which we estimate using Equation 15:

$$\Delta R_{i,t} = \alpha + \gamma \left(\widehat{\Delta \ln(s_{i,t})} \right) + a_i + b_t + \varepsilon_{i,t},$$

and our hypothesis is that $\gamma < 0$. We estimate predicted changes in the token's supply using $\Delta \ln(s_{i,t}) = \alpha + \beta_1 y_{t-1} + \beta_2 \sigma_{t-1} + \gamma' X + a_i + b_t + \varepsilon_{i,t}$, where X is a vector of controls including $s_{i,t-1}$, $\Delta \ln(s_{i,t-1})$, the log level of BTC trading volume $\ln(\text{BTC Volume}_{t-1})$, and changes in

	USDT				USDT and DAI			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta \widehat{\ln}(s_{i,t})$	-3.83** (-2.53)	-13.14*** (-3.03)	-11.54*** (-3.28)	-14.23*** (-5.09)	-2.01** (-2.18)	-9.53*** (-2.98)	-4.48** (-2.26)	-6.31*** (-3.04)
Funding Premium _t		8.96*** (3.09)	8.29*** (3.16)	9.71*** (4.05)		10.35*** (2.77)	6.30** (2.25)	7.68*** (2.61)
Bitcoin Implied Volatility _t		-8.00 (-1.41)	-6.84 (-1.41)	-7.08 (-1.40)		-10.97* (-1.90)	-3.85 (-1.10)	-4.83 (-1.22)
N	704	704	704	673	1,353	1,353	1,353	1,291
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Coin FE	n/a	n/a	n/a	n/a	Yes	Yes	Yes	Yes
Controls	No	No	Yes	Yes	No	No	Yes	Yes
Sample	Full	Full	Full	Excl. 5/22	Full	Full	Full	Excl. 5/22

Table 5: Peg Stability. Table presents regression $\Delta R_{i,t} = \alpha + \gamma(\Delta \widehat{\ln}(s_{i,t})) + a_i + b_t + \varepsilon_{i,t}$ where $R_{i,t}$ is the lending rate of stablecoin i on date t at FTX in basis points and $\Delta \widehat{\ln}(s_{i,t})$ is the expected change in the face value of the stablecoin in basis points. $\Delta \widehat{\ln}(s_{i,t})$ is estimated using the regression $\Delta \ln(s_{i,t}) = \alpha + \beta_1 y_{t-1} + \beta_2 \sigma_{t-1} + \gamma' X + a_i + b_t + \varepsilon_{i,t}$, where y is the perpetual futures funding premium, σ is Bitcoin futures' implied volatility, $s_{i,t}$ is the face value of stablecoin i , a_i is a stablecoin fixed effect, b_t is a time fixed effect, and X is a vector of controls including $s_{i,t-1}$, $\Delta \ln(s_{i,t-1})$, the log level of BTC trading volume $\ln(\text{BTC Volume}_{t-1})$, and changes in the log level of BTC trading $\Delta \ln(\text{BTC Volume}_{t-1})$. Columns with “Controls” as no exclude the X variables, and columns with “Controls” as yes include the X variables. Observation at the daily level by coin. Sample in first four columns is only USDT and in last four columns is both USDT and DAI. Columns 4 and 8 exclude the month of Terra’s collapse, May 2022. t -statistics are reported in parentheses using robust standard errors and clustered by week, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

the log level of BTC trading $\Delta \ln(\text{BTC Volume}_{t-1})$.

Table 5 shows the results. The first four columns limit the sample to USDT, and the last four include USDT and DAI. The specifications also vary which estimates include the X vector of controls when predicting $\Delta \ln(s_{i,t})$, as shown in the “Controls” row. Columns (2), (3), (4), (6), (7), and (8) also include controls for contemporaneous changes in y_t and σ_t , as BTC contemporaneous expected return and volatile should directly matter for the stablecoin lending rate over and beyond the change in the stablecoin supply. Indeed, controlling for these increases the magnitude of the effect of a change in stablecoin supply on the stablecoin lending rate. The table shows that an expected one basis point increase in token supply decreases the lending rate by 2 to 13 basis points. Again, columns (4) and (8) study the dynamics in normal times by excluding the period around TerraUSD’s collapse and find similar coefficients.

Finally, turning to step three, Table 6 shows the effect of increasing interest rates on the stablecoin’s price during Tether’s depeg episodes. We use intraday price data from Cryptocompare to identify depeg days as those with a minimum hourly intraday price below the 1st percentile or a maximum hourly intraday price above the 99th percentile. When the price depegs from \$1, either above or below, there is a positive relationship between changes in lending rates and prices, consistent with increasing lending rates pushing prices back up after a depeg below \$1 or falling lending rates pushing prices back down after a depeg above

	(1)	(2)	(3)	(4)	(5)	(6)
ΔR_t	1.69** (2.31)	1.52** (2.19)	2.70*** (2.95)	2.25*** (3.89)	2.00** (2.93)	2.57*** (5.64)
Bitcoin Implied Volatility _t			59.30** (2.33)			9.48 (0.52)
Funding Premium _t			3.82 (0.44)			20.02** (2.81)
N	22	22	22	18	18	18
R^2	0.21	0.26	0.59	0.38	0.43	0.66
Coin FE	No	Yes	Yes	No	Yes	Yes
Cond. Repeg	No	No	No	Yes	Yes	Yes

Table 6: Repegging Mechanism. Table presents regression $\Delta P_t = \alpha + \beta \Delta R_t + \varepsilon_{i,t}$ where ΔP_t is the change in the price of Tether and ΔR_t is the change in lending rate. Observations are daily with monthly time fixed effects. Sample includes periods of depegs where the previous day’s minimum intraday price is less than the 1st percentile or the maximum intraday price is above the 99th percentile. Intraday price data is hourly data from Cryptocompare. Columns (4) through (6) restrict the sample to observations where we observe a rep—defined as price, rounded to the nearest penny, equals \$1.00. ΔP_t is multiplied by 1 million for reading convenience, and R_t is in basis points. t -statistics are reported in parentheses using robust standard errors where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

\$1. There are challenges to this empirical approach: depeg events are infrequent based on our definition, which is not overly restrictive, and so the sample size is limited. Moreover, we do not have minute-by-minute lending rates, so our empirics are limited to daily observations of depeg events. Finally, factors other than a lower supply due to redemptions may drive the increase in lending rates during depeg events, which we elaborate upon in the next section.

4.3 May 2022 Stablecoin Turmoil

In May 2022, the algorithmic stablecoin TerraUSD depegged. Sentiment in crypto markets had been slugging but turned bearish after the depeg. Several prominent crypto firms failed shortly after that: 3 Arrows Capital, Voyager Digital, and Celsius, and the crypto space entered a so-called “crypto winter.” Pressure on TerraUSD spilled to other stablecoins, and Tether’s market capitalization fell from \$83 to \$73 billion in May following a stream of redemptions. Several other algorithmic stablecoins failed or teetered on the brink of viability. USDC, viewed as the highest quality stablecoin, traded at a premium to its peg and saw net inflows. Such turmoil is a natural experiment to study the model’s predictions with respect to defending the peg in times of stress when self-fulfilling runs may be at play.

Figure 6 shows the market dynamics for Tether during TerraUSD’s price collapse. The vertical line on May 10 denotes the date that TerraUSD lost its peg. The model posits that a stablecoin will lose its peg when speculative cryptocurrency risk increases or when demand for the speculative cryptocurrency falls. The top half of the figure shows that Tether lost its

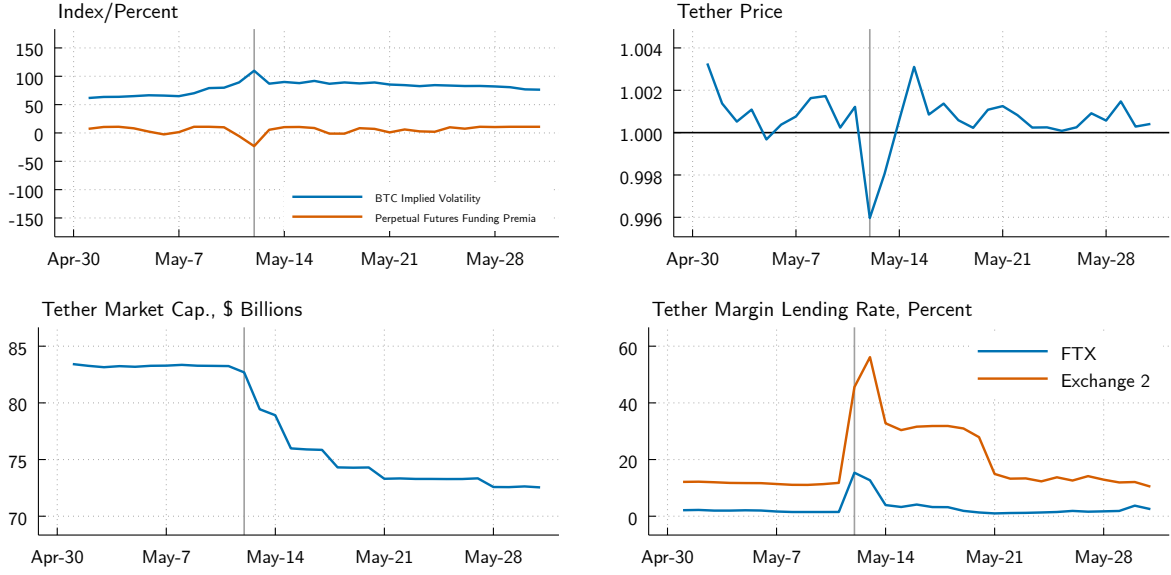


Figure 6: Tether during May 2022. Top-left figure plots BTC implied volatility and perpetual futures funding rate. Top right plots the price of Tether. Bottom left panel plots Tether’s market capitalization in billions. Bottom right panel plots Tether’s margin lending rates on two exchanges.

peg for at least two days, coinciding with spikes in BTC implied volatility and a decline in perpetual futures funding premia, a proxy for speculative demand.³⁵

The model predicts that Tether could potentially maintain its peg by decreasing the supply of tokens (s), equivalent to redeeming and burning tokens to reduce its market capitalization while holding its reserve mix (ℓ) fixed. Tether redeemed roughly \$10 billion of tokens over three weeks, consistent with the prediction, as shown in the bottom-left panel of Figure 6. The bottom-right panel shows that Tether’s lending rates spiked and remained elevated, helping stabilize the peg. We should note that the increase in the lending rate alone may not have been enough to stop the run on Tether, which could have resulted in a complete depletion of reserves following more severe shocks. We elaborated on these limitations to stabilization in section 2.4.

An additional consideration is that the higher lending rates were driven by an incentive for borrowers to bet on the collapse of Tether: their debt would be stablecoin-denominated and, thus, worth zero if Tether collapsed. In the Online Appendix in section A.4, we extend our model to encompass this motive and show that it can result in even higher lending rates as investors would require higher compensation. The rapid increase in lending rates before

³⁵The perpetual future funding rate declined slightly in May 2022, but demand does not disappear as the lender rationally expects repayment in some states; the large dip in Figure 3 reflects the Ethereum Merge in September 2022.

an actual de-peg resembles the “peso problem” in the literature studying the collapse of fixed-exchange-rate regimes (Flood and Garber, 1984, and Penati and Pennacchi, 1989). In particular, the expectation of a de-peg can result in an increase in domestic credit and in interest rates denominated in domestic currency, even prior to a collapse of a fixed exchange regime, which is akin to the mechanism we describe. However, we also show that supporting such an equilibrium requires that there are some states of the world where the stablecoin does not collapse and where speculative demand is strong enough to support the higher lending rates; this is consistent with our key mechanism.

Proposition 1 also shows that the stablecoin issuer could also maintain its peg by increasing ℓ . It is unlikely this was a primary tool used to stabilize the peg in the immediate aftermath of the TerraUSD failure, although we cannot observe ℓ during this period. Tether’s disclosure for the quarter ending March 2022 showed liquid asset holdings of \$43 billion (52 percent of its total assets), defined as the sum of its cash, reverse repos, and Treasury bills. Its liquid asset holdings would have fallen to \$33 billion (46 percent of total assets), assuming it paid for redemptions entirely out of liquid asset sales. To increase ℓ , Tether would have needed to sell \$4.4 billion of its illiquid assets and replace them with liquid assets. Such a large shift out of illiquid assets over a short period seems unlikely without material losses, so arguably, the main adjustment channel was through s during this episode.

The ℓ adjustment mechanism to maintain the peg is likely more useful over longer periods. In June 2022, rumors circulated that Tether’s commercial paper portfolio had suffered 30 percent losses. In response, Tether explicitly said that would increase ℓ in the long run:³⁶

“Tether can report that its current portfolio of commercial paper has since been further reduced to 11 billion (from 20 billion at the end of Q1 2022), and will be 8.4 billion by the end of June 2022. This will gradually decrease to zero without any incurrences of losses. All commercial papers are expiring and will be rolled into US Treasuries with a short maturity.”

Such dynamics are not limited to Tether. Dai, a decentralized and collateralized stablecoin, uses the USDC stablecoin as collateral for more than half its outstanding coins. In the summer of 2022, USDC’s issuer—Circle—began blocking wallets holding USDC that were associated with Tornado Cash.³⁷ Market participants grew concerned that DAI would be compelled to comply with the sanctions given their large USDC holdings. Rune Christianson, DAI’s co-founder, suggested that DAI should move its USDC holdings to ETH, functionally increasing

³⁶<https://web.archive.org/web/20220721170350/https://tether.to/en/tether-condemns-false-rumours-about-its-commercial-paper-holdings/>

³⁷Tornado Cash is a virtual currency mixer designed to obfuscate transaction details on the Ethereum blockchain. The U.S. Treasury sanctioned it in August 2022 for its role in money laundering.

the risk of its reserves (decreasing ℓ). In response, users redeemed roughly four percent of DAI’s outstanding tokens the next day, amounting to \$320 million.

4.4 Robustness

Our theoretical analysis suggests a causal relationship between expected returns (y) and stablecoin lending rates (R). We show the two are highly correlated in Table 2. Yet, some other unobserved variables may be driving the behavior in both variables, resulting in a high correlation. We address this concern with an instrumental variables approach using Major League Baseball (MLB) data.

In June 2021, MLB and FTX announced a five-year sponsorship deal naming FTX the “Official Cryptocurrency Exchange” of the MLB. The deal placed a prominent FTX logo on all umpire uniforms beginning July 13, 2021—previously, umpires had never worn advertising patches. Umpires wore the patch for all regular season, postseason, and spring training games. The sponsorship agreement included promotions on nationally televised MLB games, MLB.com, MLB Network (a television channel), and social media.³⁸ The deal was worth \$150 million, making it likely FTX’s largest endorsement deal in terms of annual expense.³⁹

We collect television viewership data on nationally televised MLB games from showbuzzdaily.com.⁴⁰ The data include a household rating, which measures the percentage of households watching the game. The television viewership data run from July 13, 2021, to November 5, 2022, corresponding to when umpires started wearing the logo (beginning during the 2021 All-Star game) through the end of the 2022 World Series. On many days there is only one game with a household rating. We use daily averages of the household rating as our instrument for the funding premium. Notably, the sample does not include the period of FTX’s collapse, which began on November 6, 2022.

Our identification relies on two assumptions: First, we assume that the advertising is effective, and some MLB audience members began trading cryptocurrency after viewing the advertising. FTX’s agreement to the costly sponsorship deals indicates that they believed it would lead to more customers and more trading on their platform. There is considerable evidence that advertising is effective (Guadagni and Little 1983, Ippolito and Mathios 1991, Akerberg 2003, Sethuraman et al. 2011). Bagwell (2007)’s survey of the literature on advertising’s effect on consumer behavior indicates that advertising is most effective for those without previous experience with the brand. Moreover, survey evidence shows that new retail

³⁸<https://www.mlb.com/press-release/press-release-mlb-ftx-cryptocurrency-partnership>

³⁹See <https://apnews.com/article/sam-bankmanfried-ftx-crypto-bitcoin-ab5910fe6f9c56b73dc2943848c33817>.

⁴⁰The data also includes about half a dozen nationally televised baseball events, like the home run derby, all-star game, trade deadline, and draft.

crypto traders entered the market during rapid crypto price increases.⁴¹ Second, the timing of the baseball schedule is set well in advance of the season, and it is highly improbable cryptocurrency events affect the timing or viewership of MLB games.⁴²

One concern is that the advertising might bring new customers to open accounts and begin lending stablecoins rather than speculating in other digital assets, thereby increasing the demand for stablecoins. To control for this, we include the change in the market capitalization of the stablecoin on that day to control for potential changes in the demand for the stablecoin. Moreover, were new customers to open accounts to lend stablecoins, we would expect lending rates to fall as the supply of lendable coins grew, leading to a negative relationship between the TV rating and lending rates, which would go against the result we try to establish. However, we still find a positive relationship between the two despite this potential downward bias.

Table 7 shows the IV regression results. In the first stage, we regress the daily funding premium on the average household rating; in the second stage, we regress the lending rate on the predicted funding premium. Panel A shows the second stage result. For every 1pp increase in the futures funding premium estimated using the household rating instrument, Tether’s margin lending rate is about 18bp higher on an annualized basis (column 2). Using DAI’s lending rate or different controls gives similar estimates ranging from 16bps to 22bps.

Panel B reports the first stage regression. The instrument satisfies the relevance condition, and the F -statistic indicates the instrument is largely statistically strong. Panel C shows that the instrumented regression gives similar coefficients to the OLS regression.

We provide additional robustness tests in the Online Appendix. Table A.7 shows several placebo tests. We use future household ratings as the instrument, with the columns varying by using ratings from one day, one week, or four weeks in the future. Baseball viewership in the future should not affect today’s lending rate through the funding premium, because it is unknown at date t . Using the future household rating as the instrument leads to an insignificant relationship between the funding premium and the lending rate for nearly all specifications. The F -statistic indicates the instruments are weak.

One concern is that the FTX lending and the Binance perpetual futures markets may be meaningfully segmented from one another. We use the Binance futures funding rate because it is likely the most liquid perpetual future instrument settled in USDT. Although Binance ostensibly prohibited U.S.-based traders from using its platform, the Securities and Exchange

⁴¹See <https://www.jpmorganchase.com/institute/research/financial-markets/dynamics-demographics-us-household-crypto-asset-cryptocurrency-use>.

⁴²The 2022 season schedule was modified shortly before the season began due to protracted negotiations with the players’ union.

Commission has documented that Binance still allowed U.S. clients to access its platform.⁴³

We check for additional robustness by including a proxy for the attention an individual game demands above and beyond its household rating.⁴⁴ We proxy for attention using the *Championship Leverage Index* (cLI), a standard measure used by baseball analysts to estimate the effect of an individual game’s outcomes on that team’s odds of winning the World Series championship.⁴⁵ Teams closer to earning a playoff berth and games later in the playoffs generally have higher cLI measures. The cLI is standardized, so the average game has a cLI of 1. We use data from Baseball-Reference.com for regular season cLI values, and we estimate the cLI for playoff games, and our sample includes nationally televised games. We exclude observations in a few cases when the viewership data does not provide specific team names.

In Table A.8, we estimate the IV but now use the product of the household rating and the cLI, where a date’s rating and cLI is the average of the nationally-televised games on that day. Our identifying assumption is that the audience pays attention to more important games more closely. As expected, there is a positive relationship between the household rating and the cLI, so more important games have higher viewership. Yet there remains some variation in cLI that is not entirely explained by the household rating, especially during the regular season when games have lower viewership. For two games with identical household ratings, our instrument will expect a larger effect from the more important game as measured by the cLI. The new instrument will exactly equal the previous household instrument for the average game because the average cLI is 1. The second-stage results shown in Table A.8 are similar to those without the cLI but are somewhat stronger across each specification.

Finally, we show that speculators have good reason to speculate in response to MLB advertising. In the Online Appendix Table A.9, we show that the returns for BTC, ETH, and DOGE are indeed positively related to household ratings of MLB games, consistent with our intuition that speculators may grow bullish as FTX advertises to new customers. The table regresses the daily return of BTC, ETH, or DOGE on the household rating of nationally televised MLB games on the same day. The table includes day-of-week fixed effects. While the standard errors are large, the average return increases in household rating. DOGE

⁴³For example, U.S. traders could use Virtual Private Networks (VPNs) to access the platform. See <https://www.sec.gov/files/litigation/complaints/2023/comp-pr2023-101.pdf>.

⁴⁴A large behavioral literature discusses the role of attention in traditional finance returns. Barber and Odean (2008) show that retail investors buy stocks that capture their attention. Papers such as Engelberg et al. (2012), Engelberg and Parsons (2011), and Da et al. (2011) use TV, media, or internet data to measure investor attention.

⁴⁵Specifically, Baseball Reference estimates the cLI as follows: “For each team game, we run 25,000 coin-toss simulations of the remainder of the season twice. In the first simulation, we assume the team won the game in question. In the second simulation, we assume the team lost the game in question. The difference between the team’s World Series win probabilities after a win and a loss measures the importance this game has on the team’s World Series win probability.”

coin has a positive and statistically significant relationship, perhaps because it is even more subject to animal spirits than Bitcoin and Ether.

Panel A: Second Stage		Lending Rate $R_{i,t}$		
	USDT		DAI	
	(1)	(2)	(3)	(4)
Futures $\widehat{\text{Funding Rate}}_t$	0.279*** (4.310)	0.175*** (3.069)	0.223*** (2.890)	0.155*** (3.159)
Bitcoin Implied Volatility $_t$	0.055 (0.721)	0.035 (0.540)	-0.020 (-0.463)	-0.026 (-0.879)
$\Delta \ln(s_{i,t})$	-0.006 (-1.167)	-0.004 (-0.955)	-0.008** (-2.347)	-0.006** (-2.264)
R_{t-1}		0.481*** (2.969)		0.492*** (6.576)
N	258	258	258	258
Time FE	Yes	Yes	Yes	Yes

Panel B: First Stage		Funding Premium		
	USDT		DAI	
	(1)	(2)	(3)	(4)
Rating $_t$	2.587*** (3.437)	1.941*** (2.830)	2.407*** (3.191)	2.307*** (3.359)
Bitcoin Implied Volatility $_t$	0.344* (1.730)	0.210 (1.058)	0.331** (2.303)	0.314** (2.219)
$\Delta \ln(s_{i,t})$	0.024* (1.748)	0.024* (1.750)	0.028*** (5.159)	0.028*** (5.154)
R_{t-1}		1.145** (2.570)		0.299 (0.935)
N	258	258	258	258
Time FE	Yes	Yes	Yes	Yes
F -stat	11.82	8.01	10.18	11.28

Panel C: OLS		Lending Rate $R_{i,t}$		
	USDT		DAI	
	(1)	(2)	(3)	(4)
Futures Funding Rate $_t$	0.211*** (14.232)	0.119*** (5.647)	0.153*** (7.492)	0.109*** (4.938)
Bitcoin Implied Volatility $_t$	0.013 (0.281)	0.007 (0.227)	0.067 (1.180)	0.024 (0.528)
$\Delta \ln(s_{i,t})$	0.003 (0.264)	-0.001 (-0.061)	0.004 (0.928)	0.002 (0.599)
R_{t-1}		0.489*** (6.906)		0.297*** (5.027)
N	705	704	650	649
Time FE	Yes	Yes	Yes	Yes

Table 7: Instrumental Variables Regression of Futures Funding Premia and Lending Rates.

Instrumental variables regression using the mean household rating of MLB games on a given day as an instrument to predict the perpetual futures funding premium. Panel A shows the second stage regression of the instrumented variable on margin lending rates separately for USDT and DAI. Panel B shows the first stage regression of the instrument on the perpetual futures funding premium. Panel C shows the OLS regression of the lending rate on the funding premium. Time FE indicates day of week, month of year, and year fixed effects. Kleibergen-Paap rk Wald F statistics reported. t -statistics are reported in parentheses using robust standard errors and clustered by week where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5 Conclusion

Historically, private money has often proved fragile, trading at a discount (Gorton, 2017). Stablecoins are a new form of private money that are susceptible to runs but have largely managed to maintain their peg. Stablecoins can offer compensation to their owners in several ways, including facilitating crypto trading, insuring against currency devaluation in countries with volatile currencies, or enabling illicit activities. Our analysis focuses on their novel role in crypto speculation, where stablecoin holders earn compensation by lending to traders taking levered speculative positions in other cryptocurrencies. Stablecoin lending rates are high and tightly correlated with measures of speculative demand. Fluctuations in crypto speculative motives put upward or downward pressures on stablecoin prices in secondary markets. Stablecoin issuers can maintain the peg by adjusting the total supply of stablecoin tokens outstanding or altering the liquidity of their reserves. But stablecoins can collapse quickly despite maintaining their peg most of the time.

Our analysis has placed the fragility of stablecoins at the forefront due to the structural vulnerabilities of their business model. However, we also further emphasize how stablecoins can sustain positive run-risk in equilibrium because this has been an empirically relevant feature of stablecoins during the period of the speculative crypto-trading that we investigate. This does not need to be the case in general, as stablecoin issuers can choose a run-proof structure, also established as a special case in our model. Indeed, issuers have invested in safer assets since the second half of 2022, which can be partially due to the ability to extract seigniorage from higher Treasury rates. However, the role of stablecoins in facilitating leverage for crypto speculation remains true. One example is the surge in demand for PYUSD (PayPal’s stablecoin), which invests in high-quality assets, coinciding with a promotion by Kamino Finance, a lending service on Solana (see Figure A.4 in the Online Appendix). Hence, much of our analysis—and especially our empirical work on identifying speculative crypto demand and linking it to stablecoin lending rates and stablecoin supply—continues to be informative even under zero run risk and can be useful for future research.

The mechanism we describe and test is not limited to stablecoins and the crypto-world. We expect the results to hold in other cases where privately-produced money helps satisfy leverage demand. Bankers’ acceptances in the 1920s and, more recently, tokenized money market fund shares—used to meet variation margin in derivatives and repo transactions—are two examples. Our framework and analysis could also be extended to these cases to provide valuable insights.

References

- Daniel A Akerberg. Advertising, learning, and consumer choice in experience good markets: an empirical examination. *International Economic Review*, 44(3):1007–1040, 2003.
- Kenechukwu Anadu, Pablo Azar, Catherine Huang, Marco Cipriani, Thomas M Eisenbach, Gabriele La Spada, Mattia Landoni, Marco Macchiavelli, Antoine Malfroy-Camine, and J Christina Wang. Runs and flights to safety: Are stablecoins the new money market funds? 2023.
- Pablo D. Azar, Garth Baughman, Francesca Carapella, Jacob Gerszten, Arazi Lubis, JP Perez-Sangimino, David E. Rappoport, Chiara Scotti, Nathan Swem, Alexandros Vardoulakis, and Aurite Werman. Financial stability implications of digital assets. *Federal Reserve Bank of New York Economic Policy Review*, 30, 2024.
- Kyle Bagwell. The economic analysis of advertising. *Handbook of industrial organization*, 3: 1701–1844, 2007.
- Brad M. Barber and Terrance Odean. All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors. *The Review of Financial Studies*, 21(2):785–818, 2008.
- Christoph Bertsch. Stablecoins: Adoption and fragility. *Sveriges Riksbank Working Paper Series*, 2023.
- Wilko Bolt, Jon Frost, Hyun Song Shin, and Peter Wierts. The bank of amsterdam and the limits of fiat money. 2023.
- Markus K Brunnermeier and Lasse Heje Pedersen. Market liquidity and funding liquidity. *The review of financial studies*, 22(6):2201–2238, 2009.
- Hans Carlsson and Eric van Damme. Global games and equilibrium selection. *Econometrica*, 61(5):989–1018, 1993.
- Zhi Da, Joseph Engelberg, and Pengjie Gao. In Search of Attention. *The Journal of Finance*, 66(5):1461–1499, 2011.
- Douglas W. Diamond and Philip H. Dybvig. Bank runs, deposit insurance, and liquidity. *Journal of Political Economy*, 91(3):401–419, 1983.
- Adrien d’Avernas, Vincent Maurin, and Quentin Vandeweyer. Can stablecoins be stable? *Working Paper*, 2022.
- Joseph Engelberg, Caroline Sasseville, and Jared Williams. Market Madness? The Case of “Mad Money”. *Management Science*, 58(2):351–364, 2012.
- Joseph E. Engelberg and Christopher A. Parsons. The Causal Impact of Media in Financial Markets. *The Journal of Finance*, 66(1):67–97, 2011.

- Robert P. Flood and Peter M. Garber. Collapsing exchange-rate regimes: Some linear examples. *Journal of International Economics*, 17(1):1–13, 1984.
- Ana Fostel and John Geanakoplos. Leverage cycles and the anxious economy. *American Economic Review*, 98(4):1211–44, 2008.
- John Geanakoplos. The leverage cycle. *NBER Macroeconomics Annual*, 24:1–66, 2010.
- Itay Goldstein and Ady Pauzner. Demand-deposit contracts and the probability of bank runs. *Journal of Finance*, 60(3):1293–1327, 2005.
- Gary Gorton. Inland bills of exchange: Private money production without banks+. *Explorations in Economic History*, 92:101547, 2024.
- Gary B. Gorton. The history and economics of safe assets. *Annual Review of Economics*, 9: 547–586, 2017.
- Gary B. Gorton and Jeffery Zhang. Taming wildcat stablecoins. *Available at SSRN 3888752*, 2021.
- Gary B. Gorton, Chase P. Ross, and Sharon Y. Ross. Making money. *NBER Working Paper #29710*, 2022.
- Denis Gromb and Dimitri Vayanos. Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of financial Economics*, 66(2-3):361–407, 2002.
- Peter M Guadagni and John DC Little. A logit model of brand choice calibrated on scanner data. *Marketing science*, 2(3):203–238, 1983.
- Pauline M Ippolito and Alan D Mathios. Information, advertising, and health choices: a study of the cereal market. In *Economics of food safety*, pages 211–246. Springer, 1991.
- Anil K Kashyap, Dimitrios P. Tsomocos, and Alexandros P. Vardoulakis. Optimal bank regulation in the presence of credit and run risk. *Journal of Political Economy*, 132(3): 772–823, 2024.
- JP Konig. Payments: Stablecoins vs. trading. 2023.
- Roman Kozhan and Ganesh Viswanath-Natraj. Decentralized stablecoins and collateral risk. *WBS Finance Group Research Paper Forthcoming*, 2021.
- Arvind Krishnamurthy and Annette Vissing-Jorgensen. The aggregate demand for treasury debt. *Journal of Political Economy*, 120(2):233–267, 2012.
- Alfred Lehar and Christine A Parlour. Decentralized exchanges. *Available at SSRN 3905316*, 2021.
- Gordon Liao. Macroprudential considerations for tokenized cash. *working paper*, 2022.
- Gordon Liao, Thomas Hadeed, and Ziming Zeng. Beyond speculation: Payment stablecoins for real-time gross settlements. 2023.

- Jiageng Liu, Igor Makarov, and Antoinette Schoar. Anatomy of a run: The terra luna crash. Technical report, National Bureau of Economic Research, 2023.
- Yiming Ma, Yao Zeng, and Anthony Lee Zhang. Stablecoin runs and the centralization of arbitrage. *working paper*, 2023.
- Igor Makarov and Antoinette Schoar. Cryptocurrencies and decentralized finance (defi). *Brookings Papers on Economic Activity*, 2022(1):141–215, 2022.
- Stephen Morris and Hyun Song Shin. Global Games: Theory and Applications. *In Advances in Economics and Econometrics: Theory and Applications, Eight World Congress, Vol. 1, ed. Mathias Dewatripont, Lars P. Hanse, and Stephen J. Turnovsky. Cambridge: Cambridge University Press*, 2003.
- Caitlin Ostroff and Jared Malsin. Turks pile into bitcoin and tether to escape plunging lira. *Wall Street Journal*, 2022. URL <https://www.wsj.com/articles/turks-pile-into-bitcoin-and-tether-to-escape-plunging-lira-11641982077>.
- Alessandro Penati and George Pennacchi. Optimal portfolio choice and the collapse of a fixed-exchange rate regime. *Journal of International Economics*, 27(1):1–24, 1989.
- Linda Schilling. Optimal forbearance of bank resolution. *The Journal of Finance (forthcoming)*, 2023.
- Raj Sethuraman, Gerard J Tellis, and Richard A Briesch. How well does advertising work? generalizations from meta-analysis of brand advertising elasticities. *Journal of Marketing Research*, 48(3):457–471, 2011.
- Alp Simsek. Belief disagreements and collateral constraints. *Econometrica*, 81(1):1–53, 2013.
- Harald Uhlig. A luna-tic stablecoin crash. Working Paper 30256, National Bureau of Economic Research, 2022.
- Skander Van den Heuvel. The welfare effects of bank liquidity and capital requirements. *FEDS 2022-072*, 2022.
- Cy Watsky, Jeffrey Allen, Hamzah Daud, Jochen Demuth, Daniel Little, Megan Rodden, and Amber Seira. Primary and secondary markets for stablecoins. 2024.
- Warren E. Weber. Antebellum u.s. state bank balance sheets, 2018. URL <https://researchdatabase.minneapolisfed.org/collections/wd375w37w?locale=en>.

Leverage and Stablecoin Pegs¹

Gary B. Gorton Elizabeth C. Klee Chase P. Ross
Sharon Y. Ross Alexandros P. Vardoulakis

A Online Appendix

Section A.1 includes additional derivations for statements made in Section 2 of the manuscript. Sections A.2-A.6 extend the baseline model in several dimensions mentioned in the manuscript: Section A.2 endogenizes margin requirements in stablecoin margin loans, Section A.3 incorporates benefits from additional uses of stablecoins other than their use in leveraged speculative crypto-trading, and Section A.4 incorporates the motive to borrow stablecoins to speculate on their collapse. Section A.5 derives the equilibrium token supply and stablecoin liquidity under non-observability. Section A.6 presents the choice of ℓ by the issuer when the liquid asset pays a positive interest. Section A.7 checks the robustness of using the BTC/USDT perpetual futures funding rate as a measure of speculative demand considering alternative proxies. Sections A.8 and A.9 report additional figures and tables with auxiliary results mentioned in the manuscript.

A.1 Additional Derivations

A.1.1 Derivatives of lending rate R with respect to y, m, s, λ

The effect of an increase in the cryptocurrency expected return, y , on the lending rate R , is

$$\frac{dR(\lambda, s)}{dy} = \frac{1}{1-m} > 0. \quad (\text{A.1})$$

Because $F'(e) > y$, an increase in the margin, m , yields

$$\frac{dR(\lambda, s)}{dm} = \frac{y}{(1-m)^2} - \frac{F'\left(e - \frac{m}{1-m}(1-\lambda)s\right)}{(1-m)^2} + \frac{m(1-\lambda)sF''\left(e - \frac{m}{1-m}(1-\lambda)s\right)}{(1-m)^3} < 0. \quad (\text{A.2})$$

An increase in the number of tokens, s , yields

$$\frac{dR(\lambda, s)}{ds} = \frac{m^2(1-\lambda)F''\left(e - \frac{m}{1-m}(1-\lambda)s\right)}{(1-m)^2} < 0, \quad (\text{A.3})$$

¹The views expressed in this paper are those of the authors and do not necessarily represent those of Federal Reserve Board of Governors, or anyone in the Federal Reserve System.

while an increase in redemptions, λ , yields

$$\frac{dR(\lambda, s)}{d\lambda} = -\frac{m^2 s F''(e - \frac{m}{1-m}(1-\lambda)s)}{(1-m)^2} > 0. \quad (\text{A.4})$$

A.1.2 Details steps for derivation of unique θ^* in global game

Given the private signal, an individual patient investor will update their posterior about θ , which will be uniform in $[x_i - \epsilon, x_i + \epsilon]$ and compute the expected payoff differential

$$\Delta(x_i) = \int_{x_i - \epsilon}^{x_i + \epsilon} \nu(\theta, \lambda) \frac{d\theta}{2\epsilon}. \quad (\text{A.5})$$

If $x_i \geq \bar{\theta} + \epsilon$, the individual patient investor can conclude that $\theta \geq \bar{\theta}$ and will not redeem, independent of their belief about λ ($\Delta(x_i) > 0$). Similarly, if $x_i < \underline{\theta} - \epsilon$, the individual patient investor can conclude that $\theta < \underline{\theta}$ and will redeem, independent of their belief about λ ($\Delta(x_i) < 0$). These are the *upper and lower dominance* regions for θ , where the individual action is independent of the beliefs about the actions of others.

For intermediate $x_i \in [\underline{\theta} - \epsilon, \bar{\theta} + \epsilon]$, the sign of $\Delta(x_i)$ depends on the beliefs about λ . To pin down these beliefs, we focus on a threshold strategy that all patient investors follow. We show that there exists a unique signal threshold x^* , such that every investor redeems if their private signal $x_i < x^*$ and does not redeem if $x_i > x^*$. Given this threshold, an individual investor can form well-defined beliefs about the total number of redemptions by patient investors, denoted by $\lambda^b(\theta, x^*)s$, and given by the probability that other investors receive a private signal below x^* . If $\theta > x^* + \epsilon$, all patient investors get signals $x_i > x^*$, none redeem, and $\lambda^b(\theta, x^*) = \delta$. If $\theta < x^* - \epsilon$, all patient investors get signals $x_i < x^*$, all redeem, and $\lambda^b(\theta, x^*) = 1$. If $x^* - \epsilon \leq \theta \leq x^* + \epsilon$, some patient investors get signals $x_i > x^*$, while others get signals $x_i < x^*$; thus, under the threshold strategy, $\lambda^b(\theta, x^*) = (1 - \delta)\Pr(x_i < x^*) = \delta + (1 - \delta)(x^* - \theta + \epsilon)/(2\epsilon)$. The following equation summarizes these beliefs:

$$\lambda^b(\theta, x^*) = \begin{cases} 1 & \text{if } \theta < x^* - \epsilon \\ \delta + (1 - \delta)(x^* - \theta + \epsilon)/(2\epsilon) & \text{if } x^* - \epsilon \leq \theta \leq x^* + \epsilon \\ \delta & \text{if } \theta > x^* + \epsilon \end{cases} \quad (\text{A.6})$$

Using (A.6), an investor can compute the expected payoff differential using their posterior about θ , given the signal x_i and an assumed value for x^* :

$$\Delta(x_i, x^*) = \int_{x_i - \epsilon}^{x_i + \epsilon} \nu(\theta, \lambda^b(\theta, x^*)) \frac{d\theta}{2\epsilon}. \quad (\text{A.7})$$

Unlike in (A.5), beliefs in (A.7) are uniquely determined and pin down the payoff differential.

Under a threshold strategy, a patient investor does not redeem ($\Delta(x_i, x^*) > 0$) if $x_i > x^*$ and redeems ($\Delta(x_i, x^*) < 0$) if $x_i < x^*$. By continuity, the investor that receives the threshold signal x^* is indifferent between not redeeming and redeeming, i.e.,

$$\Delta(x^*, x^*) = \int_{x^*-\epsilon}^{x^*+\epsilon} \nu(\theta, \lambda^b(\theta, x^*)) \frac{d\theta}{2\epsilon} = 0. \quad (\text{A.8})$$

A threshold strategy also implies thresholds for fundamentals $\theta_{\hat{\lambda}}$ and $\theta_{\bar{\lambda}}$ such that the issuer is solvent at $t = 2$ for $\theta \geq \theta_{\hat{\lambda}}$ and has enough liquidity at $t = 1$ for $\theta \geq \theta_{\bar{\lambda}}$ given signal threshold x^* and redemptions $\lambda^b(\theta, x^*)$. These thresholds are determined by $\hat{\lambda} = \lambda^b(\theta_{\hat{\lambda}}, x^*)$ and $\bar{\lambda} = \lambda^b(\theta_{\bar{\lambda}}, x^*)$. Using these, the threshold (A.8) can be expanded to

$$\begin{aligned} \Delta(x^*, x^*) = & - \int_{x^*-\epsilon}^{\theta_{\bar{\lambda}}} \frac{\ell + (1-\ell)\xi}{\lambda^b(\theta, x^*)} \frac{d\theta}{2\epsilon} + \int_{\theta_{\bar{\lambda}}}^{\theta_{\hat{\lambda}}} \left[\theta \frac{X(1-\ell) \left[1 - \frac{\lambda^b(\theta, x^*) - \ell}{\xi(1-\ell)} \right]}{1 - \lambda^b(\theta, x^*)} - 1 \right] \frac{d\theta}{2\epsilon} \\ & + \int_{\theta_{\hat{\lambda}}}^{x^*+\epsilon} \left[\theta R(\lambda^b(\theta, x^*), s) + (1-\theta) \max \left(\frac{\ell - \lambda^b(\theta, x^*)}{1 - \lambda^b(\theta, x^*)}, 0 \right) - 1 \right] \frac{d\theta}{2\epsilon} = 0. \end{aligned} \quad (\text{A.9})$$

As is typical in the global game literature, we focus on the limiting case where noise $\epsilon \rightarrow 0$, which also implies that $\theta_{\hat{\lambda}}, \theta_{\bar{\lambda}} \rightarrow x^*$. We will denote by θ^* this common threshold that the fundamentals' thresholds, $\theta_{\hat{\lambda}}, \theta_{\bar{\lambda}}$, and signal threshold, x^* , converge to. Expressing (A.9) in terms of θ^* and changing variables from θ to λ , such that as θ decreases from $x^* + \epsilon$ to $x^* - \epsilon$, λ uniformly increases from 0 to $1 - \delta$, we get

$$\begin{aligned} \bar{\Delta}^* = & \int_{\delta}^{\hat{\lambda}} \left[\theta^* R(\lambda, s) + (1 - \theta^*) \max \left(\frac{\ell - \lambda}{1 - \lambda}, 0 \right) - 1 \right] \frac{d\lambda}{1 - \delta} \\ & + \int_{\bar{\lambda}}^{\hat{\lambda}} \left[\theta^* \frac{X(1-\ell) \left[1 - \frac{\lambda - \ell}{\xi(1-\ell)} \right]}{1 - \lambda} - 1 \right] \frac{d\lambda}{1 - \delta} - \int_{\bar{\lambda}}^1 \frac{\ell + (1-\ell)\xi}{\lambda} \frac{d\lambda}{1 - \delta} = 0. \end{aligned} \quad (\text{A.10})$$

Existence and Uniqueness of Threshold Equilibrium. $\bar{\Delta}^*$ is continuous in θ^* because all integrands are continuous and the discontinuity in v occurs only at one discrete point, $\hat{\lambda}$. Then, from the existence of the upper and dominance regions, there exists a θ^* such that $\bar{\Delta}^* = 0$. It is, then, easy to show that the expected payoff differential is positive (negative) for an investor who receives signal $x_i > \theta^*$ ($x_i < \theta^*$), and hence the threshold strategy θ^* is indeed an equilibrium. Intuitively, observing a higher signal shifts probability from negative values of v to positive values of v as beliefs about aggregate withdrawals improve using (A.6); recall

that noise is uniformly distributed. Given that v changes sign—"crosses zero"—only once, it follows that the posterior average of v is higher (lower) for $x_i > \theta^*$ ($x_i < \theta^*$) and, hence, positive (negative); we refer the reader to Goldstein and Pauzner, 2005, and Kashyap et al. 2024 for the technical details and precise exposition.² Finally, observe that $d\bar{\Delta}^*/d\theta^* > 0$, so θ^* and, hence, the threshold equilibrium strategy are unique.

A.1.3 Derivatives of θ^* with respect to y, m, s, ℓ

Total differentiating (A.10) yields the following derivatives:

$$\frac{d\theta^*}{dx} = -\frac{d\bar{\Delta}^*}{dx} \left[\frac{d\bar{\Delta}^*}{d\theta^*} \right]^{-1} \quad \text{for } x \in \{y, m, s, \ell\}$$

Note $d\bar{\Delta}^*/dx = \int_{\delta}^{\hat{\lambda}} \theta^* dR(\lambda, s)/dx d\lambda > 0$ for $x \in \{y, m, s\}$, thus they affect ξ^* only through R . Using (A.1)–(A.3) and $d\bar{\Delta}^*/d\theta^* > 0$, we have

$$\frac{d\theta^*}{dy} < 0 \quad \& \quad \frac{d\theta^*}{dm} > 0 \quad \& \quad \frac{d\theta^*}{ds} > 0. \quad (\text{A.11})$$

Finally,

$$\begin{aligned} \frac{d\bar{\Delta}^*}{d\ell} &= \frac{d\hat{\lambda}}{d\ell} \left[\theta^* R(\hat{\lambda}, s) - 1 \right] \frac{1}{1-\delta} + \int_{\delta}^{\ell} (1-\theta^*) \frac{1}{1-\lambda} \frac{d\lambda}{1-\delta} \\ &\quad - \frac{d\hat{\lambda}}{d\ell} \left[\theta^* \frac{X(1-\ell) \left[1 - \frac{\hat{\lambda}-\ell}{\xi(1-\ell)} \right]}{1-\hat{\lambda}} - 1 \right] \frac{1}{1-\delta} + \int_{\hat{\lambda}}^{\bar{\lambda}} \frac{X(1/\xi-1)}{1-\lambda} \frac{d\lambda}{1-\delta} - \int_{\bar{\lambda}}^1 \frac{1-\xi}{\lambda} \frac{d\lambda}{1-\delta}. \end{aligned} \quad (\text{A.12})$$

Given that $d\hat{\lambda}/d\ell > 0$ from (7), all the terms in the above condition are positive apart from the last one, which means that the effect of ℓ on θ^* may be ambiguous. This is a typical property in bank-run models, and it is intuitive: It suggests that in the region of beliefs about redemptions that a run materializes, higher liquidity increases the payoff from redeeming because individuals can successfully redeem their tokens with higher probability. We derive below a (weak) sufficient—not necessary—condition for $d\bar{\Delta}^*/d\ell > 0$, which requires that the expected lending rate is below a threshold, supported by the data.

²Note that for the existence of a threshold equilibrium the strongest property of one-sided strategic complementarities is not needed and single-crossing of v suffices as Goldstein and Pauzner (2005) also point out. Given our focus on threshold equilibria, we do not make further assumptions.

Under the sufficient condition, we unambiguously obtain

$$\frac{\partial \theta^*}{\partial \ell} < 0. \quad (\text{A.13})$$

Note that (A.13) can still hold in alternative parameterizations violating the sufficient condition but may also not hold. In the latter cases, the issuer would set $\ell = 0$, which is inconsistent with observed stablecoin reserve portfolios (see Section 2.4 for issuer optimization problem).

Sufficient condition for $d\bar{\Delta}^*/d\ell > 0$. Substituting (9) in (A.12) we get that

$$\begin{aligned} \frac{d\bar{\Delta}^*}{d\ell} = & \frac{d\hat{\lambda}}{d\ell} \left[\theta^* R(\hat{\lambda}, s) - 1 \right] \frac{1}{1-\delta} + \int_{\delta}^{\ell} (1-\theta^*) \frac{1}{1-\lambda} \frac{d\lambda}{1-\delta} \\ & - \frac{1}{\ell} \int_{\delta}^{\hat{\lambda}} \left[\theta^* R(\lambda, s) + (1-\theta^*) \max\left(\frac{\ell-\lambda}{1-\lambda}, 0\right) - 1 \right] \frac{d\lambda}{1-\delta} \\ & - \frac{d\hat{\lambda}}{d\ell} \left[\theta^* \frac{X(1-\ell) \left[1 - \frac{\hat{\lambda}-\ell}{\xi(1-\ell)}\right]}{1-\hat{\lambda}} - 1 \right] \frac{1}{1-\delta} - \frac{1}{\ell} \int_{\hat{\lambda}}^{\bar{\lambda}} \left[\theta^* \frac{X(1-\ell) \left[1 - \frac{\lambda-\ell}{\xi(1-\ell)}\right]}{1-\lambda} - 1 \right] \frac{d\lambda}{1-\delta} \\ & + \int_{\hat{\lambda}}^{\bar{\lambda}} \frac{X(1/\xi - 1)}{1-\lambda} \frac{d\lambda}{1-\delta} + \frac{1}{\ell} \int_{\bar{\lambda}}^1 \frac{\xi}{\lambda} \frac{d\lambda}{1-\delta}. \end{aligned} \quad (\text{A.14})$$

Given that $d\hat{\lambda}/d\ell > 0$ from (7), the terms in the last two lines in (A.14) are all positive and, thus, we only need to sign the terms in the first two lines. Add and subtract $d\hat{\lambda}/d\ell \cdot (1-\theta^*) \cdot (\ell-\delta)/(1-\delta)^2$. Then, because (i) $dR(\lambda, s)/d\lambda > 0$ from (3), (ii) $d(\ell-\lambda)/(1-\lambda)/d\ell = -(1-\ell)/(1-\lambda)^2 < 0$, and (iii) $d(1-\lambda)^{-1}/d\lambda > 0$, the sum of the terms in the first two line is strictly higher than

$$\begin{aligned} & \left[\frac{d\hat{\lambda}}{d\ell} - \frac{\hat{\lambda}-\delta}{\ell} \right] \left[\theta^* R(\hat{\lambda}, s) + (1-\theta^*) \max\left(\frac{\ell-\delta}{1-\delta}, 0\right) - 1 \right] \frac{1}{1-\delta} \\ & + (1-\theta^*) \frac{\ell-\delta}{(1-\delta)^2} \left(1 - \frac{d\hat{\lambda}}{d\ell} \right). \end{aligned} \quad (\text{A.15})$$

The last term is strictly positive because $d\hat{\lambda}/d\ell = X(1-\xi)/(X-\xi) < 1$. Moreover,

$$\theta^* R(\hat{\lambda}, s) + (1-\theta^*) \max\left(\frac{\ell-\delta}{1-\delta}, 0\right) - 1 > \underline{\theta} R(\delta, s) + (1-\underline{\theta}) \max\left(\frac{\ell-\delta}{1-\delta}, 0\right) - 1 = 0,$$

because $\theta^* > \underline{\theta}$ and $R(\hat{\lambda}, s) > R(\delta, s)$. If $d\hat{\lambda}/d\ell - \hat{\lambda}-\delta/\ell > 0 \Rightarrow \delta < \xi(X-1)/X-\xi$, then

$d\bar{\Delta}^*/d\ell > 0$ always. For δ lower than that threshold, we can derive a sufficient condition for the lending rate such that the absolute value of the terms in the first line in (A.15) is lower than $1/\ell \int_{\lambda}^1 \frac{\xi}{\lambda} d\lambda$ and, thus, (A.14) is positive. The latter term is strictly higher than $1/\ell \xi \log \xi$, while the absolute value of the former is strictly lower than $1/\ell \cdot \xi(X-1)/(X-\xi) \cdot (\max R - 1)$, where we considered the higher possible lending rate and set $\delta = 0$. Thus, it is sufficient that $\max R \leq -\log \xi \cdot (X-\xi)/(X-1)$ for $d\bar{\Delta}^*/d\ell > 0$. This condition is easily satisfied. For example, given an expected yield of 5% for the illiquid asset, i.e., $X = 2.1$, and a liquidity discount of 25%, i.e., $\xi = 0.75$, it is sufficient that the expected lending rate is lower than 35%, which is the case in our data. As mentioned, the sufficient condition on the lending rate is rather weak and we could be relaxed further if we consider the effect of the other positive terms in (A.14).

A.1.4 Derivatives of P with respect to y, m, s, ℓ

We show how P changes with the demand and riskiness of cryptocurrencies as well as the size and liquidity of the stablecoin. We first examine the effect stemming from the cryptocurrency demand, y , and riskiness, m , as well as the size of the stablecoin s . For $x \in \{y, m, s\}$ we have

$$\begin{aligned} \frac{dP}{dx} = (1-\delta) \frac{dR(\delta, s)}{dx} \frac{1 - (\theta^*)^2}{2} \\ - \frac{d\theta^*}{dx} \left\{ (1-\delta) \left[\theta^* R(\delta, s) + (1-\theta^*) \max \left(\frac{\ell - \delta}{1-\delta}, 0 \right) \right] + \delta - (\ell + (1-\ell)\xi) \right\}. \end{aligned} \quad (\text{A.16})$$

Using (A.1)–(A.3) and (A.11), and $\theta^* R(\delta, s) + (1-\theta^*) \max((\ell - \delta)/(1-\delta), 0) > 1 > (\ell + (1-\ell)\xi)$ since $\theta^* > \underline{\theta}$, we have that

$$\frac{dP}{dy} > 0 \quad \& \quad \frac{dP}{dm} < 0 \quad \& \quad \frac{dP}{ds} < 0. \quad (\text{A.17})$$

In other words, the higher the cryptocurrency demand, the lower the risk, or the smaller the stablecoin circulation is, the higher P is for two reasons. First, a higher y , and lower m or s , increase the payoff conditional on a run not occurring (first term in (A.16)). Second, the probability that a run does not occur increases with y , and decreases with m or s , as the incentives to run are lower, all else equal (second term in (A.16)).

Finally, a change in ℓ changes P according to

$$\begin{aligned} \frac{dP}{d\ell} &= \int_{\theta^*}^1 (1 - \theta) \cdot (\ell > \delta) d\theta + \int_0^{\theta^*} (1 - \xi) d\theta \\ &\quad - \frac{d\theta^*}{d\ell} \left[\theta^* R(\delta, s) + (1 - \theta^*) \max \left(\frac{\ell - \delta}{1 - \delta}, 0 \right) - (\ell + (1 - \ell)\xi) \right] > 0. \end{aligned} \quad (\text{A.18})$$

In other words, the higher the percentage of liquid assets in stablecoin reserves is, the higher P is for two reasons. First, a higher ℓ increases the probability of being paid conditional on a run occurring (first term in (A.18)). Second, the probability that a run does not occur increases with ℓ , all else equal (second term in (A.18)).³

A.2 Model Extension: Endogenous Margin Requirements

In the baseline model, we assume that the exchange sets margin m without considering how it should be chosen optimally between traders and investors. Given that levered lending takes place after run uncertainty is resolved, m should not depend on θ^* , but can depend on R . To derive an optimal m we consider a structure—akin to Fostel and Geanakoplos (2008)—under which traders offer investors a menu of contracts $k \in K$ described by a pair (R_k, m_k) , where R_k is given, in equilibrium, by (3) for certain m_k . That is, traders offer investors a menu of contracts under all of which they break even. Given that these contracts are offered after redemptions λ s have been observed, they are only parameterized by different m_k . Investors will then choose the contract that maximizes their utility.

To introduce a trade-off, we suppose for this extension that investors face a cost c for directly holding the cryptocurrency when the trader defaults. Recall from Section 2.1 that R is the expected lending rate, incorporating the payoff when traders default, and that traders will default for cryptocurrency payoff realizations $\tilde{y} < y'_k$, i.e., y'_k is the threshold below which investors receive the collateral for margin m_k and is given by

$$y'_k = (1 - m)R_{k,c} \quad \Rightarrow \quad y'_k = \frac{\bar{y}(1 - m)R_k - y_k'^2/2}{\bar{y} - y'_k}, \quad (\text{A.19})$$

where we replaced the contractual lending rate, $R_{k,c}$, with the expected lending rate, R_k .

The probability that investors incur the cost c is equal to $\int_{\tilde{y} < y'_k} dF(\tilde{y})$. We assume $\tilde{y} \sim U[0, \bar{y}]$ for simplicity and, thus, $y = \bar{y}/2$. Investor's payoff from not redeeming is equal to $\theta \left(R_k - c \int_{\tilde{y} < y'_k} dF(\tilde{y}) \right) + (1 - \theta) \max \left(\frac{\ell - \lambda}{1 - \lambda}, 0 \right)$. In other words, the expected return from lending the stablecoin, R , is curtailed by the expected cost of holding it when the trader

³As mentioned in Section 2.2 if $d\theta^*/d\ell > 0$ and, hence, $dP/d\ell < 0$, then the issuer will choose $\ell = 0$. See the Online Appendix Section A.1.3 for a sufficient condition to exclude this case.

defaults. Among all available contracts $(R_k(m_k), m_k) \forall k \in K$, the investor will choose the one that delivers the higher payoff. Given that R_k is a function of m_k , we only need to find the m_k that maximizes the investor's payoff, which is the solution to

$$\theta \frac{\partial R_k}{\partial m_k} - \theta c \frac{\partial y'_k}{\partial m_k} \frac{1}{\bar{y}} + \underline{\psi}_k - \bar{\psi}_k = 0 \quad (\text{A.20})$$

where $\underline{\psi}_k$ and $\bar{\psi}_k$ are the Lagrange multipliers on $m_k \geq 0$ and $m_k \leq \bar{m}$, where \bar{m} is the maximum margin traders would be willing to post given by $y - \bar{m}F'(e - m/1 - m(1 - \lambda)s) = 0$. Recall that $dR_k(\lambda, s)/dm_k$ is given by (A.2) and is negative. dy'_k/dm_k is obtained by totally differentiating (A.19)

$$\frac{dy'_k}{dm_k} = \frac{\bar{y}}{\bar{y} - y'_k} \left((1 - m) \frac{dR_k}{dm_k} - R_k \right) < 0. \quad (\text{A.21})$$

For $m \rightarrow 0$, the first two terms in (A.20) converge to $y - (1 - c/(\bar{y} - y'_k))F'(e)$ and is positive only if $c \geq \bar{c} \equiv (\bar{y} - y'_k)(1 - y/F'(e)) > 0$ given the assumption $F'(e) > y$. For $m \rightarrow \bar{m}$, the sum of the first two terms converges to $dR_k/dm_k|_{m_k=\bar{m}}(1 - c/(\bar{y} - y'_k)(1 - \bar{m})) < 0$. Moreover, the sum of the first two terms is strictly decreasing for $F'''' > 0$, which is typical for widely used concave technologies such as Cobb-Douglas production function, and we will assume herein. Hence, the contract that investors choose is unique and depends on the level of c .

Case I. If $c \leq \bar{c}$, then $\underline{\psi}_k > 0$, $\bar{\psi}_k = 0$, and $m_k = 0$, such that our baseline analysis carries through in its entirety.

Case II. If $c \geq \bar{c}$, m_k is interior, i.e., $\underline{\psi}_k = \bar{\psi}_k = 0$. In this case, m_k will be a function of y , λ , and s , and we need to show that our baseline results do not change. Essentially, we need to show that the derivatives of R_k with respect to $x \in \{y, \lambda, s\}$ do not change sign. We show that this is the case for sufficiently high c . Note that

$$\frac{dR_k}{dx} = \frac{\partial R_k}{\partial x} + \frac{\partial R_k}{\partial m_k} \frac{dm_k}{dx}. \quad (\text{A.22})$$

Given that $\partial R_k/\partial x$ and $\partial R_k/\partial m_k$ are given by (A.1)-(A.4), we only need to sign dm_k/dx , which we can compute by totally differentiating (A.20):

$$\begin{aligned} \frac{dm_k}{dy} &= - \frac{\frac{\partial^2 R_k}{\partial m_k \partial y} (\bar{y} - y'_k - c(1 - m_k)) + \left(2 - \frac{\partial y'_k}{\partial y}\right) \frac{\partial R_k}{\partial m_k} + c \frac{\partial R_k}{\partial y}}{\frac{\partial^2 R_k}{\partial m_k^2} (\bar{y} - y'_k - c(1 - m_k)) + \left(2c - \frac{\partial y'_k}{\partial m_k}\right) \frac{\partial R_k}{\partial m_k}} \\ \frac{dm_k}{dx} &= - \frac{\frac{\partial^2 R_k}{\partial m_k \partial x} (\bar{y} - y'_k - c(1 - m_k)) - \frac{\partial y'_k}{\partial x} \frac{\partial R_k}{\partial m_k} + c \frac{\partial R_k}{\partial x}}{\frac{\partial^2 R_k}{\partial m_k^2} (\bar{y} - y'_k - c(1 - m_k)) + \left(2c - \frac{\partial y'_k}{\partial m_k}\right) \frac{\partial R_k}{\partial m_k}}, \quad x \in \{\lambda, s\}. \end{aligned} \quad (\text{A.23})$$

Since $F''' > 0$, $\partial^2 R_k / \partial m_k^2 < 0$, and $\partial^2 R_k / \partial m_k \partial y > 0$, $\partial^2 R_k / \partial m_k \partial \lambda > 0$, $\partial^2 R_k / \partial m_k \partial s < 0$. Moreover, by totally differentiating (A.19), we get

$$\begin{aligned} \frac{dy'_k}{dy} &= \frac{1}{\bar{y} - y'_k} \left(\bar{y}(1-m) \frac{\partial R_k}{\partial dy} + 2(1-m)R_k - 2y'_k \right) \\ \frac{dy'_k}{dx} &= \frac{\bar{y}}{\bar{y} - y'_k} (1-m) \frac{\partial R_k}{\partial x}, \quad x \in \{\lambda, s\}. \end{aligned} \quad (\text{A.24})$$

Consider $c \rightarrow \bar{y}/1 - m_k^+$. Then, using (A.21) and (A.24), we get

$$\frac{dm_k}{dy} = -\frac{\frac{(2\bar{y}-2(1-m_k)R_k)}{(\bar{y}-y'_k)}}{2c - \frac{dy'_k}{dm_k}} - \frac{c - \frac{\bar{y}}{(\bar{y}-y'_k)}(1-m_k) \frac{\partial R_k}{\partial m_k}}{2c - \frac{\bar{y}}{(\bar{y}-y'_k)}(1-m_k) \frac{\partial R_k}{\partial m_k} + \frac{\bar{y}}{(\bar{y}-y'_k)} R_k} \frac{\partial R_k / \partial y}{\partial R_k / \partial m_k} < -\frac{\partial R_k / \partial y}{\partial R_k / \partial m_k}, \quad (\text{A.25})$$

$$\frac{dm_k}{d\lambda} = -\frac{c - \frac{\bar{y}}{(\bar{y}-y'_k)}(1-m_k) \frac{\partial R_k}{\partial m_k}}{2c - \frac{\bar{y}}{(\bar{y}-y'_k)}(1-m_k) \frac{\partial R_k}{\partial m_k} + \frac{\bar{y}}{(\bar{y}-y'_k)} R_k} \frac{\partial R_k / \partial \lambda}{\partial R_k / \partial m_k} < -\frac{\partial R_k / \partial \lambda}{\partial R_k / \partial m_k}, \quad (\text{A.26})$$

$$\frac{dm_k}{ds} = -\frac{c - \frac{\bar{y}}{(\bar{y}-y'_k)}(1-m_k) \frac{\partial R_k}{\partial m_k}}{2c - \frac{\bar{y}}{(\bar{y}-y'_k)}(1-m_k) \frac{\partial R_k}{\partial m_k} + \frac{\bar{y}}{(\bar{y}-y'_k)} R_k} \frac{\partial R_k / \partial s}{\partial R_k / \partial m_k} > -\frac{\partial R_k / \partial s}{\partial R_k / \partial m_k}. \quad (\text{A.27})$$

Using the above, we get from (A.22) that $dR_k/dy > 0$, $dR_k/d\lambda > 0$, and $dR_k/ds < 0$, which means that qualitatively the derivatives of the lending rate with respect to x are the same as in the baseline analysis in Section 2.1 where the margin was constant.

It follows that away from the limit, but for $c \geq \bar{c}$, it also suffices to establish the bounds in (A.25)-(A.27). These latter expressions for the bounds hold if

$$2 \frac{\partial^2 R_k}{\partial m_k \partial x} \frac{\partial R_k}{\partial m_k} > \frac{\partial^2 R_k}{\partial m_k^2} \frac{\partial R_k}{\partial x},$$

which holds for all $x \in \{y, \lambda, s\}$ using (A.1)-(A.4) and the cross-derivatives above. Hence, endogenous margins may weaken quantitatively the effect of $\{y, \lambda, s\}$ on R_k —because $\partial R_k / \partial m_k$ —but qualitatively do not matter.

A.3 Model Extension: Payment services from stablecoins

In this section, we introduce an additional source of demand for stablecoins arising from use cases other than speculation. Such services may accrue from facilitating cross-country payments, services offered exclusively by the digital-asset ecosystem, or tax evasion and

illicit activities. The value of these services could be aggregated in a convenience yield V , which can be constant or depend on the number of stablecoins in circulation, that is $V(\lambda, s) \equiv V((1 - \lambda)s)$ with $dV/ds < 0$ and $dV/d\lambda > 0$ following Krishnamurthy and Vissing-Jorgensen (2012). Then, the stablecoin payoff from not redeeming when the issuer does not default is given by $\theta(R(\lambda, s) + V(\lambda, s)) + (1 - \theta) \max(\ell - \lambda/1 - \lambda, 0)$. A positive convenience yield increases the payoff and decreases the probability of a run *ceteris paribus*. Moreover, if the convenience yield decreases in the number of stablecoins, then the stabilization mechanism operating via the redemptions channel is strengthened. The stabilization mechanism via the liquid portfolio share continues to operate in the absence of a convenience yield.

A.4 Model Extension: Speculating on Stablecoin Collapse

In the baseline model, we considered that investors lent their stablecoins to traders, who want to take leverage on cryptocurrencies after run uncertainty has been resolved. Yet, traders may want to borrow the stablecoins before run uncertainty is resolved so that they can also speculate on the collapse of the stablecoin. The idea is that promised repayment is denominated in stablecoins and, thus, if the stablecoin price collapses to zero after a run, traders would need to repay zero without losing their pledged collateral.

We consider a very simple extension of the model to introduce this motive. There are two types of traders and investors: A and B. Both types are identical with the difference that type A traders borrow stablecoins from type A investors before $t = 1$, while type B traders borrow stablecoins from type B investors after $t = 1$. We assume that the tokens lent early are circulated back to other stablecoin investors, who want to lend them after run uncertainty is resolved at $t = 1$. Type A investors are of mass $1 - \delta$, which is equal to patient type B investors. Traders of both types have the same endowment and each type has its own, distinct, outside technology. A-traders still need to pledge collateral, thus they buy the cryptocurrency on margin as in the baseline model. Thus, the return on the outside options is given by $\rho_A = F'[e - m/(1 - m)(1 - \delta)s]$ for A-traders and $\rho_B = F'[e - m/(1 - m)(1 - \delta)s]$ for B-traders, in equilibrium when a run does not occur. The run decision for the B-investors is the same as in Section 2.2 and the stablecoin price they are willing to offer is given by (11) with the difference that the lending rate will be different. Next, we derive the lending rate and the participation decision for the A-investors.

Denote by \hat{R} the expected lending rate for borrowing before $t = 1$. As before, A-traders will break even with their outside option but in this case, they additionally do not need to repay anything when the stablecoin collapses in a run or when the issuer default conditional on a run not occurring, because the price of tokens goes to zero. Hence, their participation

constraint is

$$\int_{\theta^*}^1 [\theta(y - (1 - m)\hat{R}) + (1 - \theta)y]d\theta + \int_0^{\theta^*} yd\theta = m\rho_A, \quad (\text{A.28})$$

yielding

$$\hat{R} = \frac{y - m\rho_A}{1 - m} \frac{2}{1 - \theta^{*2}}. \quad (\text{A.29})$$

The first term in the left-hand side of (A.28) is the payoff to A-traders conditional that a run does not occur: with probability θ ($\geq \theta^*$) the issuer is solvent and A-traders need to repay their stablecoin-denominated loan, while with probability $1 - \theta$ the issuer is insolvent and tokens are worth zero, so A-traders can pocket the cryptocurrency return in its entirety. The second term in the left-hand side of (A.28) is the payoff to A-traders conditional on a run: A-traders pocket the whole cryptocurrency return because their stablecoin-denominated loan is worth zero. The right-hand side in (A.28) is the outside-option payoff.

Using (1) and (A.29), we can compare the lending rates for lending before and after $t = 1$, R and \hat{R} . It is easy to see that $\hat{R} > R$. This result is intuitive. Traders face a trade-off when borrowing early: If the run occurs, they gain a lot and are willing to offer high lending rates. But, if the run does not occur, they will pay higher lending rates with probability θ .

A.5 Token Supply and Stablecoin Liquidity Without Observability

In the paper we derive the optimal choice of s and ℓ under observability. However, as noted, the choice of ℓ may not be observable in real time contrary to s . The issuer may still use a combination of ℓ and s to maintain the peg in response to crypto-related shocks between $t = 0$ and $t = 1$ but cannot credibly commit to a certain choice of ℓ given that it is not observable. This information resembles an incomplete contract whereby the issuer may deviate from the choice of ℓ after the peg is stabilized (see Online Appendix in Kashyap et al. 2023). The issuer will maximize the profits accruing to them when choosing ℓ and s but will only internalize the effect of s and not ℓ on the peg stability condition $P = 1$. Yet, the issuer will still internalize the effect of both ℓ and s on the run threshold θ^* , since the run may happen later at $t = 1$. Then, the optimality condition with respect to ℓ is

$$\frac{1 - (\theta^*)^2}{2} \frac{d\Pi(\delta)}{d\ell} s - \theta^* \Pi(\delta) s \frac{d\theta^*}{d\ell} = 0, \quad (\text{A.30})$$

which together with $P = 1$ yields the optimal (ℓ, s) . Comparing (A.30) to (14) we see that former misses a wedge W equal to

$$W = -\frac{dP/d\ell}{dP/ds}\Pi(\delta)\left(\frac{1 - (\theta^*)^2}{2} - \theta^*\frac{d\theta^*}{ds}s\right). \quad (\text{A.31})$$

For a given s , the issuer will choose a lower (higher) ℓ if $W > 0$ ($W < 0$) when ℓ is unobservable compared to the case that it is.⁴ In turn, this means that the change in s should be higher (lower) to stabilize the peg for the same level of crypto-related shocks. Importantly, the issuer will use both stabilization mechanisms to maintain the peg even when ℓ is unobservable. The following Proposition shows that the sign of W depends on the level of run risk.

Proposition 3. *There exist a unique $\hat{\theta} \in (0, 1)$ such that W in (A.31) is positive for $\theta^* < \hat{\theta}$ and negative for $\theta^* > \hat{\theta}$.*

The proof is straightforward. Since $dP/ds < 0$ and $dP/d\ell > 0$ from (A.17) and (A.18), the sign of W depends on the sum of the terms in the parenthesis, which is continuous in θ^* , negative for $\theta^* \rightarrow 1$ and positive for $\theta^* \rightarrow 0$, while $d\theta^*/ds$ is positive and increasing in θ^* . Hence, $\hat{\theta}$ exists and is unique. This result is intuitive. When ℓ is not observable, the issuer has an incentive to deviate but at the same time still internalizes how the choice of ℓ matters for run risk and, thus, their expected profits. If run risk is low, i.e., $\theta^* < \hat{\theta}$ and $W > 0$, the issuer deviates toward a lower ℓ , and vice versa if run risk is high. Investors anticipate this deviation and respond by redeeming more or fewer tokens compared to the case of observable ℓ .

Proposition 3 also has implications for the viability of the stablecoin when the speculative demand for cryptocurrencies wanes. In particular, suppose that there is a shock pushing y below $1 + m[F'(e) - 1]$. If ℓ is observable, the issuer would set $\ell = 1$ and keep the stablecoin running with $\theta^* = 0$, i.e., no run risk (Proposition 2). But, with unobservable ℓ , the issuer will have an incentive to deviate toward $\ell < 1$. Investors would anticipate this and redeem all their tokens immediately; otherwise, they would be exposed to run risk without the proper compensation. By continuity, the same would hold for y close to, but higher than, $1 + m[F'(e) - 1]$, even though expected lending rates would be (somewhat) higher than one for this level of y . Overall, stablecoins are not viable for low enough y under non-observability of ℓ , which also provides an additional rationale why issuers may want to disclose their reserves more frequently during crypto turmoil, similar to what USDC did in May 2022.

⁴To see this, note that the solution under observable ℓ can be implemented in an environment where ℓ is not observable under a Pigouvian tax/subsidy on liquid holdings ℓ : A negative (positive) W calls for tax (subsidy), implying lower (higher) ℓ than in the unconstrained equilibrium with unobservable ℓ .

A.6 Positive interest on liquid assets

This section extends the baseline model to allow for a positive interest rate on the liquid assets. For simplicity and without loss of generality, we assume that the liquid asset pays off $r \geq 1$ from $t = 1$ to $t = 2$, while it continues to pay zero interest from $t = 0$ to $t = 1$. The case that $r = 1$ corresponds to a zero (net) interest rate in the baseline model. This extension should suffice for the purpose of studying how r matters for the $t = 2$ profits of the issuer and, hence, the choice of ℓ . To maintain a risk premium over the liquid asset we also set the illiquid asset payoff to be a function of r , i.e., $X(r)$.

The issuer's profits are then given by

$$\begin{aligned} & \int_{\theta^*}^1 \theta \left[X(r)(1 - \ell) \left(1 - \frac{(\delta - \ell)^+}{\xi(1 - \ell)} \right) + (\ell - \delta)^+ r - (1 - \delta) \right]^+ s d\theta \\ & + \int_{\theta^*}^1 (1 - \theta) \left[(\ell - \delta)^+ r - (1 - \delta) \right]^+ s d\theta. \end{aligned} \quad (\text{A.32})$$

The issuer will earn an interest rate on remaining liquid assets after repaying impatient investors, $\ell - \delta$; should the difference be positive. This may also allow the issuer to remain solvent even in the bad state that the illiquid asset pays zero. This requires that $r > \bar{r} = \frac{1 - \delta}{(\ell - \delta)^+}$. Note that for $\ell \rightarrow 1$, which implies $\theta^* \rightarrow 0$, the profits are $(1 - \delta)(r - 1)s$.

If the issuer remains solvent in the bad state of the world, then investors could lend their tokens to traders and earn the lending rate. Denote by $\tilde{\lambda}$ the maximum level of withdrawals such that the issuer remains solvent in the bad state given by $\tilde{\lambda} = \max\left(\delta, \frac{\ell r - 1}{r - 1}\right)$.

The run threshold θ^* is determined by

$$\begin{aligned} \bar{\Delta}^* &= \int_{\delta}^{\tilde{\lambda}} [R(\lambda, s) - r] \frac{d\lambda}{1 - \delta} + \int_{\tilde{\lambda}}^{\hat{\lambda}} \left[\theta^* R(\lambda, s) + (1 - \theta^*) \left(\frac{\ell - \lambda}{1 - \lambda} \right)^+ r - r \right] \frac{d\lambda}{1 - \delta} \\ &+ \int_{\hat{\lambda}}^{\bar{\lambda}} \left[\theta^* \frac{X(1 - \ell) \left[1 - \frac{\lambda - \ell}{\xi(1 - \ell)} \right]}{1 - \lambda} - r \right] \frac{d\lambda}{1 - \delta} - \int_{\bar{\lambda}}^1 \frac{\ell + (1 - \ell)\xi}{\lambda} r \frac{d\lambda}{1 - \delta} = 0. \end{aligned} \quad (\text{A.33})$$

For $\lambda \in [\delta, \tilde{\lambda})$, the issuer is solvent both in the good and bad state, and investors can lend out their tokens. For $\lambda \in [\tilde{\lambda}, \hat{\lambda})$, the issuer defaults in the bad state but may invest any remaining liquidity $(\ell - \lambda)^+$ at $t = 1$ in the liquid asset to earn r , which increases the payoff from not withdrawing when the bad state realizes. The payoffs in other regions are as in the baseline model, though note that the payoff from withdrawing at $t = 1$ increases with r in all regions. Also, the cutoffs $\hat{\lambda}$ and $\bar{\lambda}$ are functionally the same due to the simplifying assumption that the liquid asset pays interest only from $t = 1$ to $t = 2$; however, changing r affects $\hat{\lambda}$ through

$X(r)$. It is also easy to show that $\tilde{\lambda} \leq \ell < \hat{\lambda}$, i.e., the level of withdrawals needed to make the issuer insolvent in the good state is higher than the level needed for insolvency in the bad state.

The stablecoin price is given by

$$\begin{aligned}
P = & (r > \bar{r}) \cdot \int_{\theta^*}^1 \{(1 - \delta)R(\delta, s)/r + \delta\} d\theta \\
& + (r \leq \bar{r}) \cdot \int_{\theta^*}^1 \left\{ (1 - \delta) \left[\theta R(\delta, s)/r + (1 - \theta) \max\left(\frac{\ell - \delta}{1 - \delta}, 0\right) \right] + \delta \right\} d\theta \\
& + \int_0^{\theta^*} (\ell + (1 - \ell)\xi) d\theta,
\end{aligned} \tag{A.34}$$

where the $t = 2$ return from lending the token, $R(\delta, s)$, needs to be discounted by r but not the rest of the cash flow as they accrue at $t = 1$ and can thus earn r . Note that if $r > \bar{r}$, the issuer may be solvent in the bad state, and thus, investors can earn the lending rate in both states conditional on a run not occurring. Otherwise, if $r \leq \bar{r}$, investors receive pro-rata the remaining liquid resource in the bad state.

The issuer chooses ℓ and s to maximize (A.32) subject to $P = 1$, and with θ^* and P be determined by (A.33) and (A.34). Figure A.5 plots the profits of the issuer, for different levels of ℓ and r , normalized over the profits for $\ell = 0.63$; this is an illustrative parametrization of the model, which should not be taken as a realistic calibration.⁵ Nevertheless, the qualitative properties we highlight are general and do not depend on the choice of initial parameters. In particular, for both zero and positive but low-enough interest rates, the issuer optimally sets $\ell < 1$ to maximize profits, thus exposing the stablecoin to run risk. However, for a high enough interest rate, profits are maximized $\ell = 1$, alleviating any run risk.

A.7 Robustness for Measuring Expected Returns

We check that using Binance's BTC/USDT perpetual futures funding rate is a robust proxy for expected returns. One concern is that using the BTC/USDT perpetual futures as a proxy of y overweights idiosyncrasies specific to Bitcoin. In Table A.5, we show pairwise correlations of the BTC/USDT time series with several others. Binance also has perpetual futures that settle into Binance USD, another stablecoin, and we show that funding rates across perpetual futures are highly correlated regardless of which stablecoin they settle in. Another concern is that all futures funding rates on Binance reflect idiosyncrasies specific to Binance, rather

⁵We have set $y = 4$, $m = 0.1$, $e = 1$, $F(x) = \zeta x^\alpha$ with $\zeta = 2$ and $\alpha = 0.5$, $\delta = 0.55$, $\xi = 0.4$, and $X = 1.4r + 0.6$. We have considered three cases for r : $r = 1.0$ (zero interest rate), $r = 1.2$ (positive and low-enough interest rate), and $r = 1.4$ (positive and high-enough interest rate).

than aggregate expected returns for cryptocurrency beyond Binance. We compare Binance’s number with another large exchange, FTX, and find that funding rates are similar across the exchanges, confirming that the funding rates are not principally capturing exchange-specific factors. Finally, we show that perpetual futures funding rates are closely linked to expected returns embedded in crypto futures traded on the CME.

To address concerns about idiosyncrasies specific to Bitcoin, USDT, or Binance, we show correlations across several different contracts (BTC, ETH, and DOGE) settled in different types of stablecoins (USDT, BUSD, and FTX’s USD) and across both Binance and FTX. We include DOGE as it is known as a highly speculative currency and was arguably started as a joke. The last two columns are the expected return measures we infer from CME futures, which we describe below. Combined, all the series are highly correlated, indicating that variation in our main measure of y , BTC/USDT on Binance, is not principally reflecting something specific to BTC, USDT, or Binance instead of speculative expected returns. measures

We can also proxy for y using the expected return embedded in crypto futures traded on the CME. Unlike the highly levered offshore perpetual futures, these futures are vanilla futures and like equity index futures. The CME sets the rules for the derivatives, and they have standard monthly expirations. These crypto futures are widely used by U.S.-based institutional investors who want to speculate on the price of Bitcoin or Ether but are unwilling or unable to hold cryptocurrencies directly. While the futures have embedded leverage, they are considerably less levered than the offshore perpetual futures.⁶

We calculate expected returns y for Bitcoin and Ether using the futures prices. Let $F_{t,t+n}$ denote the price of a future at time t with delivery at $t + n$, and let $z_{t,t+n}$ denote the n -period discount factor implied by the risk-free rate. We can infer expected returns using a no-arbitrage argument comparing the present value of $F_{t,t+n}$ and $F_{t,t+n+1}$. The expected return is

$$\mathbb{E}_{t,t+n \rightarrow t+n+1}[y] \equiv \left(\frac{z_{t,t+n+1}}{z_{t,t+n}} \right) \frac{F_{t,t+n+1}}{F_{t,t+n}} \quad (\text{A.35})$$

We use the overnight-indexed swap curve to estimate the n -period discount factors: $z_{t,t+n} = 1/(1 + y_{t,t+n}^{\text{OIS}}/12)^{(1/12)}$ where $y_{t,t+n}^{\text{OIS}}$ is the n -month OIS yield. We prefer to use consecutive futures rather than the front-month future versus the spot because the futures include leverage which may introduce a bias relative to the spot price.

In principle, we can use the ratio of contracts with any expiration to calculate expected

⁶As of June 2022, the CME requires 50 percent (60 percent) margin for BTC (ETH) futures, allowing roughly $1 \times (0.67 \times)$ leverage. See <https://www.cmegroup.com/markets/cryptocurrencies.html>.

returns between the two contracts' expirations. We focus on the first and second front-month contracts for two reasons. First, using the shortest maturity contracts helps control for any distortions introduced by an upward-sloping term structure of risk premia. Second, the liquidity of derivative contracts falls considerably at longer terms.

Figure A.6 plots our measure of expected returns for Bitcoin and Ether. Given the tremendous bull market in cryptocurrencies over the past several years, expected returns are almost always positive, although they dipped negative in late 2018 and briefly during the 2020 pandemic. The average expected return for Bitcoin using the measure is 5.0% from December 2017 to November 2022, ranging from -10.8% in December 2018 to 23.5% in February 2021. The ETH expected return has a shorter history because the future was introduced later, but from February 2021 to November 2022 it averaged 4.8% with a standard deviation of 7.3% compared to BTC's 3.9% average and 5.3% standard deviation over the same period.

We test the model's prediction that lending rates are increasing in y by regressing Tether's lending rate on FTX on our measure of expected returns using

$$\text{USDT Lending Rate}_t = \alpha + \beta \mathbb{E}[\text{Ret}^{BTC}] + \gamma X_t + \varepsilon_t$$

where X_t is a vector of controls. Table A.6 shows the regression results. A 1pp increase in $\mathbb{E}[\text{Ret}^{BTC}]$ increases the stablecoin lending rates by between 0.8 and 1.4pp, depending on the control variables. Across all specifications, there is a positive and significant relationship between lending rates and expected returns. Figure A.7 is a scatter plot between expected returns on Bitcoin and Tether lending rates showing a positive relationship.

One concern is that we confound expected speculative returns with the term structure of risk premium. We control for this problem by including an expected return for the SPX equity index using the same logic: we compare the present value of the first and second front-month for the SPX. Including this control in column (6) does not change the statistically strong relationship between expected returns and lending rates.

A.8 Appendix Figures

A.17




Figure A.1: Speculation and USDT Market Cap. Figure plots the monthly average of the daily perpetual future funding rate and daily change in Tether's market cap.




Figure A.2: Prime Money-Market Mutual Fund Holdings of Treasuries and ONRRPs. Figure the ratio of total prime money fund assets held in Treasuries or in Treasuries plus investments at the Federal Reserve. Data from the Office of Financial Research's U.S. Money Market Fund Monitor.




Figure A.3: Tether Redemptions and Issuance. Left panel plots the daily redemptions and issuance of Tether as a percent of its face value. Redemptions are defined as the change in the face value of the stablecoin's market capitalization on date t divided by the face value on date $t - 1$ for days with net redemptions, and zero otherwise. Right panel plots the analogous measure for days with net issuance, and zero otherwise.




Figure A.4: PYUSD and Kamino. Figure plots the market capitalization of PYUSD before and after the introduction of the Kamino lending platform in Solana.




Figure A.5: Stablecoin issuer's profit for different ℓ and r .




Figure A.6: Futures-Implied Expected Returns Figure plots the one-month/one-month expected return on Bitcoin and Ether estimated using the difference in present values for one-month futures prices relative to two-month futures prices. Present values are calculated using OIS interest rates, and futures prices are CME future prices.




Figure A.7: Stablecoin Lending Rates and Futures-Implied Expected Returns. Figure plots a binscatter of the one-month/one-month expected return on Bitcoin against USDT's margin lending rate on the FTX exchange.

A.9 Appendix Tables

Haircut (%)	Coin	Ticker	FTX	Binance	Bitfinex	Kraken
<i>Major Coins</i>	Bitcoin	BTC	5	5	0	0
	Ether	ETH	10	5	0	0
	Cardano	ADA	<i>n.a.</i>	10	70	10
	Ripple	XRP	10	15	50	<i>n.a.</i>
	Solana	SOL	15	10	30	10
	Dogecoin	DOGE	10	5	80	<i>n.a.</i>
	Litecoin	LTC	10	10	0	30
	Avalanche	AVAX	15	20	80	50
	Tron	TRX	15	50	70	50
<i>Stablecoins</i>	Tether	USDT	5	0	0	10
	USD Coin	USDC	0	0	0	10
	Binance USD	BUSD	0	0	<i>n.a.</i>	<i>n.a.</i>
	Dai	DAI	15	<i>n.a.</i>	25	10
<i>Average</i>	Major Coins		11	14	42	21
	Stablecoins		5	0	8	10

Table A.1: Haircuts. The Table gives haircuts across FTX, Binance, Bitfinex, and Kraken. FTX haircut is 1 minus the initial weight; Binance haircut is 1 minus the collateral rate. Average is an unweighted average of the haircuts in the corresponding rows above. Collateral haircuts updated as of November 2022, except Binance numbers are October 2022. A lower haircut implies that a larger share of the asset’s nominal price can be used to back a levered position. While there is heterogeneity across exchanges, stablecoins have lower haircuts. Note that exchange deposits are economically equivalent to a non-tradeable stablecoin issued by the exchange and have similarly low haircuts. Suppose a trader wants to use ten times leverage to buy \$100 of BTC. The margin requirement depends on the trader’s collateral. Using Binance haircuts, if the trader posts AVAX as collateral, they must provide $\$10/(1 - 20\%) = \12.5 of AVAX. If, however, the trader posts USDT as collateral, they need to post only $\$10/(1 - 0\%) = \10 of USDT. Posting a stablecoin as collateral requires 20% less equity capital from the trader.

	1 Largest	3 Largest	5 Largest	10 Largest	All
	(1)	(2)	(3)	(4)	(5)
FTX Tether Lending Rate _t	0.05*** (23.77)	0.03*** (2.70)	0.02** (2.04)	0.03* (1.91)	0.03*** (4.39)
<i>N</i>	273	673	1,074	1,989	5,668
<i>R</i> ²	0.31	0.14	0.04	0.03	0.09
TVL Weighted	No	No	No	No	Yes
Avg. TVL (\$ millions)	183	150	104	65	15

Table A.2: FTX Lending Rates and Defi Lending Rates. Table presents regression $R_{j,t}^{Defi} = \alpha + \beta_1 R_t^{USDT} + \varepsilon_{i,t}$ where R_t^{USDT} is Tether’s margin lending rate from the FTX exchange and $R_{j,t}^{Defi}$ is the lending rate at the Defi lending platform j . Defi lending rates from DefiLlama, spanning all protocols in the lending category that include Tether. Observations are daily, and we winsorize defi lending rates at the 5 and 95 percentile to reduce the influence of outliers. Protocols are calculated using their average 2022 total value lock (TVL) in US dollars. Column 5 includes all protocols in the sample and weights the regression by the protocol’s average 2022 TVL. “Avg. TVL” row provides the average total value lock of the protocols in the given sample. Constant omitted. t -statistics are reported in parentheses using robust standard errors, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	USDT	DAI
	ΔEFFR_t	ΔEFFR_t
$\Delta R_{i,t}$	-0.007 (0.87)	0.000 (0.99)
N	486	450

Table A.3: Correlation of FTX Lending Rates and Fed Funds Rate. Table presents the correlation of FTX lending rates for stablecoin i , $R_{i,t}$, with the effective federal funds rate where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	USDT			USDT and DAI		
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \ln(s_{i,t})$	-2.65^{**} (-2.57)	-3.65^{**} (-2.55)	-4.42^{***} (-3.00)	-1.03^* (-1.89)	-1.23^{**} (-2.07)	-1.33^{**} (-2.18)
Bitcoin Implied Volatility _t			-10.66 (-1.42)			-9.25 (-1.25)
$\Delta \ln(s_{i,t-1})$			1.01 (0.69)			-0.55 (-1.22)
$\ln(s_{i,t-1})$			-4761.72^{***} (-2.69)			-806.44^{**} (-2.28)
N	704	704	704	1,353	1,353	1,353
R^2	0.01	0.02	0.04	0.00	0.01	0.01
Month FE	No	Yes	Yes	No	Yes	Yes
Coin FE	n/a	n/a	n/a	No	Yes	Yes

Table A.4: Outside Option Return and Stablecoin Volume. Table presents regression $\Delta \rho_t = \alpha + \beta_1 \Delta \ln(s_{i,t}) + \gamma' X + a_i + b_t + \varepsilon_{i,t}$ where $\Delta \rho_t$ is the change in the outside option ρ_t , $\Delta \ln(s_{i,t})$ is the change in the log change in the face value of stablecoin i (either USDT or USDT and DAI), X is a set of controls, a_i is a stablecoin fixed effect, and b_t is a time fixed effect. We define the outside option $\rho_t = y_t - (1 - m)R_t$ where y_t is proxied by the future funding rate, R_t is the FTX lending rate for the given stablecoin, and we assume $m = 0.2$. t -statistics are reported in parentheses using robust standard errors and clustered by week, where $^* p < 0.10$, $^{**} p < 0.05$, $^{***} p < 0.01$.

	BTC/USDT Binance	ETH/USDT Binance	BTC/BUSD Binance	DOGE/BUSD Binance	BTC/USD FTX	ETH/USD FTX	$\mathbb{E}[R^{BTC}]$ CME	$\mathbb{E}[R^{ETH}]$ CME
BTC/USDT, Binance	1.00							
ETH/USDT, Binance	0.89***	1.00						
BTC/BUSD, Binance	0.81***	0.70***	1.00					
DOGE/BUSD, Binance	0.64***	0.59***	0.65***	1.00				
BTC/USD, FTX	0.83***	0.79***	0.76***	0.59***	1.00			
ETH/USD, FTX	0.75***	0.87***	0.65***	0.51***	0.80***	1.00		
$\mathbb{E}[R^{BTC}]$	0.65***	0.62***	0.55***	0.50***	0.66***	0.61***	1.00	
$\mathbb{E}[R^{ETH}]$	0.65***	0.62***	0.57***	0.55***	0.64***	0.56***	0.83***	1.00

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.5: Correlation of Expected Return Proxies. Table presents the pairwise correlations of several perpetual futures funding rates and the expected return inferred using CME crypto futures. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Bitcoin		Ether		Both	
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbb{E}[Ret^{BTC}]$	1.05*** (7.64)	0.51*** (3.16)			0.55*** (2.76)	0.64*** (2.87)
Ret^{BTC}		0.23*** (2.86)				0.37*** (3.01)
$\mathbb{E}[Ret^{ETH}]$			0.78*** (8.10)	0.15 (0.94)	0.45*** (4.06)	-0.14 (-0.73)
Ret^{ETH}				0.09 (1.35)		-0.15 (-1.52)
$\mathbb{E}[Ret^{SP}]$						0.05 (0.60)
N	924	924	868	868	868	868
R^2	0.35	0.51	0.35	0.50	0.38	0.52
Month FE	No	Yes	No	Yes	No	Yes
Coin FE	No	Yes	No	Yes	No	Yes

Table A.6: Stablecoin Interest Rates and Expected Returns. Table presents regression $R_{t,i} = \alpha + \beta_1 \mathbb{E}_t[Ret^i] + \beta_2 Ret^i + a_i + b_t + \varepsilon_{i,t}$ where $R_{t,i}$ is the lending rate for stablecoin i , either USDT or DAI, $\mathbb{E}_t[Ret^j]$ is the one-month/one-month expected returns for coin j —either Bitcoin and Ether— Ret^j is the contemporaneous price returns on Bitcoin and Ether, a_i is a stablecoin fixed effect, and b_t is a time fixed effect. Observations are daily; the Bitcoin-only sample in columns (1) and (2) runs from December 2020 to November 5, 2022, and the remaining columns with Ether run from February 2021 to November 5, 2022. t -statistics are reported in parentheses using robust standard errors and clustered by week, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Lending Rate R_t					
	USDT			DAI		
	1 Day	1 Week	4 Weeks	1 Day	1 Week	4 Weeks
	(1)	(2)	(3)	(4)	(5)	(6)
Futures $\widehat{\text{Funding Rate}}_t$	0.182 (1.540)	0.207 (1.339)	0.315 (0.654)	0.178 (1.479)	0.131** (2.019)	−0.153 (−0.710)
R_{t-1}	0.606*** (4.761)	0.519*** (2.992)	0.437 (1.005)	0.500*** (6.712)	0.556*** (7.221)	0.674*** (4.719)
Bitcoin Implied Volatility $_t$	−0.033 (−0.733)	−0.039 (−0.725)	0.009 (0.181)	−0.060 (−1.193)	−0.044 (−1.506)	0.011 (0.468)
$\Delta \ln(s_{i,t})$	−0.007* (−1.773)	−0.004 (−1.332)	−0.005 (−0.569)	−0.003 (−0.922)	0.000 (0.031)	0.006 (1.256)
N	258	258	258	258	258	258
F -stat	1.88	1.25	0.33	2.45	1.47	0.96
Time FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A.7: Instrumental Variables Placebo Regression of Futures Funding Premia and Lending Rates. Instrumental variables regression using the mean household rating of MLB games on a given day in the future as an instrument to predict the perpetual futures funding premium. Table presents several placebo tests using viewership data from the future as the instrumental variable: either 1 day, 1 week, or 4 weeks in the future. Time FE indicates day of week, month of year, and year fixed effects. Kleibergen-Paap rk Wald F statistics reported. t -statistics are reported in parentheses using robust standard errors and clustered by week where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Instrument: Household Rating \times Championship Leverage Index				
	Lending Rate R_t			
	USDT		DAI	
	(1)	(2)	(3)	(4)
Futures $\widehat{\text{Funding Rate}}_t$	0.342*** (10.872)	0.217*** (3.280)	0.252*** (4.019)	0.162*** (3.660)
Bitcoin Implied Volatility $_t$	0.032 (0.366)	0.030 (0.424)	-0.028 (-0.597)	-0.027 (-0.860)
$\Delta \ln(s_{i,t})$	-0.008 (-1.246)	-0.005 (-1.149)	-0.009*** (-2.673)	-0.006** (-2.331)
R_{t-1}		0.406* (1.925)		0.512*** (6.302)
N	245	245	245	245
F -stat	24.84	7.11	22.48	21.01
Time FE	Yes	Yes	Yes	Yes

Table A.8: Instrumental Variables Regression of Futures Funding Premia and Lending Rates with Championship Leverage Index. Instrumental variables regression using the mean household rating of MLB games on a given day as an instrument to predict the perpetual futures funding premium. Instrument is the product of the Household Rating and the Championship Leverage Index. The Championship Leverage Index (cLI) is a common sabermetrics estimate of the importance of a game to a team's chances of winning the World Series. cLI data provided by Baseball Reference for the regular season, and we manually calculate it for playoff games. The cLI is standardized so that its value is 1 for the average game. Time FE indicates day of week, month of year, and year fixed effects. Kleibergen-Paap rk Wald F statistics reported. t -statistics are reported in parentheses using robust standard errors and clustered by week where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	BTC		ETH		DOGE	
	(1)	(2)	(3)	(4)	(5)	(6)
Rating	0.14 (1.18)	0.15 (1.27)	0.23 (1.29)	0.23 (1.29)	0.36* (1.93)	0.40** (2.25)
Constant	−0.14 (−0.56)	−0.13 (−0.43)	−0.11 (−0.34)	−0.01 (−0.01)	−0.27 (−0.81)	0.26 (0.34)
N	258	258	258	258	258	258
R^2	0.00	0.04	0.00	0.04	0.01	0.04
Day-of-Week FE	No	Yes	No	Yes	No	Yes

Table A.9: Speculative Returns and Household Rating. Table presents regression $Ret_{i,t} = \alpha + \beta \text{Household Rating}_t + b_t + \varepsilon_{i,t}$ where $Ret_{i,t}$ is the price return of coin i —where i is Bitcoin, Ether, or Dogecoin—Household Rating $_t$ is the household rating of nationally televised MLB games on date t , and b_t are day of week fixed effects. Observations are daily. t -statistics are reported in parentheses using robust standard errors and clustered by week, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.