BIS

BIS Working Papers
No 1355

Making stablecoins stable(r): can regulation help?

by Tirupam Goel, Ulf Lewrick and Isha Agarwal

Monetary and Economic Department

June 2026

JEL classification: G2, G28, C6

Keywords: capital regulation, liquidity regulation, stablecoins, crypto, money market funds, financial stability, buffer usability

This publication is available on the BIS website (www.bis.org).

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Making stablecoins stable(r): can regulation help?*

Tirupam Goel, Ulf Lewrick, and Isha Agarwal†

June 2026

Abstract

We model a stablecoin issuer that optimises capital, cash and bond holdings under persistent stablecoin flows. Absent regulation, the issuer holds little capital and favours interest-bearing but less-liquid bonds over cash. This exposes coin-holders to default risks and poses systemic spillovers via price impact of bond fire-sales. How can regulation mitigate these risks? We consider capital and liquidity thresholds as usable buffers. They can be breached in stress but discipline issuers by triggering additional redemptions, thus endogenising stablecoin flows. The thresholds work through asymmetric channels. While the liquidity threshold only raises cash holdings, the capital threshold increases both capital and cash. Both thresholds mitigate default and spillover risks, suggesting they are substitutes. However, they are complements for regulators targeting *both* risks. Using stablecoin flows and US Treasury market depth, we calibrate a two-way mapping that enables regulators to recover capital-liquidity threshold combinations implied by chosen risk targets (and vice-versa).

JEL Codes: G2, G28, C6

Keywords: Capital regulation; liquidity regulation; stablecoins; crypto; money market funds; financial stability; buffer usability.

*The authors thank Juliane Begenau, Ulrich Bindseil, Christopher LC Boedihardjo, Paolo Cavallino, Francesca Carapella (discussant), Jon Frost, Vera Lubbersen (discussant), Marcus Opp (editor of the special issue of the Review of Finance on *The Future of Payments*), Linda Schilling (discussant), Yizhou Xiao (discussant) and seminar participants at the 3rd International Conference on Payments and Securities (Bank of Canada), Asian Finance Association conference 2025, EEA Congress 2025, BAFFI-CEPR-ECB Future of Payments conference 2025, a Reserve Bank of India seminar, IGIDR Macroeconomics conference 2025 (where the paper won the third prize), ISI Delhi Conference 2025, IESE Workshop on Digital Money 2026, and a research meeting at the Bank for International Settlements (BIS) for helpful comments.

†Tirupam Goel (tirupam.goel@bis.org; corresponding author) and Ulf Lewrick (ulf.lewrick@bis.org) are at the BIS. Isha Agarwal (isha.agarwal@sauder.ubc.ca) is at the University of British Columbia, Sauder School of Business. The authors are grateful to Jimmy Shek, Rudraksh Kansal and Ilaria Mattei for their help in procuring the data for this project. The views expressed in this paper are those of the authors and not necessarily those of the Bank for International Settlements or its member central banks.

1 Introduction

Stablecoins have grown rapidly over the past five years, with the two largest issuers now rivalling major US Treasury money market funds in their holdings of short-dated government securities. Their deepening footprint in traditional financial markets has placed a spotlight on one of their main vulnerabilities: liquidity mismatch in reserve management. Fiat-backed stablecoin issuers hold demandable liabilities against interest-bearing but less liquid assets (e.g. government bonds) while maintaining small cash buffers. Large redemptions can therefore force bond sales at discounts, erode capital and transmit stress to money markets through the price impact channel. This combination of fragility at the issuer level and spillovers at the financial market level motivates the central question of this paper: how can prudential regulation protect coin-holders while limiting market spillovers?

Major jurisdictions are currently exploring a regulatory architecture that combines liquidity and capital thresholds for stablecoin issuers.[1] The design and calibration of these regulatory thresholds, however, is evolving differently across jurisdictions, reflecting a lack of consensus among policymakers. Two questions are unresolved. First, how do liquidity and capital thresholds *interact* when applied to a non-bank issuer whose balance sheet adjusts mechanically with the demand for its liabilities? Second, what numerical thresholds are needed to meet specific micro- and macro-prudential targets? This paper provides a framework to answer both. In a nutshell, we model a stablecoin issuer and study its response to the regulatory thresholds. The micro- and macro-prudential risks posed depend on the issuer's behaviour and how the thresholds interact. Ultimately, we calibrate the model and derive a two-way mapping between regulatory thresholds and policy targets that can provide concrete calibration guidance to policymakers.

Our dynamic model features a fiat-backed stablecoin issuer that faces serially correlated flow shocks. The issuer responds by choosing the privately optimal cash-bond mix of its reserve assets and raises capital to absorb losses. Bonds earn a spread but are costly to unwind. Bond sales depress prices and mark down remaining holdings, with the price impact increasing in sale size (Shleifer and Vishny [2011]). Cash, by contrast, is risk-free but earns no

[1] Examples include the European Union (Markets in Crypto-Assets Regulation ie MiCAR), the United States (Guiding and Establishing National Innovation for U.S. Stablecoins ie GENIUS Act), Hong Kong Special Administrative Region, and Singapore.

interest. The unregulated model admits a closed-form solution under mild assumptions. The optimal cash policy as a function of flows mimics a "shark-fin" that sequences cash use across four flow phases: (i) for large subscriptions, which render future outflows unlikely, the issuer holds no cash and only invests in bonds; (ii) for small to moderate subscriptions, it maintains a precautionary cash buffer; (iii) for small to moderate redemptions, it uses cash first to avoid costly bond liquidations; and (iv) for large redemptions, it must sell bonds, with extreme outflows leading to large bond sale losses, pushing capital below zero and triggering default.[2] Several extensions and comparative statics, such as introducing asymmetric persistence in stablecoin flows, regime-dependent flows (run risk), and variations in the risk-return profiles of bonds and cash, help strengthen intuition and affirm robustness of the issuer's balance-sheet strategy.

Regulation seeks to address two wedges that a stablecoin issuer does not internalise in its privately optimal choices. First, limited liability induces risk-shifting. The issuer holds little capital and favours interest-bearing but less liquid bonds over cash. It thus underweights tail losses borne by coin holders (e.g. Li and Mayer [2026]). Second, forced liquidations transmit stress to money markets by depressing the prices of bonds and close substitutes (e.g. Ahmed and Aldasoro [2025]; d'Avernas and Vandeweyer [2024]). These externalities motivate a dual regulatory objective comprising a microprudential target—the issuer's probability of default (PD)—and a macroprudential target—the expected price impact (EPI) from bond sales. This dual objective is in line with international recommendations on stablecoin regulatory frameworks (e.g. Financial Stability Board [2023]). To address these risks, we consider two regulatory tools—a liquidity-ratio (LR) threshold and a capital-ratio (CR) threshold—in line with ongoing policy discussions.

A novel feature of our framework is that regulatory thresholds operate as *usable buffers*, which makes stablecoin flows endogenous. Imposing regulatory thresholds as hard constraints is common practice in the literature and is analytically convenient but can be self-defeating: for example, cash that cannot be drawn down when redemptions arrive is of little use.[3] We model regulatory thresholds as reference points that can be breached in stress. However, this invokes coin-holder discipline by triggering additional redemptions, making stablecoin flows

[2]The cash-first-bonds-next result mirrors evidence on liquidity management in open-ended funds (Jiang et al. [2021]; Ma et al. [2022]; Chernenko and Sunderam [2016]).

[3]In fact, in Annex I we show that when the LR threshold is imposed as a hard constraint, the issuer's default probability increases.

endogenously linked to the issuer's balance-sheet strength. This design creates incentives for the issuer to build buffers in normal times while preserving their usability during periods of stress, thereby balancing regulatory stringency and effectiveness.

Our model delivers three main results. First, the LR and CR thresholds shape issuer behaviour through *asymmetric* channels. The LR threshold induces the issuer to hold more cash but has no direct effect on capital. By contrast, the CR threshold not only raises capital but also cash because additional cash reduces the bond sales that erode capital.

The second result concerns the ultimate impact of thresholds on policy targets. More cash reduces both PD and EPI, while more capital reduces PD but has no direct effect on EPI. Combined with the first result, we obtain a key insight: both thresholds help achieve both policy objectives. This suggests that the thresholds are substitutes. However, the thresholds are complements — i.e. both are necessary — when the regulator targets PD and EPI jointly.

Our third result is quantitative. The calibrated model provides a *two-way mapping* between the regulatory thresholds and the policy targets. For a given threshold pair, the mapping provides the resulting PD and EPI. Conversely, for a joint PD and EPI target, the mapping points to suitable threshold pairs that can achieve the target. For instance, a plausible low-end calibration of the LR and CR thresholds yields a weekly issuer PD of around 0.7 bps and an EPI of 2.7 bps during stressed conditions, notably lower than the PD of more than 15 bps and an EPI of around 4 bps in the unregulated case. Higher LR and CR thresholds further reduce PD and EPI.

Our calibration relies on observed stablecoin flow dynamics and US money market depth to ensure that the mapping we derive is relevant for the design of stablecoin policies in practice. The weekly persistence in stablecoin flows is estimated using a panel of major stablecoins. We find economically and statistically significant AR(1) dynamics. The price impact of forced bond sales is estimated using redemption-induced sales by US Treasury money market funds whose portfolios closely resemble those of leading stablecoin issuers. The remaining model parameters are disciplined via a method of moments to match observed liquidity and capital ratios of a major stablecoin issuer.

Finally, we examine the normative implications of regulation. We compare the welfare benefits from lower PD and EPI against the constraints imposed on stablecoin issuers and the attendant welfare costs for users. Any such assessment is necessarily stylised as it depends on how much value society places on stablecoins. Nonetheless, the welfare analysis captures the regulatory trade-off and helps identify the optimal thresholds under alternative assumptions about the social value of stablecoins.

While some of our findings echo insights from the banking literature (Cecchetti and Kashyap [2018]; Carletti et al. [2020]; Kara and Ozsoy [2020]), two key differences justify a model tailored to stablecoins. The first concerns balance-sheet rigidities. Banks can elastically adjust their balance-sheet size by lending and simultaneously creating deposits, and by borrowing to compensate for deposit outflows. By contrast, stablecoin issuers' balance-sheet size adjusts mechanically with demand for their coins, constraining their degrees of freedom when responding to shocks. These rigidities place liquidity management at the centre of the issuer's optimisation problem, as our model demonstrates. Second, risk exposures differ substantially. For banks, exposures to loans and other risky assets place the emphasis on credit risks and capital requirements. Bank liabilities, by comparison, tend to be more stable because banks rely on a mix of funding sources, benefit from deposit insurance, and have access to central bank backstops. For stablecoins, the opposite holds. Their main vulnerability lies instead on the liability side. Their liabilities are demand-able, uninsured, and not supported by a public safety net, while their assets (predominantly short-dated government bonds and cash) carry little risk.[4] Vulnerabilities therefore stem from more volatile redemptions, placing the regulatory focus on liquidity requirements. At the same time, because large redemptions can force bond sales and liquidation costs that erode capital, liquidity regulation should be complemented by capital requirements.[5]

Related literature The academic literature on stablecoins has expanded rapidly, reflecting concerns about their stability and systemic spillovers. This paper connects to three strands of literature.

The first focuses on the link between stablecoin flow dynamics—especially episodes in

[4]In other words, stablecoins are information-sensitive assets while bank deposits tend to be information-insensitive assets (during normal times), as discussed in Gorton and Pennacchi [1990] and Dang et al. [2020].

[5]Stablecoins also share some similarities with money market funds (Jiang et al. [2021]; Anadu et al. [2023]). However, fund investors earn a share of the fund's returns, aligning their incentives more closely with the fund's risk taking. By contrast, stablecoin holders are creditors who earn no interest and have no upside from greater risk-taking by the issuer. These differences affect not only the risk-taking incentives of fund managers and stablecoin issuers, but may also make stablecoin flows more sensitive to the strength of the issuer's balance sheet.

which users rush to redeem their holdings—and par convertibility. A foundation of this literature is the seminal work by Morris and Shin [1998], Rochet and Vives [2004] and He and Xiong [2012] on global games and endogenous runs in the case of banks and similar financial institutions. For instance, Gorton and Zhang [2023] draw parallels to nineteenth-century private banknotes, highlighting the susceptibility of stablecoins to runs. Liu et al. [2023] analyse the Terra-Luna collapse, showing how opacity and weak disclosure accelerated redemptions, while Routledge and Zetlin-Jones [2021] emphasise vulnerabilities arising from insufficient reserves and mechanical redemption protocols. Other studies, such as Gorton et al. [2025]; d'Avernas et al. [2026]; Bertsch [2023], examine how flow dynamics impact stablecoin prices – especially the parity. Still others, including Ma et al. [2025], Ahmed et al. [2024] and Lyons and Viswanath-Natraj [2023], explore how arbitrage efficiency, disclosures, and market microstructure may affect parity and exacerbate run risk by amplifying first-mover advantages. This work has greatly advanced our understanding of why stablecoin flows are volatile, but it typically abstracts from the liquidity-management problem that the issuer solves in response to volatile flows. By contrast, we model the issuer's problem explicitly, while adopting a parsimonious setup for endogenous stablecoin flows. In particular, we characterise the issuer's privately optimal sequencing of cash use and bond purchase/sales following redemption and subscription shocks, including when these shocks depend on the strength of the issuer's balance sheet relative to regulatory thresholds.

A second strand examines regulatory tools to mitigate stablecoin risks. Li and Mayer [2026] study capital requirements and reserve-asset risk limits and show that capital requirements reduce the issuer's reliance on procyclical seigniorage. Liao et al. [2024] proposes a capital adequacy framework to derive the amount of capital needed to absorb losses stemming operational risks. Carapella [2025] develops a limited-commitment framework with banks and stablecoin issuers, proposing a loss mutualisation mechanism with costly membership to discipline risk-taking and improve welfare in the absence of regulation. By comparison, our focus is on the interplay between capital and liquidity thresholds for stablecoin issuers. We link these thresholds with micro- and macro-prudential objectives to provide a tractable framework for policy calibration.

The third strand concerns the design and use of regulatory buffers. Empirical work documents how banks (e.g. Couaillier et al. [2025], Berrospide et al. [2024]) and money market funds (e.g. Cipriani and La Spada [2020]) build and draw buffers relative to regulatory thresholds. We complement by providing a model where such buffers arise endogenously. Regulatory thresholds are cast not as hard minima but as reference points that can be temporarily breached. This mechanism reflects evolving regulatory practice in which liquidity is meant to be usable in stress (see e.g. Bank of England [2025]) while recognising that disclosure combined with investor discipline shape flow dynamics.

Overall, our paper provides fresh insights into how capital and liquidity thresholds shape stablecoin issuer behaviour and financial stability outcomes, and it delivers a practical mapping between regulatory tools and policy objectives. The rest of the paper is organised as follows. Section 2 introduces the model and presents its closed-form solution under mild assumptions. Section 3 outlines the calibration strategy. The quantitative analysis is presented in Section 4, with extensions of the model and a concise welfare analysis provided in Section 5. Section 6 concludes.

2 Model

We develop a three-period model of a stablecoin issuer. The dates are denoted $t = 0, 1, 2$. On date 0, the issuer begins with a given volume of stablecoins s_0. It must decide how much capital k_0 to raise and how to back the stablecoins with two types of assets: cash c_0 and bonds b_0. In doing so, the issuer faces the following considerations.

First is the balance-sheet constraint $c_0 + b_0 = s_0 + k_0$. Second is a trade-off with regard to its asset allocation. Cash bears no interest and can be used to meet redemptions without any liquidation costs. By contrast, bonds provide interest μ, can be bought at par, but are subject to a liquidation cost when sold.[6] We allow this cost to increase with the amount sold, reflecting a bigger fire-sale discount when selling larger amounts. Specifically, if quantity x is sold, bond prices decline by $g\,(x)$, where $g'(x) \geq 0$. Therefore, the unit price received per bond sold is $1 - g\,(x)$ and the value of bonds that remain on the issuer's balance sheet is marked down to reflect the new market price. Finally, while capital helps absorb fire-sale losses and reduces the risk of insolvency, raising capital incurs a quadratic cost $\lambda\, k_0^2$.

On date 1, the date-0 balance-sheet configuration (c_0, b_0, s_0, k_0) is taken as given. A stablecoin flow shock r_1 is realised so that the date-1 stablecoins outstanding is given as

[6] The stablecoin issuer's trade-off is reminiscent of the dynamic cash management decision of open-ended mutual funds (e.g. Zeng [2017], Morris et al. [2017]).

$s_1 = s_0 (1 + r_1)$. In this period, the issuer must decide the cash and bond allocations c_1 and b_1 subject to the balance-sheet identity $c_1 + b_1 = s_1 + k_1$ and the capital evolution identity:

$$k_1 \quad = \quad \underbrace{k_0}_{\text{Previous capital}} \quad + \quad \underbrace{b_0 \mu}_{\text{Revenue from bonds}} \quad - \underbrace{b_0 \, g\Big(\max\{0, b_0 - b_1\}\Big)}_{\text{Bond sale cost \& valuation loss}} \tag{1}$$

The equation states that the source of capital growth on date 1 is retained earnings.[7] Also note that the bond-sale cost term accounts for mark-to-market losses. That is, sales depress the market price of bonds and therefore losses are incurred on bonds that are sold as well as those that the issuer continues to hold.

On date 2, the stablecoin flow shock r_2 implies $s_2 = s_1 (1 + r_2)$ and the date-1 balance-sheet configuration is taken as given. There are no decisions taken on date 2—only the payoff is realised and capital is given by:

$$k_2 \; = \; k_1 + \mu \, b_1 - b_1 \, g\Big(\max\{0, -r_2 s_1 - c_1\}\Big). \tag{2}$$

Since there is no future period, cash has no continuation value. Cash is therefore used first to meet redemptions, minimising bond sales and associated costs to the extent possible.

The objective of the issuer is to maximise the present discounted value of its capital net of capital raising costs (as also in, for example, Gertler and Kiyotaki [2010]):

$$V_0 = \max_{k_0, c_0, b_0, c_1, b_1} \Big\{ - \lambda k_0^2 + \beta \, \mathbb{E}_{r_1} \Big[k_1 \Big] + \beta^2 \, \mathbb{E}_{r_2} \Big[k_2 \Big] \Big\}. \tag{3}$$

The issuer defaults whenever capital falls below zero. This may arise at date 1 or date 2 if the bond sales necessary to satisfy large redemptions impose liquidation costs and valuation losses that exceed available capital. We assume limited liability: in default, coin-holders have no recourse beyond the issuer's assets, and owners are not liable for residual claims.

We assume that stablecoin flow shocks are serially correlated. In the baseline version of our model, this correlation is taken as exogenous. The distribution of r_t is conditional on the realisation of r_{t-1} (see equation (4)). Our specification captures the idea that what ultimately matters for the stablecoin issuer's decisions is the distribution of flow shocks next

[7] Following Myers and Majluf [1984], we do not allow for raising capital on date 1. The intuition is that if the issuer is in stress and needs more capital to avoid insolvency, shareholders are unlikely to commit more capital to the issuer, making any capital raising prohibitively costly.

period. We motivate this assumption and estimate the serial correlation in flow shocks using stablecoin data in Section 3.1.

In the workhorse version of our model, we allow stablecoin flow shocks to be endogenous. It is natural to expect that stablecoin holders are more likely to redeem from an issuer with a weaker balance sheet. Two key indicators of balance-sheet resilience are the liquidity ratio, c/s, and the capital ratio, k/s. A higher liquidity ratio implies that the issuer holds more cash, which helps meet redemptions more easily. Indeed, the benefit of holding cash is that, unlike bonds, it can be used to meet redemptions without any liquidation costs.

In line with the above intuition, we consider two regulatory thresholds that serve as *anchors* for coin-holder behaviour. The first one is the liquidity ratio (LR) threshold, $\theta \geq 0$, which is measured against the issuer's liquidity ratio (c/s). The second one is the capital ratio (CR) threshold, $\chi \geq 0$, which is benchmarked against the capital ratio (k/s). The CR threshold is equivalent to expecting that each stablecoin should be backed by $(1 + \chi)$ units of reserve assets, i.e. enforcing overcollateralisation. The thresholds are inspired by evolving policy frameworks.[8]

We treat the thresholds as usable buffers rather than hard minima. The issuer may operate above, at, or below the thresholds. Indeed, there is no point requiring a stablecoin issuer to hold 30 percent cash *at all times* if cash cannot be used as an insurance during times of stress.[9] When the issuer dips below the threshold, coin-holders become concerned and are more likely to redeem. Effectively, the issuer faces the wrath of stablecoin holders—so-called coin-holder discipline. We assume that the deeper the breach, the higher the likelihood of

[8]For instance, regulation in the European Union (MiCAR) imposes specific liquidity ratio thresholds under which issuers must hold at least 30% of their reserves as deposits with commercial banks. Implementation of US regulation (GENIUS Act) may establish thresholds in due course. In terms of total reserve asset requirements, a minimum 100% backing is the norm. Regulation in the European Union also establishes capital requirements alongside additional overcollateralisation requirements based on a historical look-back approach that takes into account the size, complexity and nature of the reserve assets (European Banking Authority [2024]).

[9]There is precedence for buffer usability in regulatory practice. For example, money market funds may draw down liquidity buffers and fall below cash thresholds when facing large redemptions; such use typically triggers business constraints and/or investor discipline (e.g. Cipriani and La Spada [2020]). More generally, breaches of regulatory thresholds do not lead to immediate closure: supervisors impose remedial measures (e.g. payout restrictions) or require the execution of pre-arranged recovery plans, as foreseen for stablecoins under MiCAR in the European Union. The notion of coin-holder discipline is likewise informed by Pillar 3 of Basel III, under which enhanced disclosures enable stakeholders to assess institutional soundness and to discipline excessive risk-taking by reallocating funds elsewhere (i.e. stakeholder discipline).

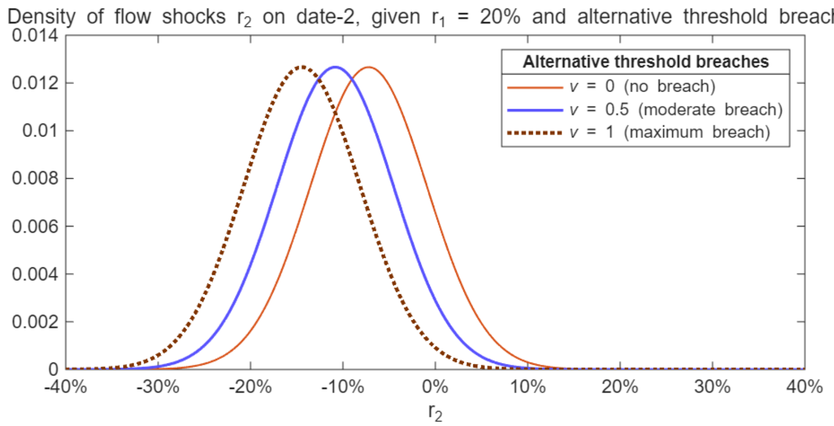

Figure 1: Illustration of how the distribution of flow shocks on date 2, r_2 depends on the severity of threshold breaches, conditional on a 20% flow shock on date 1, i.e. $r_1 = -0.2$.

redemptions (also see e.g. Ahmed et al. [2024]). In contrast to much of the prudential regulation literature which imposes always-binding requirements, our approach places buffer usability at the centre of the analysis.

Turning to the formal specification, we define the degree of violation for each type of threshold as $\nu^{LR} = \min\Big\{1, \max\{0, 1-(c/s)/\theta\}\Big\}$ and $\nu^{CR} = \min\Big\{1, \max\{0, 1-(k/s)/\chi\}\Big\}$. By construction, $\nu^{\text{LR}}, \nu^{\text{CR}} \in [0,1]$. Thus, $\nu^{\text{LR}} = 0$ when $c/s \geq \theta$ (no breach) and $\nu^{\text{LR}} = 1$ when $c = 0$ (the maximal breach, for $\theta > 0$). The capital measure is analogous: $\nu^{\text{CR}} = 0$ when $k/s \geq \chi$ and $\nu^{\text{CR}} = 1$ when $k = 0$ (for $\chi > 0$). These normalised measures bound the severity of any breach and ensure comparability across thresholds. Moreover, we treat the two thresholds symmetrically in how breaches affect coin-holder behaviour, thereby avoiding artefactual asymmetries that arise purely from specification. We aggregate violations via the maximum, capturing the most severe breach at a point in time: $\nu = \max\{\nu^{LR}, \nu^{CR}\}$. Formally, the distribution of flow shocks as a function of the breach is given as follows:

$$f_t(r_t|r_{t-1}, \nu) = \begin{cases} \mathcal{N}\Big(\gamma r_{t-1}, \sigma^2\Big) & \text{if } \nu = 0, \\ \mathcal{N}\Big(\gamma r_{t-1} - \nu\delta, \sigma^2\Big) & \text{if } \nu > 0. \end{cases} \tag{4}$$

Here ν determines the degree to which the distribution of shocks shifts to the left, thus resulting in more redemptions. The factor $\delta = \gamma(r_{t-1} - r_{min})$ simply ensures that the mean of the distribution stays within the range $[\gamma r_{min}, \gamma r_{t-1}]$.[10] In particular, for maximum violation of regulation i.e. $\nu = 1$ the mean of the distribution is γr_{min}. See Figure 1 for an illustration of how breaches of the regulatory threshold on a given date affect the distribution of flows on the next date.

Before proceeding, a few comments on how our approach to endogenous flows relates to the literature are in order. Both coin-holders in our paper and investors/depositors in papers on endogenous runs such as Rochet and Vives [2004] and He and Xiong [2012] use threshold strategies. The key difference is that, in our setting, the threshold is set by the regulator, whereas those papers use belief-driven thresholds—i.e. thresholds shaped by investors' beliefs about balance-sheet health of the financial institution. We focus on regulatory thresholds on three grounds. First is policy relevance: emerging frameworks such as MiCAR and the GENIUS Act contemplate explicit thresholds intended to serve as tangible reference points for stablecoin issuers and coin-holders; our model is designed to capture precisely this mechanism. Second, disclosure and lower information asymmetry: policy frameworks mandate frequent disclosures and permit only low-risk balance-sheet assets (e.g. government bonds and cash), reducing information asymmetry that is central to a model with belief-driven thresholds. Third is calibration: regulatory thresholds deliver a clean, two-way mapping from instruments (LR/CR) to policy targets (PD/EPI), which we expect to be useful for policy calibration.

Formal problem statement The issuer's problem can be stated in terms of backward induction. The solution to the date-2 problem is straightforward (see Section 2.1). As such we focus on specifying the date-1 problem statement. The date-0 problem statement follows from equation (3) subject to the analogous balance-sheet constraint and the corresponding conditional flow distribution $r_1 \sim f_1(r_1|r_0, \nu_0)$.[11]

Let V_1 be the value of the issuer on date 1, as a function of the shock r_1, and the state vector (b_0, k_0).[12] The date-1 problem statement then reads:

[10] Here r_{min} is the minimum value that r can take on either date. Conceptually, it can take a value of -1 corresponding to a 100% redemption. See Annex B for more details.

[11] We note that without loss of generality the starting value of stablecoins at date 0, s_0, already subsumes any prior flow shocks so that $r_0 = 0$ by construction.

[12] While the full balance-sheet configuration on date 0 is (c_0, k_0, b_0, s_0), (b_0, k_0) is a sufficient state vector given that s_0 is fixed and c_0 follows from the balance-sheet identity.

$$V_1(r_1, b_0, k_0) = \max_{b_1} \left\{ k_1 + \beta\, \mathbb{E}_{r_2|r_1}\Big[k_2\Big] \right\} \quad s.t. \quad s_1 = s_0(1+r_1), \quad c_1 + b_1 = s_1 + k_1$$

$$k_1 = k_0 + b_0\mu - \; b_0\, g\Big(\max\{0,\, b_0 - b_1\}\Big), \quad k_2 = k_1 + b_1\mu - b_1 g\left(\max\{0,\, -r_2 s_1 - c_1\}\right)$$

$$r_2 \sim f_2(r_2|r_1, \nu_1), \quad \nu_1 = \max\left\{\nu_1^{LR}, \nu_1^{CR}\right\}$$

$$\text{with} \quad \nu_1^{LR} = \min\left\{1, \max\left\{0,\, 1 - (c_1/s_1)/\theta\right\}\right\}, \quad \nu_1^{CR} = \min\left\{1,\, \max\left\{0, 1 - (k_1/s_1)/\chi\right\}\right\}.$$

Discussion of the issuer's trade-offs The issuer's liquidity-management problem entails a trade-off between near-term viability and longer-term profitability. When confronted with a redemption shock at date 1, the issuer must choose whether to meet outflows by drawing down cash or by liquidating bonds. Using cash avoids fire-sale discounts and preserves the stock of interest-bearing bonds. However, it depletes the liquidity buffer that insures against future redemptions. Coin-holder discipline complicates the trade-off: using cash can push the issuer closer to, or below, the liquidity threshold (LR), thereby heightening coin-holder discipline and the likelihood of future redemptions. By contrast, selling bonds immediately helps meet redemptions without reducing the cash buffer, but incurs liquidation costs that depress market prices (and hence mark-to-market valuations), erode capital and may precipitate default if losses exceed available capital.

These trade-offs are shaped by the balance-sheet configuration with which the issuer enters date 1. Accordingly, date-0 choices must balance risk reduction (via higher cash and capital buffers) against return maximisation (via larger holdings of interest-bearing bonds), while anticipating the feedback effects of coin-holder discipline from breaches of regulatory reference points (LR and CR). This intertemporal calculus can make it privately optimal to maintain precautionary buffers above the thresholds, thereby lowering the frequency and severity of buffer breaches under stress and improving resilience to outflows.

2.1 Closed-form solution

To build intuition on the issuer's liquidity management trade-off, we consider a simplified version of the workhorse model that admits a closed-form solution. We make two simplifying assumptions: first, bond sales incur a constant liquidation cost, $\kappa \in (0, 1)$, for every unit of bonds sold, with $\kappa > \mu$; and second, flows are exogenous. The issuer's objective remains as in equation (3), and we solve the model by backward induction.

Date-2 problem At date 2, there are no further decisions besides meeting redemptions and liquidating bonds if necessary. If $r_2 \geq 0$ (net subscriptions), no rebalancing of cash and bonds is required. If $r_2 < 0$ (net redemptions), there is a payout $-s_1 r_2 > 0$ that must be met. If $c_1 < -s_1 r_2$, the issuer sells the minimum amount of bonds required to meet the shortfall. Because the liquidation cost κx reduces net worth and there is no continuation beyond date 2, cash retention has no value. Hence, the issuer always uses cash before selling bonds at date 2. In analogy to the general case in (2), we have $V_2(r_2, b_1, k_1) = k_2 = k_1 + \mu b_1 - \kappa \max\{0, -s_1 r_2 - c_1\}$.

Date-1 problem At date 1, the issuer observes r_1 and chooses b_1 (and implicitly c_1) to solve $V_1(r_1, b_0, k_0) = \max \left\{ k_1 + \beta \, \mathbb{E}_{r_2|r_1} \left[V_2(r_2) \right] \right\}$ where:

$$k_1 = \begin{cases} k_0 + \mu b_0, & \text{if } b_1 \geq b_0 \quad \text{(buying or holding)}, \\ k_0 + \mu b_0 - \kappa(b_0 - b_1), & \text{if } b_1 < b_0 \quad \text{(selling)}. \end{cases} \tag{5}$$

Anticipating date 2 and substituting the balance-sheet identity into $V_2(r_2)$, the continuation value conditional on r_1 reads:

$$\mathbb{E}_{r_2|r_1}[V_2(r_2)] = k_1 + \mu b_1 - \kappa \, \mathbb{E}_{r_2|r_1}[\max\{0, \ -s_1 r_2 - c_1\}] \, . \tag{6}$$

Recalling that $\mathbb{E}[r_2|r_1] = \gamma r_1$ and $r_2|\{r_1\} \sim \mathcal{N}(\gamma r_1, \sigma^2)$, we note that the term inside the expectation is a "hinge" (positive-part) function of a normal variable. Applying Lemma 1 in Annex A, the continuation term becomes:

$$\mathbb{E}_{r_2|r_1}[\max\{0, \ -s_1 r_2 - c_1\}] = -(c_1 + s_1 \gamma \, r_1)\Big(1 - \Phi(z)\Big) + s_1 \sigma \, \phi(z), \quad z \equiv \frac{c_1 + s_1 \gamma \, r_1}{s_1 \sigma}, \tag{7}$$

where z measures the *distance-to-shortfall* in standard deviation units at date 2.[13] Combining (5) and (7), we can write the date-1 objective (up to constants) as a function of the choice b_1 through its effects on k_1 and c_1.

To solve for the date-1 policy, we work piecewise across whether the issuer is buying/holding bonds ($b_1 \geq b_0$) or whether it is selling bonds ($b_1 < b_0$).

Case B (buying/holding) We have $k_1 = k_0 + \mu b_0$. Differentiating $V_1(r_1)$ w.r.t. b_1, while noting that k_1 does not depend on b_1 in this case, gives:

$$\frac{\partial}{\partial b_1} V_1(r_1) = \mu - \kappa \frac{\partial}{\partial c_1} \mathbb{E}_{r_2|r_1}[\max\{0, \ -s_1 r_2 - c_1\}] \frac{\partial c_1}{\partial b_1} = \mu - \kappa \, (1 - \Phi(z)). \tag{8}$$

where we make use of (21) in Annex A and $\partial c_1 / \partial b_1 = -1$ in Case B. An interior optimum must satisfy the first-order condition $\mu = \kappa \, (1 - \Phi(z))$, which states that the marginal value of holding an extra unit of bonds must equal the marginal expected liquidation cost avoided by holding an extra unit of cash.[14] This pins down z at $z^\star = \Phi^{-1}(1 - \mu/\kappa)$, with $1 - \Phi(z^\star)$ representing the issuer's target shortfall probability. Recalling $z = (c_1 + s_1 \gamma \, r_1)/(s_1 \sigma)$ from (7), the target cash policy in Case B is:

$$c_1^\star(r_1) \ = \ s_0(1 + r_1)(\sigma z^\star - \gamma r_1), \tag{9}$$

from which the bond policy follows directly as $b_1^\star = s_1 + k_0 + \mu b_0 - c_1^\star$.

Several comparative statics arise from (9): first, lower liquidation costs, κ, or higher bond returns, μ, raise the attractiveness of bonds, leading to a reduction in optimal cash holdings, $c_1^\star$, and a corresponding increase in the target shortfall probability (lower $z^\star$); second, higher persistence, γ, reduces $c_1^\star$ for $r_1 > 0$ and increases it for $r_1 < 0$ as contemporaneous flows become more indicative of next-period flows; and third, increasing volatility, σ, scales $c_1^\star$ up proportionally, raising precautionary cash since higher uncertainty about future flows raises the option value of cash.

[13]Recall that a cash shortfall occurs when $c_1 < -s_1 r_2$. Therefore $1 - \Phi(z) = \mathbb{P}(r_2 < -c_1/s_1 \mid r_1)$ is the conditional probability of a liquidity shortfall that would necessitate bond sales at date 1. Larger z means a lower shortfall probability and a lower expected liquidation need.

[14]An interior solution exists for $1 - \Phi(z) \in (0, 1)$ since $\kappa > \mu$.

Feasibility of Case B requires $b_1 \geq b_0$, or equivalently $c_1 \leq c_{\text{cap}}(r_1)$ where

$$c_{\text{cap}}(r_1) = s_1 + k_1 - b_0 = s_0(1 + r_1) + k_0 - (1 - \mu)\, b_0. \tag{10}$$

If $c_1^{\star}(r_1) > c_{\text{cap}}(r_1)$, the Case B interior solution in (9) would imply a discretionary bond sale (i.e. selling bonds although cash is still available). Optimality will then place b_1 at the boundary $b_1 = b_0$ and set cash to its no-discretionary-bond-selling cap $c_1 = c_{\text{cap}}(r_1)$, provided that $c_1 \geq 0$ holds.

For sufficiently large redemptions, however, the issuer is forced to sell bonds as cash cannot drop into negative territory. The threshold flow below which $b_1 = b_0$ becomes infeasible since $c_1 \geq 0$ would be violated follows from (10) as:

$$r_{\text{FS}} = \frac{(1 - \mu)b_0 - k_0}{s_0} - 1, \tag{11}$$

Case S (selling). In this alternative case, we have $k_1 = k_0 + \mu b_0 - \kappa(b_0 - b_1)$ and $c_1 = s_1 + k_0 + \mu b_0 - \kappa b_0 - (1 - \kappa)b_1$. Differentiating $V_1(r_1)$ w.r.t. b_1 in Case S yields two terms: the direct effect through k_1 and the continuation effect:

$$\begin{aligned} \frac{\partial}{\partial b_1} V_1(r_1) &= \kappa + \beta\Big\{\kappa + \mu - \kappa \frac{\partial}{\partial b_1} \mathbb{E}_{r_2|r_1}[\max\{0, \; -s_1 r_2 - c_1\}]\Big\} \\ &= \kappa + \beta\Big\{\kappa + \mu + \kappa(1 - \kappa)\Big(1 - \Phi(z)\Big)\Big\} > 0. \end{aligned} \tag{12}$$

where we make use of $\partial k_1/\partial b_1 = \kappa$ as well as $\partial c_1/\partial b_1 = -(1-\kappa)$ in Case S and (21) in Annex A. Since $\kappa > 0$ ensures strict inequality for non-trivial choices of parameters, the objective is strictly increasing in b_1 when $b_1 < b_0$ so that any optimum cannot lie strictly inside Case S. In other words, the issuer never sells bonds if not forced to do so by redemptions.

State-space partition for the optimal $c_1(r_1)$ We now collect the interior solution and feasibility boundaries to obtain the piecewise optimal policy for $c_1(r_1)$. We define four phases based on whether the interior $c_1^{\star}(r_1)$ hits zero (cash non-negativity), exceeds the no-discretionary-selling cap $c_{\text{cap}}(r_1)$, or is feasible inside Case B:

$$c_1(r_1) = \begin{cases} 0, & \text{Phase I: } r_1 \geq \sigma z^\star/\gamma \\ c_1^\star(r_1), & \text{Phase II: } 0 \leq c_1^\star(r_1) \leq c_{\text{cap}}(r_1) \text{ and } r_1 \geq r_{\text{FS}}, \\ c_{\text{cap}}(r_1), & \text{Phase III: } c_1^\star(r_1) > c_{\text{cap}}(r_1) \text{ and } r_1 \geq r_{\text{FS}}, \\ 0, & \text{Phase IV: } r_1 < r_{\text{FS}}. \end{cases} \tag{13}$$

This yields a "shark-fin" profile for the optimal c_1 (Figure 2, left panel), with b_1 determined by the balance-sheet identity (right panel). The policy tallies closely with the workhorse model's optimal response, as we will show in Figure 4 in Section 4.1.

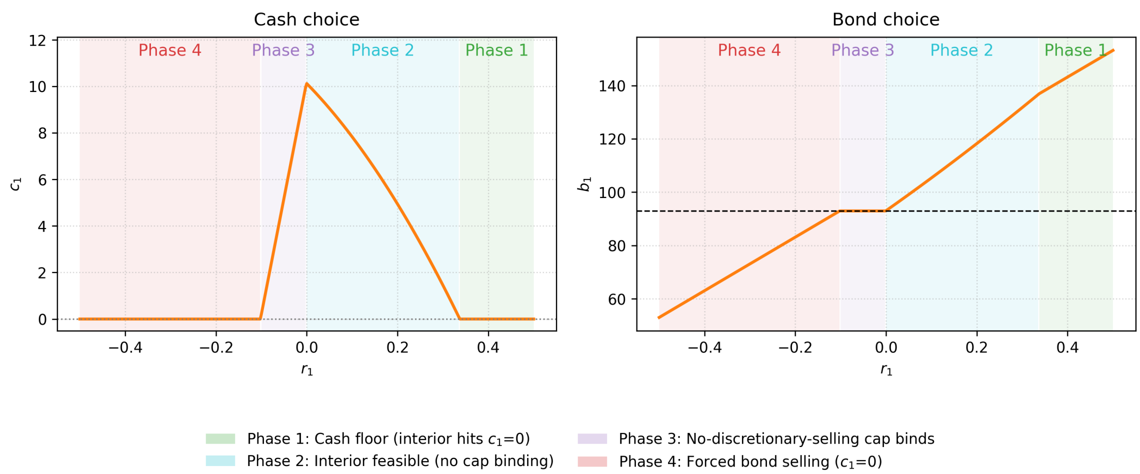

Figure 2: Illustrative example of optimal cash and bond choices at date 1 as a function of stablecoin flow shock r_1 (horizontal axis). The dashed black line in the right-hand panel indicates the bond choice, b_0, at date 0 for reference.

The main economic intuition is that the issuer's optimal cash policy reflects the option value of liquidity and a sequencing of adjustments across flow regimes: for large subscriptions (Phase 1), the expected continuation of inflows renders the insurance value of cash negligible, driving cash to zero; in the interior phase (Phase 2), cash is feasibly chosen and varies monotonically and concavely with flows, declining as inflows rise (recall equation (9)); for small to moderate redemptions (Phase 3), the no-discretionary-bond-selling cap binds and the issuer avoids bond liquidation costs by holding bonds at their date-0 level and uses cash to meet redemptions, letting cash adjust linearly at the cap (recall equation (10)); finally,

Calibrated parameters	Symbol	Calibrated value
Discount factor	β	0.99
Return on bonds	μ	7.67 bps
Stablecoins at date 0	s_0	100
Flow shock: persistence	γ	0.361
Flow shock: standard deviation	σ	0.063
Price impact: curvature	α	0.037
Price impact: maximum impact	m	41 bps
Capital cost factor	λ	9.23
Target moments	**Value in data**	**Value in model**
Capital ratio	0.101%	0.108%
Liquidity ratio	11.8%	11.7%

Table 1: Calibrated parameters and moments.

under large redemptions (Phase 4), feasibility constraints force bond sales and cash again hits zero, with the issuer liquidating bonds to restore non-negativity of cash. We discuss the phases and their boundaries in more detail in Annex A.

3 Calibration

Given that our workhorse model does not admit a closed-form solution, we solve it numerically. To this end, we adopt a parsimonious calibration approach that matches key aspects of the stablecoin ecosystem. In the subsections below, we first estimate the process governing stablecoin flows. Then, we identify the price impact function relevant for flow-induced sales of short-term government securities. Finally we discipline the residual parameters via a method of moments that aligns the model's balance-sheet ratios with those observed for a leading stablecoin issuer.

The list of parameters and their calibrated values is provided in Table 1. We align one period in the model with one week in the data and express the issuer's balance sheet in terms of billions of US dollars.

To calibrate the first two parameters, we assume an annual risk-free rate of 4%, a standard value in the literature. This translates into a weekly discount factor of roughly 0.99. Relatedly, the corresponding weekly yield on bonds is 7.67 bps (assuming compounding). Furthermore, we set the starting amount of stablecoins to $s_0 = 100$ at date 0, roughly equivalent to the average market capitalisation (in USD billions) of the two prominent stablecoins,

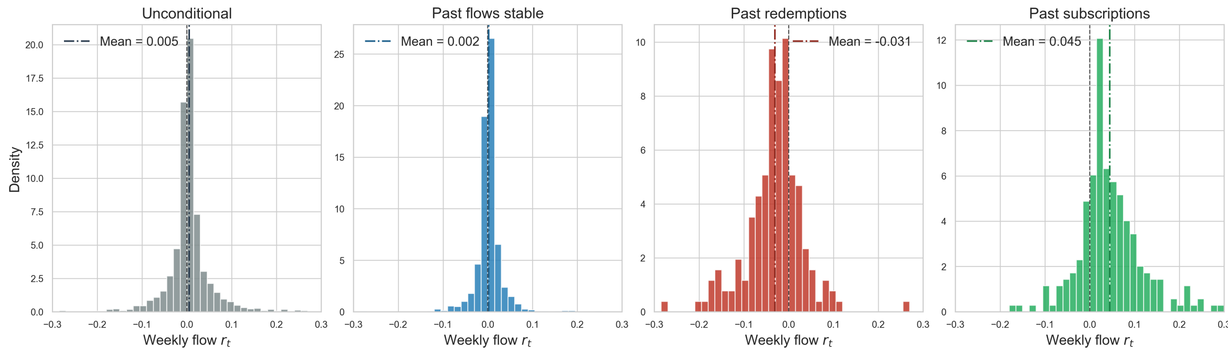

Figure 3: Pooled distribution of weekly (net) stablecoin flows for five issuers from 2020 to 2025. The first panel shows the unconditional distribution. The subsequent three panels display the conditional distributions based on net flows in the previous week being within ± 0.03 (stable), less than -0.03 (redemptions) and more than $+0.03$ (subscriptions). Flows of less than -0.3 and more than 0.3 have been dropped for expositional clarity. Dashed vertical lines indicate the mean of the distribution. Source: CoinGecko.

USDC and USDT, over the course of the year 2025. The estimation of the remaining parameters is described in the following sections.

3.1 Estimation of persistence in stablecoin redemptions

We assess whether a parsimonious AR(1) process, as assumed by our model, provides a suitable empirical characterisation of weekly stablecoin flows.

A visual inspection of weekly flows indicates substantial inertia. Figure 3 illustrates the pooled distribution of weekly changes (first panel) alongside the conditional distributions (second to fourth panels) categorised by the previous week's net flows having remained within ± 0.03 (stable; blue), falling below -0.03 (redemption; red) and exceeding $+0.03$ (subscriptions; green), respectively. When conditioned on a large redemption (subscription) in the previous week, the distribution of redemptions becomes skewed to the left (right). To assess persistence formally, Table 2 reports estimates from the dynamic panel regression

$$r_{i,t} = \gamma \, r_{i,t-1} + \eta_i + \eta_t + \varepsilon_{i,t}, \tag{14}$$

where $r_{i,t}$ denotes net flows of stablecoin i in week t, η_i captures time-invariant stablecoin fixed effects, η_t captures week fixed effects, and standard errors are clustered by week. We use a balanced sample of five US dollar-backed stablecoins over 2020–2025 for which continuous data are available.[15] The estimates indicate economically and statistically significant

Stablecoins	All (1)	All (2)	All (3)	Top2 (4)	All (5)	All (6)
Flows, r_{t-1} ($\hat{\gamma}$)	0.432***	0.426***	0.361***	0.462***	0.265***	0.362***
	(0.063)	(0.063)	(0.065)	(0.078)	(0.052)	(0.051)
Outflows, $r_{t-1} \times \mathbf{1}\{r_{t-1} < 0\}$					0.297	
					(0.211)	
Flows, r_{t-2}						−0.038
						(0.029)
Flows, r_{t-3}						−0.001
						(0.047)
Flows, r_{t-4}						0.049
						(0.034)
SD residuals ($\hat{\sigma}$)	0.065	0.065	0.063	0.028	0.062	0.061
	(0.010)	(0.010)	(0.009)	(0.004)	(0.009)	(0.010)
N	1,455	1,455	1,455	588	1,455	1,365
Adj. R^2	0.192	0.193	0.250	0.353	0.262	0.241
Stablecoin FE		Yes	Yes	Yes	Yes	Yes
Week FE			Yes	Yes	Yes	Yes

Table 2: AR(1) coefficients ($\hat{\gamma}$) based on equation (14). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Robust standard errors, clustered by week, in parentheses. The table reports pooled regressions for a balanced sample of five stablecoins based on weekly flows from 2020 to 2025. Column (4) comprises only the two largest stablecoins by market capitalisation (Top2), USDT and USDC. The table also reports the standard deviation (SD), $\hat{\sigma}$, of the estimated residuals, which we use as a proxy for the standard deviation of the redemption shock, σ, in the calibration.

persistence. Across specifications that vary the set of fixed effects (columns (1) to (3)), the AR(1) coefficient, $\hat{\gamma}$, is around 0.4. Persistence appears somewhat stronger for the two largest stablecoins by market capitalisation (USDT and USDC; column (4)), suggesting more predictable flow dynamics at scale. Likewise, distinguishing the effect of past outflows from inflows (column (5)) suggests somewhat higher persistence for outflows (0.562 versus 0.265), although we note that the difference is not statistically significant. Augmenting the specification with additional lags (column (6)) does not improve fit, supporting the use of a simple AR(1) structure. For calibration of the flow process, we therefore set $\gamma = 0.361$ and $\sigma = 0.063$ based on the specification with the richest set of fixed effects (column (3)), where $\hat{\sigma}$ is the standard deviation of residuals used to proxy the variance of redemption shocks.

[15] Consistent with the model, net flows in week t are computed as $r_t = s_t/s_{t-1} - 1$, based on weekly averages of coins outstanding s_t. The sample comprises Tether (USDT), Circle (USDC), Binance-USD (BUSD), Paxos-Standard (USDP) and TrueUSD (TUSD).

3.2 Calibration of the price impact function

Estimating the price impact of bond sales is central to our quantitative analysis but presents several identification challenges. Detailed, high-frequency data on stablecoin issuers' security-level holdings and transactions are not publicly available, and naive regressions of three-month T-bill returns on stablecoin flows are vulnerable to endogeneity from omitted variables as well as reverse causality, since T-bill returns can themselves influence stablecoin demand (e.g. Ahmed and Aldasoro [2025]). These data and identification frictions necessitate a proxy strategy that isolates mechanically forced sales due to redemptions.

Our identification strategy exploits the close alignment between the reserve portfolios of leading fiat-backed stablecoins and the asset holdings of US Treasury money market funds (MMFs). Both hold predominantly short-dated government securities and overnight repos, and both issue money-like, demandable liabilities to investor bases that are prone to withdrawals in stress (e.g. Anadu et al. [2023]).[16] Price dislocations observed during MMF redemptions are therefore informative about the sale-induced price impact associated with stablecoin outflows. The parallel is reinforced by scale: the two largest stablecoin issuers now rival major MMFs, with Tether among the top holders of short-dated US Treasuries (e.g. Ahmed and Aldasoro [2025]) and Circle managing reserves via a dedicated MMF.

We construct a weekly measure of total outflow-induced bond sales by MMFs that approximates the mechanically forced component of bond liquidations. Following the mutual fund literature (e.g. Coval and Stafford [2007]; Jiang et al. [2021]; Ma et al. [2022]), we combine weekly fund outflows with lagged portfolio weights to form

$$x_t \;=\; \sum_{f=1}^{F} \sum_{j=1}^{J} \Big(\mathit{Outflow}_{f,t} \times \mathit{Holdings\ share}_{j,f,t(m-1)} \Big) \;\geq 0, \tag{15}$$

where $\mathit{Outflow}_{f,t}$ is the outflow of fund $f \in F$ in week t and $\mathit{Holdings\ share}_{j,f,t(m-1)}$ is the share of security $j \in J$ in fund f's portfolio in the previous month $m-1$. To align with the profile of stablecoin reserves, we restrict our sample to US Treasury MMFs and consider only securities with a remaining maturity of 9 to 12 weeks as of week t to obtain total weekly flow-induced sales from January 2020 to December 2025 (see Annex C).[17]

[16] Although MMF shares confer ownership claims on the underlying portfolio, they are redeemable on demand and thus closely resemble demandable liabilities.

[17] Holdings are observed monthly; we therefore follow standard practice and combine weekly flows with

	(1)	(2)	(3)
Flow-induced MMF sales ($\hat{\alpha}$)	0.051***	0.034***	0.037***
	(0.014)	(0.012)	(0.012)
$-\ R_t^{1m}$ (bps)		1.156***	1.151***
		(0.298)	(0.354)
Aggregate MMF outflows ($\Delta \ln$)			−0.081
			(0.349)
MOVE ($\Delta \ln$)			−1.746
			(1.745)
RRP ($\Delta \ln$)			−0.029
			(0.033)
N	312	312	308
Adj. R2	0.013	0.351	0.351

Table 3: Price impact estimates. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Bootstrapped standard errors in parentheses. The cap m (see (16)) is set to the maximum weekly loss on three-month Treasury bills observed over 2020–2025. Data sources and variable construction are detailed in Annex C.

Our measure x_t is designed to capture exogenous, flow-driven selling pressure in the segment of the Treasury bill market that represent the preferred habitat of major stablecoin issuers. Aggregating across funds averages out idiosyncratic cash management decisions, yielding a more robust proxy for the net supply shock faced by the market. Aggregating across securities within a narrow remaining-maturity window ensures comparable duration and liquidity, tightly linking x_t to three-month Treasury bill pricing and limiting confounding factors.

We then estimate a concave price impact function that accommodates an upper bound on price dislocations (e.g. Fukker et al. [2022]). Specifically, we relate the negative of the weekly return on three-month T-bills (so that higher values correspond to lower prices) to outflow-induced sales,

$$\textit{Price impact}_t \ = \ m\Big(1 - e^{-\alpha x_t}\Big) \ - \ \vartheta_r \, R_t^{1m} \ + \ \vartheta_z \, Z_t \ + \ \varepsilon_t, \qquad (16)$$

where $m > 0$ governs the maximum (asymptotic) price impact, $\alpha > 0$ controls the curvature (concavity) of the impact function, R_t^{1m} is the weekly return on one-month T-bills that nets out common rate movements and close-substitute co-movements, and Z_t is a vector

lagged (monthly) portfolio weights.

of controls. The cap m reflects the entry of arbitrageurs if prices deviate materially from fundamentals; in the baseline we set m to the maximum weekly loss observed over our sample (9.52 bps), noting that a conservative risk-management stance would justify higher values (see Section 3.3 below). The control vector Z_t includes changes in market volatility (log changes in MOVE), quantity-based shifts in aggregate MMF investor demand (log changes in gross MMF outflows), and changes in MMF liquidity conditions (log changes in the Federal Reserve's overnight reverse repurchase agreements, RRP), consistent with theoretical work on the role of MMFs in Treasury bill markets (e.g. Stein and Wallen [2025]; d'Avernas and Vandeweyer [2024]).

Table 3 reports estimates for increasingly rich specifications. Our preferred specification (column (3)) suggests economically meaningful price impact: $\hat{\alpha} = 0.037$ implies that \$1 billion of outflow-induced sales lowers weekly three-month T-bill returns by about 0.3 bps, while \$10 billion (\$30 billion) reduces them by roughly 2.9 bps (6.4 bps).[18]

3.3 Matching balance-sheet ratios

We discipline the remaining two parameters in Table 1, the maximum price impact factor m and the capital cost factor λ, by matching the model's liquidity and capital ratios to those observed for a major stablecoin issuer. The idea is as follows. On the one hand, m is interpreted as the issuer's perceived upper bound on price dislocations under forced sales, which can rationally exceed the maximum realised impact in the sample owing to prudent risk management. On the other hand, λ determines the cost of raising capital and thus shapes the issuer's privately optimal capital choice.

Data constraints shape our calibration strategy. Stablecoin balance-sheet disclosures are heterogeneous across issuers and time: not all issuers report, definitions change, and consistent monthly histories are only available in recent years. To map the model's cash–bond structure to observables, we focus on issuers whose reserve assets closely match the ones in our model and whose disclosures are regular, sufficiently granular and externally attested.

Over our sample period, this criterion selects USDC, a dollar stablecoin with a market share of about 30% in terms of end-2025 market capitalisation. Using monthly observations for the year 2025 (see also Figure 14 in Annex D), we obtain an average liquidity ratio (cash-

[18] These magnitudes accord with complementary evidence on price effects from large purchases by stablecoin issuers as estimated by Ahmed and Aldasoro [2025].

to-stablecoins) of 11.8% and an average capital ratio (capital-to-stablecoins) of 0.101%.[19] These moments provide informative targets for the model's endogenous choices of cash and capital. Formally, we choose (m, λ) to minimise the distance between the empirically observed means of the liquidity and capital ratios in the year 2025 and the model-implied ratios on date 0.

The method of moments delivers a close fit, despite the model's parsimony. The calibrated (m, λ) result in a liquidity ratio of 11.7% and a capital ratio of 0.108%, close to their empirical counterparts, lending support to the subsequent quantitative analysis.

4 Quantitative analysis

In this section, we present the quantitative findings from our model based on the calibration described in the previous section. We first focus on the issuer's optimisation problem and then on regulatory considerations. The issuer's problem is solved using numerical methods (see Annex B for details).

4.1 Exogenous stablecoin flows

We begin by analysing the issuer's optimal strategy when stablecoin flows are exogenous. The four phases of the closed-form solution in Section 2.1 provide the intuition for understanding the quantitative findings.

At date 0, the issuer chooses a mix of cash and bonds to back its stablecoin liabilities. The optimality of an interior solution—as opposed to a corner solution involving only bonds or only cash—reflects the issuer's trade-off: cash avoids liquidation costs when meeting future redemptions, while bonds earn a spread but are costly to unwind under stress (Shleifer and Vishny [2011]). Quantitatively, the issuer allocates a small proportion of reserve assets to cash (12%) and the majority to interest-bearing bonds (88%), consistent with the

[19] When computing the liquidity ratio, we focus on reserve assets reported as cash (primarily bank deposits). Other reserve assets, such as outright bond holdings and overnight reverse repos, are typically exposed to settlement frictions, collateral haircuts, dealer dependence and reduced market depth during stress. In our model, these assets are therefore treated as *bonds* rather than cash. Our approach is consistent with evidence on liquidity management in open-ended funds, where cash is the primary buffer used for meeting outflows, while securities (including repo positions) are liquidated once cash buffers are drawn down (Ma et al. [2022]; Jiang et al. [2021]; Chernenko and Sunderam [2016]).

calibration's target moments.

Additionally, the issuer raises capital to absorb potential liquidation and valuation losses that could arise from future bond sales. The amount of capital raised depends not only on the cost of raising capital, but also reflects an optimal balance between raising the capacity to absorb losses due to future bond sales and maintaining sufficient cash as insurance against future redemptions. Quantitatively, the issuer's privately optimal choice of date 0 capital is 0.108, translating into a capital-to-stablecoin ratio of around 0.11%. This aligns, by design, with the calibration target and protects the issuer against a broad, though not exhaustive, range of redemption shocks.

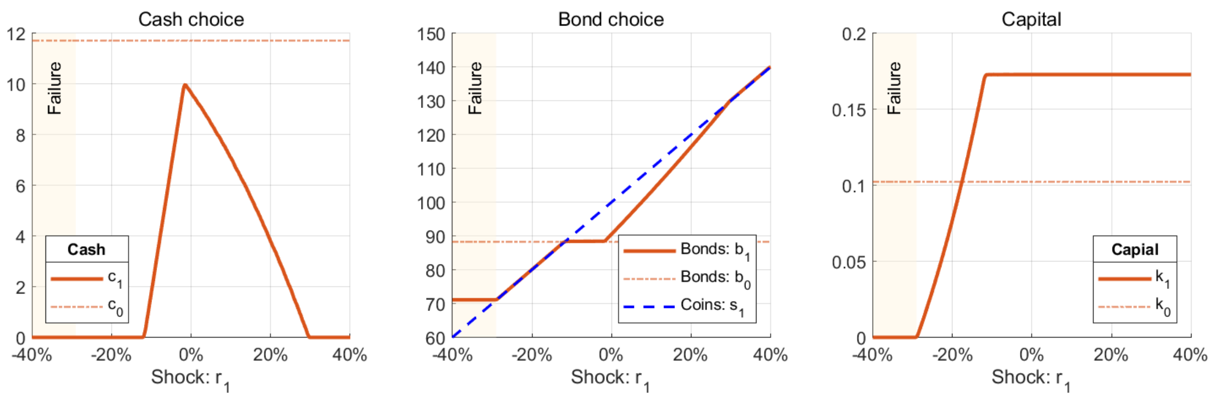

Figure 4: Solution to the date-1 problem as a function of stablecoin flow shock r_1, conditional on the issuer's optimal date-0 decisions. The date-0 stablecoin supply is fixed at $s_0 = 100$.

On date 1, the optimal cash use and bond sale response to flow shocks goes through four phases (Figure 4), consistent with the "shark-fin" pattern for cash use derived in Section 2.1.

For *small to moderate subscriptions*, the issuer reduces its cash holding with inflows (left-hand panel). Due to flow persistence, higher inflows at date 1 increase the likelihood of continued inflows at date 2, reducing the need to preserve cash. Instead, the issuer purchases additional bonds ($b_1 > b_0$). Meanwhile, capital rises above its date-0 level due to the accumulation of bond revenue and in the absence of losses associated with bond sales (right-hand panel).

For sufficiently *large subscriptions*, the issuer opts for a corner solution: the option value of cash becomes negligible since facing outflows at date 2 becomes highly unlikely. Cash is therefore driven to its floor and the issuer tilts further towards bonds.

In the phase of *small to moderate redemptions*, the issuer relies on cash to meet outflows. By maintaining the bond stock above or at its date-0 level, the issuer avoids liquidation costs and valuation losses, which helps preserve capital. However, this comes at the expense of entering date 2 with less cash.

Under *large redemptions*, the issuer enters the final phase, and feasibility constraints bind. Cash hits zero and bonds have to be liquidated to meet redemption requests. Fire-sale costs and valuation losses on remaining bond holdings start depleting the issuer's capital position (Figure 4, right-hand panel). As outflows grow to extreme levels, losses exceed the available capital and the issuer defaults (left shaded area). Given limited liability, capital cannot be negative. We assume that in such a situation the issuer is placed into receivership and bond sales are halted.[20]

The above sequencing of cash use and bond sales is consistent with empirical evidence on liquidity management in open-ended funds, where managers first use cash ("horizontal" cuts) and resort to selling securities once cash buffers are exhausted (Jiang et al. [2021]; Ma et al. [2022]; Chernenko and Sunderam [2016]).

Robustness A host of robustness checks help strengthen the intuition behind the mechanisms within our model and affirm the robustness of our findings.

In Annex E, we show that higher bond returns (as compared to the baseline) reduce cash holdings on both dates 0 and 1. By contrast, making bonds riskier by increasing the price impact of bond sales or making bond returns uncertain (Annex F) increases cash holdings.[21] Meanwhile, a risk that cash is not fully available makes cash a less effective buffer.[22] This induces the issuer to hold less cash and more bonds at date 1 (Annex F).[23]

[20] While outside of our model and not relevant for the solution to the issuer's problem, this scenario could imply that the receiver guarantees par redemption to all remaining coin-holders and covers losses through a dedicated fund (e.g. Carapella [2025]).

[21] The baseline model assumes a fixed bond return μ, consistent with the near risk-free nature of short-dated US Treasury bills. However, in practice, changes in short-term interest rates can introduce modest fluctuations in μ (market risk). A mean-preserving return risk increases the weight of costly low-return states, lowers the certainty-equivalent of bond payoffs and raises the value of cash as insurance.

[22] Stablecoin issuers typically hold their "cash" as deposits with commercial banks, which exposes them to credit risk. As noted in Section 5.1, the collapse of Silicon Valley Bank briefly exposed Circle (USDC) to such risk.

[23] Given our calibration, we do not find the issuer to hold more cash on date 0 for any given level of cash riskiness. For sufficiently high probability of cash being unavailable or sufficiently low availability when that happens, the date-0 cash holding can fall below its baseline value. At date 1, it is possible that the issuer

Next, we find that higher cost of raising capital implies less capital is raised on date 0 and also leads to lower capital on date 1 (Annex G). Further, we study the role of persistence in stablecoin flows. Lower persistence compared to the baseline leads to somewhat lower cash holdings on date-0 and in the case of redemptions on date-1 (Annex G). Indeed, starting from redemptions on date-1, the likelihood of date-2 redemptions is now lower, reducing the need for cash as insurance. By contrast, in the case of subscriptions on date 1, the likelihood of future redemptions is now higher, justifying an *increase* in date-1 cash holdings.

Finally, to ensure that our qualitative findings are not tied to the limited foresight implied by a three-period model, we consider an infinite-horizon version in Annex H. In this model, the issuer exhibits similar behaviour in the case of redemptions—which is when it is under pressure to meet the outflows. By contrast, in the case of subscriptions, it raises its cash buffers well above the levels observed in the three-period model as insurance against redemptions in the many periods that follow. Overall, however, the "shark-fin"-shaped cash response continues to hold.

4.2 Endogenous flows and regulatory thresholds

This section introduces endogenous flows and quantifies how prudential thresholds and coinholder discipline shape the issuer's liquidity management.

LR threshold We consider an LR threshold of $\theta = 5\%$, roughly equivalent to one standard deviation of redemptions, σ, below the issuer's date-0 liquidity ratio in the unregulated baseline. This represents a plausible low-end calibration. While this threshold is set below the liquidity ratio chosen by the issuer at date 0 in the absence of regulation, it induces significant adjustments by the issuer on both dates.

On date 0, the issuer increases its cash holdings, as shown by the dashed blue line above the dashed red line in left-hand panel of Figure 5. This adjustment is driven entirely by date-1 considerations: the issuer anticipates the potential need for more cash on date 1 to avoid breaching the LR threshold and positions itself accordingly.

On date 1, the issuer holds the same or a larger amount of cash (solid blue line) compared to the baseline case (solid red line). Specifically, in the case of *redemptions*, cash holdings are

holds more cash if cash is only slightly risky.

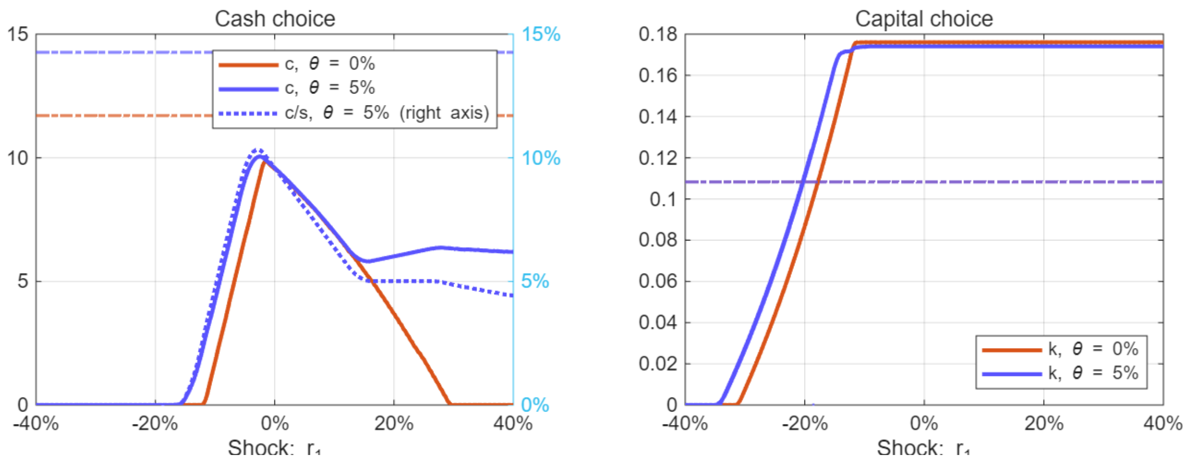

Figure 5: Comparing the baseline model where stablecoin flows are exogenous (shown in red colour) with the case where an LR threshold $\theta = 5\%$ gives rise to endogenous flows (shown in blue colour). The left-hand panel shows the cash choice while the right-hand panel shows the capital choice. Date-0 choices are shown via dashed lines and date-1 choices are shown via solid lines. Notably, the dotted blue line in the left-hand panel shows the cash-to-stablecoin ratio on date 1.

marginally higher relative to the baseline, and remain the primary instrument for meeting outflows.

In contrast, under subscriptions, the regulatory bite is more evident. The LR threshold induces materially higher cash holdings. The issuer maintains a liquidity ratio at or above $\theta = 5\%$ to avoid triggering coin-holder discipline, as shown by the cash-to-stablecoin ratio (dotted blue line). For sufficiently large subscriptions, however, small breaches of the LR threshold become privately optimal. For instance, when facing a subscription as large as 40%, a minor breach of the threshold may still result in subscriptions on date 2. This implies a low insurance value of cash and allows the issuer to dip below θ and hold more bonds instead.

The issuer's cash management strategy exemplifies *buffer building* on date 0 and *buffer usability* on date 1. On date-0, even though issuer's liquidity ratio exceeds θ in the unregulated case, it pre-emptively raises its cash holdings in the regulated case to reduce the likelihood and severity of threshold breaches on date-1. On date 1, the issuer tries to remain close to the LR threshold when this is feasible (i.e. during subscriptions and small redemptions), while using cash and breaching the threshold when necessary (i.e. in the case of large redemptions).

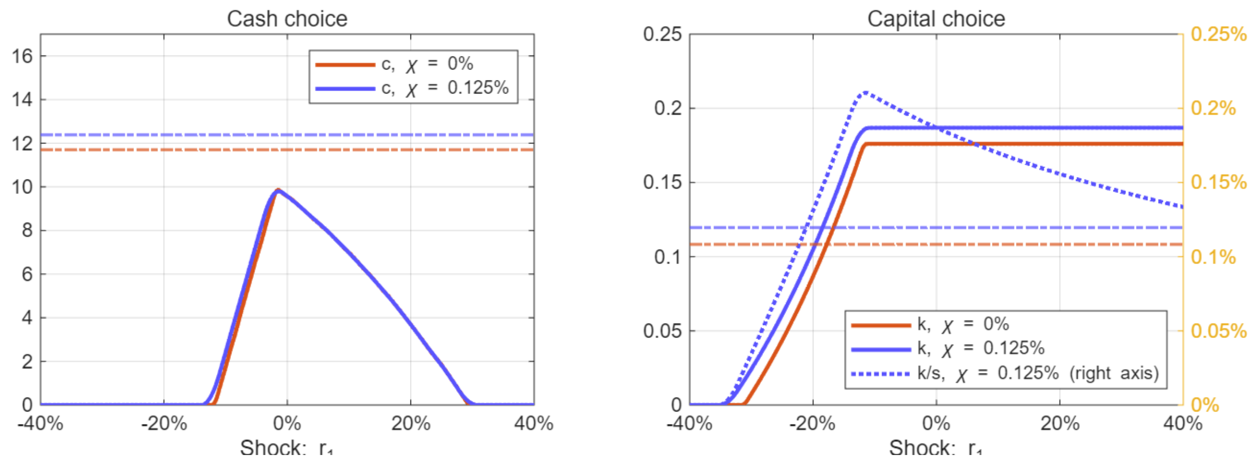

Figure 6: Comparing the baseline model where stablecoin flows are exogenous (shown in red colour) with the case where an CR threshold $\chi = 0.125\%$ gives rise to endogenous flows (shown in blue colour). The left-hand panel shows the cash choice while the right-hand panel shows the capital choice. Date-0 choices are shown via dashed lines and date-1 choices are shown via solid lines. Notably, the dotted blue line in the right-hand panel shows the capital-to-stablecoin ratio on date 1.

By comparison, the LR threshold has no impact on the capital choice of the issuer on date 0 (right-hand panel of Figure 5). On date 1, in the absence of bond sales, the issuer has somewhat less capital. This is because the issuer holds more cash and fewer bonds on date 0, resulting in lower bond interest income to add to capital on date 1. However, because of the extra cash on date 0, bonds sales are pushed out ie *begin* at higher levels of redemptions. Therefore, the kink in capital also shifts left and so does the point of default.

CR threshold Next, we consider a capital-ratio threshold of $\chi = 0.125\%$, a value modestly above the calibrated date-0 capital-to-stablecoins ratio of the issuer.[24] On date 0, the issuer reacts by raising capital and thus positions itself closer to the threshold (dotted blue line in the right-hand panel of Figure 6). This helps avoid coin-holder discipline. Higher capital levels on date 0 translate into greater capital buffers for any flow shock r_1 on date 1, reducing the likelihood of breaching the capital-ratio threshold and ultimately lowering insolvency risk.

We further note that the introduction of the CR threshold induces the issuer also to raise

[24]The calibration of the CR threshold varies across jurisdictions. Some specify no minimum (e.g. Hong Kong SAR, Singapore). MiCAR in the European Union, by contrast, requires a level of 2-3%, and additional overcollateralisation requirements may apply, to account for both liquidity risk and a range of operational risks. Since our model abstracts from operational risks, the relevant CR threshold is correspondingly lower.

its cash holding on date 0. This finding underscores the complementarity between cash and capital: higher cash holdings reduce the need to liquidate bonds, thereby mitigating losses that undermine capital. This, in turn, lowers the probability of breaching the CR threshold.

Alternative forms of coin-holder discipline So far, we have assumed that coin-holder discipline kicks in if *either* threshold is violated, and it is proportional to the *maximum* violation across the two thresholds. We consider two variations of this mechanism in Annex I. First, we assume coin-holders assess the two thresholds holistically. Buffers relative to one threshold compensate for breaches relative to the other (i.e. a liquidity ratio above the LR threshold compensates for a capital ratio below the CR threshold and vice versa). This additional flexibility, or weakened coin-holder discipline, induces the issuer to hold less cash and capital on both dates 0 and 1. Second, we consider the case of hard thresholds. Recall the argument that if regulatory thresholds are imposed as minimum *requirements*, they may be self-defeating. We confirm this insight in Annex I by showing that, while a hard LR threshold elicits greater cash holdings, it increases the likelihood of default by forcing the issuer to sell bonds when it breaches the hard threshold.

Varying the regulatory thresholds We now examine the issuer's sensitivity to changes in the regulatory thresholds. In Figure 7, we partition the response by LR threshold (top row) and CR threshold (bottom row) and highlight an asymmetry.

As the LR threshold rises, the issuer responds by monotonically raising its cash holdings on date 0 (top-left panel), while date-0 capital remains unchanged (top-right). This pattern reflects the sequencing of cash versus bond use on date 1. Cash serves as the issuer's first line of defence against redemptions, meaning that higher cash holdings reduce both the frequency and severity of breaching the LR threshold. In contrast, capital plays a loss-absorbing role only when bond sales are pursued. As we noted in the context of Figure 5, the LR threshold is already breached by the time bond sales begin. As such, capital does not directly affect the probability of breaching the LR threshold on date 1, which explains the muted response of capital to changes in the LR threshold.

As the CR threshold rises, the issuer initially holds its cash level constant (bottom-left panel) while raising more capital on date 0 (bottom-right panel). Beyond a certain point, however, the marginal cost of raising capital begins to exceed the marginal effectiveness

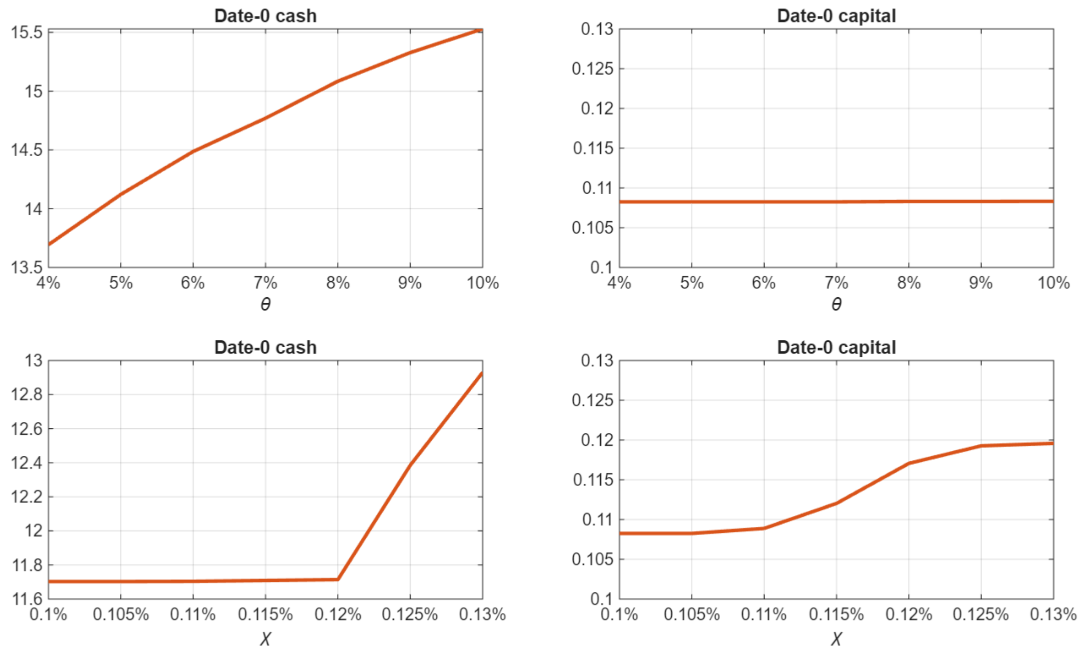

Figure 7: The impact of alternative regulatory thresholds on the capital and cash choices of the issuer.

of capital in preventing breaches of the CR threshold. Instead, increasing cash holdings becomes attractive because cash reduces the need to liquidate bonds and thus lowers the losses that erode capital. Consequently, beyond this point, the issuer increasingly relies on cash to comply with the CR threshold (bottom-left panel), while capital becomes relatively insensitive to further increases in the CR threshold.

The asymmetry in how the issuer responds to changes in the LR threshold versus the CR threshold constitutes a key result of the paper. It underscores that the LR threshold only operates through the cash margin, whereas the CR threshold can elicit a response by the issuer along both cash and capital margins. We study the policy implications of this asymmetric response to regulatory threshold in the next section.

4.3 Policy objectives and calibration of regulatory thresholds

In this section we link the prudential thresholds and policy objectives. Consistent with international recommendations on stablecoin arrangements (e.g. Financial Stability Board

[2023]), the regulator faces two wedges that a private issuer does not fully internalise: first, limited liability leads the issuer to underweight tail losses that are borne by coin-holders (e.g. Li and Mayer [2026]); second, forced liquidations transmit stress to money markets by depressing the prices of reserve assets and close substitutes (e.g. Ahmed and Aldasoro [2025]; d'Avernas and Vandeweyer [2024]).[25] These wedges motivate a dual objective comprising a micro-prudential target—default risk—and a macro-prudential target—the expected market impact from sales.

We operationalise these objectives via the issuer's probability of default (PD) and the expected price impact (EPI) stemming from bond sales by the issuer. Both of these policy targets are evaluated for date 1 from the vantage point of date 0. Specifically, PD is the probability that capital becomes negative on date 1 given the balance-sheet choice at date 0. EPI is the expected decline in the price of bonds induced by the issuer's sales at date 1, irrespective of default.[26]

In line with stress-testing practice, we compute both metrics under a stressed flow distribution centred at ζ:

$$PD(\chi, \theta, \zeta; c_0, k_0, b_0) = \int \mathbb{1}\Big(k_1(r_1) < 0\Big) f\Big(r_1 \,|\, \mathbb{E}[r_1]{=}\zeta\Big) dr_1. \tag{17}$$

$$EPI(\chi, \theta, \zeta; c_0, k_0, b_0) = \int g\Big(b_0 - b_1(r_1)\Big) f\Big(r_1 \,|\, \mathbb{E}[r_1]{=}\zeta\Big) dr_1. \tag{18}$$

We find that regulatory thresholds combined with coin-holder discipline help reduce both PD and EPI irrespective of the hypothetical stress level under which these policy targets are computed. Changes in the issuer's date-0 and date-1 cash and capital choices (discussed in Section 4.2) underpin this finding. Notably, regulation is more beneficial when stress is more severe (Figure 8).

In what follows, we fix $\zeta = -0.126$, which corresponds to a two standard-deviation redemption shock, and examine the impact of alternative regulatory thresholds. At this level, the unregulated case yields a PD of more than 15 bps and an EPI of close to 4 bps (red lines, Figure 8).

[25]Depressed prices can adversely affect other asset holders and, in extreme cases, initiate cascades of defaults (Cantú et al. [2024]).

[26]Alternative formulations of the policy targets, such as constraints on Value-at-Risk or market liquidity indicators, can be accommodated by the model and are likely to yield qualitatively similar trade-offs.

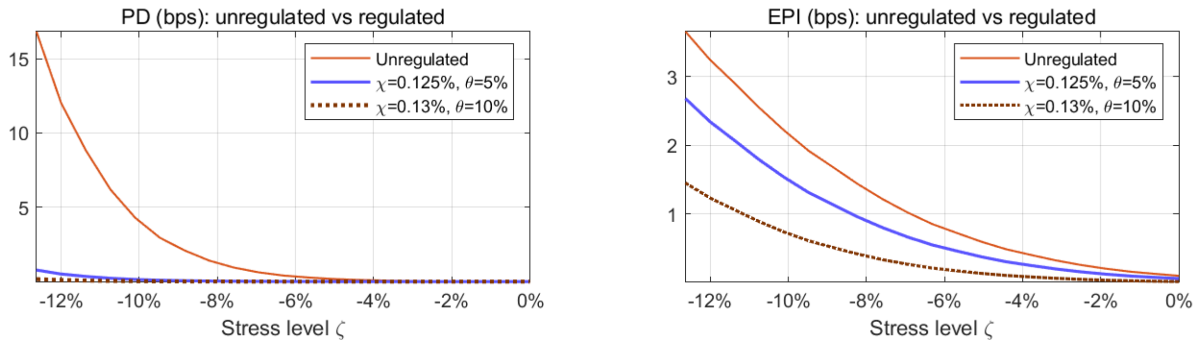

Figure 8: The figure plots probability of default (PD) and expected price impact (EPI) across stress levels ζ—defined as the assumed expectation of the date-1 flow—contrasting the unregulated baseline with selected regulated regimes.

Increases in both the LR and the CR thresholds can reduce PD and EPI (Figure 9). This suggests that the two thresholds exhibit some degree of substitutability if one were to focus only on either the PD or the EPI objective.[27] However, if the policy objective is to meet both PD and EPI targets, then the joint use of both prudential thresholds becomes necessary, in line with the Tinbergen principle that attaining multiple independent policy targets requires at least as many independent instruments.

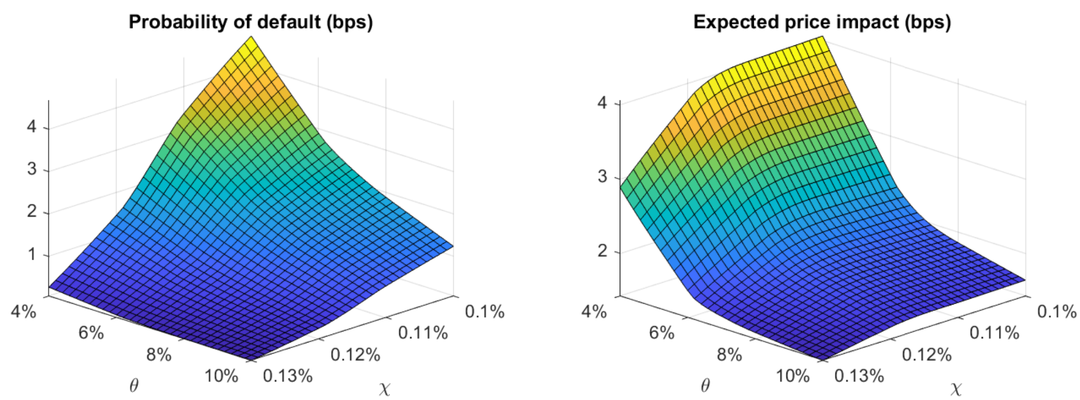

Figure 9: The figure shows how regulatory thresholds affect the issuer's probability of default (PD) and expected price impact (EPI), with the date-1 flow distribution's mean fixed at the stress level $\zeta = -0.126$.

We further note that thresholds act through distinct channels. The LR threshold, θ, operates along the cash margin, inducing higher cash holdings but no additional capital.

[27] It is worth emphasising the partial substitutability of the two thresholds: depending on the baseline level of χ, raising the CR threshold may have no effect on EPI or may reduce it.

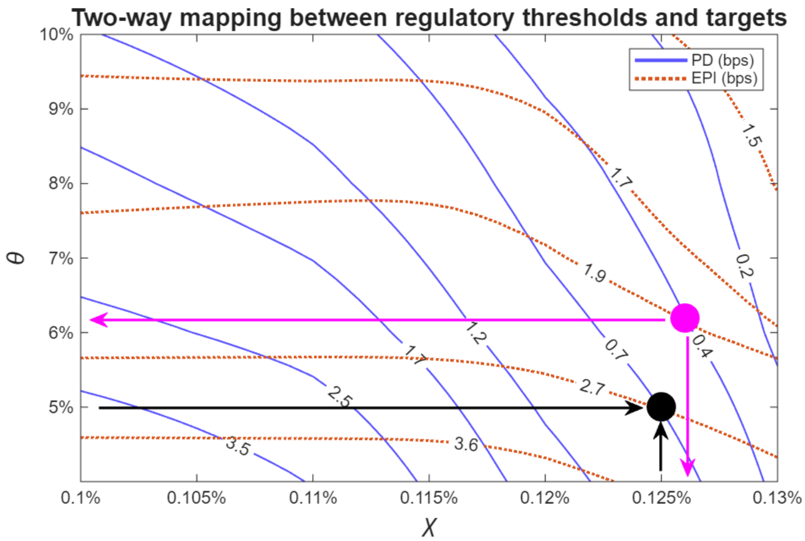

Figure 10: Mapping between regulatory thresholds (χ, θ) and targets (PD, EPI). The dotted lines are the iso-EPI curves and the solid lines are iso-PD curves—i.e. PD and EPI are constant along these curves respectively. The constant value of PD and EPI is indicated along each curve.

By contrast, the CR threshold, χ, primarily induces more capital and, at sufficiently high levels, also induces more cash holdings. This asymmetry matters for joint calibration. Cash has a first-order impact on EPI by reducing the incidence and scale of bond sales, and a secondary impact on PD by limiting sale-induced losses. Capital has a first-order impact on PD by absorbing liquidation and valuation losses, but no direct impact on EPI. These insights highlight that ignoring a stablecoin issuer's privately optimal response to prudential thresholds—including how buffers are built ex ante and used in stress—risks systematic miscalibration of the thresholds.

Ultimately, our analysis provides a two-way mapping between regulatory thresholds and targets that can inform policy calibration. The *iso-plot* in Figure 10 shows this mapping by overlaying the two panels from Figure 9. It allows one to project $(PD,\ EPI)$ combinations for a given choice of thresholds (χ, θ). Alternatively, the iso-plot can be used to trace combinations of $(\chi,\ \theta)$ that attain a given $(PD,\ EPI)$ target.

To illustrate this two-way mapping, consider the regulatory pair $\theta = 5\%$ and $\chi = 0.125\%$ examined above: it results in a weekly PD of around 0.7 bps and an EPI of 2.7 bps (black annotations in Figure 10). If regulation aims for a PD of less than 0.5 bps and an EPI of less than 2 bps, then moderately higher thresholds are needed (magenta annotations).

5 Extensions

This section presents extensions of the model to affirm the robustness of our main findings and also build further intuition. First, we assess how the issuer adapts its balance sheet in response to asymmetric persistence in stablecoin flows or even run-like flows. We then present a life-cycle simulation that tracks the cash, bond and capital choices of the issuer under varying flow scenarios. We conclude with a welfare assessment of the regulatory thresholds.

5.1 Asymmetric flow persistence and run risk

A potential policy concern is that outflows may exhibit significantly higher persistence compared to inflows, particularly during periods of stress. In fact, in the context of creditor coordination failures—as explored in the global games literature (e.g. Goldstein and Pauzner [2005]; Rochet and Vives [2004]; Morris and Shin [1998])—adverse shocks could precipitate a regime shift in flows, and even culminate in a run.

To explore these possibilities, we first consider a scenario where the persistence of redemptions exceeds the persistence of subscriptions. We focus on exogenous flows to ensure that any effects of changes in persistence are not conflated with the role of regulatory thresholds. Our findings in Section 3.1 (see Table 2, column (5)) guide our calibration in this case, suggesting an AR(1) coefficient of $\gamma_{r<0} = 0.562$ for outflows and $\gamma_{r>0} = 0.265$ for inflows.[28]

Under more persistent outflows, the issuer holds more cash on date 0. This reflects the increase in downside risks against which cash serves as an insurance (Figure 11, left-hand panel). On date 1, the issuer holds more cash when facing subscriptions. This is because the persistence of subscriptions is smaller than that in the baseline, implying a higher probability of facing redemptions at date 2.

Next, we assess the impact of run risk on the issuer's liquidity management. While our previous analysis assumes mean-reversion of flows, particularly large outflows could prompt panic among remaining coin-holders, leading to further aggravation of flows.

The experience of USDC in the context of Silicon Valley Bank's failure in March 2023 provides an illustrative example (e.g. Ahmed et al. [2024]). When news spread that Circle,

[28] Recall that the difference between the two coefficients lacks statistical significance, which is why we use a common coefficient for our workhorse model.

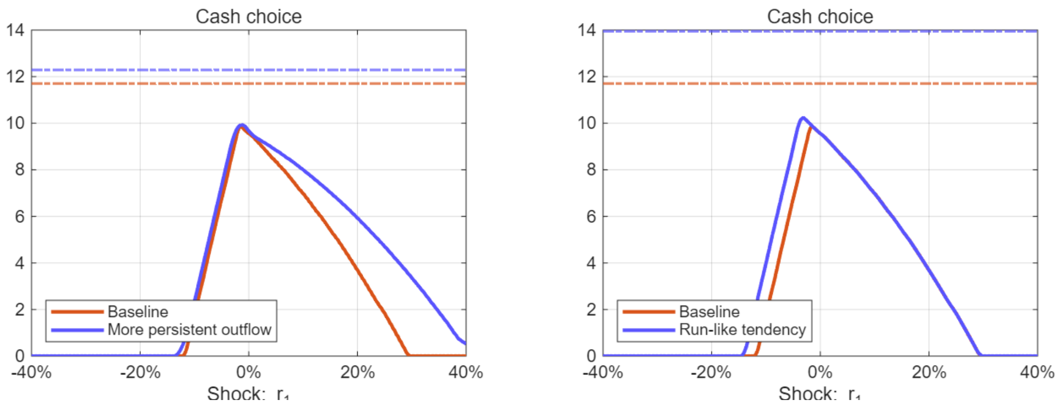

Figure 11: Left-hand panel: Cash choice conditional on flows being more persistent in the case of outflows ($\gamma_{r<0} = 0.562$) than for inflows ($\gamma_{r\geq 0} = 0.265$). Right-hand panel: Cash choice conditional on expected future redemptions increasing to twice the standard deviation of flows (instead of reverting to mean) as soon as current redemptions exceed one standard deviation of flows. Solid lines show the date-0 cash choice while dashed lines show the date-1 cash choice.

the issuer of USDC, held significant deposits with the bank, *daily* flows intensified on impact (see Figure 21 in Annex J). We posit that there might have been a regime shift that was quickly reigned in by public backstops to stabilise the US banking sector. While the event period was too short to estimate a regime switching process at weekly frequency, we take the spirit of the March 2023 episode to illustrate the impact of run-like tendencies in flows on issuer behaviour in Figure 11 (right-hand panel). Here, we assume that whenever redemptions exceed one standard deviation of the distribution on date 1 (i.e. $r_1 < -\sigma$), the expected flow on date 2 falls to negative of twice the standard deviation (i.e. $\mathbb{E}_{r_1}[r_2 \,|\, r_1 < -\sigma] = -2\sigma$).

Intuitively, the issuer responds by holding more cash at date 0 given that the distribution of flows has shifted towards higher expected redemptions. Moreover, when facing date-1 redemptions, run risk induces the issuer to hold more cash as insurance against the elevated risk of date-2 redemptions (unless the issuer has no option but to run down cash). In the case of subscriptions, cash holdings at date 1 coincide with the baseline choice since the persistence of inflows is unchanged.

5.2 Simulating a stablecoin issuer's life cycle

Simulating the life cycle of a stablecoin issuer helps shed light on their default and growth trajectories. To this end, we generate randomly drawn sequences of shocks from the condi-

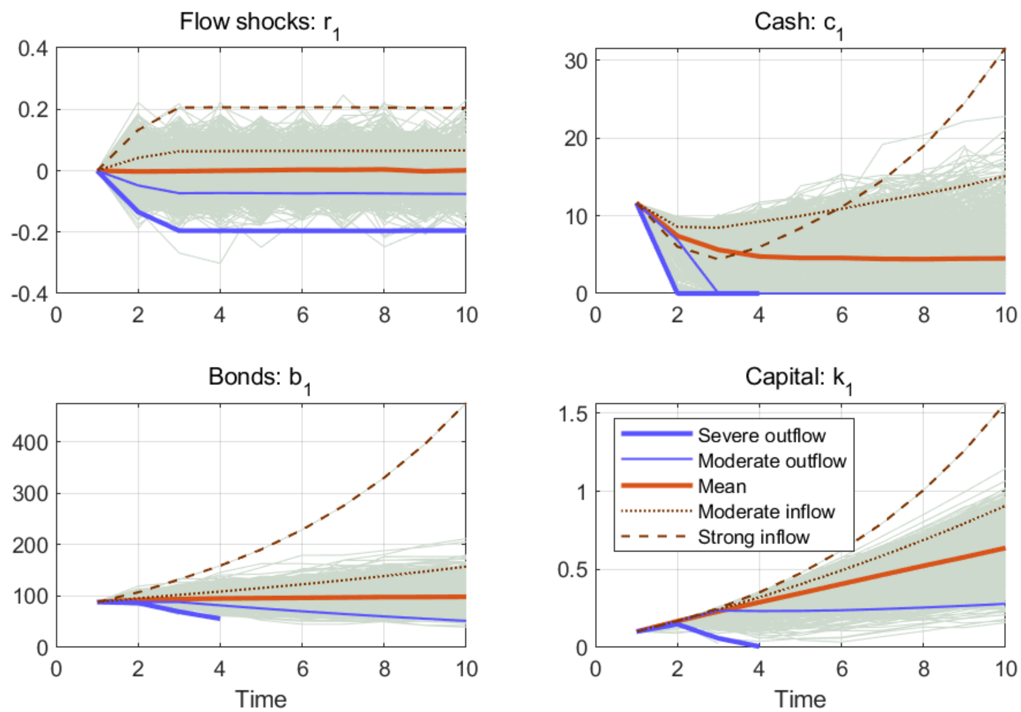

Figure 12: Simulating the issuer's balance sheet for randomly drawn sequences of flow shocks over 10 weeks. The bold red line shows the mean across simulations. The other lines (blue and brown) indicate manually created special cases of sustained weekly inflows or outflows.

tional distribution of flows and track the cash, bond, and capital positions of the issuer. We assume that in each period the issuer responds to the shock as per its optimal strategy at date 1 in our workhorse model.

The following insights emerge (Figure 12). First, the range of possible outcomes is wide. Strong and moderate inflows lead to rapid growth (dashed and dotted brown lines), reminiscent of the rise in stablecoins' market capitalisation over the past years. Meanwhile, the bold blue line—which ceases to exist by the fourth period in the fourth panel—shows how sustained severe outflows can drive capital below zero and lead to default. Moderate outflows, by comparison, can lead to a scenario where the issuer is left with no cash but still has positive amounts of capital (thin blue line). Second, there is greater variation in cash (relative to its mean) as compared to bonds. Specifically, the coefficient of variation in the case of cash (measured as the ratio of standard deviation and mean) is more than twice that in the case of bonds. This underscores the issuer's preference to maintain a steady holding of interest-bearing bonds and use unremunerated cash as a shock absorber.

5.3 Towards optimal regulation

Thus far, our analysis has been positive: we take regulatory thresholds as given and study how they impact issuer behaviour and ultimately PD and EPI. Regulation, however, also constrains stablecoin issuers. If set too stringently, it can render the business model unsustainable. This can translate into fewer entrants and ultimately fewer stablecoins in circulation, hurting users. A normative assessment must therefore balance the welfare benefits from lowering PD and EPI against the constraint imposed on issuers and the welfare costs borne by users. Any such assessment is necessarily stylised, as it depends on how much value society places on stablecoins and on the outside options of issuers and users. This subjectivity limits the insights from any welfare analysis, motivating our paper's focus on a positive analysis. This caveat notwithstanding, we sketch a simple objective function to illustrate how our mapping can be embedded in welfare-like comparisons.

We posit a regulator's objective in reduced form that trades off PD and EPI against the issuer's franchise value (which in turn should shape the provision of stablecoin services):

$$W(\theta, \chi) = \omega V_0 - PD - EPI, \tag{19}$$

where (θ, χ) are the LR and CR thresholds, respectively, V_0 is the issuer's date-0 charter value as defined in equation (3), and PD and EPI are computed under the stressed flow distribution used in Section 4.3 i.e. expected redemptions on date 1 equal $\zeta = -0.126$. The weight $\omega > 0$ captures how much the regulator values the sustainability of stablecoins relative to prudential outcomes. Higher thresholds reduce PD and EPI but also lower V_0, giving rise to the regulatory trade-off.[29]

We first vary the LR threshold θ, keeping the CR threshold χ at zero. The objective $W(\theta, \chi)$ is hump-shaped in θ for empirically plausible values of ω, yielding an interior optimum (Figure 13, left-hand panel). When ω is low (i.e. limited weight on issuer franchise value), the regulator prefers a more stringent LR threshold. As ω rises, the optimum shifts

[29] V_0 declines because greater cash holdings imply foregone bond returns and greater capital implies higher raising costs. Alternative scaling factor or weights on PD and EPI could be used. Our aim is illustrative rather than definitive.

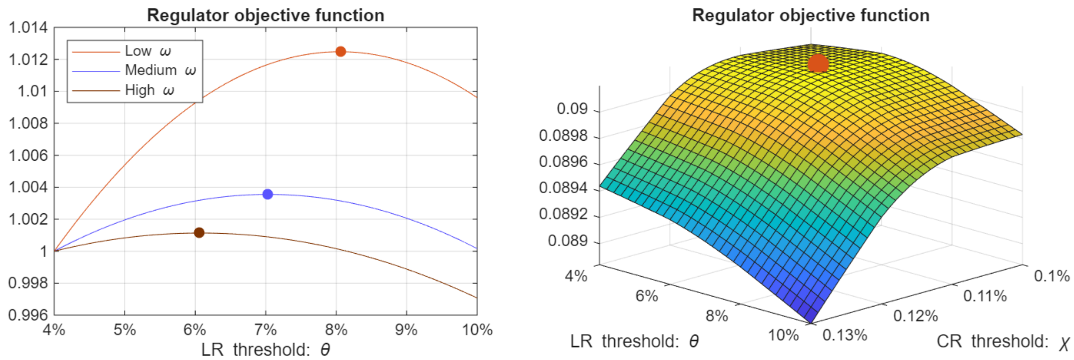

Figure 13: Left-hand panel: Regulator's objective $W(\theta, \chi)$ as a function of LR threshold θ for $\omega = 0.1$ (low), $\omega = 0.2$ (medium) and $\omega = 0.3$ (high), holding the CR threshold $\chi = 0$. Each objective function is normalised to 1 at $\theta = 4\%$. Right-hand panel: Contour plot of $W(\theta, \chi)$ over varying LR and CR thresholds given a medium value of $\omega = 0.2$, illustrating an interior optimum (red dot).

towards looser LR thresholds because the marginal cost in terms of V_0 dominates the diminishing marginal gains in PD and EPI. For extreme values of ω, corner solutions emerge: a high LR threshold when ω is close to zero and a very low threshold when ω is large.

Finally we optimise jointly over the LR and CR thresholds (θ, χ). To fix ideas, we assume a medium value of $\omega = 0.2$. The objective exhibits an interior maximum, reflecting diminishing returns to either individual or joint use of the two thresholds (Figure 13, right-hand panel). We note complementarities across the instruments. If used individually the optimal LR threshold is around 7% (recall left-hand panel). If used jointly, the optimal LR threshold declines to around 6% and welfare increases. In other words, when both thresholds are at the regulator's disposal, welfare optimisation implies using each threshold, but less intensely than if used in isolation.

6 Conclusion

This paper develops a dynamic model of a fiat-backed stablecoin issuer that jointly optimises capital, cash and bond holdings, while facing serially correlated stablecoin flows and bond sale discounts. In the unregulated benchmark, the issuer holds little capital and favours interest-bearing but less liquid bonds over cash. This behaviour raises default risk as well as the risk of systemic spillovers via forced bond sales. These risks justify the use of liquidity and capital thresholds. We model the thresholds as usable buffers—i.e. they can be breached, but at the cost of coin-holder discipline—therefore creating incentives for the issuer to hold buffers. The thresholds affect the issuer through asymmetric channels. The liquidity threshold operates through the cash margin, whereas the capital threshold operates through both capital and cash margins.

Translated into policy goals, capital helps absorb liquidation and valuation losses, thus reducing default risk. Cash helps reduce bond sales and associated losses, thereby lowering default risk as well as systemic spillovers. Taken together, both regulatory thresholds help achieve each policy objective, suggesting that the thresholds are substitutes. However, they are complements if the regulator targets default and spillover risks jointly. We provide a two-way mapping between liquidity and capital threshold pairs to dual policy targets—PD and EPI. Calibrated to stablecoin flows and US Treasury market depth, plausible thresholds reduce both PD and EPI related risks (e.g. PD from above 15 bps to 0.7 bps and EPI from about 4 bps to 2.7 bps).

Our findings carry concrete policy implications. For one, we show that hard minimum requirements that must always bind can increase default risk. Instead, *usable* thresholds combined with coin-holder discipline can improve the alignment of private incentives with prudential objectives. Further, the complementarities between the liquidity and capital thresholds justify their joint calibration in order to meet both micro- and macro-prudential targets. The two-way mapping between (PD, EPI) and (CR, LR) thresholds serves as a useful basis for the design of prudential policies for stablecoins.

Several aspects of our analysis can be expanded in future research. First, while we focus on liquidity and market risks, operational risks are also relevant for stablecoins. These risks tend to entail low-probability but high-impact events (e.g. blockchain forks or 51% attacks) that may require tailored modelling techniques. Relatedly, operational risks might justify higher capital thresholds relative to the levels studied in the paper. Second, while we study a stablecoin issuer that invests in only two types of assets, in practice, issuers may invest in a wider range of assets with more diverse risk-return profiles than cash and government bonds. How issuers respond to regulation in these cases may offer additional insights. Third, as stablecoin issuers grow further, their stress could become systemic. This could justify studying concentration limits as another regulatory tool (e.g. a requirement that cash deposits must be spread across banks and money market funds). Our framework offers a starting point for analysing these extensions.

References

Ahmed, R. and I. Aldasoro (2025). Stablecoins and safe asset prices. *BIS Working Papers 1270.*

Ahmed, R., I. Aldasoro, and C. Duley (2024). Public information and stablecoin runs. *BIS Working Papers 1164.*

Anadu, K., P. Azar, C. Huang, M. Cipriani, T. Eisenbach, G. La Spada, M. Landoni, M. Macchiavelli, A. Malfroy-Camine, and J. C. Wang (2023). Runs and flights to safety: are stablecoins the new money market funds? *Federal Reserve Bank of Boston Research Department Working Papers 23-11.*

Bank of England (2025). Proposed regulatory regime for sterling-denominated systemic stablecoins. Consultation paper, November.

Berrospide, J., A. Gupta, and M. Seay (2024). The usability of bank capital buffers and credit supply shocks at SMEs during the pandemic. *International Journal of Central Banking 20*(3), 185–255.

Bertsch, C. (2023). Stablecoins: Adoption and fragility. *Sveriges Riksbank Working Papers 423.*

Cantú, C., J. Frost, and T. M. Noone (2024). Some perspectives on the regulation of stablecoins. *Regulating Financial Innovation: FinTech, Crypto-assets, DeFi, and Beyond (Edited by Steven Schwarcz).*

Carapella, F. (2025, May). Stablecoin self-regulation. Mimeo.

Carletti, E., I. Goldstein, and A. Leonello (2020). The interdependence of bank capital and liquidity. *BAFFI CAREFIN Centre Research Paper* (2020-128).

Cecchetti, S. and A. Kashyap (2018). What binds? Interactions between bank capital and liquidity regulations. *The Changing Fortunes of Central Banking*, 192–202.

Chernenko, S. and A. Sunderam (2016, July). Liquidity transformation in asset management: Evidence from the cash holdings of mutual funds. *NBER Working Papers 22391.*

Cipriani, M. and G. La Spada (2020). Sophisticated and unsophisticated runs. *Federal Reserve Bank of New York Staff Reports 956*. Revised April 2024.

Couaillier, C., M. Lo Duca, A. Reghezza, and C. Rodriguez D'Acri (2025). Caution: do not cross! Distance to regulatory capital buffers and corporate lending in a downturn. *Journal of Money, Credit and Banking 57*(4), 833–862.

Coval, J. and E. Stafford (2007). Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics 86*(2), 479–512.

Dang, T. V., G. Gorton, and B. Holmström (2020). The information view of financial crises. *Annual Review of Financial Economics 12*(1), 39–65.

d'Avernas, A., V. Maurin, and Q. Vandeweyer (2026). Can stablecoins be stable? *Management Science*. Forthcoming.

d'Avernas, A. and Q. Vandeweyer (2024). Treasury bill shortages and the pricing of short-term assets. *The Journal of Finance 79*(6), 4083–4141.

European Banking Authority (2024). Draft regulatory technical standards to further specify the liquidity requirements of the reserve of assets under Article 36(4) of Regulation (EU) 2023/1114. Final report.

Financial Stability Board (2023). High-level recommendations for the regulation, supervision and oversight of global stablecoin arrangements: final report.

Fukker, G., M. Kaijser, L. Mingarelli, and M. Sydow (2022). Contagion from market price impact: a price-at-risk perspective. *ECB Working Paper Series 2692*.

Gertler, M. and N. Kiyotaki (2010). Financial intermediation and credit policy in business cycle analysis. *Handbook of Monetary Economics 3*, 547–599.

Goldstein, I. and A. Pauzner (2005). Demand–deposit contracts and the probability of bank runs. *Journal of Finance 60*(3), 1293–1327.

Gorton, G. and G. Pennacchi (1990). Financial intermediaries and liquidity creation. *The Journal of Finance 45*(1), 49–71.

Gorton, G. B., E. C. Klee, C. P. Ross, S. Y. Ross, and A. P. Vardoulakis (2025). Leverage and stablecoin pegs. *Journal of Financial and Quantitative Analysis*, 1–65.

Gorton, G. B. and J. Y. Zhang (2023). Taming wildcat stablecoins. *U. Chi. L. Rev. 90*, 909.

He, Z. and W. Xiong (2012). Rollover risk and credit risk. *Journal of Finance 67*(2), 391–430.

Jiang, H., D. Li, and A. Wang (2021). Dynamic liquidity management by corporate bond mutual funds. *Journal of Financial and Quantitative Analysis 56*(5), 1622–1652.

Kara, G. I. and S. M. Ozsoy (2020, June). Bank regulation under fire sale externalities. *The Review of Financial Studies 33*, 2554–2584.

Li, Y. and S. Mayer (2026). Financial intermediation with fragile seigniorage: a dynamic model of stablecoin issuers. Mimeo.

Liao, G., D. Fishman, and J. Fox-Geen (2024). Risk-based capital for stable value tokens. Mimeo.

Liu, J., I. Makarov, and A. Schoar (2023, April). Anatomy of a run: the Terra Luna crash. *NBER Working Papers 31160*.

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

Ma, Y., K. Xiao, and Y. Zeng (2022). Mutual fund liquidity transformation and reverse flight to liquidity. *The Review of Financial Studies 35*(10), 4674–4711.

Ma, Y., Y. Zeng, and A. L. Zhang (2025). Stablecoin runs and the centralization of arbitrage. *NBER Working Papers 33882*.

Morris, S., I. Shim, and H. S. Shin (2017). Redemption risk and cash hoarding by asset managers. *Journal of Monetary Economics 89*, 71–87.

Morris, S. and H. S. Shin (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review 88*(3), 587–597.

Myers, S. C. and N. S. Majluf (1984). Corporate financing and investment decisions when firms have information that investors do not have. *Journal of Financial Economics 13*(2), 187–221.

Rochet, J.-C. and X. Vives (2004). Coordination failures and the lender of last resort: Was bagehot right after all? *Journal of the European Economic Association 2*(6), 1116–1147.

Routledge, B. and A. Zetlin-Jones (2021). Currency stability using blockchain technology. *Journal of Economic Dynamics and Control 128*, 104155.

Shleifer, A. and R. Vishny (2011). Fire sales in finance and macroeconomics. *Journal of Economic Perspectives 25*(1), 29–48.

Stein, J. C. and J. Wallen (2025). The imperfect intermediation of money-like assets. *The Journal of Finance 80*(6), 3185–3221.

Zeng, Y. (2017). A dynamic theory of mutual fund runs and liquidity management. *European Systemic Risk Board Working Papers Series 42*.

Annex A Closed-form model

This annex provides additional derivations and results for the closed-form model presented in Section 2.1.

Lemma 1. *Let* $r_t \sim \mathcal{N}(\gamma r_{t-1}, \sigma^2)$ *and fix* $s_t > 0$, $c_t \in \mathbb{R}$. *Then*

$$\mathbb{E}_{r_2|r_1}\Big[\max\{0,\ -s_1\, r_2 - c_1\}\Big] = -(c_1 + s_1\, \gamma\, r_1)\Big(1 - \Phi(z)\Big) + s_1\, \sigma\, \phi(z), \qquad z \equiv \frac{c_1 + s_1\, \gamma\, r_1}{s_1\, \sigma}, \tag{20}$$

where Φ *and* ϕ *are the standard normal CDF and PDF. Moreover,*

$$\frac{\partial}{\partial c_1}\, \mathbb{E}_{r_2|r_1}\Big[\max\{0,\ -s_1\, r_2 - c_1\}\Big] = \Phi(z) - 1\ \in\ (-1, 0). \tag{21}$$

Proof. Define $W = s\, r + c \sim \mathcal{N}(\mu_w, \sigma_w^2)$ with $\mu_w = s\, r + c$, $\sigma_w^2 = s^2\sigma^2$. Then

$$\mathbb{E}\Big[\max\{0, -W\}\Big] = \int_{-\infty}^{0} (-w) f_W(w)\, dw = -\,\mathbb{E}\Big[W\, \mathbf{1}\{W \leq 0\}\Big].$$

Using truncated normal moments, $\mathbb{E}\Big[W\, \mathbf{1}\{W \leq d\}\Big] = \mu_w \Phi\Big(\frac{d-\mu_w}{\sigma_w}\Big) - \sigma_w\, \phi\Big(\frac{d-\mu_w}{\sigma_w}\Big)$ for any cutoff d, hence with $d = 0$,

$$-\mathbb{E}\Big[W\, \mathbf{1}\{W \leq 0\}\Big] = -\mu_w\, \Phi\Big(-\frac{\mu_w}{\sigma_w}\Big) + \sigma_w\, \phi\Big(-\frac{\mu_w}{\sigma_w}\Big)\,.$$

Symmetry gives $\Phi(-z) = 1 - \Phi(z)$ and $\phi(-z) = \phi(z)$, yielding (20) with $z = \mu_w/\sigma_w = (c_1 + s_1\gamma r_1)/(s_1\sigma)$. Differentiating w.r.t. c_1 and using the chain rule gives (21). ■

State-space partition for the optimal c_1 **given by** (13) Figure 2 visualises the issuer's optimal response to flows at date 1. In this part of the annex, we derive the four phases that give rise to the "shark fin" approach for cash.

Phase I: Interior hits $c_1 = 0$ **(large subscriptions)** By (9), the interior target cash truncates at the non-negativity boundary $c_1 = 0$ when

$$s_0(1 + r_1)\, (\sigma z^\star - \gamma r_1) = 0 \quad \Rightarrow \quad r_1 = \frac{\sigma z^\star}{\gamma}. \tag{22}$$

Since $s_0(1+r_1) > 0$ requires $r_1 > -1$, the economically relevant crossing is $r_1 = \sigma z^\star/\gamma$. For $r_1 \geq \sigma z^\star/\gamma$, the interior rule sets $c_1 = 0$. Intuitively, when r_1 is large and $\gamma > 0$, the issuer expects subscriptions in the next period, so the option value of cash is low, driving c_1 down to zero.

Phase II: Interior feasible (no cap binding) When $c_1^\star(r_1)$ is strictly between 0 and $c_{\text{cap}}(r_1)$, the Case B interior solution in (9) is implementable. Monotonicity follows from $dc_1^\star/dr_1 = s_0\,(\sigma z^\star - \gamma - 2\gamma r_1)$. Thus, $c_1^\star(r_1)$ decreases in r_1 for $r_1 > (\sigma z^\star - \gamma)/(2\gamma)$ and, with $\gamma > 0$, it is concave with second derivative $d^2 c_1^\star/dr_1^2 = -2s_0\gamma < 0$.

Phase III: No-discretionary-selling cap binds (small to moderate redemptions) The interior $c_1^\star(r_1)$ exceeds the cap when:

$$c_1^\star(r_1) - c_{\text{cap}}(r_1) = s_0(1+r_1)\,(\sigma z^\star - \gamma r_1 - 1) - (k_0 + (\mu - 1)b_0) > 0. \tag{23}$$

Expanding the quadratic in r_1 leads to the expression for the real roots:

$$r_1^{\text{cap},\pm} = \frac{(\sigma z^\star - 1 - \gamma) \pm \sqrt{(\sigma z^\star - 1 - \gamma)^2 + 4\gamma\left[(\sigma z^\star - 1) - (k_0 - (1-\mu)b_0)/s_0\right]}}{2\gamma}. \tag{24}$$

For $\gamma > 0$ and a positive discriminant, $c_1^\star(r_1) - c_{\text{cap}}(r_1) > 0$ on the open interval $(r_1^{\text{cap},-},\, r_1^{\text{cap},+})$. In this phase, optimality places b_1 at the boundary $b_1 = b_0$ (no discretionary selling), and sets cash at the cap $c_1^\star(r_1) = c_{\text{cap}}(r_1)$ for redemptions not exceeding the forced-selling boundary, $r_1 \geq r_{\text{FS}}$, as defined in (11). In this phase, the precautionary motive (high c_1) would require $b_1 < b_0$. But since discretionary selling (i.e. sales of bonds when cash is still available) is suboptimal, the issuer stays at the boundary $b_1 = b_0$ and uses cash immediately to meet date 1 redemptions, letting c_1 adjust linearly with r_1 via $c_{\text{cap}}(r_1)$.

Phase IV: Forced selling, cash hits zero (large redemptions) If $b_1 = b_0$ violates $c_1 \geq 0$, i.e. for $r_1 < r_{\text{FS}}$ in (11), feasibility forces sales to restore $c_1 = 0$ so that $c_1(r_1) = 0$ for $r_1 < r_{\text{FS}}$. Extreme outflows at date 1 deplete balance-sheet room and the issuer must sell bonds to meet cash non-negativity, resulting in $c_1 = 0$ and maximal feasible b_1.

Annex B Numerical solution algorithm

We solve the issuer's problem using backward induction. We first solve the date-2 problem where—by construction—any redemption is first met with cash and then with bonds. This gives a date-2 capital schedule $k_2(r_2)$ that depends on date-1 decisions. Next, we discretise the flow shock r_2 using 501 equally spaced grid points in the range $[-0.4, 0.4]$. We then solve the date-1 problem for every possible r_1 shock and state vector $[k_0, c_0]$. While r_1 is discretised similarly as r_2, the state vector $[k_0, c_0]$ is discretised using 101 points along each dimension. The mid-points of these grids are chosen to match the calibration targets—i.e. the k_0 and c_0 grids are centered around 0.1 and 11.7 respectively. Choosing more focused grids (instead of one that spans the entire feasibility set) helps one arrive at an accurate solution more quickly. All computations are done in MATLAB using parallel processing.

Finally, we solve the date-0 problem by computing the total three-period value corresponding to each possible choice of k_0 and c_0. This involves computing the expected values on date 1 and date 2. To compute expectations, we discretise the probability distributions using the grids for r_1 and r_2. Integrals of the density function over symmetric intervals around the grid points are used to estimate probability weights of each grid point (we ensure that the weights add up to 1). The weights are properly adjusted whenever the balance-sheet choices entail a breach of the regulatory thresholds. When computing the expected values, we set capital to zero in case of issuer default, and assume that the balance sheet is frozen i.e. no further bond sales take place. Finally, we solve for the optimal $[k_0, c_0]$, one that delivers the maximum value for the issuer. Subsequently, we trace the issuer's date-1 cash decision as a function of the flow shock: $c_1(r_1)$.

Annex C Price impact estimation

This annex documents the data sources, measurement choices, and aggregation steps underpinning our construction of total flow-induced sales by US Treasury money market funds (MMFs), x_t, as used in Section 3.2.

We assemble a sample comprising all 56 MMFs (including Institutional US Treasury MMFs) in the LSEG Lipper Fund Research Database that were active between January 2020 and December 2025, with aggregate total net assets (TNA) amounting to USD 2.3

trillion at end-2025. Funds in this classification allocate at least 99.5% to cash, Treasury repurchase agreements and short-term US Treasury instruments, which closely mirrors the asset composition of major fiat-backed stablecoin reserves. For each fund, we collect daily TNA and monthly security-level holdings across 1,967 securities. Fund-level weekly net flows are inferred from daily returns and changes in TNA, consistent with standard practice in the mutual fund literature.

We construct weekly outflow-induced sales at the security–fund level by combining observed outflows with lagged portfolio weights (e.g. Coval and Stafford [2007]; Jiang et al. [2021]; Ma et al. [2022]). Specifically, for security j held by fund f in week t, we define

$$\textit{Outflow-induced sale}_{j,f,t} \;=\; \textit{Outflow}_{f,t} \times \textit{Holdings share}_{j,f,t(m-1)},$$

where $\textit{Outflow}_{f,t}$ is the dollar outflow of fund f in week t and $\textit{Holdings share}_{j,f,t(m-1)}$ is the share of security j in fund f's portfolio in the previous month $m-1$ (the highest available frequency for holdings). The implicit assumption is that MMFs meet one dollar worth of outflows by selling securities in proportion to their portfolio holdings. To sharpen our analysis, we include only securities with a remaining maturity of 9 to 12 weeks as of week t (in total, 892 securities held by 48 MMFs), consistent with our focus on measuring the price impact for 3-month Treasury bills. Aggregating across all funds that held a given security in $m-1$ yields total outflow-induced sales of that security in week t, and summing across all such securities produces total weekly flow-induced sales, x_t, as defined in (15).

We complement x_t with several weekly variables used in the price impact estimation reported in Table 3 and provide the corresponding summary statistics and variable definitions in Table 4.

	Mean	Stdev	Min	P10	P25	P50	P75	P90	Max	N
Outflow-induced MMF sales ($ bn)	0.39	0.33	0.00	0.07	0.16	0.30	0.52	0.80	1.84	312
−1×Return on 3-month T-bills (bps)	0.16	1.87	−18.00	−0.76	−0.26	0.04	0.42	1.94	9.52	312
−1×Return on 1-month T-bills (bps)	0.05	0.95	−5.24	−0.36	−0.09	0.00	0.10	0.73	7.35	312
MOVE (Δ ln)	0.00	0.08	−0.39	−0.08	−0.05	−0.00	0.04	0.10	0.30	312
Reverse repo agreements RRP (Δ ln)	−0.01	1.12	−6.62	−0.54	−0.10	−0.00	0.07	0.61	5.12	308
Aggr. MMF outflows (Δ ln)	0.00	0.29	−0.84	−0.39	−0.19	0.00	0.21	0.34	0.76	312

Table 4: Summary statistics of the variables used in the price impact estimation (see Table 3). Outflow-induced MMF sales are defined as in equation (15) and based on the imputed sales of 48 MMFs using data from LSEG Lipper Fund Research Database and Bloomberg. The return on three-month (one-month) Treasury bills (T-bills) is the week-on-week change in the average weekly price of synthetic three-month (one-month) T-bills inferred from corresponding yields (St. Louis FRED). MOVE is the log change in the weekly average of the ICE BofA US Bond Market Option Volatility Estimate Index. RRP is the log change in the weekly average of the Federal Reserve's outstanding overnight reverse repurchase agreements. Aggregate US MMF outflows is the log change in weekly gross outflows of all 3,224 US MMFs in the EPFR iMoneyNet database.

Annex D Stablecoin liquidity and capital ratios

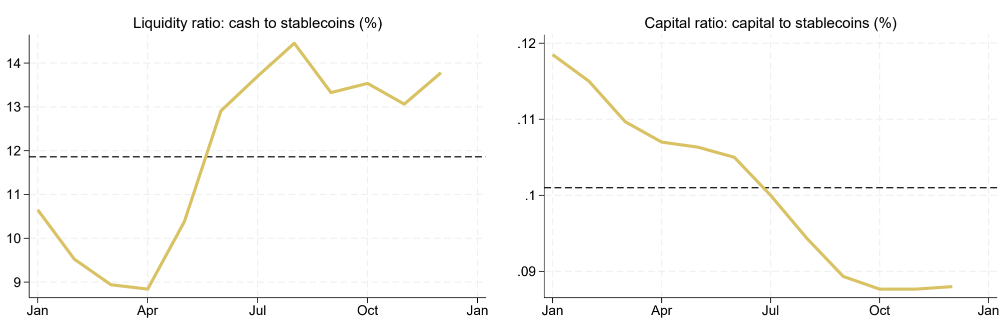

Figure 14: The three-month moving average liquidity and capital ratios of the target stablecoin issuer, namely USDC, during 2025. The dashed lines show the mean ratios during 2025. Source: Bloomberg and company disclosures.

Annex E Alternative bond returns and sale costs

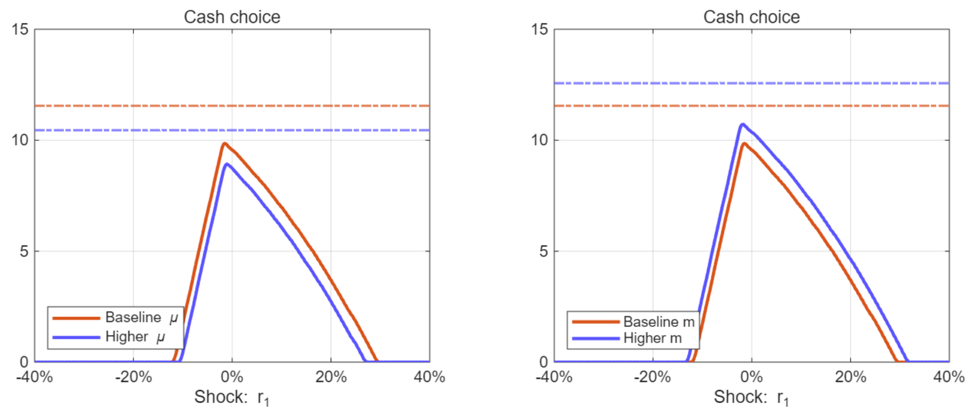

Figure 15: Left-hand panel: The impact of higher bond return on the issuer's cash choice. Baseline $\mu = 7.67$bps. Higher $\mu = 10$bps. Right-hand panel: The impact of higher maximum price impact of bond sales on the issuer's cash choice. Baseline $m = 41$bps. Higher $m = 50$bps. Solid lines show the date-0 cash choice while dashed lines show the date-1 cash choice.

Annex F Risky bond returns and cash deposits

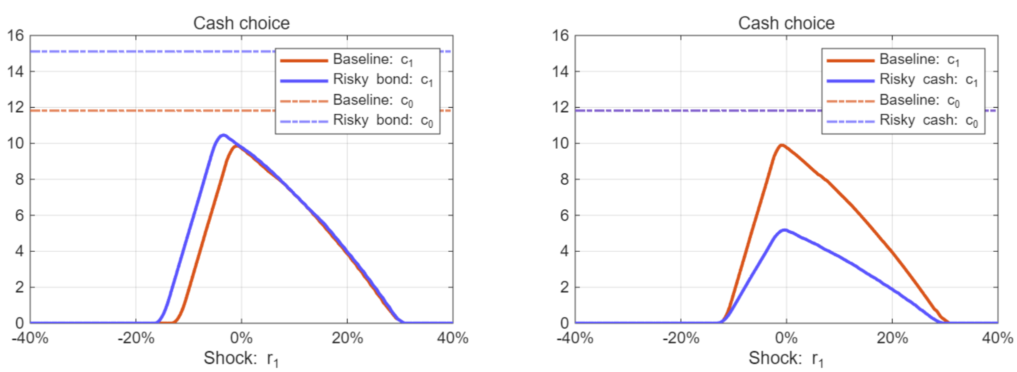

Figure 16: Left-hand panel: The impact of risky date-2 bond returns on the issuer's cash choice. Compared to a fixed baseline $\mu = 7.67$bps, we consider a mean-preserving spread, allowing μ to take two values $\mu_L = 0$ and $\mu_H = 15.3$bps with equal probability. Right-hand panel: The impact of risky cash availability on date 2 on the issuer's cash choice. Specifically, compared to the baseline where cash deposits are always fully available, in the counterfactual we assume a 20% probability that only 50% of cash is usable on date 2. Solid lines show date-0 cash choice while dashed lines show date-1 cash choice.

Annex G Alternative capital costs and flows

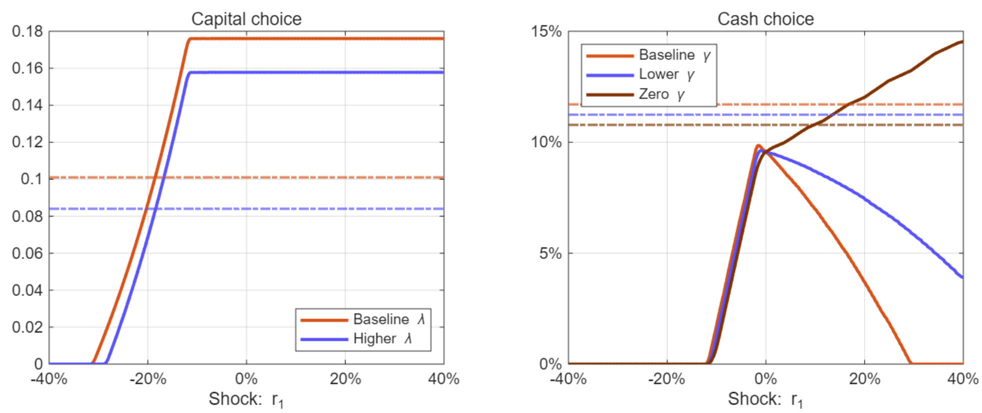

Figure 17: Left-hand panel: The impact of higher cost of raising capital on the issuer's capital choice. Baseline $\lambda = 9.23$. Higher $\lambda = 12$. Right-hand panel: The impact of lower persistence in flows on the issuer's cash choice. Baseline $\gamma = 0.36$. Lower $\gamma = 0.2$. Solid lines show date-0 cash choice while dashed lines show date-1 cash choice.

Annex H Infinite-horizon model

A three-period model implies a limited planning horizon. This makes the stablecoin issuer less forward-looking as compared to a scenario with more periods. To assess the implications of such a setting, we extend the three-period model to an infinite-horizon one. We cast the issuer's problem recursively, with W being the value of the stablecoin franchise as a function of four state variables: the current-period redemption/subscription shock, r; the previous-period values of the cash-to-stablecoin ratio, c_{s-1}; the capital-to-stablecoin ratio, k_{s-1}; and the stablecoins outstanding, s_{-1}:[30]

$$W(r, c_{s_{-1}}, k_{s_{-1}}, s_{-1}) = \max_{b} \quad \mathcal{H}(k) + \beta \mathbb{E} W(r', c_s, k_s, s)$$

$$\text{s.t.} \quad s = s_{-1}(1+r), \quad k = k_{-1} + \mu b_{-1} - b_{-1}\, g\Big(\max\{0, b_{-1} - b\}\Big), \quad c = s + k - b,$$

$$\text{where} \quad k_s = k/s, \quad c_s = c/s, \quad r' \sim f(r'|r).$$

[30]Compared to the three-period model where the state variables are expressed in levels, in the infinite-horizon model it helps express cash and capital as ratios relative to stablecoins. This is because the relevant grids for cash and capital in ratio terms do not depend on the level of stablecoins, while in level terms they do.

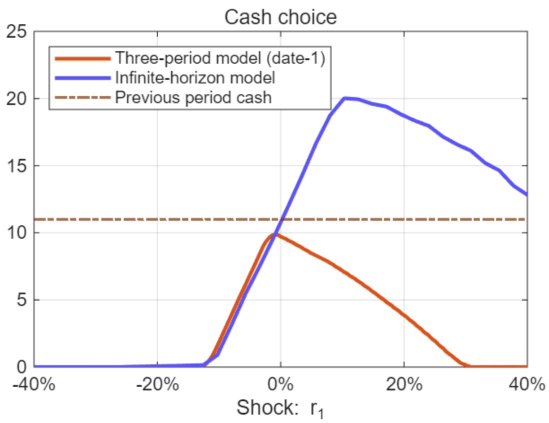

Figure 18: Comparing the cash policy of the stablecoin issuer on date 1 of the three-period model with that in the infinite-horizon model for a starting state that is the same as the date-0 cash and capital ratios in the three-period model.

For ease of comparison, we maintain the same assumptions as in the three-period model, with the exception that capital cannot be raised externally—neither at date 0 as in the three-period model, nor at any other date—and can only grow via retained bond revenues. We solve the infinite-horizon problem using a global solution method (value function iteration) and assess the cash policy of an issuer that has the same cash and capital ratios as on date 0 in the three-period model.

We find that the issuer holds more cash in the infinite-horizon model in the case of subscriptions but similar amounts of cash during redemptions. During redemptions, the issuer is under pressure to use the cash to meet the outflows in order to avoid costs associated with bond sales. By contrast, during subscriptions, the issuer has more flexibility to choose the share of cash versus bonds. It uses this flexibility to build a cash buffer as an insurance against future redemptions. Notably, cash holdings peak for intermediate subscriptions, after which the incentive to hold cash begins to decline as the likelihood of subscriptions in the next period becomes materially small. Overall, while the cash profile is qualitatively similar to the three-period case, a longer horizon induces the forward-looking issuer to secure more insurance.

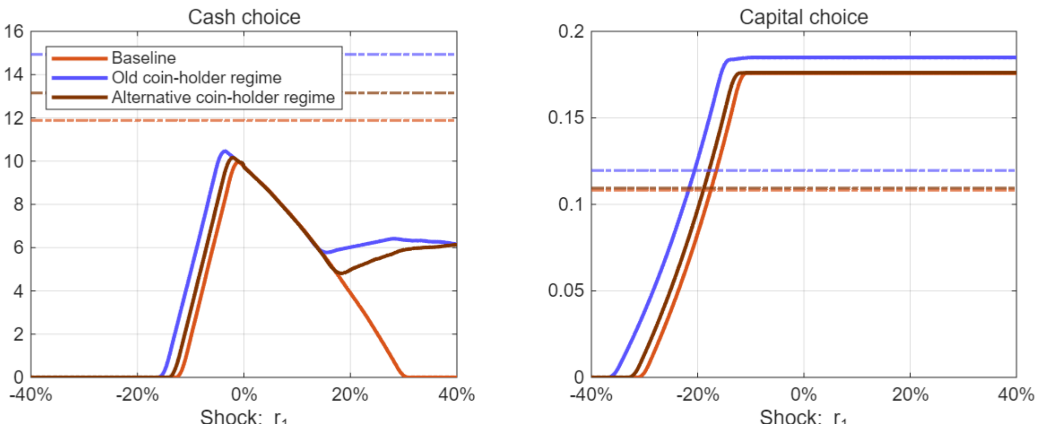

Figure 19: Left-hand panel: Cash choice in the baseline versus different coin-holder disciplining regimes. Right-hand panel: Capital choice. In both panels, the LR threshold is set at 5% and the CR threshold at 0.125%.

Annex I Variations in coin-holder discipline

This Annex discusses the impact of variations in how coin-holders respond to breaches in the regulatory thresholds.

First, we consider a scenario where the coin-holders assess the two thresholds jointly rather than separately. For example, an issuer with 100 percent cash reserves may not need any capital. If coin-holders appreciate this complementarity, they should not penalise an issuer holding only cash but no capital, even if this implies a violation of the capital threshold. To this end, we allow violations of one threshold to be compensated by buffers relative to the other. That is, the effective violation is set to $\nu = \max\{\nu^{LR}, \nu^{CR}\}$ in case both thresholds are breached (similar to the workhorse model) and to $\nu = min\{1, max\{0, \nu^{LR}+\nu^{CR}\}\}$ otherwise. The amended regime loosens coin-holder discipline and leads the issuer to hold less cash and capital on both dates 0 and 1 as compared to the baseline model (Figure 19).

Second, we consider the case of hard thresholds: the issuer is considered defunct upon violating a regulatory threshold. To illustrate the implications of this regime, we consider an LR threshold of 15% and assume that the regulator imposes a significant monetary penalty if the issuer's cash ratio falls below this threshold on either date 0 or date 1. In response, the issuer raises its cash holding on date 0 (relative to the baseline) and holds a buffer relative to the LR threshold (Figure 20). On date 1, the issuer positions itself exactly at the threshold (unless doing this leads to default, as reflected in the discontinuous drop of cash to zero in

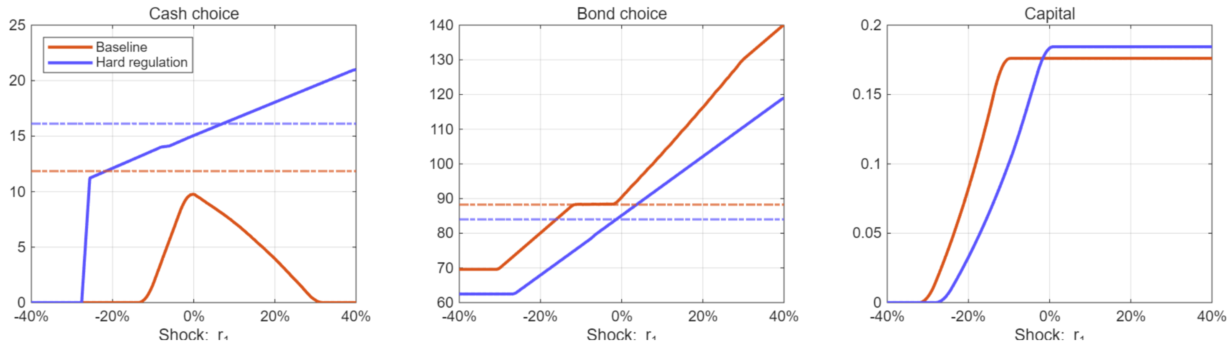

Figure 20: Cash, bond and capital choices in the baseline regime compared with one where an LR threshold of 15% is imposed as a hard minimum requirement.

the first panel). The middle panel shows that the cash-first-bonds-next principle no longer applies. The issuer has no cash buffer on date 1, so it must resort to bond sales immediately when facing a redemption (blue line has no flat part in the case of small redemptions, unlike the red line). This leads to a higher probability of default on date 1 as the issuer runs out of capital faster (third panel). This is despite the issuer raising more capital on date 0. Overall, hard thresholds may not improve the resilience of the issuer.

Annex J Outflow from USDC in March 2023

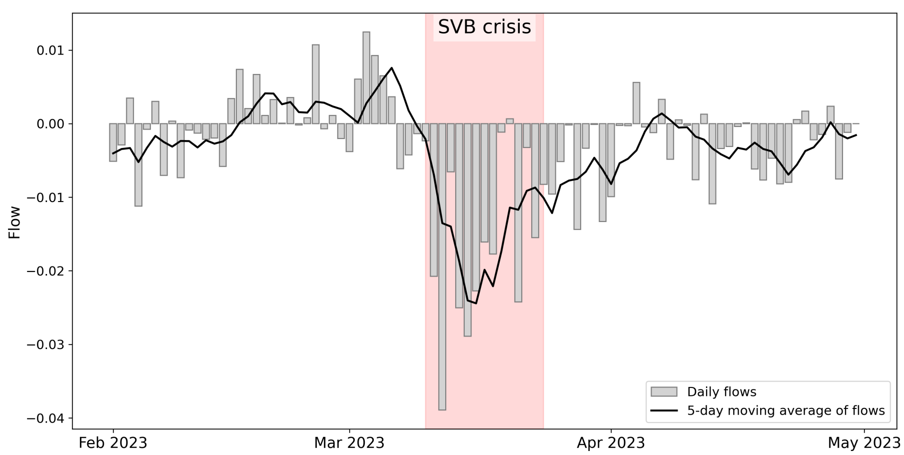

Figure 21: Daily USDC flows, as a share of the previous day's coins outstanding, during the Silicon Valley Bank (SVB) crisis.