Making stablecoins stable(r): Can regulation help?

Tirupam Goel, Ulf Lewrick, and Isha Agarwal*

This version: June 13, 2025

Abstract

Stablecoin issuers typically allocate their reserve assets between interest-bearing bonds and cash, with the latter serving as insurance against redemptions. How can issuers balance the inherent trade-off of having to meet current redemptions versus preserving liquidity for future redemptions, while also remaining sufficiently capitalised and profitable? Our model demonstrates that for small-scale redemptions, issuers primarily rely on cash reserves to minimise bond-sale costs. Conversely, larger redemptions necessitate increasing reliance on bond sales to conserve cash, though such sales can exert downward pressure on bond prices and cause broader market disruptions. What regulatory measures can mitigate these risks? We find that both liquidity and capital requirements can reduce the likelihood of stablecoin defaults and their spillover effects. However, their comparative advantages differ: liquidity requirements are more effective in curbing market spillovers, while capital requirements are better suited to lowering default probabilities. Using data on stablecoin redemptions and U.S. Treasury market price dynamics, our calibration suggests that a regulatory framework combining a 1% capital requirement with a 30% liquidity requirement could limit an issuer's default probability to around 0.05 basis points per year and restrict the expected price impact of bond sales in the Treasury market to around 65 basis points.

JEL Codes: G2, G28, C6

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

*Tirupam Goel (tirupam.goel@bis.org; corresponding author) is at the Bank for International Settlements (BIS). Ulf Lewrick (ulf.lewrick@bis.org) is at the BIS and the University of Basel. 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 and Ilaria Mattei for their help in procuring the data for this project. The authors thank Juliane Begenau, Christopher LC Boedihardjo, Paolo Cavallino, Jon Frost and seminar participants at the BIS for helpful comments. The views expressed in this paper are those of the authors and not necessarily those of the Bank for International Settlements.

1 Introduction

Fiat-backed stablecoins strive to be money-like assets that can be redeemed at par on demand. However, to generate revenue, issuers not only hold cash but also back their stablecoins with interest-bearing assets. The inherent liquidity mismatch exposes issuers to the risk that large redemptions trigger fire-sale of reserve assets (e.g. U.S. Treasuries), depress prices in core financial markets, and lead to knock-on effects. The rapid rise in their global market capitalisation—from $120 billion to $200 billion during 2024 alone—indicates that stablecoins have been deepening their footprint in the underlying asset markets, most notably those for short-dated U.S. Treasuries and associated repo transactions. These developments have placed a spotlight on the financial stability risks posed by stablecoins and raise a natural question: how to regulate stablecoins?

Several jurisdictions have pushed forward dedicated regulatory frameworks, such as the European Union's Markets in Crypto-Asset Regulation (MiCAR). However, the lack of analytical frameworks to inform the specific design of prudential regulation has challenged policymakers.[1] This paper aims to help fill this gap.

We develop a dynamic model of a stablecoin issuer to assess its liquidity management in response to redemption and subscription risks. We then study how the issuer would react to the introduction of prudential requirements. Finally, we assess how different prudential requirements interact and how they map to regulatory objectives, such as a tolerance level for the issuer's probability of default (PD).

The model considers a stablecoin issuer managing reserves composed of cash and interest-bearing bonds, the latter subject to liquidation costs (e.g. Shleifer and Vishny [2011]).

[1]MiCAR provides a high-level regulatory framework for stablecoins, with many regulatory technical standards still in the process of implementation and expected to evolve over time. Similar frameworks have been introduced in Singapore and Hong Kong, among others. Meanwhile, legislative efforts in the United States are underway.

The issuer faces persistent subscription and redemption shocks, leading to several trade-offs. The first concerns the allocation of reserves: how much cash versus bonds to *hold* ex-ante. While cash provides insurance against redemptions, it comes at the cost of forfeiting bond revenues. The second trade-off arises during redemptions: how much cash to *use* to meet current redemptions while preserving liquidity for future needs. Additionally, the issuer holds capital to absorb potential losses, mitigating default risk when liquidation costs deplete reserves.

The stablecoin issuer is subject to liquidity and capital requirements, which draw inspiration from ongoing regulatory discussions.[2] The liquidity requirement (LR) is defined as the ratio of cash to stablecoins issued, while the capital requirement (CR) is defined as the ratio of capital to stablecoins issued. The CR effectively serves as a minimum overcollateralisation requirement, stipulating the extent to which stablecoins must be backed by more than 100% of their value.

The objectives of these requirements are both micro-prudential and macro-prudential. On the micro-prudential side, they aim to ensure the stability and safety of stablecoins for their users. On the macro-prudential side, they seek to safeguard the broader financial system from potential spillovers.[3]

The issuer is subject to coin-holder scrutiny, akin to investor discipline. Specifically, the further the issuer falls short of the regulatory requirements, the greater the redemption risk due to a loss of confidence among coin holders. This dynamic introduces an additional layer of market discipline, as issuers must weigh the potential consequences

[2]For instance, in Europe, liquidity requirements for stablecoins are set between 30% and 60%, depending on the significance of the issuer. In Hong Kong, a capital requirement ranging from 1% to 2% has been considered.

[3]These objectives are analogous to those underlying prudential regulation for banks (e.g., Van den Heuvel [2008]). However, prudential regulation for stablecoins cannot fully resolve the "singleness" issue, i.e. the risk that stablecoins may not always trade at par value (Garratt and Shin [2023]). Nonetheless, by reducing liquidity risk, prudential regulation can enhance coin holders' confidence and mitigate the frequency and severity of deviations from parity.

of eroding trust among their coin holders.[4]

In addition to market discipline, regulatory violations attract fines. The two mechanisms combined pose a decision-making problem for the issuer: how to position itself relative to the requirements during *good* times and how to utilise its buffers (or even dip below the requirements) during *bad* times.

Our model demonstrates that the issuer seeks to build LR buffers during periods of stablecoin subscriptions. The size of the buffer increases with the magnitude of subscriptions, peaking at moderate subscription levels. For very large subscriptions, however, the need for a substantial buffer diminishes due to the expected persistence of subscriptions. In other words, the likelihood that an issuer experiencing a large subscription today will face a correspondingly large redemption tomorrow—justifying a substantial buffer—is relatively low. As a result, the issuer opts for a smaller LR buffer and allocates more reserves to bonds to benefit from higher returns.

The issuer's considerations become more intricate in the case of redemptions. For small redemption shocks, the issuer primarily relies on cash to fulfill redemption requests, thereby avoiding the liquidation costs associated with bond sales. This strategy is supported by the expectation that small redemptions today will likely turn into subscriptions in the near future, allowing the issuer to replenish its cash position.

However, for larger redemptions, the issuer must adopt a mixed approach, using both cash and bond sales equitably. While this incurs bond-sale costs, it is necessary to preserve some cash for future needs. In extreme cases, depending on the issuer's ex-ante balance sheet strength, it may be forced to dip below the LR—especially if it lacked a buffer to begin with—or even run out of cash entirely, ultimately defaulting

[4]This assumption is inspired by the design of the Basel III Liquidity Coverage Ratio (LCR), which allows banks to fall below the LCR threshold temporarily. However, there has been significant debate about whether banks actually utilise this flexibility (e.g. Bank of England [2022]). Our emphasis on coin-holder scrutiny is informed by this discussion.

due to large bond-sale losses.[5]

A key advantage of the model, relative to a purely empirical analysis, is its ability to conduct counterfactual policy experiments. We show that both the CR and LR can reduce the issuer's PD and the expected price impact (EPI) of forced bond sales in the next period. While the two requirements act as substitutes, they are jointly necessary to achieve a dual PD-EPI objective. Moreover, each requirement has a distinct comparative advantage: the LR primarily reduces the EPI by limiting the need for bond sales, whereas the CR is more effective at lowering the PD by absorbing bond-sale losses.

The paper further contributes to the policy discussion by deriving a quantitative mapping from policy objectives—expressed in terms of PD and EPI targets—to the corresponding CR and LR levels. To achieve this, we calibrate the redemption process using data on major dollar-pegged stablecoins and estimate the price impact of forced sales of U.S. Treasury securities using money market fund (MMF) data.[6] Our estimates suggest that a 1% CR and a 30% LR correspond to a PD of around 0.05 basis points (bps) per year and an EPI in the U.S. Treasury market of around 65 bps.

Various comparative statics and robustness checks confirm the intuition underlying these results. For instance, we find that larger issuers rely more heavily on cash to meet the same redemption amount, as bond sales become proportionally more costly at larger scales. A higher price impact of bond sales has a similar effect. Conversely, having a larger ex-ante LR buffer means that the issuer has more cash for future needs, thereby reducing both PD and EPI. Furthermore, greater coin-holder scrutiny or higher regulatory fines incentivise issuers to maintain larger LR buffers, further reducing PD and EPI.

[5] While testing the model's predictions in the context of stablecoins is challenging due to the lack of detailed public disclosure of issuers' holdings, similar dynamics are well-documented in the literature on mutual funds. See, for instance, Jiang et al., 2021; Ma et al., 2022.

[6] This approach addresses the limitation that stablecoin asset holdings are typically not disclosed with the same granularity as MMF holdings. It leverages the similarity in balance sheet structures between the two entities, as highlighted by Anadu et al. [2023b].

Our paper contributes to the nascent but vibrant literature on stablecoins, with a particular focus on the risks associated with their operation and regulation. A central theme in this literature is the risk of runs, which poses a significant challenge to the stability of stablecoins.

Ma et al. [2023] examine the interplay between run risk and arbitrage efficiency. They find that while more efficient arbitrage enhances price stability in secondary markets (see also Lyons and Viswanath-Natraj [2023]), it exacerbates run risk by reducing the price impact of coin sales, thereby amplifying coin holders' first-mover advantage during periods of stress. We complement their findings on secondary markets with our analysis of the importance of price dynamics in stablecoin stability through the channel of the underlying reserve asset markets, which accords with recent empirical work by Ahmed and Aldasoro [2025], who document this channel in T-bill markets.

Gorton and Zhang [2023] compare stablecoins to the private banknotes issued by banks in the nineteenth century, arguing that both are prone to run risk and represent an inferior medium of exchange. Similarly, Liu et al. [2023] analyse the collapse of Terra in 2022, attributing the event to growing concerns about the system's sustainability, with public transaction visibility accelerating the run. Routledge and Zetlin-Jones [2021] focus on speculative attacks and coordination failures in under-collateralised stablecoins, highlighting the vulnerabilities stemming from insufficient reserves. Gorton et al. [2025], in turn, investigate how stablecoins maintain a constant price of one dollar despite significant run risks. Bertsch [2023] explores the factors influencing the adoption and fragility of stablecoins. Unlike these studies, our paper emphasises the interaction between regulatory compliance and coin-holder discipline, showing how regulatory requirements can shape redemption risks and issuer behaviour.

Relatively few studies have examined the role of regulation in mitigating stablecoin risks. Li and Mayer [2022] analyse stablecoin management within a general equilibrium framework, finding that capital requirements improve stability but that constraining the riskiness of reserve assets could destabilise prices if issuers reduce leverage to compensate for reduced income. Liao et al. [2024] propose a capital framework to mitigate non-financial risks arising from the underlying technology and operations of stablecoins. Ahmed et al. [2024] point to the ambiguous effect of public disclosures on coin holders' run incentives. Our contribution to this literature lies in the joint assessment of capital and liquidity requirements. By incorporating endogenously determined regulatory buffers and occasionally binding constraints, we show how coin-holder discipline can interact with regulation to influence issuer behaviour. This connection underscores the role of regulation as a coordination device, akin to mechanisms observed in other financial intermediaries, such as banks or MMFs (e.g. Cipriani and La Spada [2020]).

More broadly, our paper is part of the literature on the vulnerabilities of financial intermediaries that create money-like assets while engaging in liquidity transformation. Stablecoin issuers share many similarities with MMFs, which are subject to run risk as investors can redeem shares at a fixed price Kacperczyk and Schnabl, 2013; Schmidt et al., 2016; Sunderam, 2015; Parlatore, 2016. Existing studies on open-ended mutual funds have shown that run-like situations can arise when investors rush to redeem shares in funds holding illiquid assets Chen et al., 2010; Goldstein et al., 2017. Similarly, bank runs have been modeled extensively in Diamond and Dybvig [1983], Goldstein and Pauzner [2005], and Bernardo and Welch [2004], among others.

Building on this literature, our paper develops a framework to analyse the regulatory design for stablecoins, emphasising the interplay between prudential requirements, issuer behaviour, and spillovers into the markets of issuers' reserve assets.

The rest of the paper is organised as follows. Section 2 introduces the model. Section 3 outlines the calibration strategy. The results of the quantitative analysis are presented in Section 4. Finally, Section 5 concludes.

2 Model

2.1 Stablecoin issuer's optimisation

We develop a dynamic model of a centralised stablecoin issuer. Each period, the issuer starts with a balance sheet that consists of two types of assets, a_{-1}: cash, c_{-1}, and bonds, b_{-1}. The issuer's liabilities comprise the issued stablecoins, s_{-1}. The difference between assets and liabilities represents the issuer's capital, k_{-1}, so that the balance sheet identity is given by $b_{-1} + c_{-1} = s_{-1} + k_{-1}$ (Figure 1).[7]

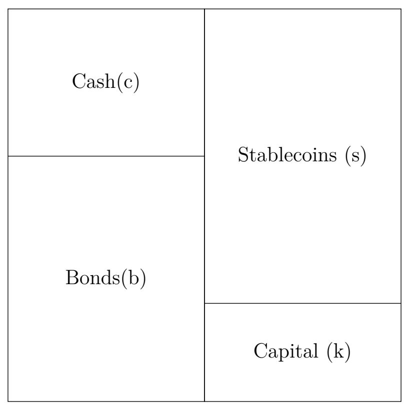

Figure 1: A stylised stablecoin issuer's balance sheet

The issuer faces a trade-off with regard to its asset allocation.[8] While cash can be used to meet redemptions without any liquidation costs, it bears no interest. Bonds, by contrast, yield a stochastic return μ, which takes value $\mu_L < 0$ with probability p_L and $\mu_H > 0$ with probability $p_H = (1 - p_L)$. Bonds can be bought at par. However, selling bonds to obtain cash is subject to a liquidation cost. This cost increases with the amount sold, reflecting a fire-sale discount when selling larger amounts. Specifically, if x assets are sold, the discount is $m(x)$, where $m'(x) > 0$. Therefore, the unit price

[7] We focus on centralised stablecoin issuers, which represent the vast majority of stablecoin market capitalisation (see e.g. Li and Mayer [2022] for a discussion of alternative protocols, referred to as decentralised autonomous organisations (DAOs)). In the context of stablecoins, the issuer's capital (i.e. equity share) is also referred to as "governance tokens" or "secondary units".

[8] In this regard, the stablecoin issuer's dynamic liquidity risk management is reminiscent of the cash management decisions of open-ended mutual funds (e.g. Zeng [2017], Morris et al. [2017]).

received per asset sold is $1 - m(x)$. The overall loss incurred when selling a quantity x is given as $g(x) = m(x)x$, where $g''(x) > 0$.

On the liability side, the demand for stablecoins is exogenous. Each period, a random fraction, r, of the stablecoins are redeemed or subscribed to. We assume this random demand shock follows an AR(1) process: $f(r'|r) \sim \mathcal{N}(\gamma r, \sigma)$, where γ represents the persistence in the shock process and σ is the standard deviation. The serial correlation helps capture the observed momentum in stablecoin redemptions.[9]

Issued stablecoins also represent a source of revenue for the issuer. Unit revenue per stablecoin is denoted by τ.[10]

Finally, the role of capital is to absorb losses and thus protect stablecoin holders. The main source of capital growth is retained earnings. Meanwhile, we also allow the issuer to raise capital externally, k^e, when it is facing subscriptions, but not when facing redemptions.[11] We assume that the issuer defaults when capital falls below zero. Default occurs when, for instance, a particularly large redemption forces the issuer to draw down all its cash and sell bonds to an extent that the associated fire-sale cost cannot be absorbed by the available capital.

Regulation Regulation intends to reduce the issuer's probability of default (PD) and limit the expected price impact (EPI) of bond sales in the face of redemptions

[9]Potential micro-foundations for such momentum have been studied in the case of stablecoins and money market funds. Such momentum is understood to be driven by herding, coordination failures, and incentives to run *first*. Global games are typically used to model such mechanisms. In this paper, instead, we adopt the perspective that what ultimately matters for the stablecoin issuer's decisions is the conditional distribution of the demand shock in the next period given the shock this period, rather than a precise understanding of the origin or underpinnings of the shocks. As such, we focus on estimating the serial correlation in the distribution of demand shocks using stablecoin data rather that modeling the shock process (see Section 3.3).

[10]As in Li and Mayer [2022], this revenue can be thought of as the pecuniary benefit extracted from the transactions data obtained from users of the stablecoin.

[11]The intuition is related to that provided in Myers and Majluf [1984]. When facing subscriptions, the issuer is able to send a signal to its shareholders and ask for more funding for expanding the business. However, it cannot credibly send such a signal when facing redemptions.

(see also Section 4.2). To achieve these objectives, two types of regulatory constraints are imposed on the issuer. The first one is a minimum capital requirement (CR), $k/s \geq \chi$ with $\chi \geq 0$. The requirement ensures overcollateralisation: each stablecoin in circulation is backed by at least $1 + \chi$ units of assets (cash and/or bonds).

The second one is a minimum liquidity requirement (LR), $c/s \geq \theta$ with $\theta \geq 0$. The LR seeks to ensure that the issuer holds at least some cash, given that cash, unlike bonds, can be used to meet redemptions without any liquidation costs.

If the stablecoin issuer fails to meet either requirement, it faces the wrath of stablecoin holders. As cash and/or capital decline below the regulatory requirements, coin-holders impose market discipline on the issuer (e.g. Ahmed et al. [2024]) and the distribution of redemptions shifts to the left. We assume that the intensity of redemptions depends on the degree to which the issuer violates either requirement. For instance, if the issuer violates the LR by choosing a cash-to-stablecoin ratio less than θ, i.e. $c/s < \theta$, then the degree of violation ν_{LR} is given as $1 - c/s/\theta$. Overall violation across the two requirements is given as $\nu = max\{\nu_{LR}, \nu_{CR}\}$. The distribution of demand shock r' for the next period is then given as:

$$f(r'|r,\nu) = \begin{cases} \mathcal{N}(\gamma r, \sigma) & \text{if } \nu = 0, \\ \mathcal{N}\left(\gamma\left(\min(0,r) - \lambda\nu\right), \sigma\right) & \text{if } \nu > 0. \end{cases} \tag{1}$$

Here, λ is the penalty parameter that governs the relationship between regulatory violations and coin-holder discipline. For instance, when $c \geq s\theta$ (i.e. no violation of the LR), no penalty arises. However, when $c = 0$, representing the most extreme violation of the LR, expected redemptions increase by at least $\gamma\lambda\nu$.[12] Additionally, the issuer

[12]Calibrating λ is inherently challenging. To address this, we rely on an informed estimate for λ, as detailed in Section 3.4. To ensure robustness, we also perform comparative statics in Section 4 to evaluate the sensitivity of our results to variations in λ.

incurs a regulatory fine, proportional to both the degree of violation and the size of the issuer, given by $\nu\, s\, \Omega$.[13]

Issuer's decision problem The problem of the issuer is to structure its asset allocation in a way that maximises the current value of its expected return on capital (ROC) in the future. ROC is computed as profits per unit of capital, with profits in the current period being equal to the transaction revenues minus the costs of selling bonds:[14]

$$\Delta \quad = \quad \underbrace{\tau s_{-1}}_{\text{stablecoin revenue}} \quad - \quad \underbrace{\mathbb{1}(b < b_{-1}) g(b_{-1} - b)}_{\text{bond fire-sale cost}}.$$

We cast the issuer's problem recursively, with V being the value of the stablecoin franchise as a function of four state variables: the balance sheet size, a, in the previous period; the cash to stablecoin ratio *relative* to the LR, $c_{s\theta}$; the capital to stablecoin ratio *relative* to the CR, $k_{s\chi}$, and the current period realisation of the demand shock, r.

$$V(a_{-1}, c_{s\theta_{-1}}, k_{s\chi_{-1}}, r) = \max_{c,\, k^e} \quad \mathcal{H}(\Delta/k_{-1}) + \beta \mathbb{E} V(a, c_{s\theta}, k_{s\chi}, r')$$

$$s.t. \quad LR, \quad CR, \quad s = s_{-1}(1+r), \quad b = s + k - c, \quad k = k_{-1} + \Delta + k^e \mathbb{1}(r > 0)$$

$$a = c + b + b\mu' - \Omega\nu s, \quad k_{s\chi} = (k + b\mu' - \Omega\nu s)/s/\chi, \quad c_{s\theta} = (c + b\mu' - \Omega\nu s)/s/\theta$$

$$\mu' \sim (\mu_L, \mu_H; p_L, p_H), \quad r' \sim f(r'|r, \nu)$$

Here $\mathcal{H}$ is a concave function that captures the utility that the issuer derives from a given level of ROC. $\beta\mathbb{E}V(\cdot)$ denotes the discounted expected continuation value. This

[13] Setting the regulatory fine to zero does not alter the qualitative insights of the model as long as coin-holder discipline ensures that the regulatory requirements have traction. We include both coin-holder discipline and regulatory fine in the model for realism, but only one would also suffice.

[14] The return on bonds accrues at the beginning of the next period and is therefore included in the continuation value as shown in the expressions for a, $k_{s\chi}$ and $c_{s\theta}$ in the optimisation problem.

takes account of any regulatory fine, stablecoin demand shocks, stochastic bond returns, as well as the risk of default. We note that a, $k_{s\chi}$ and $c_{s\theta}$ depend on μ' and Ω: while a higher return on bonds improves the cash and capital position of the issuer, a larger fine worsens it.

Discussion of the issuer's trade-offs Each period, for any given ex-ante balance sheet configuration, the issuer needs to decide how to react to redemption and subscription shocks. A large redemption, for instance, raises the following question: to what extent should the redemption be met by drawing down cash or by selling bonds? On the one hand, using cash avoids the liquidation costs associated with using bond sales in the current period, and also helps save the opportunity cost associated with foregone interest on bonds sold. On the other hand, depleting cash not only makes the issuer more vulnerable in case of future redemptions, it can also raise the likelihood of redemptions if the LR is violated in the process. In extreme cases, the issuer may not be able to meet the redemptions even after using all the cash. This can happen when capital is insufficient to cover bond liquidations costs, leading to default.

Past decisions also affect the issuer's considerations. While holding more cash during normal times reduces default risk, it implies forgoing the return on bonds. As such, the issuer needs to constantly balance a trade-off between viability and profitability.

The regulatory requirements introduce additional trade-offs. These requirements clearly incentivise issuers to build regulatory buffers during normal times, ensuring their availability during periods of stress. Such buffers can help mitigate regulatory breaches and, consequently, reduce the likelihood of an increasingly unfavourable redemption scenario. However, determining the optimal size of these buffers and the extent to which they should be utilised remains less straightforward. While buffers offer significant benefits, they also entail costs. For example, maintaining a larger buffer above the LR necessitates holding more cash, which can adversely impact profitability, as previously discussed. At the same time, operating too close to the regulatory limit may elevate the risk of future breaches.

3 Model calibration

In this section, we discuss the calibration of the model. We start with a description of the data sources. We then discuss how the standalone parameters are calibrated. Next, we estimate the stablecoin demand shock process. We conclude this section by estimating the price impact of bond sales.

3.1 Data

We employ three data sources. The first consists of daily market capitalisation and prices for the nine largest fiat-backed stablecoins by circulation as of end-2023, sourced from CoinGecko. This dataset spans the period from 2020 to 2023. By dividing the market capitalisation by the price, we derive the supply of coins in circulation.

Second, we collect stablecoin balance sheet data. This process presents several challenges. For one, not all stablecoins publicly disclose such information on their websites. Among those that do, the reporting practices vary significantly. For example, the data typically begin in 2022 or 2023, with only few extending back to 2020. Additionally, the level of detail in the disclosures varies not only across issuers but also over time for the same issuer. Furthermore, while some issuers release their balance sheets on a monthly basis, others do so quarterly. Given these limitations, we are able to obtain reliable disclosures from only the two largest stablecoins, USD Coin (USDC) and Tether.

Third, we gather data on MMF flows and their holdings of U.S. Treasury securities to estimate the price impact of large outflows. Our analysis focuses on MMFs with

Calibrated parameters	Symbol	Calibrated value
Discount factor	β	0.98
Return on assets (weekly)	μ_L, μ_H	-0.586 bps, 3.668 bps
Probability of low return	p_L	0.500
Stablecoin turnover fee	τ_r	0.001
Demand shock: persistence	γ	0.801
Demand shock: standard deviation	σ	0.059
Coin-holder penalty	λ	0.050
Price impact: $m(x) = m\,x$	m	0.002
Regulatory penalty	Ω	0.100
Regulatory parameters	**Symbol**	**Baseline value**
Capital regulation	χ	1%
Liquidity regulation	θ	30%

Table 1: The top panel lists the calibrated parameters whose values are set using counterparts in the data. The bottom panel lists the baseline value of the regulatory parameters.

portfolio compositions similar to those of major stablecoin issuers, for which detailed security-level data is unavailable. Specifically, we collect data on funds classified under the Lipper Objective code UST. These are US-domiciled MMFs that primarily invest in short-term US government securities, mirroring the investment strategies of major stablecoin issuers. UST funds cater to retail investors and allocate 99.5% of their assets to cash and U.S. Treasury securities with maturities of less than 60 days. We include all funds within this category that were active for at least part of the period from January 2016 to March 2024, resulting in a sample of 35 funds. For each fund, we obtain daily total net assets (TNA) data from Lipper based on which we approximate weekly net flows. Next, we compile the MMFs' monthly holdings of U.S. Treasury securities, resulting in a total of 1,274 securities. Of these, 764 are U.S. Treasury bills, while the remaining 520 are notes and bonds.[15] All securities in the sample have a residual maturity of less than 60 days when held by the MMFs. Finally, for each security, we retrieve daily bid and ask prices or yields from Bloomberg and calculate weekly returns (see Annex A).

[15]U.S. Treasury bills have an original maturity of up to one year, notes have an original maturity of 2 to 10 years, and bonds have an original maturity of more than 10 years.

3.2 Calibration of the standalone parameters

The full list of parameters, along with their calibrated or estimated values, is provided in Table 1. In the model, we assume that one period corresponds to one week in the data, and one unit of asset represents $1 billion in the data.

For certain model parameters, standard values are available in the literature or direct counterparts can be identified in the data, allowing us to set their values accordingly. We refer to these as standalone parameters. Specifically, we set the discount factor β to 0.98. The interest rate parameters μ_L and μ_H are determined using the 25th and 75th percentiles of the distribution of weekly returns on the U.S. Treasury securities in our sample. Based on this, we assume a symmetric distribution, and that each outcome is equally likely.

Next, we calibrate the subscription/redemption fee parameter τ_r. This fee can vary in terms of the level and structure across issuers, which makes their calibration more challenging. For instance, some charge for redemptions beyond a certain limit or apply staggered fees. Some also impose non-price constraints, such as a lag with which redemption requests can be met.[16] As an approximate average of the various levels of the fee, we consider a uniform all-in-fee of 0.01% for τ_r. After all, this fee are not the main source of revenue for stablecoin issuers, and we account for it because fees exist in reality.

3.3 Estimation of persistence in stablecoin redemptions

We estimate the persistence of demand shocks, γ, based on the empirically observed inertia in stablecoin flows. Our calibration sets one period in the model to a week in practice. Accordingly, we estimate the persistence parameter in the conditional

[16] For instance, USDC has no fee for redemptions that are typically fulfilled in two business days. Near-instant redemptions are free up to a certain limit, but there is a fee that increases in steps of 0.03%, 0.06% and 0.1% for higher amounts.

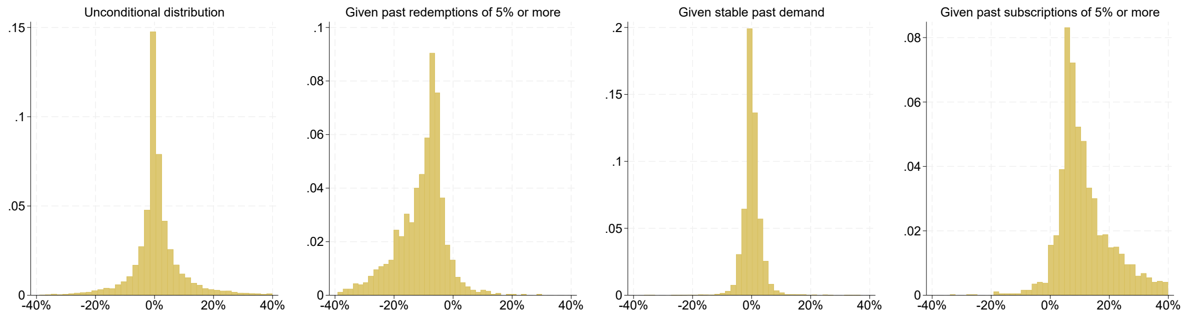

Figure 2: Pooled distribution of the weekly percentage change in the outstanding supply of coins for 9 issuers for the period from the start of 2020 to the end of 2023. The first panel shows the unconditional distribution. The subsequent three panels display the conditional distributions based on net redemptions exceeding 5 percent in the previous week, net redemptions or net subscriptions within 5 percent, and net subscriptions exceeding 5 percent. Source: CoinGecko.

distribution of r using net weekly changes in the amount of stablecoins outstanding.

Weekly flows exhibit substantial inertia. Figure 2 illustrates the pooled distribution of weekly changes (first panel) alongside the conditional distributions (second to fourth panels) categorised by net redemptions exceeding 5 percent in the previous week, net redemptions or net subscriptions within 5 percent, and net subscriptions exceeding 5 percent. When conditioned on a large redemption (subscription) in the previous week, the distribution of redemptions becomes notably skewed to the left (right). To account for this formally, we estimate the following dynamic panel regression:

$$flow_{i,t} = \gamma\, flow_{i,t-1} + \alpha_t + \mu_i + \epsilon_{i,t}, \tag{2}$$

based on the two-step GMM developed in Arellano and Bover [1995] and Blundell and Bond [1998]. $flow_{i,t}$ is the percentage change in the outstanding supply of stablecoin i between t and $t-1$. μ_i and α_t are stablecoin and time fixed effect, respectively.

Flows are highly persistent at a weekly frequency, with $\gamma = 0.801$, which is significant at all usual confidence levels. Persistence is also observable at higher frequencies, with 3-day flows suggesting a value of $\gamma = 0.583$, and it is broadly comparable across the

Model type	Coefficient	Standard error
Pooled, weekly flows (γ)	0.801	(0.023)
Pooled, weekly flows (st.dev.($\hat{\varepsilon}$))	0.059	(0.011)
Pooled, 3-day flows (γ)	0.583	(0.038)
BUSD	0.935	(0.006)
FdUSD	0.872	(0.032)
GeminiD	0.846	(0.009)
HUSD	0.888	(0.005)
PAXD	0.901	(0.008)
USDT	0.891	(0.005)
TrueUSD	0.914	(0.003)
USDC	0.950	(0.005)
ZUSD	0.910	(0.010)

Table 2: AR(1) coefficients (γ) based on Equation (2) for the pooled sample and for regressions based on individual stablecoins, where the coefficients are based on estimating $flow_t = \gamma flow_{t-1} + \mu + \epsilon_t$. The second row reports the standard deviation of the estimated residuals, $\hat{\varepsilon}$, of the pooled regression. This measure is used as a proxy for the standard deviation of the demand shock, σ.

different stablecoins (Table 2).

Role of balance sheet characteristics As discussed in the model section, a deterioration in the liquidity profile of the stablecoin issuer's balance sheet can increase the likelihood of redemptions. This hypothesis has been verified in the case of mutual funds. For instance, Ma et al. [2022] show that outflows are larger at more illiquid fixed-income mutual funds. Cipriani and La Spada [2020] find a similar result in the case of institutional money market mutual funds but less so in the case of retail money market funds as retail investors are less attentive to fund asset compositions. Relatedly, Chen et al. [2010] and Goldstein et al. [2017] find that for more illiquid mutual funds, bad performance is more likely to lead to outflows, especially when aggregate illiquidity is high.

Testing this hypothesis formally for stablecoins is challenging because, as mentioned in the previous section, most issuers do not disclose detailed balance sheet data. For the few that do, the frequency is low or historical coverage is insufficient. These limitations prevent a rigorous estimation of the penalty parameter, λ. As such, we use the collapse of Silicon Valley Bank and its impact on USDC as a guide. The collapse led to severe

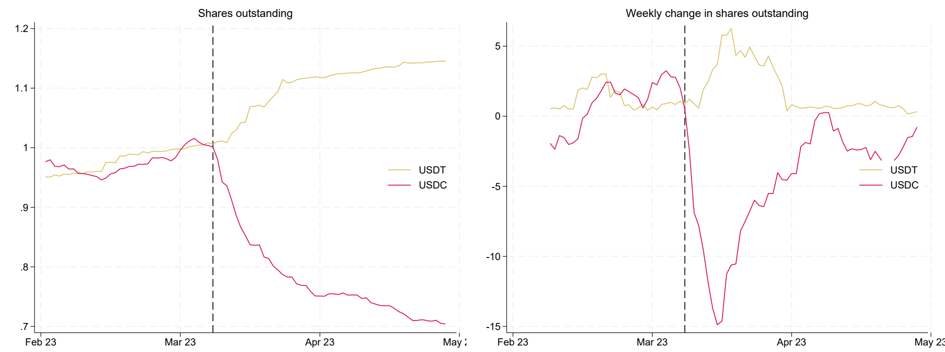

Figure 3: First panel: Shares outstanding before and after the collapse of Silicon Valley Bank. Average shares outstanding in the first week of March is normalised to 1. Second panel: Weekly change in shares outstanding, in percent. Source: CoinGecko.

outflows from USDC while other stablecoins were less affected. Indeed, Tether experienced inflows possibly reflecting a "run to safety" by USDC coin holders (Figure 3). Coin-holders' concerns about blocked USDC's deposits at Silicon Valley Bank and thus Circle's ability to service USDC redemptions likely triggered the redemptions. The episode underscores that perceived balance sheet strength shapes demand shocks (e.g. Ahmed et al. [2024]). USDC faced a redemption of close to 10 percent in the first week and a cumulative redemption of close to 30 percent in the first two months. Based on this observation, we set λ such that redemptions increase by up to 5 percentage points depending on the degree to which the regulatory requirements are violated.

3.4 Calibration of the price impact function

We examine the impact of large, flow-induced sales at MMFs to calibrate the model's price impact function. Stablecoins and MMFs have notable commonalities in their balance sheets and also in the distribution of demand for their respective liabilities (e.g. Anadu et al. [2023a]), suggesting that price effects induced by MMF investors' redemptions can serve as a useful proxy of the corresponding impact of redemptions by stablecoin investors. Moreover, the two leading stablecoin issuers have grown to sizes comparable to major US MMFs, with Tether featuring among the top-3 holders of short-dated U.S. Treasury securities.

To draw comparisons to stablecoins, we estimate how redemption pressure on MMFs transmits into pressure on the prices of the securities sold by these funds. We focus on U.S. Treasury securities, which represent the bulk of the largest stablecoin issuers' asset holdings, using weekly fund data and returns for the period from January 2016 to March 2024. For robustness, we consider two variants of measuring returns: one based on weekly averages of security prices and one based on using prices for the last business day of each week.[17]

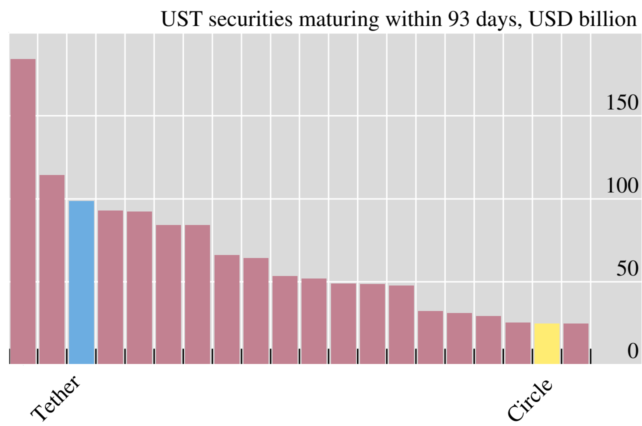

Figure 4: The figure compares the holdings of U.S. Treasury securities with remaining maturity of up to 93 days, the currently proposed maximum maturity in draft US legislative proposals, for the largest US money market funds (by holdings) and stablecoin issuers as of March 2025. Source: EPFR iMoneyNet.

In line with the extant literature on MMFs (e.g. Coval and Stafford, 2007; Jiang et al., 2021; Ma et al., 2022), the selling pressure is assumed to be proportional to the fund redemptions, giving rise to the following measure of flow-induced sales (see

[17]For consistency, we measure flows and T-bill returns based on the analogous approach in the respective regressions.

Annex A for details) :

$$Flow\text{-}induced\ sale_{i,t} = \sum_{f=1}^{F} (Outflow_{f,t} \times Holdings_{i,f,t(m-1)}), \tag{3}$$

where $Outflow_{f,t}$ is the outflow experienced by MMF $f \in F$ in week t and $Holdings_{i,f,t(m-1)}$ is the asset share of security i in fund f's portfolio in the previous month, $m - 1$.[18] Accordingly, *Flow-induced sale* approximates MMFs' total sales of security i in week t that were induced by outflows experienced by all funds that held security i.

Following the literature, we estimate the impact of flow-induced sales on the returns of the securities as follows:

$$Return_{i,t} = \alpha + \beta Flow\text{-}induced\ sale_{i,t} + \delta_i + Return_{4\text{-}week,t} + \varepsilon_{i,t}, \tag{4}$$

where the dependent variable $Return_{i,t}$ is the weekly return on security i at date t. δ_i accounts for security fixed effects, whereas $Return_{4\text{-}week,t}$ controls for market-wide changes in the returns on short-dated US Treasuries (4-week T-bills).

Our estimates suggest that flow-induced sales of \$100 million reduce weekly returns by about 1 to 2 basis points (Table 3). For the calibration, we set the price impact to the upper range, consistent with regulation taking a conservative approach to calibration.[19]

[18] As is standard practice in the literature, we estimate weekly flow-induced sales using weekly data on fund outflows and monthly data on fund holdings since holdings are not available at a higher frequency.

[19] Specifically, we assume that the price impact is linear in the size of the sale, i.e. $m(x) = mx$, and set m equal to 0.002 with $x = 1$ equivalent to selling \$1 billion. The estimated price impact is roughly equivalent to 40% of the standard deviation of weekly returns on U.S. Treasury securities in our sample (see also Table 4 in Annex A). We discuss an alternative specification in Annex C. Ahmed and Aldasoro [2025] estimate the price impact of stablecoin issuers' T-bill purchases. Their findings suggest a five basis point increase in yields for sales of \$3.5 billion over a 5-day horizon. For foreign official flows in US notes, Ahmed and Rebucci [2024] find an impact of 100 basis points on yields for flows of \$100 billion over 30 days.

	Weekly averages			End-of-week prices		
	(1)	(2)	(3)	(4)	(5)	(6)
Flow-induced sale	-1.786^{**}	-0.953^{**}	-1.086^{***}	-1.268^{**}	-0.872^{***}	-0.573^{**}
	(0.736)	(0.390)	(0.341)	(0.596)	(0.295)	(0.262)
Return on UST 4-week bills			1.220^{***}			0.852^{***}
			(0.124)			(0.088)
Observations	27,222	27,212	27,212	27,222	27,212	27,212
Adj. R2	0.001	0.622	0.654	0.000	0.599	0.622
Security fixed effects		Yes	Yes		Yes	Yes

Table 3: This table presents the estimated results from Equation 4. The dependent variable is the weekly return of security i in week t. The independent variable is the flow-induced sale of security i at time t. Returns are measured based on weekly average prices (columns (1) to (3)) or prices on the last day of the week (columns (4) and (6)), respectively. 1-month return on UST is the return on a generic 4-week U.S. Treasury bill. See Annex A for additional information on the definition of the variables. Robust standard errors, clustered at the security level, in parentheses. The significance of the coefficient estimate is indicated by * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

3.5 Regulatory parameters

The regulatory framework for stablecoins is evolving. In general, mentions of liquidity requirements are more common and well-defined as compared to discussions of capital requirements. Across jurisdictions, the European Union has made the most progress in developing a framework for regulating stablecoins. We thus use MiCAR as the main basis for calibrating the regulatory parameters.

In terms of liquidity, MiCAR requires stablecoin issuers to hold at least 30% of their reserves as deposits with commercial banks.[20] In terms of capital, while all jurisdictions specify a minimum of 100 % backing of stablecoins with reserve assets, any overcollateralisation or capital requirements are yet to be finalised. MiCAR does establish basic capital requirements for cryptoasset service providers. Meanwhile, the European Banking Authority's draft regulatory technical standards consider a mandatory overcollateralisation in the case of stablecoins based on a historical look-back approach, which takes into account the size, complexity and nature of the reserve of assets (Euro-

[20]Systemically important stablecoins, referred to as "significant e-money tokens" in MiCAR, need to hold as least 60% of their reserves as deposits.

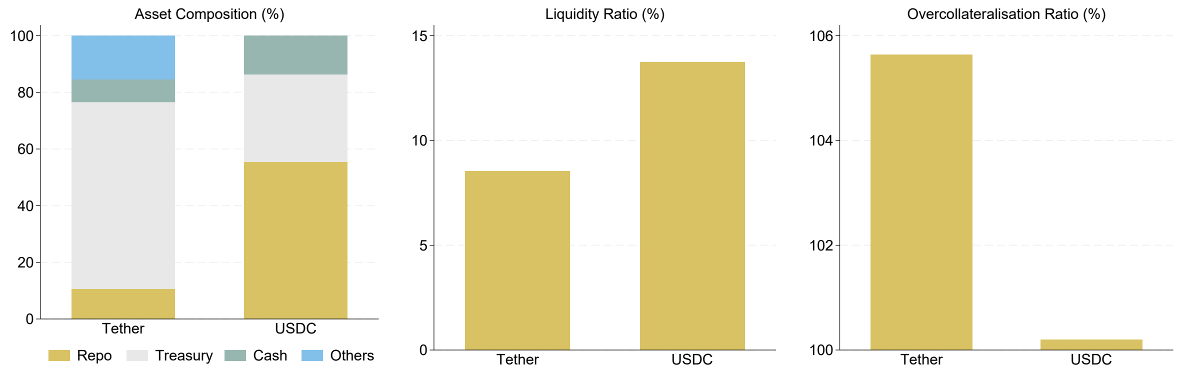

Figure 5: First panel: Asset mix of the two largest stablecoin issuers. Second panel: Liquidity ratio defined as the ratio of cash to coins in circulation. Third panel: Overcollateralisation defined as the ratio of reserve assets to coins in circulation. All panels report average values based on data from Q4 2023 and Q1 2024. Source: company disclosures.

pean Banking Authority, 2024). Its initial impact assessment based on historical data indicates that this overcollateralisation requirement would have been equivalent to a capital requirement in a range of 0.2% to 6.3% for stablecoin issuers of significant size and of 0.3% to 10.3% for the rest. Meanwhile, in Hong Kong, a 1-2% CR was considered initially.

Guided by these policy calibrations, we start from an LR benchmark of 30% and a CR benchmark of 1% (as outlined in Table 1). In Section 4, we perform comparative static experiments to evaluate the impact of changing these requirements, thus providing insights to inform policy calibrations.

The joint calibration of liquidity and capital requirements is a key aspect of ongoing regulatory deliberations. The balance sheets of USD Coin and Tether offer some indications of differences in their capital and liquidity strategies (see Figure 5). For both coin issuers, U.S. Treasuries (held outright or as part of a repo) account for most of the reserve assets (first panel), with cash holdings ranging in the 5-15 percent range (second panel). Interestingly, the liquidity ratios appear to be inversely related to the capital ratios (i.e. the degree of overcollateralisation, third panel). This suggests that liquidity and capital have some degree of substitutability. In other words, issuers adopt different approaches when balancing the need to hold cash (as an insurance against redemptions) against the need to hold capital (as an absorber of fire-sale losses). We shed light on the relative roles and potency of the two regulatory requirements in the next section.

4 Quantitative analysis

In the first part of our quantitative analysis, we take capital and liquidity requirements as given and focus on the issuer's problem. In the second part of the quantitative analysis, we conduct comparative static experiments, especially with respect to the regulatory requirements.

4.1 Liquidity management by the stablecoin issuer

In this section, we focus on the *baseline issuer*. Its balance sheet configuration in the previous period is as follows: it has a balance sheet size of $a_{-1} = 100$ (corresponding to \$100bn, approximately the average size of USDC and Tether), and it satisfies the two regulatory requirements exactly, i.e., $c_{s\theta_{-1}} = 1 = k_{s\chi_{-1}}$. We then trace how this baseline issuer responds to various demand shocks, r. We solve the issuer's problem using global solution methods (discussed in Annex B). A key advantage of this approach is that it allows us to study the privately optimal behaviour of the issuer for any ex-ante balance sheet configuration and demand shock.

When facing no subscriptions or redemptions, i.e. $r = 0$, the issuer seeks to build an LR buffer (see Figure 6, first panel). It achieves this by selling some of its bonds to raise cash. While this entails some bond liquidation costs, the revenues are sufficient to cover these costs. Consequently, profits remain positive and are added to the capital

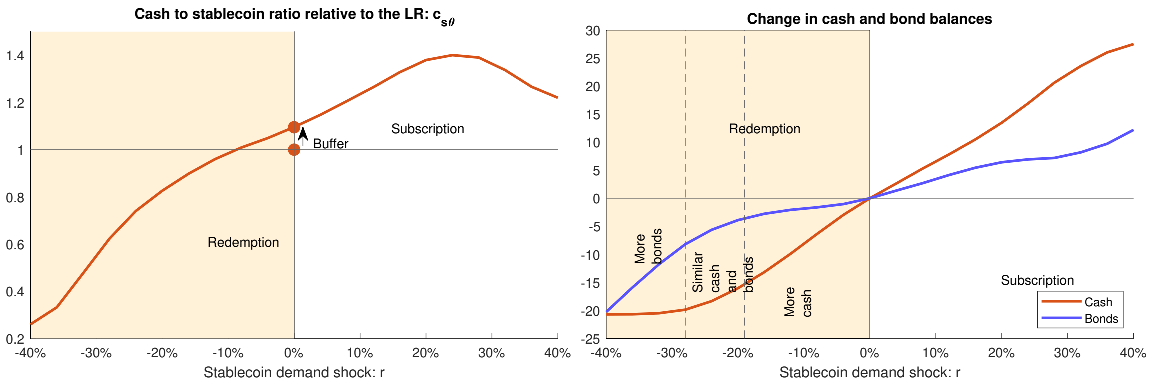

Figure 6: Left-hand panel: The liquidity management strategy as a function of the demand shock r for the baseline issuer, i.e., an issuer with $a_{-1} = 100; c_{s\theta_{-1}} = 1 = k_{s\chi_{-1}}$, meaning that the issuer is exactly at the CR and LR. Right-hand panel: The cash and bond balances of the issuer relative to what it chooses when $r = 0$. as a function of the demand shock r for the baseline issuer, i.e., an issuer with $a_{-1} = 100; c_{s\theta_{-1}} = 0.66 = k_{s\chi_{-1}}$, meaning that the issuer is exactly at the CR but below the LR.

stock as retained earnings. This also enables the issuer to build a CR buffer. Notably, the incentive to build an LR buffer is even stronger if the issuer violated the LR in the previous period.

In the case of small- to medium-sized subscriptions, i.e. $r > 0$, the issuer takes advantage of the inflow of funds to build a larger cash buffer. This behaviour reflects the issuer's ability to strengthen its liquidity position under favourable conditions.

For particularly large subscriptions, i.e., $r \gg 0$, the issuer opts to maintain a smaller buffer. This decision is driven by the persistence of demand shocks – when subscriptions are large, the issuer anticipates continued inflows in the near future, which reduces the relative benefits of holding a larger buffer compared to a scenario where $r = 0$.

The issuer's privately optimal behaviour becomes more intricate in the case of redemptions. This can be better understood through the concept of 'cuts'. For small redemption shocks, the issuer primarily relies on its cash holdings (see Figure 6, second panel). This strategy minimises bond liquidation costs. However, by predominantly using cash, the issuer may temporarily violate the LR. This is a manageable risk because current redemptions are small, and the persistence of demand shocks suggests that redemptions are likely to remain small or even turn into subscriptions in the near future, allowing the issuer to replenish its cash position.

This strategy, however, becomes infeasible for large redemption shocks. Large redemptions require the sale of more assets, and relying solely on cash may not be practical due to limited cash reserves. Even if sufficient cash were available, fully meeting redemptions with cash could result in a significant violation of the LR. Such a violation, combined with already large redemptions in the current period, would severely worsen the distribution of redemptions in the subsequent period. Moreover, insufficient cash reserves for future periods would heighten the risk of default.

To manage larger redemptions, the issuer employs a balanced approach, using both cash and bond sales in roughly equal proportions. While this strategy incurs liquidation costs, it is necessary to maintain future viability, particularly given the persistence of demand shocks and the likelihood of continued large redemptions in the next period. These model-based findings are consistent with empirical evidence on how money market funds manage their liquidity Jiang et al., 2021; Ma et al., 2022.

Issuers with alternative balance sheet configuration We consider alternative issuers, those with a different balance sheet than the baseline issuer. For instance, a larger issuer uses cash more intensely in order to meet redemptions. This is because the same bond sale – in ratio terms – would be more costly given that the sale would be larger in absolute value (first panel, Figure 7).

For an issuer that has an LR buffer of 15 percentage points (i.e. a liquidity ratio at 1.5 times the LR of 30%) to begin with, the incentive to widen the buffer when facing no redemptions/subscriptions is smaller (second panel). This is because the issuer has a relatively high liquidity ratio to begin with. For very large subscriptions, such an issuer even prefers to reduce the size of the buffer. The benefit of having an ex-ante buffer is

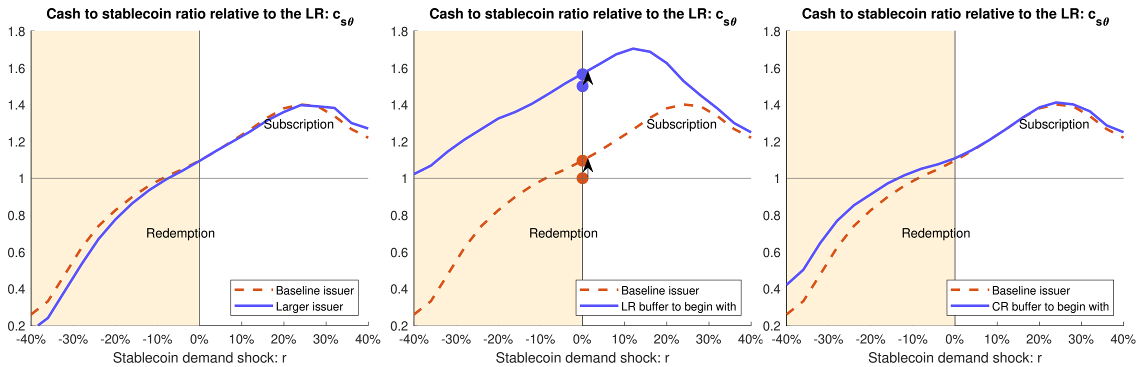

Figure 7: The liquidity management strategy of alternative issuers: (i) one with a larger balance sheet (ii) with an LR buffer (iii) with a CR buffer.

evident in the case of redemptions – the issuer does not need to go below the LR for redemptions even as large as 40%. This issuer is thus highly likely to avoid regulatory fines and the so-called coin-holder discipline.

For an issuer that has a CR buffer (i.e. a capital ratio at 1.5 times the CR of 1%) but no LR buffer also benefits. When facing redemptions, it can rely more on bond sales and preserve liquidity for the future (relative to the baseline issuer). This is because the capital cushion helps absorb more of the bond-sale costs.

Comparative statics with respect to key parameters In this subsection, we present how the behaviour of the issuer changes with three key parameters of the model. First we assess the implications of a change in the price impact parameter m. A higher price impact of bond sales leads the issuer to use cash more intensively during redemptions (first panel, Figure 8). This is expected given that using bond sales to the same extent as in the baseline would take the issuer closer to default. Second, we study the role of coin-holder scrutiny i.e. the degree to which redemptions increase when the issuer violates a regulatory requirement. We find that in this case the issuer holds a somewhat larger liquidity buffer in order to reduce the probability of violating the requirements. It also uses cash less intensively during redemptions to reduce the

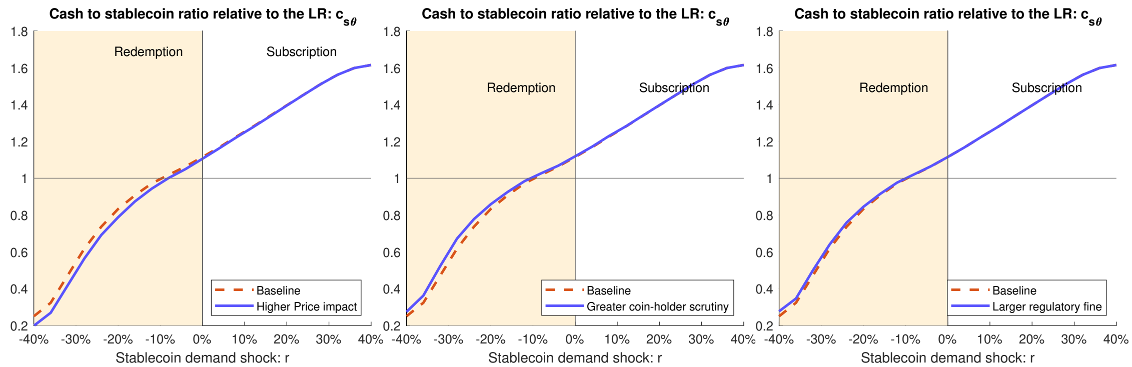

Figure 8: The liquidity management strategy of the baseline issuers under different parameter settings: (i) price impact m increases from 2bps to 2.5bps (ii) coin-holder scrutiny parameter λ goes from 5 to 10 percentage points rise in redemptions, and (iii) the regulatory fine Ω increases from 10 to 15 percent of capital.

degree with which it violates LR. A similar effect is noted when the regulatory fine increases, confirming the intuition that coin-holder scrutiny and stringency of regulatory implementation serve as substitutes.

4.2 Counterfactual policy experiments

The regulator has two objectives. First is to limit the probability of default (PD) of the issuer. This micro-prudential objective focuses on strengthening the stablecoin arrangement and reducing the risks for the immediate stakeholders, especially the coin holders. The second objective is to limit the expected price impact (EPI) of bond sales by the issuer, typically in response to redemption pressures. This macro-prudential objective has to do with limiting the spillovers from stablecoins to the broader financial system, irrespective of whether the issuer defaults.[21] The two prudential objectives justify the application of two regulatory requirements, in line with the Tinbergen Principle.[22]

[21] For instance, a negative impact on bond prices can hurt other entities in the financial system that hold the same or similar bonds and have to mark losses on their balance sheets. In extreme cases, the price impact can trigger a cascade of defaults (see e.g. Cantú et al., 2024 for a discussion).

[22] Prudential objectives could be framed differently. Regulators might aim to limit the expected volatility of the market price of reserve assets or focus on minimising losses to coin-holders in the event of an issuer's default. These alternative formulations of prudential objectives are likely to result in qualitatively similar regulatory trade-offs as those observed in the current model.

The issuer's PD in the next period depends on the current state vector as well as the regulatory requirements. With the indicator function capturing the event of default when capital becomes negative, the PD is given by:

$$PD(a_{-1}, c_{s\theta_{-1}}, k_{s\chi_{-1}}, r; \chi, \theta) = \sum_{i \in [L,H]} \int \mathbb{1}\Big(k(a, c_{s\theta}, k_{s\chi}, r') < 0\Big) p_i f(r'|r) dr$$

The expected price impact (EPI), in turn, is the expected amount by which the bond price declines in the next period due to the expected bond sales by the issuer:

$$EPI(a_{-1}, c_{s\theta_{-1}}, k_{s\chi_{-1}}, r; \chi, \theta) = \sum_{i \in [L,H]} \int m\Big(b - b'(a, c_{s\theta}, k_{s\chi}, r')\Big) p_i f(r'|r) dr$$

These equations underscore that there exists a mapping from (χ, θ) to (PD, EPI) for any stablecoin issuer with an ex-ante balance sheet composition $(a_{-1}, c_{s\theta_{-1}}, k_{s\chi_{-1}})$ and facing a demand shock r. This mapping offers regulators a framework to calibrate the capital and liquidity requirements according to their PD and EPI objectives while taking account of how stablecoin issuers optimise their balance sheet composition in the presence of regulation and redemption risks.[23]

To this end, we perturb the regulatory requirements χ and θ around their baseline values and examine how the PD and EPI of a representative issuer facing a stable demand respond.[24] We find that both LR and CR contribute to reducing PD and EPI (Figure 9), making them substitutes. However, if the regulator has a joint PD-EPI objective in sight, then a LR-CR combination becomes necessary. Crucially, each requirement affects PD and EPI in different ways. The main benefit of an LR is that

[23]Our analysis is positive in nature and does not take a position on the optimal level of regulation. A normative assessment would require taking a stance on the utility, if any, that stablecoins provide to the economy – a topic on which opinions differ significantly among academics and policymakers. For a discussion, see, e.g., Gorton and Zhang [2023] or Aldasoro et al. [2024].

[24]By stable demand, we mean r = 0, ie the issuer is not facing any subscription or redemption in the current period.

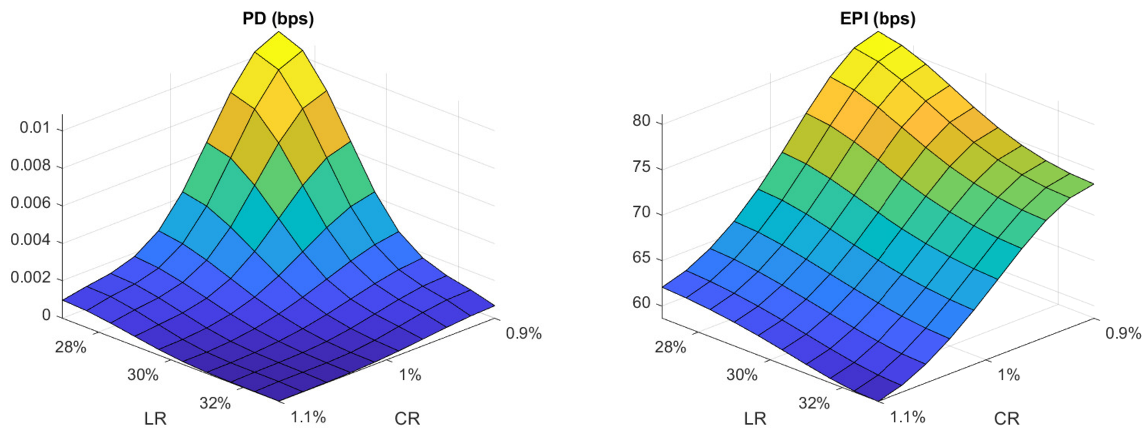

Figure 9: The impact of varying regulatory requirements on the probability of default (PD) and the expected price impact (EPI) of a representative issuer that meets the requirements exactly and is facing no redemption or subscription in the current period (r=0).

it helps reduce EPI by reducing the need to sell bonds ceteris paribus. In turn, the LR helps reduce PD. The CR primarily works to reduce PD as it absorbs bond-sale costs, but can also reduce EPI via its effect on the cash ratio of the issuer.

The contour plot in Figure 10 depicts the combinations of PD and EPI that can be achieved under different configurations of LR and CR. For instance, a 1% CR combined with a 30% LR yields a weekly PD of approximately 0.001 bps, equivalent to 0.05 bps on an annual basis, and an EPI of 65 bps.[25] The contour plot also serves as a tool to map regulatory objectives to specific requirements. For example, if the regulatory goal is a PD of 0.001 bps, the plot identifies the minimum achievable EPI and the corresponding LR-CR combination (and vice versa).

5 Conclusion

The paper has two main contributions—to provide a model of stablecoin issuer's reserve management strategy, and a framework that helps design prudential regulation for

[25] Setting low PD targets is a common feature of regulatory frameworks for payment service providers.

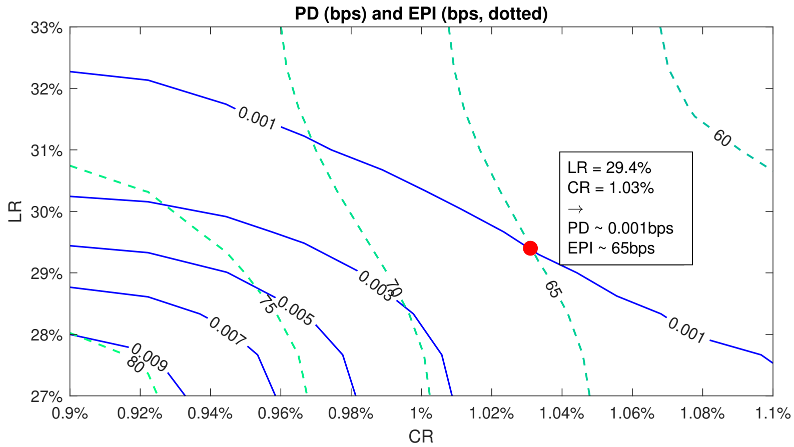

Figure 10: The mapping from regulatory objectives to requirements and vice-versa.

stablecoins.

The model helps analyse how stablecoin issuers manage their balance sheet, in particular their liquidity, in response to stochastic demand by coin holders, especially redemption shocks. It highlights the trade-offs the issuer faces in choosing between cash and interest-bearing securities. The issuer exhibits strategic balance sheet adjustments. For small redemption shocks, it primarily relies on its cash holdings to minimise the liquidation costs associated with selling securities. In contrast, large redemption shocks prompt the issuer to reduce its holdings of cash and bonds in roughly equal proportions, with the hope of replenishing its cash buffer in the future. However, if the redemption shock persists, the issuer may face difficulties in replenishing cash reserves, ultimately encountering fire-sale losses and potentially defaulting. Meanwhile, subscriptions (i.e. inflows) provide the issuer an opportunity to build buffers relative to the regulatory requirements. That said, very large subscriptions obviate the need for very large buffers.

The framework further enables an exploration of how regulation can make stablecoins more stable. We find that both capital and liquidity requirements are effective in reducing the PD of a stablecoin issuer, making them substitutes from a micro-prudential perspective. However, when the policy objective extends to macro-prudential concerns—such as mitigating spillovers from stablecoins to financial markets during periods of stress—the use of both capital and liquidity requirements becomes necessary. We highlight that each requirement has its own comparative advantage. The main benefit of an LR is that it helps reduce EPI by reducing the need to sell bonds, while the CR is more potent for PD as it helps absorb bond-sale sales. Finally, we provide a quantitative mapping from policy targets to regulatory requirements. This mapping can inform policy deliberations on how capital and liquidity requirements for stablecoins should be calibrated.

The objective of this paper is to provide a practical framework to help address key policy considerations. This inevitably involves simplifications but also opens up several avenues for future research. These include, for example, investigating the benefits of concentration limits in containing the EPI within individual market segments, mitigating risks associated with the failure of banks holding stablecoin issuers' cash deposits, and addressing broader financial stability concerns arising from tighter interconnections between stablecoin issuers and the banking sector (e.g., Bank for International Settlements [2023]).

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*.

Ahmed, R. and A. Rebucci (2024). Dollar reserves and u.s. yields: Identifying the price impact of official flows. *Journal of International Economics 152*, 103974.

Aldasoro, I., P. Mehrling, and D. Neilson (2024). On par: a money view of stablecoins. *Journal of Financial Market Infrastructures 11*(4), 47–64.

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

Anadu, K., P. Azar, C. Huang, M. Cipriani, T. M. Eisenbach, G. La Spada, M. Landoni, M. Macchiavelli, A. Malfroy-Camine, and J. C. Wang (2023b). Runs and flights to safety: Are stablecoins the new money market funds? *FRB of Boston Working Paper*.

Arellano, M. and O. Bover (1995). Another look at the instrumental variable estimation of error-components models. *Journal of Econometrics 68*, 29–51.

Bank for International Settlements (2023). The crypto ecosystem: key elements and risks. Report submitted to the G20 Finance Ministers and Central Bank Governors.

Bank of England (2022). The prudential liquidity framework: supporting liquid asset usability. Discussion Paper 1/22.

Bernardo, A. E. and I. Welch (2004). Liquidity and financial market runs. *The Quarterly Journal of Economics 119*(1), 135–158.

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

Blundell, R. and S. Bond (1998). Initial conditions and moment restrictions in dynamic panel data models. *Journal of Econometrics 87*, 115–143.

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).*

Chen, Q., I. Goldstein, and W. Jiang (2010). Payoff complementarities and financial fragility: Evidence from mutual fund outflows. *Journal of Financial Economics 97*(2), 239–262.

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

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

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

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.

Garratt, R. and H. S. Shin (2023). Stablecoins versus tokenised deposits: implications for the singleness of money. *BIS Bulletin* (73).

Goldstein, I., H. Jiang, and D. T. Ng (2017). Investor flows and fragility in corporate bond funds. *Journal of Financial Economics 126*(3), 592–613.

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

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.

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.

Kacperczyk, M. and P. Schnabl (2013). How safe are money market funds? *The Quarterly Journal of Economics 128*(3), 1073–1122.

Li, Y. and S. Mayer (2022). Money creation in decentralized finance: a dynamic model of stablecoin and crypto shadow banking. 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. Working Paper 31160, National Bureau of Economic Research.

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 (2023). Stablecoin runs and the centralization of arbitrage. mimeo.

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

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.

Parlatore, C. (2016). Fragility in money market funds: Sponsor support and regulation. *The Review of Financial Studies 29*(7), 1866–1913.

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

Schmidt, L., A. Timmermann, and R. Wermers (2016). Runs on money market mutual funds. *American Economic Review 106*(9), 2625–2657.

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

Sunderam, A. (2015). Money creation and the shadow banking system. *The Review of Financial Studies 28*(4), 939–977.

Van den Heuvel, S. J. (2008). The welfare cost of bank capital requirements. *Journal of Monetary Economics 55*(2), 298–320.

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

A Price impact estimation

This section provides additional details on the construction of the main variables used to estimate the price impact due to flow-induced sales by money market funds presented in Section 3.4.

We start by describing the calculation of the weekly returns on government securities, $Return_{i,t}$, the dependent variable in Equation 4. To construct this variable, we gather daily data on U.S. Treasury bills, securities with maturity of up to one year, and on U.S. Treasury notes and bonds, securities with original maturity of more than one year for the period, from January 2016 to March 2024.

For Treasury bills, the available data is in terms of daily annualised yields, from which we recover daily prices in two steps. First, we calculate the daily yield, $yield_{i,d}$, of security i at day d based on its daily annualised yield, $yield^{a}_{i,d}$ as:

$$(1 + yield_{i,d}) = (1 + yield^{a}_{i,d})^{(1/365)}. \tag{5}$$

Next, we obtain the daily price, $price_{i,d}$, based on:

$$price_{i,d} = \frac{100}{(1 + yield_{i,d})^{m_{i,d}}} \tag{6}$$

where 100 is the face value of the Treasury bill i and $m_{i,t}$ is the number of days to maturity for security i at day d.

We construct weekly returns for all securities as follows:

$$Return_{i,t} = \frac{price_{i,t} - price_{i,t-1}}{price_{i,t-1}}. \tag{7}$$

where $price_{i,t}$ is the weekly price of security i in week t. We consider two measures of weekly prices: the first one is based on using the weekly average price of the securities, whereas the second one uses the last observation in each week.

Our final dataset, once matched with the fund flows as discussed below, consists of 1,098 U.S. Treasury securities – 761 bills and 337 notes and bonds.

Our measure of selling pressure, $\textit{Flow-induced sale}_{i,t}$, the main independent variable in Equation 4, approximates the dollar volume of sales of security i in week t that are due to money market funds (MMFs) liquidating their government bond holdings to meet redemption requests. Its construction proceeds in two steps.

First, we obtain data on MMFs' daily (net) outflows in US dollar billions over the period from January 2016 to March 2024. For consistency with the our return measures, we consider two corresponding measures of flows: one based on changes in the weekly averages of TNA and the other one based on changes in the end-of-week values of TNA.

Second, we compute the weekly total flow-induced sales of security i at fund f in week t as the product of the outflows experienced by fund f at time t, $Outflow_{f,t}$, and the share of security i in fund f's portfolio in the previous month, $m-1$, which we denote as $Holdings_{i,f,t(m-1)}$. We infer the latter from monthly asset holdings, the highest available frequency. Following the literature studying the price impact of asset sales by mutual funds (e.g. Coval and Stafford, 2007; Jiang et al., 2021; Ma et al., 2022) we assume that holdings are constant at the fund-level within a given month and approximate flow-induced sales as:

$$\textit{Flow-induced sale}_{i,f,t} = Outflow_{f,t} \times Holdings_{i,f,t(m-1)}. \tag{8}$$

The implicit assumption is that MMFs meet one dollar worth of outflows by selling the various securities in proportion to their portfolio holdings.

Aggregating the amount of sales across all MMFs, $f \in F$, that held security i in their portfolio in the previous month yields the total volume of flow-induced sales by

MMFs for bond i at time t as defined in Equation 3.

We provide summary statistics of weekly returns and flow-induced sales (expressed in $ millions for enhanced visibility) based on weekly average values and end-of week values, respectively, in Table 4.

Table 4: Summary statistics

Weekly average values	Mean	Stdev	P10	P25	P50	P75	P90	Obs.
Returns (basis points)	1.289	4.972	–3.867	–0.508	0.382	3.628	8.461	27,222
Flow-induced sales (US$ mn)	2.982	6.513	0.000	0.055	0.731	3.014	8.329	27,222
End-of week values								
Returns (basis points)	1.301	5.085	–3.894	–0.586	0.395	3.668	8.534	27,222
Flow-induced sale (US$ mn)	3.669	7.930	0.000	0.066	0.934	3.670	10.223	27,222

The two top rows report summary statistics based on using weekly average security prices and TNA to measure weekly returns and flow-induced sales, respectively. The two bottom rows depict the corresponding measures based on using the last observation in each week.

B Solving the model using global methods

We solve the issuer's problem using Value Function Iteration. The state-space is made discrete using 21 linearly spaced grid points in each dimension. For size, we choose a range of $[0.1, 200]$ where the mid-point, i.e. 100, is assumed to be the size of the baseline issuer. For the regulatory ratios, we consider a range of $[0.1, 2]$, and for the demand shock, a range of $[-0.4, 0.4]$. To compute the value and policy function at off-grid points, we use cubic-spline interpolation. Starting with an initial guess for the value function, we iterate on the solution to the bank's problem. A key aspect of solving the problem of the issuer in each iteration is that the various constraints must be accounted for. In the case of default, we impose a steep penalty on the value function which is proportional to the amount by which capital is below zero. This creates incentives for the issuer to minimise default risk when choosing its balance sheet components. After each iteration, we update the value function until the maximum difference (at any grid point) between the old and updated value functions is smaller than a threshold.

C Alternative price impact specifications

In the main model, we consider the price impact to be a linear function of the size of the bond sale. In this Appendix, we consider a concave price impact function that increases with the size of the sale but the maximum impact is capped. The latter setup assumes that as the size of the sale becomes larger, the seller engages with an increasingly deeper liquidity pool where a larger number of buyers are ready to buy at the already discounted price, thus limiting further price discounts.

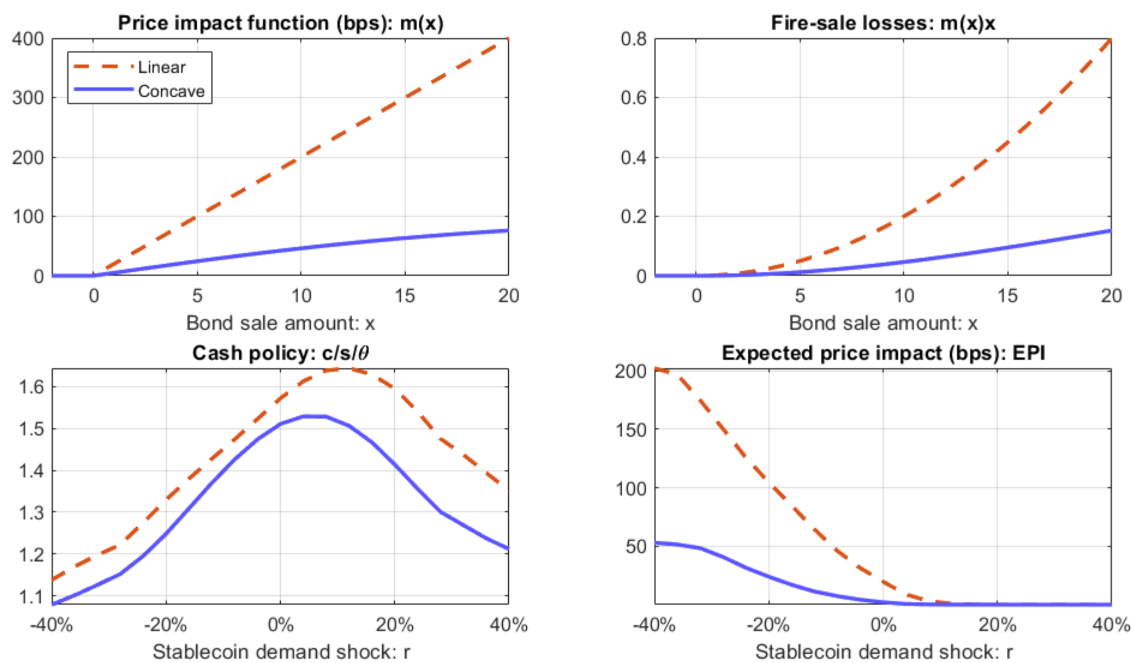

Figure 11: In case of the linear price impact (dashed red lines), the marginal impact is 2bps per \$100 million sold (the unit of the x-axis is USD billions). In case of the concave price impact (solid blue lines), the maximum impact is capped at 100 basis points (bps). The top left panel compares the two price impact functions while the top right compares the associated fire-sale losses. The bottom left and right panels show the cash to stablecoin ratio relative to θ that is chosen by the issuer and the expected price impact in the two regimes, respectively, assuming that the issuer's starting cash ratio is 45%, i.e. 1.5 times the benchmark LR of 30%.

Figure 11 depicts the two price impact functions (top left-hand panel) and the associated fire-sale losses (top right-hand panel) as a function of the bond sales (horizontal axis). Figure 11 also illustrates the issuer's cash buffer choices (bottom left-hand) and the expected price impact (bottom right-hand) for a range of redemption shocks (horizontal axis), assuming the issuer starts with an LR buffer of 15 percentage points (i.e. a liquidity ratio at 1.5 times the LR of 30%) as in the case of Figure 7 (centre panel) in Section 4.

Intuitively, when the price impact of large sales is significantly smaller (blue line), the issuer prefers to maintain a smaller cash buffer relative to the requirement as the loss, and hence the risk of default, associated with bond sales declines. Yet despite the issuer choosing to hold less cash, the expected price impact for all possible demand shocks is lower under this calibration (blue vs. red line in the bottom right-hand panel). This underscores the key role of market depth of the underlying reserve assets for the regulatory risk assessment.