Bank of England

A simulation framework for sterling money market funds: estimating redemption capacity and evaluating liquidity requirements

Staff Working Paper No. 1,177

March 2026

Rishabh Kumar

Bank of England

Staff Working Paper No. 1,177

A simulation framework for sterling money market funds: estimating redemption capacity and evaluating liquidity requirements

Rishabh Kumar[1]

Abstract

Money market funds (MMFs) aim to provide near-on-demand liquidity yet often hold assets that become hard to sell under stress, leaving them vulnerable to run-like redemptions. I build a simulation framework for sterling MMFs to estimate redemption capacity and failure probability across alternative redemption profiles and market-liquidity scenarios. Resilience of funds depends on both the timing of outflows and the effective liquidity of weekly liquid assets (WLA): front-loaded redemptions are most destabilising, and the benefit of asset sales shrinks as market depth thins. Removing the 30% WLA threshold effect – under which managers must consider measures to deter further redemptions – yields sizeable resilience gains by reducing cliff-edge behaviour. Under historically extreme shocks and without threshold effects, most resilience improvements come from holding WLA above the 30% regulatory minimum; in my simulations, gains concentrate around 40% WLA, with diminishing returns beyond.

Key words: Money market funds, liquidity, simulations.

JEL classification: G23, C63, G28.

---

(1) Bank of England. Email: rishabh.kumar@bankofengland.co.uk

The views expressed in this paper are those of the author, and not necessarily those of the Bank of England or its committees. I would like to thank Nick Vause, Gerardo Ferrara, Steven Dodkins, David Baumslag, Adam Brinley Codd, Arzu Uluc and Pelagia Neocleous for their valuable comments and suggestions.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH
Email: enquiries@bankofengland.co.uk

ISSN 1749-9135 (on-line)

1. Introduction

Money market funds (MMFs) aim to provide investors on-demand liquidity, often with same-day or next-day access, while investing in instruments that can be difficult to liquidate quickly at par. That mismatch is a feature rather than a flaw: it is precisely what allows MMFs to earn a competitive yield. But it also makes them vulnerable to run dynamics especially during stress periods where investors might prefer cash. When redemptions rise abruptly, funds may be forced to generate cash by selling assets into stressed markets, potentially reinforcing price moves and tightening funding conditions. Because MMFs sit at the junction of short-term funding markets and cash management for non-financial corporations and financial institutions, disruptions can propagate beyond the sector and become a financial stability concern (Bank of England, 2021).

Past episodes have illustrated these vulnerabilities in practice. During the March 2020 "dash for cash" (DFC) episode, funds with larger exposures to non-government assets, including U.S. Prime MMFs and EU Low Volatility Net Asset Value (LVNAV) funds, experienced substantial outflows. In the UK, the DFC episode triggered redemptions of roughly £25 billion (around 11% of sterling MMF assets), concentrated in EU-domiciled LVNAV[1] funds (Financial Conduct Authority (FCA), 2023). Investors in need of liquidity made large MMF redemptions to raise cash quickly (Czech et al., 2021). The pattern echoes earlier stress episodes, including the Global Financial Crisis (GFC), when the Reserve Primary Fund "broke the buck" after Lehman Brothers' default and the sector required extraordinary public backstops (McCabe, 2010; Bouveret, Martin, and McCabe, 2022).

Regulatory reforms have tried to address these fragilities. The United States Securities and Exchange Commission's (SEC) 2010 MMF reforms and the European Union's (EU) 2017 Money Market Fund Regulation (MMFR) introduced liquidity management tools, minimum liquid-asset requirements, and constraints on portfolio composition (Bouveret, Martin, and McCabe, 2022). Yet recurring stress events suggest that liquidity mismatches remain a source of vulnerability for the sector. Moreover, the increase in interest rates from the 2020 pandemic-era lows introduced a separate tension: higher rates can reduce the market value of fixed-income holdings and interact with liquidity needs in stress. The Bank of England's *Financial Stability in Focus* highlights the potential for rate-driven valuation changes to intensify liquidity risks across non-bank financial institutions, including MMFs (Bank of England, 2023).

A growing empirical literature documents MMF outflows and transmission channels during stress events. However, empirical work is necessarily backward-looking: it reflects realised

[1] LVNAV funds maintain a stable net asset value (NAV) within a ±20bp collar and use limited amortised-cost pricing, unlike Variable Net Asset Value (VNAV) funds which fully mark-to-market and float continuously. LVNAV and Constant Net Asset Value (CNAV) funds face stricter daily/weekly liquidity requirements and may apply fees or gates when liquidity falls, whereas VNAV funds are subject to looser liquidity rules and no fees-and-gates framework.

outcomes under the institutional and policy conditions of a particular episode. In practice, those conditions may include use of emergency lending facilities and other central bank intervention that alter both investor behaviour and the liquidity of assets. This limits what can be inferred about MMF resilience under alternative market environments, different redemption paths, or counterfactual policy designs.

This paper takes a complementary approach. I develop a forward-looking simulation framework for sterling-denominated MMFs that allows me to trace fund-level liquidity dynamics under a wide range of redemption shocks and market-liquidity assumptions. The framework delivers two outcomes that are directly relevant for policy. The first is *redemption capacity*: the maximum cumulative redemptions as a share of assets that a fund can meet over a specified horizon without breaching a requirement to hold a certain share of its assets in financial instruments that can be liquidated within one day. The second is *failure probability*: the share of simulated paths in which this daily liquidity constraint is breached under a given shock, redemption profile and time horizon. Using these metrics, I can compare in a unified way the resilience of MMFs to redemption stresses over different time horizons and profiles, depending on their initial share of WLA and a specified failure level of 10% of assets under management (AUM).

The model draws on behavioural mechanisms proposed in the recent literature (Bouveret, Martin, and McCabe, 2022; Dunne and Giuliana, 2022; Baes, Bouveret, and Schaanning, 2023) and explicitly varies the market environment. Specifically, I consider three liquidity settings: (i) a conservative benchmark in which secondary markets are effectively frozen and funds rely only on cash and maturing assets; (ii) a setting in which government securities can be sold to help meet redemptions; and (iii) an extension in which certificates of deposit (CDs) can also be monetised prior to maturity, under explicit limits intended to capture market depth. These market counterfactuals are chosen to reflect the different conditions seen in the past. I also study a policy counterfactual motivated by evidence of threshold effects near the regulatory minimum (Dunne and Giuliana, 2022). Specifically, I switch off threshold-driven redemption acceleration and quantify the resulting change in redemption capacity.

I have three main findings. First, in the frozen-market benchmark, redemption capacity rises monotonically with initial WLA and with the length of the stress horizon over which redemptions are spread, and it is strongly shaped by the *timing* of redemptions: front-loaded (left-skewed) paths are substantially more destabilising than historically observed unimodal paths. Second, allowing asset sales shifts the redemption-capacity frontier favourably, but the size of this improvement is highly sensitive to market depth. When secondary-market liquidity deteriorates, funds cannot reliably liquidate the assets counted as WLA at scale, so the effective buffer is smaller than the headline WLA measure and resilience is correspondingly weaker. Third, failure probabilities imply a non-linear trade-off in the design of liquidity buffers. Under historically extreme shocks, the largest reduction in failure risk is generated by increasing the current regulatory minimum to a higher buffer that lifts funds away from the region in which constraints tend to bind; in my simulations, that improvement is concentrated around a move to approximately 40% WLA. Beyond 40%, the marginal reduction in failure risk tends to diminish, especially under unimodal (historically observed) redemption paths. This pattern suggests that raising the minimum WLA above the current floor can materially improve resilience (especially for historically observed outflows), while also indicating where additional increases yield smaller incremental benefits.

These findings speak directly to ongoing policy debates, including the FCA review of sterling MMFs (FCA, 2023) and the SEC's 2023 reforms (SEC, 2023). The central issue is calibration: buffers must be large enough to reduce the likelihood of destabilising liquidity spirals, yet not so high that they materially impair intermediation. Crucially, the required buffer must also be credible in *periods of stress.*

Related literature: The paper sits at the intersection of two literatures: MMF fragility and the design of liquidity regulation. The GFC spurred a large literature on run dynamics, showing how liquidity transformation amplifies stress when investors demand cash against portfolios that cannot be liquidated at par (Pedersen, 2009; Gorton and Metrick, 2012; Chernenko and Sunderam, 2014; Ivashina, Scharfstein, and Stein, 2015). In the European context, European Systemic Risk Board (ESRB) (2021), reviews structural vulnerabilities of EU MMFs and documents how these vulnerabilities manifested during the COVID-19 pandemic.

Building on this foundation, more recent empirical work has turned to the interaction between investor behaviour and regulatory thresholds. Using U.S. data, Li et al. (2021) shows that prime MMFs—which are more likely to be exposed to fees and gates given lower liquidity buffers—experienced larger outflows during stress. Using both U.S. and European data, Cipriani and La Spada (2020) find that runs were more severe among institutional-share-class funds and among funds closer to binding regulatory constraints, consistent with pre-emptive redemption incentives. Related evidence for March 2020 suggests that as WLA buffers approached the regulatory threshold, outflows accelerated (Dunne and Giuliana, 2021; Li et al., 2021; Cipriani and La Spada, 2020).

The closest paper to mine is Baes, Bouveret, and Schaanning (2023), who develop a stylised model in which the ability of MMFs, especially LVNAVs, to satisfy NAV and liquidity constraints depends on both the shape of outflows and the market liquidity of the underlying assets. They then quantify how liquidity management tools such as fees, gates, and liquidity requirements affect the maximum redemptions a fund can meet. My contribution differs in two main respects. First, I implement a simulation framework calibrated on a broad panel of sterling MMF portfolios over multiple stress episodes, rather than focusing on a small number of portfolios drawn from a single pre-stress window. Second, I use the framework to compare resilience across market-liquidity counterfactuals (no asset sales, government asset sales under varying depth, and an extension with limited CD monetisation) and a policy counterfactual that removes threshold-driven redemption acceleration. This allows me to connect portfolio composition, market functioning, and threshold effects to both redemption capacity and failure probabilities.

The policy relevance of threshold effects is also emphasised by contemporaneous industry commentary during March 2020. In a client note, BlackRock (2020) described the WLA ratio as an "amber flashing light" and argued that fear of fees and gates effectively transformed the 30% WLA threshold into a new run focal point. Fitch Ratings made similar observations about investor attention to WLA during the episode (McLaughlin, 2020). Taken together, these studies motivate treating WLA not only as a regulatory metric but also as a behavioural coordination device, which I model directly through a threshold-driven acceleration channel.

The remainder of the paper is organised as follows. Section 2 defines the data and simulation design. Section 3 reports the results, structured around redemption capacity, market counterfactuals, the policy counterfactual, and failure probabilities. Section 4 concludes.

2. Data and simulation design

2.1 Data

I use the Crane[2] MMF holdings and daily flows dataset for sterling-denominated MMFs. The sample covers January 2020 to January 2023, spanning both the March 2020 DFC episode and the September 2022 Liability-Driven Investment (LDI) event.[3] Over this period, the number of funds ranged from 22 to 16 and aggregate AUM fluctuated between approximately £125bn and £225bn.

For each instrument in the holdings data, Crane reports a range of fields, including the security name, issuer and issuer type, maturity date, coupon, amount outstanding, issuer country, and credit ratings. Crane classifies holdings into four broad asset types. The first, labelled *Commercial Paper*, includes short-term instruments issued by non-government entities (asset-backed commercial paper, financial company commercial paper, non-financial company commercial pa-

[2]Crane Data LLC (Crane Data) is a private vendor of MMF statistics and holdings. Through subscription datasets they provide fund-level yields, assets/flows, monthly portfolio holdings, and other metrics with daily snapshots and historical archives.

[3]The LDI crisis of September–October 2022 saw gilt yields rise sharply. Pension funds running LDI strategies faced large margin calls, forcing rapid gilt sales that amplified the yield spike and created severe dysfunction in UK government bond markets (Bank of England, 2022).

per, and certificates of deposit). The second, labelled *Deposits*, includes non-negotiable time deposits and repurchase agreements, including government agency repo, Treasury repo, and other repo. The third, labelled *Government Debt*, comprises government-related instruments such as Treasury debt and other government debt. For the purposes of this paper, I treat these instruments as UK government securities.[4] The fourth category, labelled *Other*, includes the remaining instruments (e.g., longer-dated holdings and investment company shares). Figure 1 shows average asset allocation by fund type.

The dataset is not without limitations, so I apply several cleaning steps. First, I exclude observations in which reported daily liquid assets (DLA) are below 10% of AUM.[5] Second, I recompute WLA and DLA from the underlying holdings to mitigate reporting errors that could affect the simulation results. I define WLA as the sum of: (i) instruments with residual maturity of five business days or less and (ii) UK government-related instruments with residual maturity of 190 days or less. Under the EU/UK MMFR, LVNAV funds may include such government-related assets in WLA up to 17.5 percentage points.

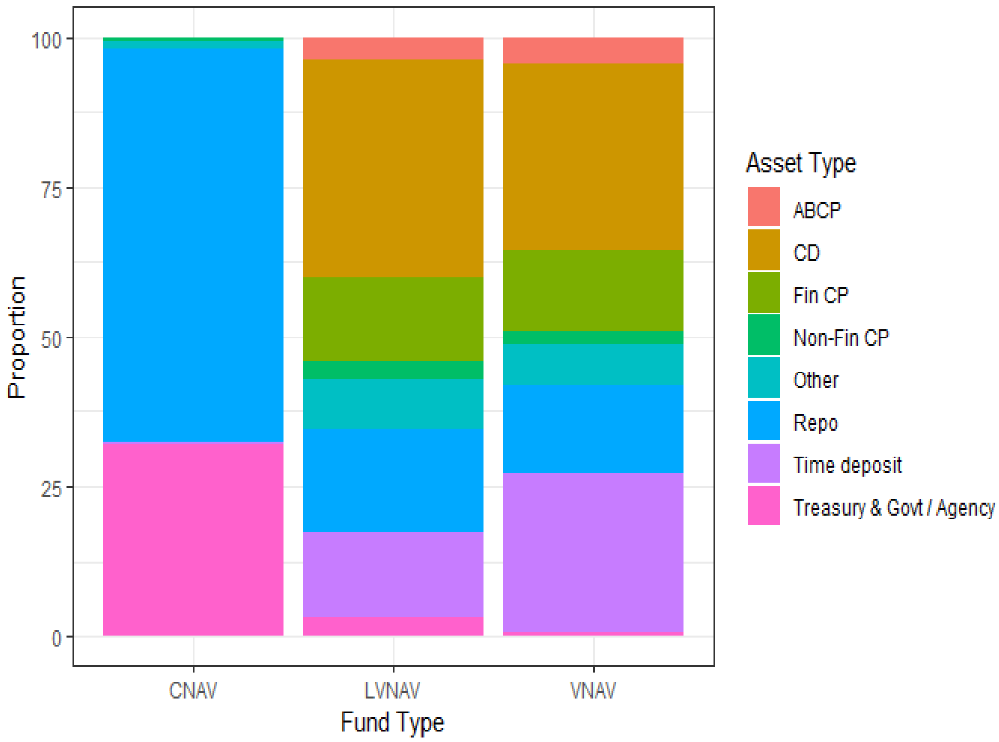

Figure 1: Average asset allocation by fund type

[4]If the dataset contains non-UK sovereign instruments within this category, I abstract from cross-sovereign differences and treat them as government securities for liquidity-classification purposes.

[5]I do this to remove observations that are inconsistent with the regulatory constraint and to reduce the influence of reporting anomalies.

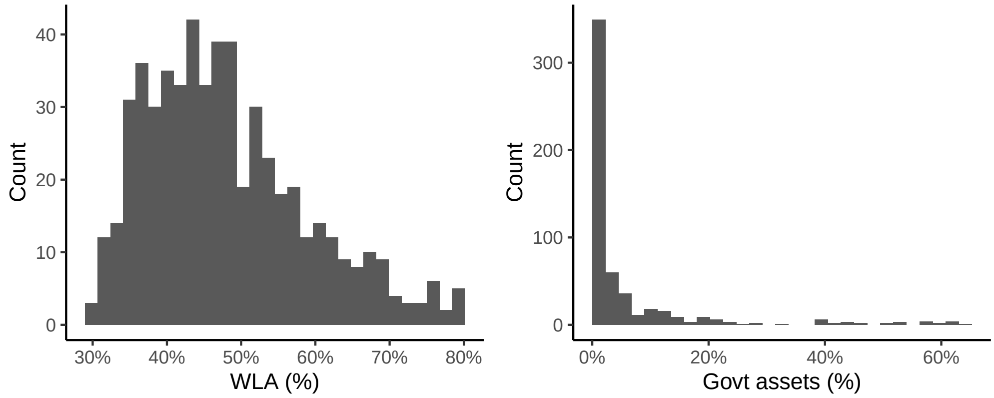

Figure 2: Distribution of weekly liquid assets (WLA) and government assets across fund-date observations.

2.2 Simulation Design

I develop a fund-level simulation framework to study run dynamics in sterling MMFs under stressed redemptions. I focus on LVNAV funds because they dominate the sterling MMF universe: in my sample, more than 75% of funds—and close to 90% of sterling MMF assets—are EU-domiciled (primarily Ireland and Luxembourg) and structured as LVNAVs. This universe also exhibits substantial cross-sectional variation in WLA, which is central to identifying how balance-sheet liquidity maps into resilience.

The simulation combines: (i) regulatory-style liquidity constraints, and (ii) behavioural dynamics that can amplify redemptions near binding thresholds. For each fund-date portfolio observed in the data, I simulate daily cash flows and redemptions over horizons of up to 30 business days and compute the maximum cumulative redemption share the portfolio can withstand without breaching the failure condition. I repeat the simulation across alternative redemption profiles and market-liquidity environments to quantify how resilience depends on both investor behaviour and the usability of assets classified as liquid. To capture variation in how usable liquid assets actually are during stress, I consider three market-liquidity environments:

1. **No asset sales (frozen secondary market).** The fund cannot sell assets to meet redemptions. It may meet outflows using cash and assets maturing on the day, including deposits and reverse repo. In particular, the fund may also choose not to roll overnight reverse-repo positions.

2. **Government asset sales.** In addition to maturing assets, the fund can sell government-issued assets. In sterling MMFs these holdings are predominantly Treasury bills (with limited gilt exposure). This regime captures the practical role of government assets in liquidity management and is aligned with the fact that government instruments can enter WLA up to the regulatory cap of 17.5%.

3. **CD sales (extension).** In addition to maturing assets, the fund can sell or exchange certificates of deposit (CDs). This regime is motivated by evidence from recent stress episodes in which sterling MMFs reduced CD holdings (see Table 5). Because CD secondary-market liquidity is more state-dependent than government debt market, I treat this regime as an extension (see Appendix A2) that also allows saleability to vary with market depth.

Within regimes 2 and 3, I also allow market depth to vary by restricting the fraction of eligible assets that can be monetised (0%, 50%, or 100%), which provides a reduced-form way to capture stressed-market illiquidity.

The simulation is built on five design choices. They govern asset classification, redemption dynamics, liquidation policy, and the failure condition. Before stating these, I introduce notation. I represent the portfolio as a set of positions indexed by residual maturity. Let $a_{i,t}$ denote the market value of position i with t business days remaining to maturity, where $t = 0$ denotes instruments that mature today. Let n_t be the number of positions with residual maturity t, and let T be the maximum residual maturity in the portfolio. Define total portfolio value:

$$V = \sum_{t=0}^{T} \sum_{i=1}^{n_t} a_{i,t}. \tag{1}$$

Let C_t denote cash and other daily liquid assets (DLA-eligible) available at the start of day t.

Defining weekly liquid assets. Regulation defines WLA as instruments maturing within five business days (including cash), plus a capped contribution from government assets for LVNAV/CNAV funds.

Let $w^W_{i,t} = \mathbf{1}\{t \leq 5\}$ indicate assets maturing within five business days, and let $w^G_{i,t}$ indicate government-issued assets. I cap the contribution of government assets to WLA at 17.5 percentage points of total assets, consistent with MMF regulations (FCA, 2023):

$$G = \min\left\{ \sum_{t=0}^{T} \sum_{i=1}^{n_t} w^G_{i,t} a_{i,t},\ 0.175\, V \right\}. \tag{2}$$

Then WLA at the start of the simulation (seed date) is:

$$\text{WLA} = \frac{\sum_{t=0}^{T} \sum_{i=1}^{n_t} w^W_{i,t} a_{i,t} + G}{V}. \tag{3}$$

Redemption shocks and redemption profiles. For each simulation, a cumulative fraction $p \in (0,1)$ of initial AUM is redeemed over a horizon of $n \in \{5, 10, 20, 30\}$ business days. Daily redemptions follow a profile $R = (R_1, \dots, R_n)$ satisfying

$$\sum_{t=1}^{n} R_t \;=\; p \cdot \mathrm{AUM}_0. \tag{4}$$

I consider seven stylised profiles (Bimodal, Exponential Decay, Left-skewed, Random, Right-skewed, Uniform, and Unimodal), illustrated in Figure 3. In the main results I focus on unimodal redemptions as observed historically, e.g. in the 2020 DFC episode and left-skewed redemptions as an extreme, front-loaded stress case.

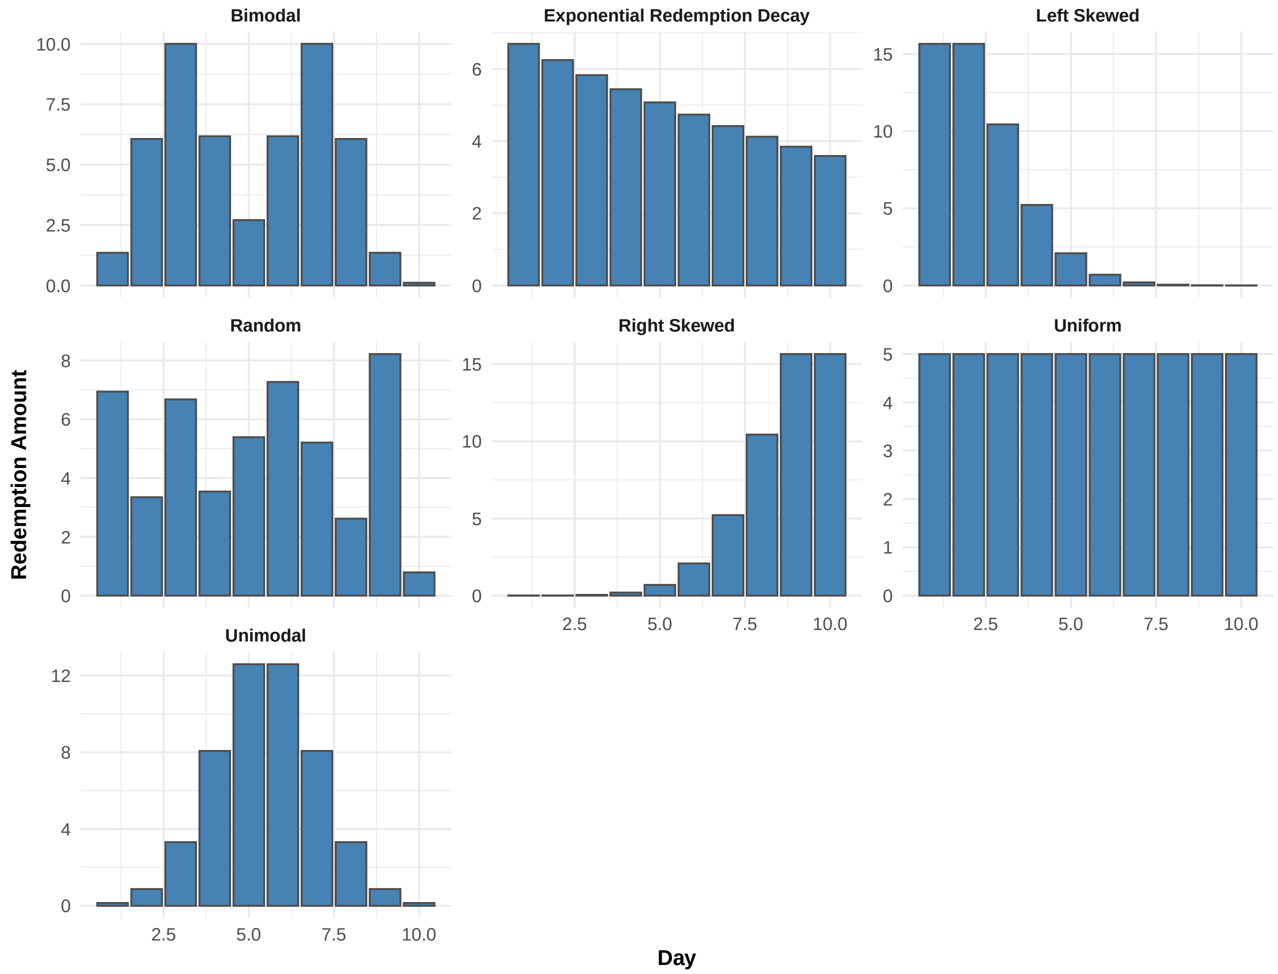

Figure 3: Illustrative redemption profiles over a 10-day horizon, scaled to a total cumulative redemption of 50% of initial AUM.

Threshold-driven redemption acceleration. Redemptions may accelerate as liquidity buffers approach the regulatory minimum. Motivated by evidence of threshold effects near 30% WLA (Dunne and Giuliana, 2022), I introduce an acceleration component that activates once WLA falls within 10 percentage points of the 30% threshold. Concretely, I set an activation threshold $\delta = 0.40$ and increase daily redemptions by 1 percentage point of AUM for each 5 percentage point decline in WLA below 40% (as shown in Figure 4). Total redemptions are capped by the previous day's AUM.

Let r_t denote the non-accelerated redemptions as a share of total assets implied by the chosen profile. The accelerated redemption share is:

$$r_t^* = \min\left\{ r_t + \max\left(0,\ 0.02 + 0.2\delta - 0.2\,\mathrm{WLA}_t\right),\ 1 \right\}. \tag{5}$$

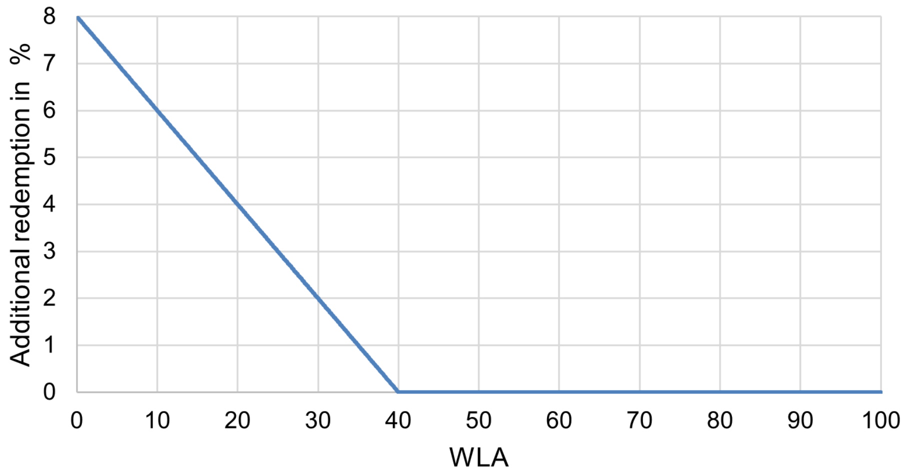

Figure 4: Illustration of threshold-driven redemption acceleration as WLA approaches regulatory thresholds.

Liquidation policy. Funds meet redemptions using available cash and maturing assets first. If asset sales are permitted, the fund sells eligible assets from the long end of the maturity spectrum to preserve short-dated liquidity and maintain WLA. Let $w_{i,t}^S$ indicate whether a position is sellable under the regime, and let $\eta \in \{0, 0.5, 1\}$ denote the maximum sellable fraction (market depth). I also apply an optional proportional haircut $\gamma \in [0, 0.0028)$ to sale proceeds, calibrated to the spread estimates reported in Baes, Bouveret, and Schaanning (2023). In practice, I found this haircut to have negligible effect on simulated redemption capacity given its small magnitude, and results are robust to setting $\gamma = 0$.

Given required cash outflow R_t, the fund can raise cash through:

$$\mathrm{CashAvailable}_t = C_t + \sum_{t=0}^{T} \sum_{i=1}^{n_t} \left(w_{i,k}^S \cdot \eta \cdot (1-\gamma)\, a_{i,t} \right), \tag{6}$$

where liquidation is implemented from longest to shortest maturities until the redemption is met or eligible assets are exhausted.

Failure condition. A portfolio is deemed to fail if daily liquid assets fall below 10% of total assets:

$$\frac{C_t}{V_t} \;<\; q, \qquad q = 0.10, \tag{7}$$

where V_t is total portfolio value at day t. Although in the UK MMFR, the dual-trigger framework is defined using the 30% WLA threshold (combined with 10% daily outflows), I adopt a 10% DLA threshold because DLA provides a more accurate measure of immediately realisable liquidity. Funds can satisfy the WLA requirement while holding substantial volumes of WLA-eligible but difficult-to-sell government securities, whereas DLA captures genuinely same-day-usable liquid resources. For this reason, 10% DLA serves as the relevant operational benchmark for a severe daily liquidity shortfall.

Given these assumptions, the fund's daily mechanics follow a fixed sequence, summarised in Figure 5. For a given portfolio, horizon, market regime, and redemption profile, I estimate redemption capacity via reverse stress testing. I search for the largest cumulative redemption share p such that the fund survives the full horizon without breaching Equation 7. To do so, I use a binary-search procedure (Knuth, 1998) over $p \in [0, 1]$, evaluating survival at each candidate p and iterating until the interval width falls below 0.01 percentage points. This yields a high-precision estimate of redemption capacity for each portfolio-profile-regime combination.

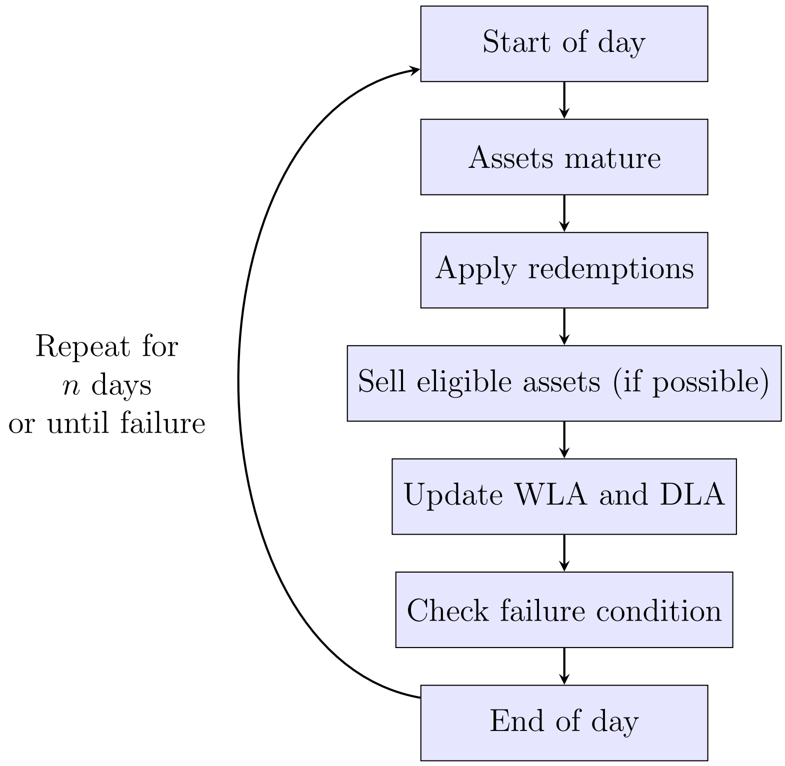

Figure 5: Simulation loop used for reverse stress testing.

I run a full-factorial grid across fund-date portfolios, liquidation regimes (none, government, CDs), market depth (0%, 50%, 100%), acceleration (on/off), horizons (5, 10, 20, 30 days), and redemption profiles. For each combination, I compute the maximum deliverable redemption share. The full grid implies 322,560 runs per redemption-profile shape.

Resilience metrics. Redemption capacity gives the maximum shock a fund can absorb; failure rates complement this by summarising how often observed portfolios would breach the DLA constraint under a fixed shock of size θ. To summarise the relationship between initial liquidity and redemption capacity across portfolios, I estimate:

$$\text{MaxOutflow} \;=\; \alpha \;+\; \beta \cdot \text{InitialWLA}. \tag{8}$$

To compute failure rates, let (f, t) index a fund and the date its portfolio is observed (the *seed date*), and let $h \in \{5, 10, 20, 30\}$ denote the horizon. For a given shock threshold θ, I define redemption capacity, survival, and failure as:

$$\begin{aligned} \mathit{max\_share}_{f,t,h} &= \text{estimated redemption capacity for } (f, t, h), &(9)\\ \text{survival}_{f,t,h} &= \mathbf{1}\{\mathit{max\_share}_{f,t,h} > \theta\}, &(10)\\ \text{failure}_{f,t,h} &= 1 - \text{survival}_{f,t,h}. &(11) \end{aligned}$$

I summarise failure rates in 5 percentage point WLA bins, using the UK business-day calendar throughout. The failure probability for a given WLA bin b, horizon h, and shock size θ is computed as the sample mean of the failure indicator:

$$p_{\text{fail}}(b, h, \theta) = E[\text{failure}_{f,t,h}(\theta) \,|\, \text{WLA}_{f,t} \in b]\,. \tag{12}$$

Unless otherwise noted, I restrict attention to seed-date portfolios with AUM $\geq$ £1 billion.

3. Results

In this section, I present simulation evidence on the resilience of UK sterling MMFs under stressed redemptions. I organise the results around two outcome measures. First, *redemption capacity*: the maximum cumulative redemption share a fund can meet over a given horizon without breaching the daily liquidity constraint. Second, *failure probability*: the fraction of simulated paths in which the fund breaches the constraint under a given shock size, redemption profile and shock horizon. I use these outcomes to: (i) characterise a baseline resilience frontier as a function of initial WLA, (ii) quantify how secondary-market liquidity assumptions shift this frontier, (iii) evaluate a policy counterfactual that removes threshold-driven acceleration near the regulatory minimum, and (iv) translate these mechanisms into failure risk under alternative redemption profiles.

Key definitions. To keep the results manageable, I group funds into initial WLA buckets (30–40%, 40–50%, etc.) and evaluate stress horizons of 5, 10, 20, and 30 business days. Unless otherwise stated, results are averaged across redemption-profile assumptions; where applicable, 95% confidence intervals are reported. This discretisation can introduce measurement error: for example, funds with starting WLA of 30.01% and 39.99% are treated as equivalent within the 30–40% bucket, even though their mean redemption capacity may differ.

3.1 Redemption capacity under a frozen secondary market (baseline)

I begin with a conservative benchmark in which secondary markets are effectively frozen and funds cannot sell assets to meet redemptions. In this setting, funds meet redemptions using only cash and assets maturing on the day (including overnight maturities). Figure 6 summarises average redemption capacity across initial WLA buckets and redemption-profile design choices, faceted by the stress horizon (5, 10, 20, and 30 business days). These estimates offer a direct mapping from initial WLA to stress absorption in the absence of asset sales and provide a lower-bound benchmark for resilience.

Two patterns stand out. First, redemption capacity increases with initial WLA and with horizon length, consistent with a larger liquidity buffer and more time for contractual cash flows to arrive. Second, the *timing* of redemptions materially affects resilience. Front-loaded profiles (e.g., left-skewed) reduce capacity disproportionately, particularly for funds close to the regulatory minimum, whereas more gradual or back-loaded profiles yield higher capacity. This highlights a key advantage of the simulation framework: it can evaluate not only shock size but also the shape of redemption pressure.

3.2 Market counterfactuals: how secondary-market liquidity shifts capacity

I next evaluate how redemption capacity changes when funds can monetise assets in secondary markets. This step matters because the regulatory definition of WLA implicitly assumes that assets counted as “liquid” remain saleable during stress, while market depth can deteriorate precisely when many funds attempt to sell simultaneously. I therefore treat market liquidity as a counterfactual dimension rather than a fixed assumption.

Liquid government market. In Figure 7, I allow funds to sell government assets to meet redemptions. Relative to the frozen-market baseline, the capacity frontier shifts outward

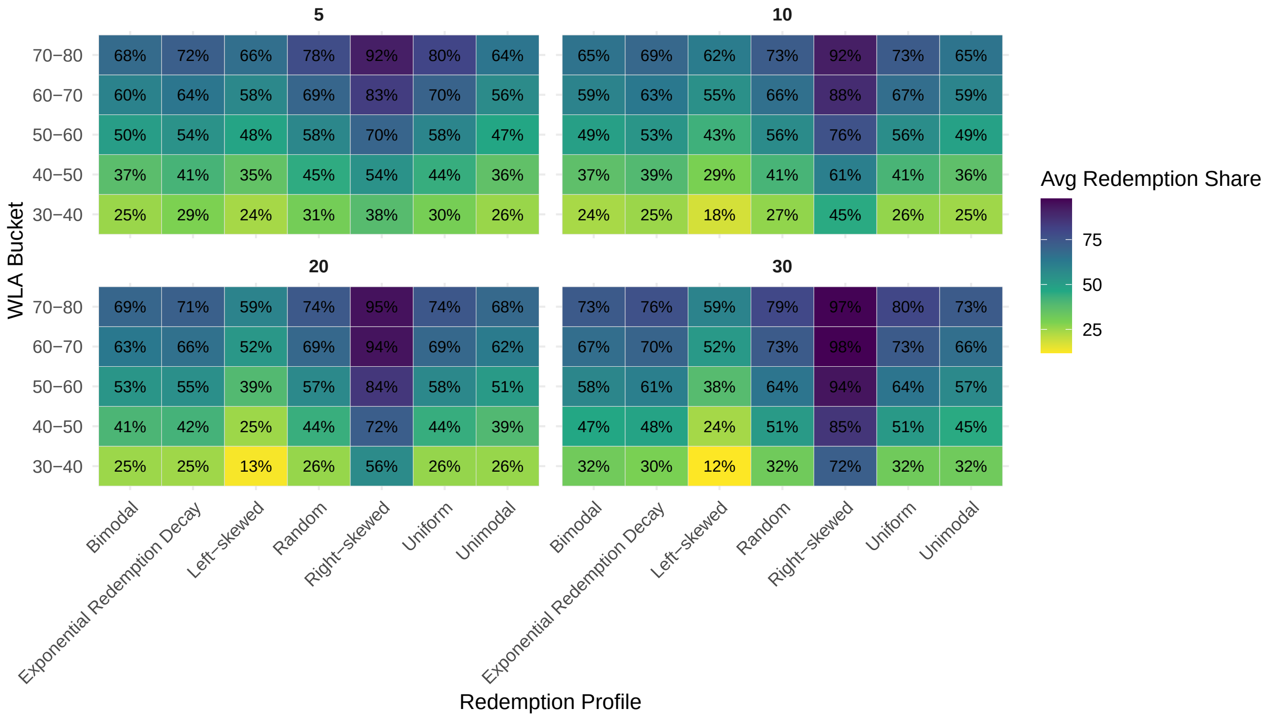

Figure 6: Heatmap of average redemption capacity (baseline: no asset sales). Average maximum cumulative redemption share (%) by initial WLA bucket and redemption profile, faceted by stress horizon. Darker tiles indicate higher redemption capacity.

across WLA buckets and redemption profiles, with the largest improvements for lower-WLA funds and for front-loaded redemption profiles. Intuitively, the ability to sell government assets provides the funds a bigger pool of assets that can be monetised which is particularly helpful when redemptions come quickly at the start of the stress window.

Tight government market. In Figure 8, I impose tighter market conditions by restricting the fraction of government assets that can be sold. The capacity gains shrink and dispersion across profiles widens, demonstrating that resilience is sensitive not only to measured WLA but also to whether WLA assets remain saleable at scale. Taken together, these two figures show that the same initial WLA level can imply very different resilience depending on market functioning.

Outward shift relative to frozen benchmark. To make the mechanism explicit, Figure 9 plots the change in capacity (Δ, in percentage points) relative to the frozen benchmark under the tight and liquid government-market counterfactuals. The comparison clarifies that the stabilising role of WLA depends on the stress-realised sellability of the assets that contribute to measured WLA (including the capped government component). When saleability is high, the frontier shifts outward; when saleability is constrained, the marginal resilience value of measured WLA is smaller and more profile-dependent.

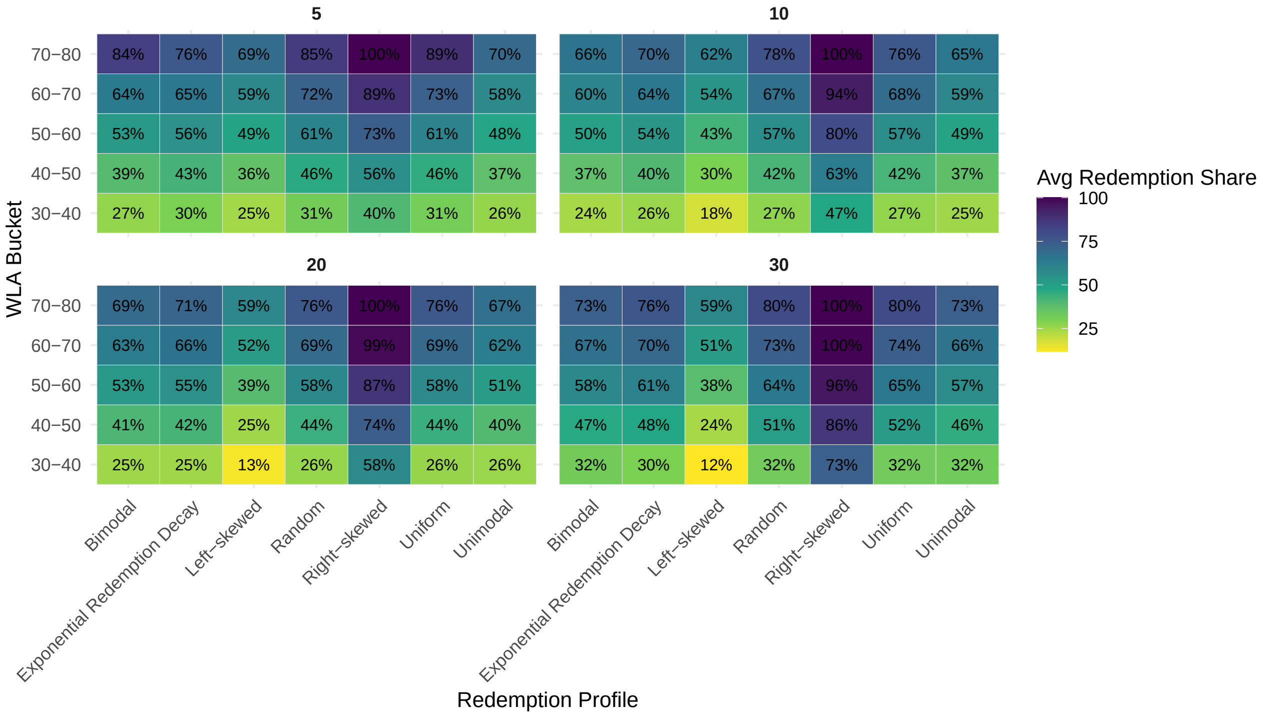

Figure 7: Heatmap of average redemption capacity (market counterfactual: government asset sales). Average redemption capacity (%) when funds may sell government assets, by WLA bucket and redemption profile, faceted by horizon.

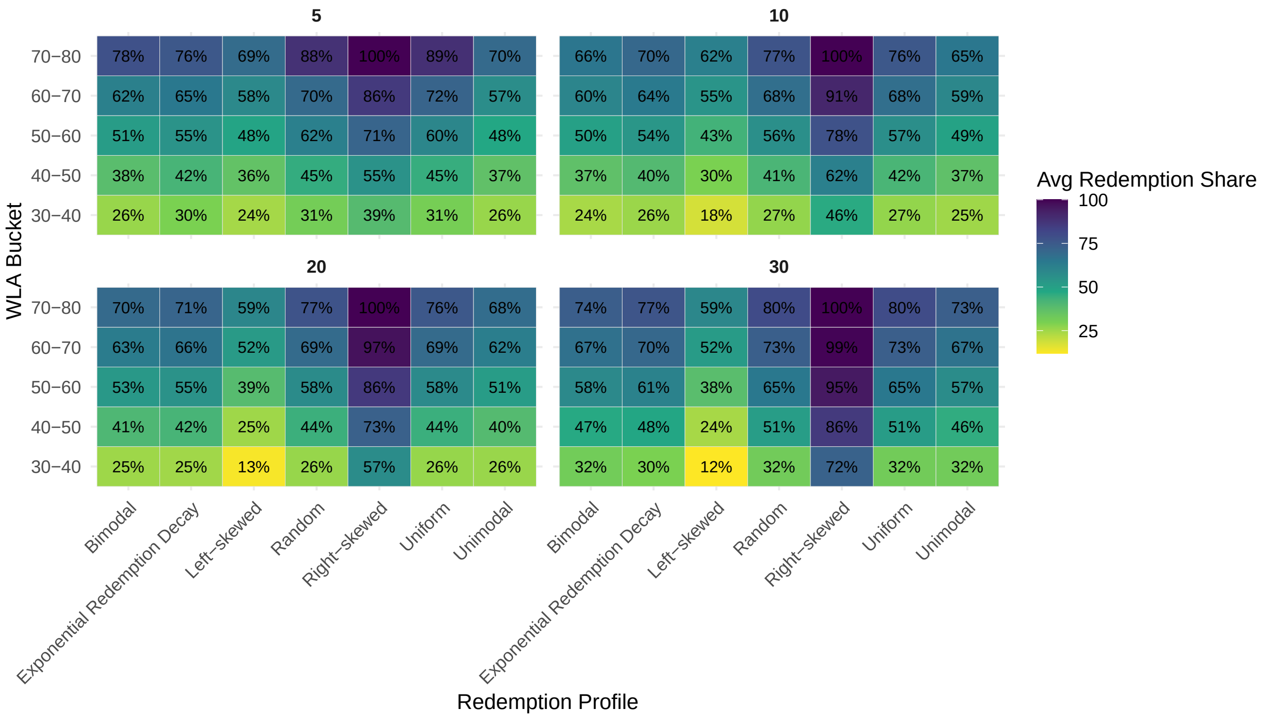

Figure 8: Heatmap of average redemption capacity (market counterfactual: tight government market). Average redemption capacity (%) when secondary-market depth is reduced (restricted government asset sales), by WLA bucket and redemption profile, faceted by horizon.

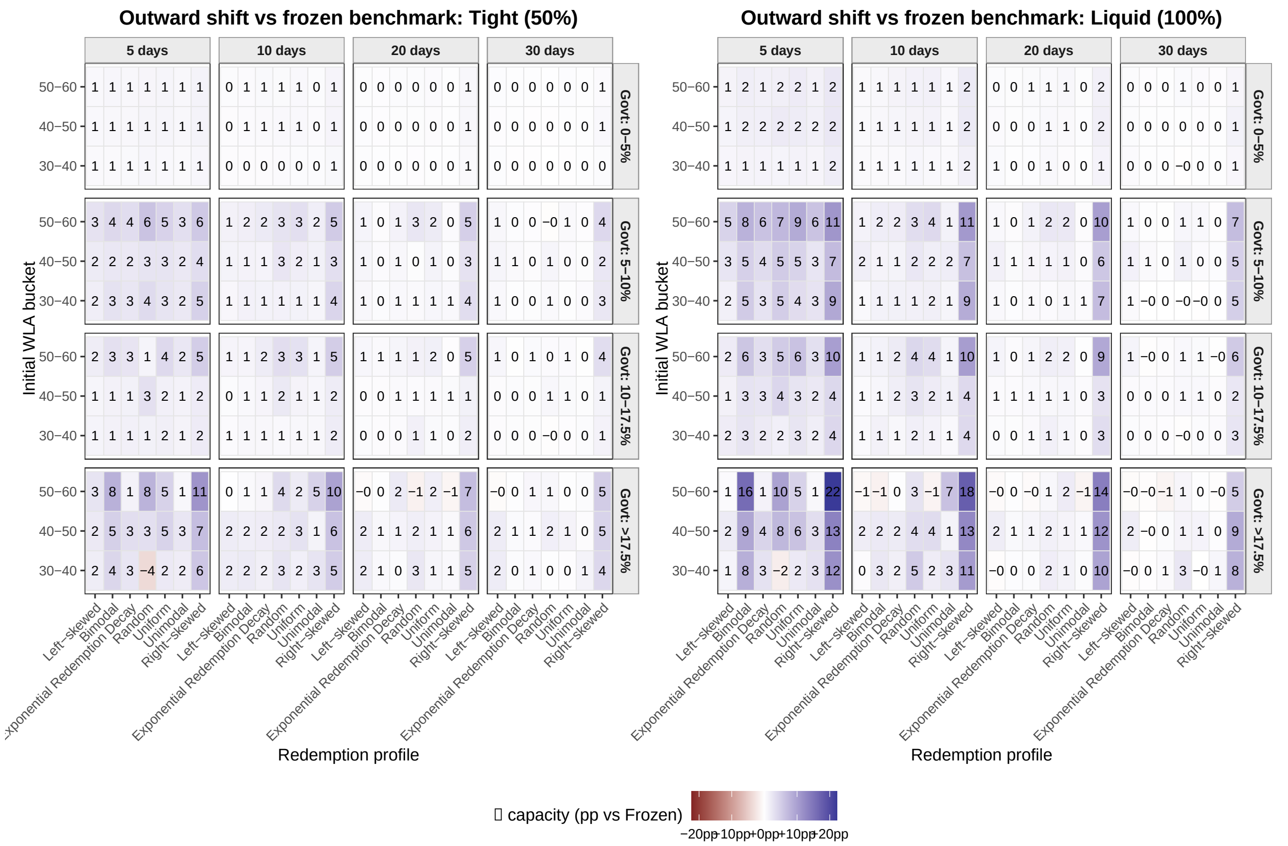

Figure 9: Outward shift in redemption capacity relative to the frozen-market benchmark. Heatmaps report Δ redemption capacity (percentage points) relative to the frozen-market benchmark, comparing a tight government market (50% of government holdings sellable) to a liquid government market (100% sellable). Results are shown by initial WLA bucket and redemption profile, faceted by shock horizon and initial government-share bucket.

3.3 Policy counterfactual: removing threshold-driven acceleration

LVNAV and CNAV MMFs with WLA close to the regulatory minimum experienced disproportionately large outflows during the DFC episode (Dunne and Giuliana, 2022). A leading explanation is that threshold-linked measures (fees/gates) can induce pre-emptive redemptions, generating a discontinuity in investor behaviour near the threshold. In my framework, I capture this channel through redemption acceleration design choice, which increases redemption intensity as funds approach the regulatory minimum. Since recent policy discussions consider removing the linkage between breaching the WLA threshold and the imposition of fees/gates, I repeat the simulations with acceleration switched off and compare outcomes.

Figure 10 reports the change in redemption capacity at 0% asset sales (or frozen secondary market), defined as $\Delta = \text{no acceleration} - \text{acceleration}$. Two results stand out. First, removing acceleration increases simulated redemption capacity across all WLA buckets and horizons, with the largest gains concentrated among funds closest to the regulatory minimum. Second, the gains increase with the stress horizon, indicating that threshold-driven acceleration compounds over longer stress windows.

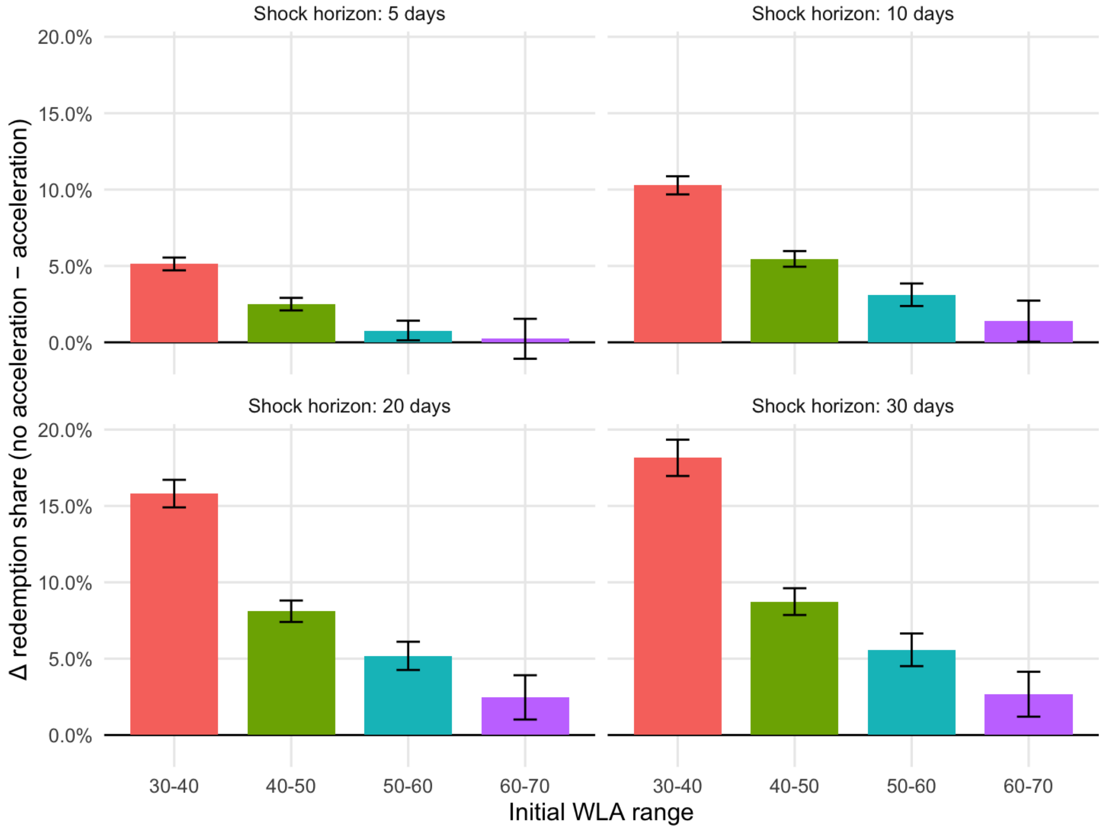

Figure 10: Policy counterfactual: change in redemption capacity when threshold-driven acceleration is removed. Bars show $\Delta =$ no acceleration $-$ acceleration at 0% asset sales, by initial WLA range and shock horizon; error bars denote 95% confidence intervals. Positive values indicate higher simulated redemption capacity under the no-acceleration specification. I computed the difference in means and used a conservative standard error treating group means as independent. In the group where Δ was close to zero, this interval crossed zero and the lower bound could be negative, indicating no statistically clear difference in that subgroup.

3.4 Failure risk under alternative redemption patterns

While redemption capacity provides a continuous measure of resilience, policy is often concerned with tail outcomes. I therefore report *failure probabilities* defined as the share of simulated paths in which a fund breaches the daily liquidity constraint. I focus on two redemption profiles that bracket plausible investor behaviour in stress. The **unimodal** profile represents an "expected-case" pattern consistent with historically observed episodes (see Falato, Goldstein, and Hortaçsu (2021)), whereas the **left-skewed** profile represents a more adverse, front-loaded pattern that concentrates redemptions early in the stress window, before some assets mature, and is intended to proxy an *extreme* run dynamic. In both cases, acceleration of redemptions has been switched off.

Note rows in Table 1 to 4 report the lower bound of each 5 percentage point initial WLA bucket (e.g., "30%" denotes portfolios with $30\% \leq$ WLA $< 35\%$, "35%" denotes $35\% \leq$ WLA $< 40\%$, and so on). Columns report stress horizons (5, 10, 20, and 30 business days), with each horizon split into two shock sizes. I report both *observed but unlikely* shocks (20%, 25%) and *extreme/historically unobserved* shocks (of 30% and 40%).

Historically observed behaviour: unimodal outflows. Tables 1–2 report failure probabilities under the unimodal profile. Two patterns emerge. First, for moderate shocks (of 20% and 25%), failure is concentrated in low-WLA buckets and short horizons; at longer horizons, failure probabilities fall sharply as maturing assets replenish liquidity. Second, under larger shocks (of 30% and 40%), failure rates remain sensitive to initial WLA, but the improvement with horizon length is more pronounced than under front-loaded profiles (discussed below), reflecting the fact that a unimodal pattern does not exhaust liquidity immediately at the start of stress.

Table 1: Unimodal redemptions: Average failure probability under observed but unlikely shocks (20%, 25%)

	5-day		**10-day**		**20-day**		**30-day**	
WLA	20%	25%	20%	25%	20%	25%	20%	25%
30%	8.6%	37.1%	0.0%	8.6%	0.0%	0.0%	0.0%	0.0%
35%	0.0%	10.4%	0.0%	1.0%	0.0%	0.0%	0.0%	0.0%
40%	0.0%	2.8%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
45%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
50%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%

Table 2: Unimodal redemptions: Average failure probability under extreme/unobserved shocks (30%, 40%)

	5-day		**10-day**		**20-day**		**30-day**	
WLA	30%	40%	30%	40%	30%	40%	30%	40%
30%	97.1%	100.0%	45.7%	100.0%	2.9%	60.0%	0.0%	22.9%
35%	52.1%	100.0%	11.5%	92.7%	3.1%	37.5%	0.0%	9.4%
40%	10.4%	96.2%	0.9%	66.0%	0.0%	9.4%	0.0%	0.9%
45%	1.1%	44.1%	0.0%	12.9%	0.0%	0.0%	0.0%	0.0%
50%	0.0%	4.8%	0.0%	2.4%	0.0%	0.0%	0.0%	0.0%

Extreme behaviour: left-skewed (front-loaded) outflows. Tables 3–4 report failure probabilities under left-skewed outflows. Relative to unimodal redemptions, front-loading materially increases failure risk at low WLA because the constraint is reached before asset matures and other contractual cash flows can help meet redemptions. This is especially visible under unobserved shocks (30–40%), where failure probabilities remain elevated across horizons for low-WLA buckets. Put differently, the same cumulative shock can be survivable under a smoother redemption path but destabilising when redemptions are concentrated early.

Table 3: Left-skewed redemptions: Average failure probability under observed but unlikely shocks (20% and 25%)

	5-day		**10-day**		**20-day**		**30-day**	
WLA	20%	25%	20%	25%	20%	25%	20%	25%
30%	8.6%	34.3%	8.6%	34.3%	8.6%	34.3%	8.6%	34.3%
35%	0.0%	5.2%	0.0%	3.1%	0.0%	3.1%	0.0%	3.1%
40%	0.0%	2.8%	0.0%	2.8%	0.0%	2.8%	0.0%	2.8%
45%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
50%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%

Table 4: Left-skewed redemptions: Average failure probability under extreme/unobserved shocks (30% and 40%)

	5-day		**10-day**		**20-day**		**30-day**	
WLA	30%	40%	30%	40%	30%	40%	30%	40%
30%	97.1%	100.0%	97.1%	100.0%	97.1%	100.0%	97.1%	100.0%
35%	41.7%	100.0%	42.7%	100.0%	42.7%	100.0%	42.7%	100.0%
40%	10.4%	90.6%	10.4%	91.5%	10.4%	91.5%	10.4%	91.5%
45%	1.1%	28.0%	1.1%	37.6%	1.1%	37.6%	1.1%	37.6%
50%	0.0%	4.8%	0.0%	4.8%	0.0%	4.8%	0.0%	4.8%

The identical values across horizons reflect a structural feature of the left-skewed shock design. When redemptions are heavily front-loaded, funds either breach the liquidity threshold within the first few days—producing the same failure outcome regardless of horizon—or remain comfortably above the threshold throughout the stress window, again yielding minimal variation across horizons. Only when initial WLA is close to the failure region do longer horizons meaningfully increase risk. Thus, the horizon effects collapse in cases where failure (or survival) is effectively determined by the early-day portion of the shock.

Taken together, the unimodal and left-skewed patterns indicate that failure risk is jointly shaped by: (i) the *level* of initial liquidity (WLA), (ii) the *size* of the redemption shock, and (iii) the *timing* of redemptions within the stress window. In my simulations, increasing WLA buffers from the current regulatory minimum of 30% to moderately higher levels delivers the largest reductions in failure probability, especially under front-loaded redemption scenarios (except in the most extreme cases). Beyond this range, further increases in liquidity provide diminishing marginal benefits. Overall, the results underscore the motivation for the framework: assessing resilience requires accounting not only for shock size, but also for investor behaviour and market dynamics that can amplify stress.

Results should be read in light of the framework's simplifying assumptions. Shocks are applied to portfolios; asset sales are constrained by exogenous saleable-fraction limits over fixed liquidation horizons; and prices do not adjust endogenously to aggregate sales. Investor flows, policy backstops, and trading frictions are imposed rather than modelled, and cross-fund interactions or liquidity-spiral feedbacks are outside scope. For robustness, I vary shock size and path, saleable-fraction caps, shock styles (front-loaded vs smoothed), and shock time horizons.

4. Conclusion

In this paper, I develop and apply a simulation framework to quantify the resilience of sterling MMFs under stressed redemptions. The framework is designed to answer four related questions: (i) how do initial WLA map into *redemption capacity* over different stress horizons, (ii) how this mapping shifts when market liquidity allows funds to monetise assets, (iii) how does the policy design, in particular, threshold-driven redemption acceleration near the regulatory WLA minimum affect resilience, and (iv) how do these mechanisms translate into *failure probabilities* under alternative redemption profiles that represent historically observed and extreme investor behaviour.

I have three main findings. First, redemption capacity is non-linear in initial WLA, and the timing of outflows matters independently of their size. In the frozen-market baseline (no asset sales), redemption capacity increases monotonically with initial WLA and with the length of the stress horizon—though this monotonicity holds on average: a fund with 35% WLA composed entirely of short-dated instruments can, in some configurations, prove more resilient than one with 40% WLA that exhausts its 17.5% government-asset allowance. More importantly, the relationship is strongly shaped by the *time profile* of redemptions. Front-loaded outflows materially reduce capacity, particularly for funds close to the regulatory minimum, whereas more gradual (unimodal) redemptions are less destabilising for a given cumulative shock. This highlights a central contribution of the framework: assessing MMF resilience requires modelling not only shock size but also the time profile of redemptions.

Second, the stabilising role of WLA depends on whether WLA assets remain liquid in periods of stress. When I allow funds to sell government assets, the redemption-capacity frontier shifts favourably outward relative to the frozen-market benchmark, especially for funds in low WLA buckets and for front-loaded redemption patterns. However, when market conditions are tight and only part of those government holdings can be sold, the improvement, due to WLA increases, is much smaller, and outcomes become more uneven across redemption profiles (some outflow patterns become far harder to withstand than others). These results highlight a key policy nuance: a higher measured WLA buffer improves resilience most reliably when it is achieved through assets that generate cash flows under stress (short maturities), rather than through assets that are often assumed – but cannot be guaranteed – to be liquid in secondary markets precisely when market depth could be impaired. Because our WLA measure allows government securities to contribute up to 17.5% of the total assets, funds can appear to hold large liquidity buffers even though a substantial part of that buffer can depend on secondary-market depth. CD-sale scenarios reported in the Appendix reinforce this point: when private-instrument liquidity is assumed to be abundant, the marginal resilience value of additional WLA diminishes; when it is constrained, dependence on initial WLA strengthens.

Third, motivated by evidence of threshold effects around that minimum—triggered when rules require funds to consider fees, gates, or suspending redemptions once WLA falls closer to the threshold—I implement a counterfactual that removes threshold-driven redemption acceleration. Under a no-asset-sales assumption, switching off acceleration increases simulated redemption capacity across WLA buckets and redemption horizons, with the largest gains concentrated in funds closest to the minimum and at longer horizons.

These mechanisms translate directly into tail risk. Comparing failure probabilities under unimodal (historically observed) and left-skewed (extreme, front-loaded) redemption profiles shows that failure risk is jointly determined by the *level* of initial liquidity, the *size* of the shock, and the *timing* of redemptions. Failure rates are materially higher under left-skewed outflows, particularly at low WLA and under larger shocks, because funds run short of liquidity before assets mature and other contractual cash flows can provide relief.

A further implication of the failure-rate results is that for a given shock resilience gains exhibit diminishing returns as initial WLA increases. Under historically extreme shocks without accelerated redemptions, moving from the regulatory minimum into moderately higher WLA buckets delivers the largest reductions in failure probability, particularly for front-loaded redemption. In my simulations, the reduction in failure risk is concentrated in the transition from 30% to 40% WLA; beyond 40%, additional increases in WLA typically generate comparatively smaller marginal improvements in survival across horizons. This non-linearity suggests that the primary resilience benefits come from lifting funds away from the region where shocks are more likely to bind rather than from increasing buffers at already high WLA levels.

Taken together, the results yield three principal policy implications. First, they provide strong support for reforms aimed at eliminating threshold-driven amplification effects in MMFs. These threshold dynamics disproportionately weaken resilience for funds operating near minimum liquidity requirements, and removing them materially reduces failure risk across a broad range of redemption profiles. Second, conditional on the removal of such threshold effects, the simulations indicate that raising the minimum WLA requirement above the current 30% floor would significantly strengthen resilience. Increasing the requirement to around 40% delivers a sizeable reduction in failure probabilities across both historically observed extreme redemption scenarios, while further increases toward 50% generate comparatively modest incremental gains. Third, although secondary to these structural reforms, the results also show that the effective liquidity of assets counted toward WLA—not only their quantity—can influence outcomes in stress. This underscores that, while threshold effects and higher baseline WLA levels are the primary levers, ensuring that WLA is composed of reliably liquid assets can further enhance resilience.

The simulation framework necessarily relies on simplifying assumptions. In particular, asset sales are modelled without endogenous price impact: market frictions are introduced through exogenous limits on saleable fractions rather than through a dynamic price discount linked to aggregate selling pressure. Future work could incorporate endogenous pricing and discounts, allowing market depth to respond to sales volumes and macro-financial conditions. The framework could also be extended to model cross-fund interactions and contagion channels, such as common-asset fire sales and liquidity spirals. It could further be applied to other open-ended fund sectors with similar liquidity transformation profiles. These extensions would further strengthen the framework's usefulness for assessing system-wide vulnerabilities and the effectiveness of alternative policy designs.

References

Baes, Michel, Antoine Bouveret, and Eric Schaanning. 2023. "Bang for (breaking) the buck: Regulatory constraints and money market funds reforms." *ESMA Working Paper* (2).

Bank of England. 2021. "The Role of Non-Bank Financial Intermediaries in the Dash for Cash in Sterling Markets." *Financial Stability Paper* 47.

———. 2022. "Bank of England announces gilt market operation." Press release. URL `https://www.bankofengland.co.uk/news/2022/september/bank-of-england-announces-gilt-market-operation`. Accessed: 2026-03-17.

———. 2023. "Financial Stability in Focus, July 2023." URL `https://www.bankofengland.co.uk/financial-stability-in-focus/2023/july-2023`. Accessed: 2024-10-07.

BlackRock. 2020. "Lessons from COVID-19: The Experience of European Money Market Funds in Short-Term Markets." Tech. rep.

Bouveret, Antoine, Antoine Martin, and Patrick E. McCabe. 2022. "Money Market Fund Vulnerabilities: A Global Perspective." Tech. Rep. 2022-012, Board of Governors of the Federal Reserve System, Washington, D.C. URL `https://doi.org/10.17016/FEDS.2022.012`.

Chernenko, Sergey and Adi Sunderam. 2014. "Frictions in shadow banking: Evidence from the lending behavior of money market mutual funds." *The Review of Financial Studies* 27 (6):1717–1750.

Cipriani, Marco and Gabriele La Spada. 2020. "Sophisticated and unsophisticated runs." *FRB of New York Staff Report* (956).

Czech, Robert, Bernat Gual-Ricart, Joshua Lillis, and Jack Worlidge. 2021. "The Role of Non-Bank Financial Intermediaries in the 'Dash for Cash' in Sterling Markets." *Bank of England Financial Stability Paper* (47).

Dunne, Peter G. and Raffaele Giuliana. 2021. "Do liquidity limits amplify money market fund redemptions during the COVID crisis?" ESRB Working Paper Series 127, European Systemic Risk Board.

———. 2022. "Do redemptions increase as money market funds approach regulatory liquidity thresholds?" *Central Bank of Ireland Research Technical Paper* 2022.

European Systemic Risk Board (ESRB). 2021. "Issues note on systemic vulnerabilities of and preliminary macroprudential policy considerations for leveraged loan markets and collateralised loan obligations." Report, European Systemic Risk

Board. URL https://www.esrb.europa.eu/pub/pdf/reports/esrb.report.210701_Issues_note_on_systemic_vulnerabilities~db0345a618.en.pdf. Accessed: 2024-08-12.

Falato, Antonio, Itay Goldstein, and Ali Hortaçsu. 2021. "Financial Fragility in the COVID-19 Crisis: The Case of Investment Funds in Corporate Bond Markets." *Journal of Monetary Economics* 123:35–52. URL https://finance.wharton.upenn.edu/~itayg/Files/covid19fragility-published.pdf.

FCA. 2023. "Financial Conduct Authority CP23/28: Updating the Regime for Money Market Funds." URL https://www.fca.org.uk/publications/consultation-papers/cp23-28-updating-regime-money-market-funds. Accessed: 2024-10-07.

Financial Conduct Authority (FCA). 2023. "Consultation Paper: CP23/28." URL https://www.fca.org.uk/publication/consultation/cp23-28.pdf. Accessed: 2024-08-06.

Gorton, Gary and Andrew Metrick. 2012. "Securitized banking and the run on repo." *Journal of Financial economics* 104 (3):425–451.

Ivashina, Victoria, David S Scharfstein, and Jeremy C Stein. 2015. "Dollar funding and the lending behavior of global banks." *The Quarterly Journal of Economics* 130 (3):1241–1281.

Knuth, Donald E. 1998. *The Art of Computer Programming, Volume 3: Sorting and Searching*. Reading, MA: Addison-Wesley, 2 ed.

Li, Lei, Yi Li, Marco Macchiavelli, and Xing Zhou. 2021. "Liquidity restrictions, runs, and central bank interventions: Evidence from money market funds." *The Review of Financial Studies* 34 (11):5402–5437.

McCabe, Patrick. 2010. "The Cross Section of Money Market Fund Risks and Financial Crises." Tech. Rep. 2010-51, Board of Governors of the Federal Reserve System. URL https://www.federalreserve.gov/pubs/feds/2010/201051/index.html.

McLaughlin, Tim. 2020. "Fed's money market move lifts Northern Trust fund above key threshold." Reuters. URL https://finance.yahoo.com/news/feds-money-market-move-lifts-233101457.html. Syndicated Reuters copy.

Pedersen, Lasse Heje. 2009. "When everyone runs for the exit." Tech. rep., National Bureau of Economic Research.

SEC. 2023. "U.S. Securities and Exchange Commission Final Rule: Release No. 33-11211, 2023." URL https://www.sec.gov/files/rules/final/2023/33-11211.pdf. Accessed: 2024-10-07.

Appendix

A1 Historical Outflows

I calculate cumulative outflow as the percentage change in asset under management (AUM) between an initial seed date and a subsequent date $t + h$ (where $h \in \{5, 10, 20\}$ business days), as defined in Equation 13.

$$\text{Cumulative Outflow}_{i,t} = \left(\frac{\text{AUM}_{i,t+h}}{\text{AUM}_{i,t}} - 1 \right) \times 100 \tag{13}$$

Table 5: Change in asset allocation pre and post LDI

Asset Type	31/08/2022	30/09/2022	Proportion Chng
Certificate of Deposit	66 B	43 B	-12%
Government Agency Repurchase Agreement	21 B	13 B	-5%
Other Note	17 B	11 B	-4%
Financial Company Commercial Paper	20 B	14 B	-3%
Treasury Repurchase Agreement	9 B	8 B	-1%
Asset Backed Commercial Paper	8 B	7 B	-1%
Treasury Debt	3 B	2 B	0%
Other Instrument	4 B	4 B	0%
Other Repurchase Agreement	3 B	4 B	0%
Government Agency Debt	1 B	1 B	0%
Non-Negotiable Time Deposit	25 B	25 B	0%
Non-Financial Company Commercial Paper	3 B	4 B	0%

A2 Extension: Certificates of Deposit (CD) sales

This appendix reports additional market-liquidity counterfactuals in which funds can meet redemptions by selling certificates of deposit (CDs). I treat these scenarios as an extension rather than a main-text result because secondary-market liquidity for CDs is more state-dependent than for government bills and therefore requires stronger assumptions about market depth, haircuts, and the fraction of the position that can be liquidated during stress. The purpose of this appendix is twofold: (i) to illustrate how allowing liquidation of private money-market instruments can shift the redemption-capacity frontier, and (ii) to assess whether the main conclusions about the role of initial weekly liquid assets (WLA) remain robust under alternative assumptions about CD market functioning.

Setup and interpretation

The simulations follow the baseline framework and the market-liquidity counterfactuals in Section 2.1. Funds face exogenous redemption shocks over horizons of 5, 10, 20, and 30 business days and may meet redemptions using (a) cash and assets maturing on the day and, in the counterfactuals below, (b) CD sales up to a specified fraction of holdings. As in the main text, *redemption capacity* is defined as the maximum cumulative redemption share the fund can absorb without breaching the daily liquidity constraint. Results are summarised by initial WLA bucket and redemption profile; where relevant, I also distinguish the least resilient funds within each WLA bucket.

A2.1 High CD market liquidity: sellable fraction = 100%

I first consider a stylised case in which funds can liquidate their CD holdings without binding market-depth constraints (sellable fraction = 100%). Figure 11 summarises average redemption capacity by WLA bucket and redemption profile, faceted by horizon. Two patterns emerge. First, redemption capacity is substantially higher than in the frozen-market baseline, reflecting the additional liquidation channel provided by CDs. Second, when CD liquidity is abundant, the relationship between initial WLA and redemption capacity becomes flatter: once funds can reliably monetise private instruments, marginal increases in WLA provide smaller incremental resilience gains, particularly at longer horizons. This result highlights that the mapping from WLA to resilience depends on whether liquidity buffers can be converted into cash under stress.

A2.2 Tight CD market liquidity: sellable fraction = 50%

I next restrict CD market depth by limiting sales to 50% of CD holdings. This counterfactual captures a stress environment in which dealer intermediation is limited or haircuts/market impact effectively constrain liquidation volumes. Figure 12 shows that once CD market liquidity is constrained, the positive relationship between initial WLA and redemption capacity re-emerges more clearly. In particular, low-WLA funds remain substantially more vulnerable under front-loaded redemption profiles, while higher-WLA buckets provide a materially larger buffer. This sensitivity illustrates that "liquidity" is conditional: the resilience contribution of private instruments depends on their saleability precisely when many funds may attempt to sell concurrently.

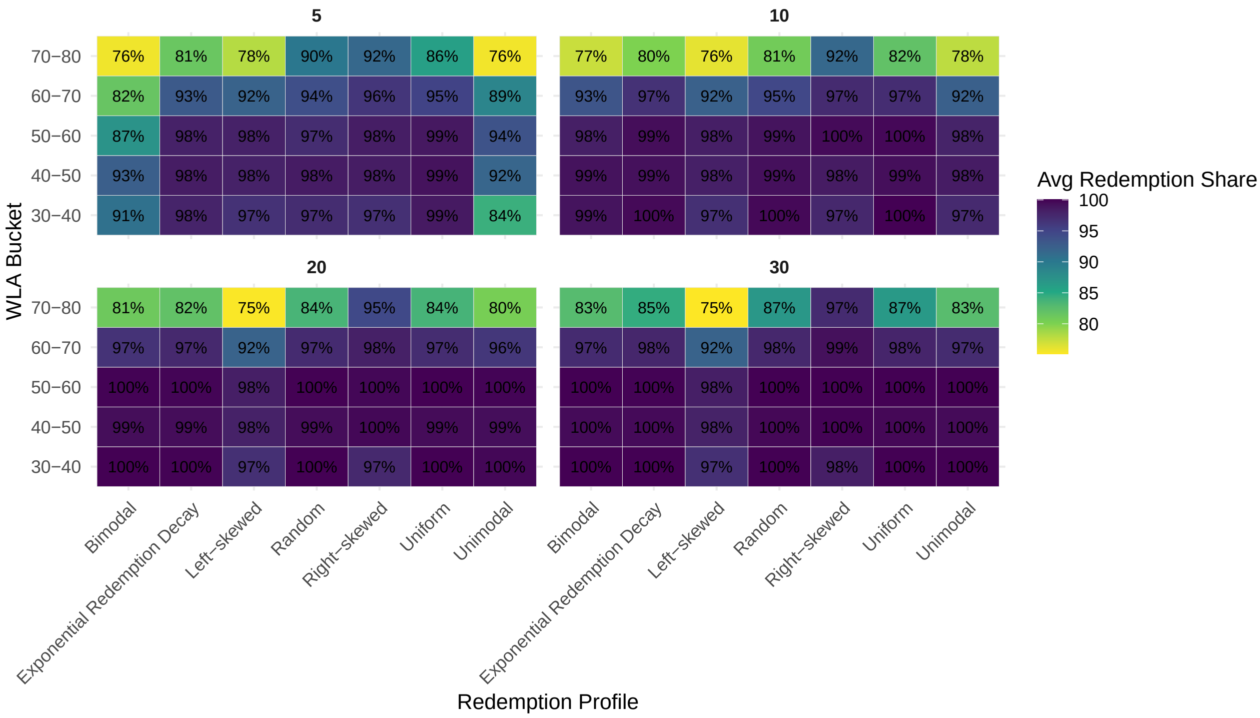

Figure 11: Heatmap of redemption capacity (extension: CD sales, high-liquidity case). Average redemption capacity (%) by initial WLA bucket and redemption profile when funds can sell up to 100% of CD holdings, faceted by stress horizon. Darker tiles indicate higher capacity.

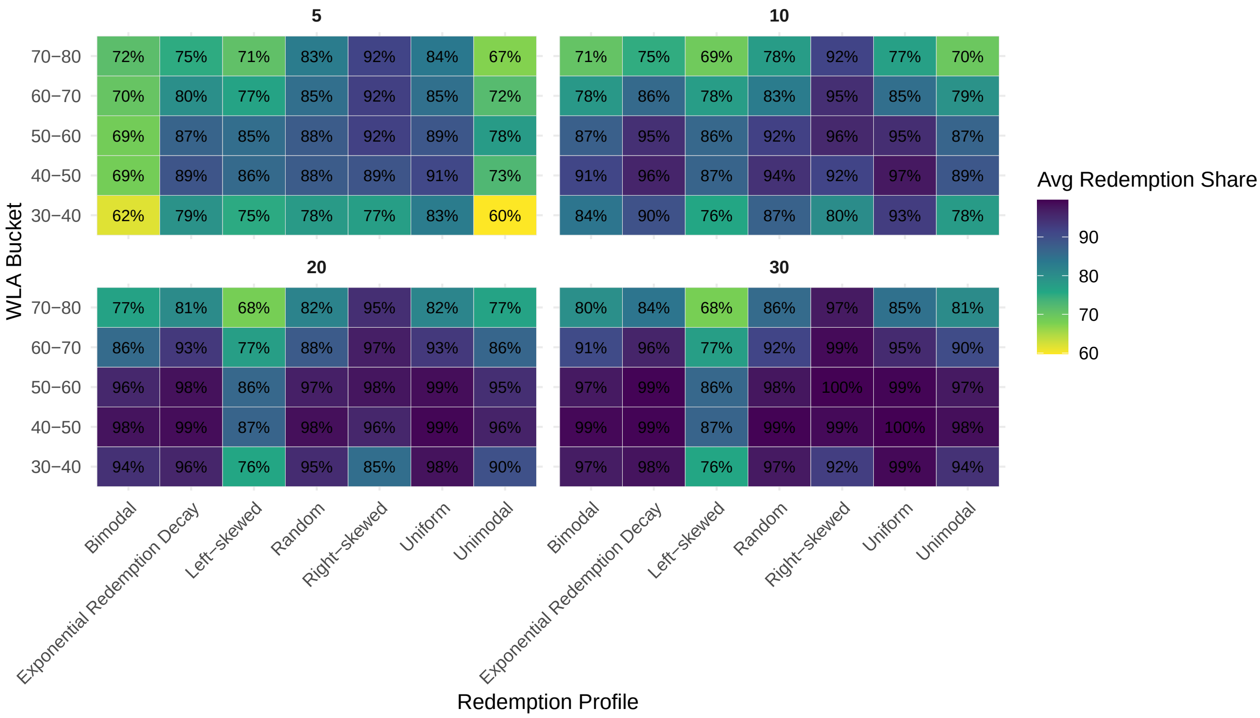

Figure 12: Heatmap of redemption capacity (extension: CD sales, tight-market case). Average redemption capacity (%) by initial WLA bucket and redemption profile when funds can sell up to 50% of CD holdings, faceted by stress horizon. Darker tiles indicate higher capacity.

A2.3 Summary and link to the main results

Overall, the CD-sale extension reinforces two key messages from the main text. First, liquidity buffers matter most when liquidation channels are impaired: restricting CD saleability restores a strong dependence of resilience on initial WLA, especially under front-loaded redemption profiles. Second, resilience assessments are sensitive to market-depth assumptions for the assets that provide liquidity in stress. For this reason, the main text emphasises counterfactuals centred on government asset sales, which are directly tied to WLA definitions and regulatory design, while CD-sale scenarios are presented here as a robustness exercise under alternative assumptions about private-instrument liquidity.

A3 Asset holdings by WLA buckets

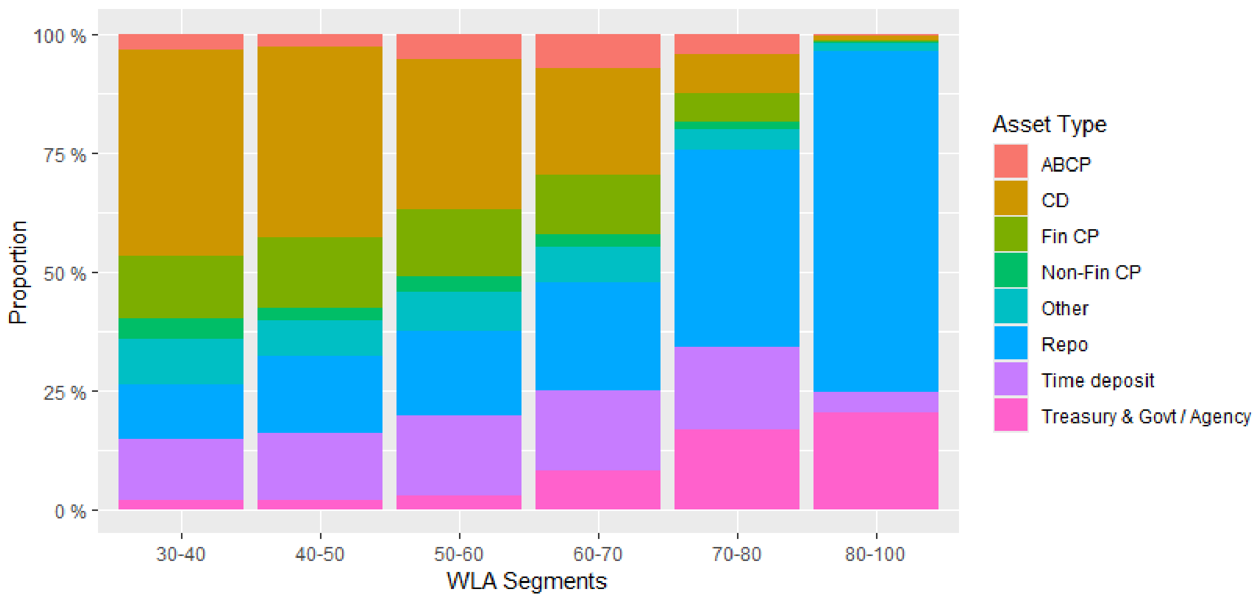

Figure 13: Asset holdings by WLA buckets. Stacked bar chart showing the average portfolio composition (proportion of AUM) for each asset type across WLA segments, averaged over the sample period.