STAFF REPORTS

NO. 1160
AUGUST 2025

Can Redemption Fees Prevent Runs on Funds?

Xuesong Huang | Todd Keister

FEDERAL RESERVE BANK of NEW YORK

Can Redemption Fees Prevent Runs on Funds?
Xuesong Huang and Todd Keister
Federal Reserve Bank of New York Staff Reports, no. 1160
August 2025
https://doi.org/10.59576/sr.1160

Abstract

We ask whether imposing fees on redeeming investors can prevent runs on money market mutual funds (MMFs) and related intermediation arrangements. We first show that imposing a fee only in extraordinary times often leaves the fund susceptible to a *preemptive* run where investors rush to redeem before the fee applies. We then show how a policy that imposes a fee when current redemption demand is above a threshold, even in normal times, can make the fund run proof. We characterize the best policy of this type, which is immune to a run of any size. We show that the reform adopted in the U.S. in 2023 leaves funds vulnerable to runs in some market conditions and imposes an inefficiently large fee in others.

JEL classification: G28, G23, D82
Key words: financial stability policy, preemptive runs, shadow banking

---

Keister: Federal Reserve Bank of New York, Rutgers University (email: todd.keister@ny.frb.org).
Huang: Lingnan College, Sun Yat-sen University (email: xuesong.huangll@gmail.com). The authors thank Huberto Ennis, Kinda Hachem, Chao He, Cecilia Paratore, Bo Wang, David Skeie, Yizhou Xiao, Shengxing Zhang, Zhen Zhou, and participants at many conferences and seminars for helpful comments and suggestions. Ziyu Xing provided excellent research assistance. Part of this research was conducted when Keister was a short-term visitor at the CIRJE, University of Tokyo, whose hospitality and support are gratefully acknowledged. Huang gratefully acknowledges financial support from the National Science Foundation of China (Project No. 72303257).

1 Introduction

The failure of Lehman Brothers in September 2008 sparked a run on prime money market mutual funds (MMFs) in the U.S, with over $400 billion withdrawn in a two-week period. Because these funds play an important role in short-term funding markets, the U.S. Treasury and Federal Reserve responded with extraordinary liquidity facilities and guarantees for MMF investors. In 2014, the Securities and Exchange Commission (SEC) introduced a reform designed to prevent a repeat of this experience. The new rules allowed a fund to limit redemptions and impose a redemption fee when its liquid assets fell below a threshold. This reform proved ineffective: prime MMFs experienced heavy outflows at the onset of the Covid crisis in March 2020, and the Federal Reserve again responded by providing extraordinary liquidity facilities. In July 2023, the SEC adopted a second reform aimed at preventing runs on MMFs during periods of financial stress. The new rules replace the regime of discretionary redemption limits/fees based on a threshold for liquid assets with mandatory fees based on current redemption demand.[1] How effective this new reform will be remains to be seen. At a conceptual level, however, the effectiveness of redemption fees as a financial stability tool, and the principles that should guide their use, are not well understood.

We study how redemption fees can best be used to prevent runs. Our goal is to provide a framework for evaluating reform proposals and for understanding the principles that should govern the design of a redemption-fee policy. Such fees could be used in any intermediation arrangement where runs may arise. MMFs provide a useful laboratory for studying fee policies because of their relatively simple structure and their recent history of both runs and policy changes. While we focus our analysis and policy conclusions on MMFs, we also uncover principles that are likely to apply more broadly.

Our model builds on the approach in Engineer (1989), which adds an additional time period to the well-known framework of Diamond and Dybvig (1983). This additional period allows for the possibility that investors will run on the fund *preemptively* if they anticipate fees may be applied or redemptions may otherwise be restricted in the future. Such preemptive reasoning is believed to have played an important role in the run on prime MMFs in March 2020.[2] We modify the framework in several ways, in part to reflect the operation of a mutual fund. In particular, there is no first-come-first-served (or *sequential service*) constraint within a period. Instead, the fund is able to observe total redemption requests in each period before setting a redemption value for that period.

[1] For the full text of each reform, see SEC (2014) and SEC (2023).

[2] See, for example, the discussions in the reports of the President's Working Group on Financial Markets (2020) and Government Accountability Office (2023).

We first use our model to study a policy that captures the spirit of the 2014 reform. In this regime, the fund redeems shares at par unless doing so would require liquidating invesment, in which case a fee is imposed. We require the fee policy to be *time-consistent* to prevent the fund from using non-credible threats to influence investor behavior. We show that this approach often admits a run equilibrium. In this equilibrium, all investors with an opportunity to redeem in the first period do so because (*i*) they may need to redeem in the next period and (*ii*) they correctly anticipate that a redemption fee may be imposed at that point. In other words, when fees are imposed only if redemption demand is large enough to indicate a run is underway, investors may have an incentive to run preemptively on the fund. Policy discussions of the failure of the 2014 reform have focused largely on how restricting redemptions (i.e., imposing "gates") can lead to a preemptive run. Our results show that redemption fees suffer from the same problem.

Our analysis of this first policy regime highlights an interesting non-monotonicity. Runs in our framework are partial, with only some investors participating. A large run, in which nearly all investors participate, would cause the fee to be imposed immediately, before any redemptions have been processed. If the fee is applied to all redemptions, however, investors are unable to redeem preemptively and have no incentive to join the run. With a smaller run, in contrast, redemption demand may not exhaust the fund's liquid assets in the first period, which would allow some investors to redeem before a fee is applied. If the run is very small, however, few assets would need to be liquidated and investors would again have no incentive to run. The danger in our model, therefore, comes from the possibility of a *moderate-sized* run: one that is small enough that redemption demand may initially be below the fee threshold, but large enough to eventually require costly liquidation of assets. This non-monotonicity plays an important role throughout our analysis.

We next study policies that impose a redemption fee based on current redemption demand, as in the 2023 reform. To be effective in preventing runs, the fee must sometimes apply even when redemptions are in the normal range. Such policies are costly because they reduce the liquidity-provision role of the fund and impose risk on investors even in the absence of a run. We derive the best run-proof policy when the fee can be an arbitrary function of redemption demand. We show that the optimal fee for a given level of redemption demand depends on the relative likelihood of that demand arising during a run compared to normal times *divided by* the level of redemption demand. The first part of this result is intuitive: the fee should be larger in situations that are more likely to occur in the event of a run and smaller in situations that are more likely in normal times. The second is counterintuitive: holding this likelihood ratio constant, the fee should be a *decreasing* function of redemption demand. In other words, the optimal fee is often larger when fewer investors redeem and smaller when more investors redeem. This pattern arises because, when the policy is run-proof, large redemption demand implies many investors have a true liquidity need. Charging a high fee in such states is costly in welfare terms. The best policy, therefore, keeps the fee relatively low in states where fundamental redemption demand is large, while still imposing a fee in enough states to remove investors' incentive to run.

The optimal fee schedule can be complex. To provide more concrete policy advice, we specialize to simple policies characterized by two numbers: (i) a threshold for redemptions below which there is no fee and (ii) a constant fee that applies whenever redemptions are above the threshold. The 2023 reforms have this form, with the threshold set to 5% of a fund's total assets and a fee equal to the costs that would be associated with liquidating a pro-rata share of each asset in the fund's portfolio. We show that the best simple run-proof policy often takes a very intuitive form: the fee is imposed whenever redemption demand is potentially consistent with a run. By applying the fee in all relevant states, this approach keeps the fee as low as possible in states where many investors have a true liquidity need.

The optimal fee schedule in either general or simple form depends on model parameters, including the size of a potential run and the probability distribution of future liquidation costs. These parameters may be particularly difficult to measure and monitor in real time. To address this concern, we study policies that are *robust* in the sense of preventing runs for a range of these parameters. We characterize the best simple, robust policy. The fee in this policy is determined by the need to prevent large runs, which would impose substantial liquidation costs. The threshold, in contrast, is set to prevent smaller runs, which are more likely to initially go undetected.

Finally, we compare the best simple, robust policy with the 2023 reform, which sets the fee equal to the cost of liquidating a "vertical slice" of the fund's portfolio. Our analysis highlights two weaknesses of the new rule. In situations where liquidity conditions may deteriorate, a robust fee policy takes potential future liquidation costs into account. The current policy is not forward-looking, however, which can leave funds vulnerable to a run. When liquidity conditions are stable, in contrast, the best policy has a smaller fee than the vertical slice rule and a lower threshold, meaning the fee will be applied more often in equilibrium. Intuitively, the new rule imposes a large fee in states where many investors have a true liquidity need. Our analysis shows that MMFs will be more attractive to investors if the fee is smaller in these states, even though the fee must then be imposed more often.

Related literature. Our paper is related to several strands of the broad literature that studies runs on financial intermediaries. We follow the approach advocated in Wallace (1990) of allowing the bank/fund to choose from a large set of contracts beyond the simple demand-deposit contract studied by Diamond and Dybvig (1983) and many others. In a framework with two consumption periods and sequential service, efficient ways to prevent runs with flexible contracts have been identified by Green and Lin (2003) and Huang (2024) using direct mechanisms and by Cavalcanti and Monteiro (2016) and Andolfatto et al. (2017) using indirect mechanisms.[3] We expand the model to include a third consumption period, which introduces the possibility of a preemptive run, where investors rush to redeem because they worry future redemptions may be restricted rather than over concern that the intermediary may fail. This approach follows Engineer (1989) and more recent work by Cipriani et al. (2014) and Voellmy (2021). We deviate from all of these papers by removing the sequential service constraint, which is inappropriate for mutual funds and other shadow banking arrangements, as argued by Andolfatto and Nosal (2024). Ours is the first paper to show that a self-fulfilling run equilibrium can exist when payment contracts are fully flexible, there is no sequential service constraint, and investment technologies have constant returns to scale.

Our paper is also closely related to the literature that focuses specifically on runs on MMFs. For example, Ennis (2013) studies the desirability of different ways of computing a fund's net asset value (NAV). McCabe et al. (2013) propose a policy that allows investors to redeem only a fraction of their balance at once, with the rest remaining in the fund. Parlatore (2016) shows how sponsors' decisions to support their funds in periods of stress can be a source of fragility. Ennis et al. (2023) propose a remedy that requires MMFs to have contractual commitments of liquidity support. We focus instead on redemption fees, as introduced in the 2023 reform. Our contribution is to derive the most efficient fee policy, which can then be compared to these and other alternative approaches to reforming MMFs.

We also contribute to the literature on runs on mutual funds more broadly. Zeng (2017) studies how a fund's desire to rebuild its cash buffer following heavy redemptions can create an incentive for investors to run under a flexible NAV rule. A group of recent papers study *swing pricing* policies, which adjust the price of a mutual fund share in response to redemption demand. See, for example, Lewrick and Schanz (2017b), Capponi et al. (2020), and Ma et al. (2025).[4] The redemption-fee policies we study are economically equivalent to swing pricing: both set the amount received by redeeming investors as a function of

[3] For settings where the best direct mechanism does not prevent runs, see Peck and Shell (2003), Ennis and Keister (2009b) and Sultanum (2014), among others.

[4] See Capponi et al. (2023) for a survey of the literature on swing pricing.

redemption demand. In contrast to these papers, however, we derive the best policy within a broad class. This additional flexibility eliminates the standard incentive to run based on concern that the fund will fail and shifts the focus to preventing preemptive runs of the type observed in March 2020.

Finally, our paper draws motivation from the growing empirical literature on mutual fund runs and patterns of redemptions. Chen et al. (2010) provide evidence of strategic complementarities in investor redemptions from mutual funds using data from 1995 - 2005. Schmidt et al. (2016) study the 2008 runs on prime MMFs and also provide evidence of strategic complementarities leading to increased redemptions. Cipriani and La Spada (2020) and Li et al. (2021) study the March 2020 episode and show that concern about future restrictions on redemptions exacerbated the runs on prime MMFs. These papers provide evidence that key features of our model are important in practice. Lewrick and Schanz (2017a) and Jin et al. (2022) document the effectiveness of swing pricing in removing the first-mover advantage in open-end mutual funds using unique data sets from Luxembourg and the U.K., respectively. This evidence is consistent with our results showing that fees based on current redemption demand can eliminate the incentive for investors to join a run.

2 The model

In this section, we describe our environment and the operation of a money market fund in this setting. We introduce an asset portfolio and fundamental uncertainty into the framework of Engineer (1989), and we modify the information structure to reflect the operation of a mutual fund rather than a bank.

2.1 The environment

There are three periods ($t = 1, 2, 3$) and a single consumption good in each period. Two technologies exist for transforming goods across periods, *storage* and *investment*. Stored goods earn a gross return of 1 between any two periods. Invested goods yield a return of $R > 1$ if held to maturity in period 3, but $r_t \leq 1$ if liquidated in period $t \in \{1, 2\}$. The value of r_1 is known, but r_2 is initially uncertain and has support $[\underline{r}, \overline{r}]$ with $\overline{r} \leq r_1$.

Each of a continuum of investors, indexed by $i \in [0,1]$, has preferences given by

$$u_i(c_1, c_2, c_3; \omega_i) = \begin{cases} u(c_1) & \text{if } \omega_i = 1 \\ u(c_1 + c_2) & \text{if } \omega_i = 2 \\ u(c_1 + c_2 + c_3) & \text{if } \omega_i = 3, \end{cases}$$

where u is the natural log function, c_t is consumption in period t, and ω_i is the investor's liquidity-preference type. Both type-1 and type-2 investors are *impatient* in the sense that they need to consume before investment matures, while type-3 investors are *patient*. Investors' types are private information, and each investor learns her own type gradually. In period 1, an investor discovers only whether she is type 1 or not. In period 2, the remaining investors discover whether they are type 2 or type 3. A known fraction $\pi \in (0,1)$ of investors will be impatient, but the distribution of these investors between type 1 and type 2 is random. In other words, there is no uncertainty about total early consumption demand, but there is uncertainty about its timing. Let π_1 denote the fraction of type 1 investors, which has a density function $f(\pi_1)$ with full support on $[0,\pi]$. The fraction of type 2 investors is then $\pi - \pi_1$.

Investors are endowed with equal shares in a *fund*. The fund owns one unit of the good per investor; a fraction s of the good is in storage and the remaining $1-s$ is invested. Investors are isolated from each other and can only interact with the fund.[5] In periods 1 and 2, investors receive information about their own type and may have an opportunity to submit a redemption request to the fund. We assume all type 1 investors can redeem in period 1, but only a fraction $\delta \in (0,1]$ of non-type-1 investors can do so. The remaining fraction $1-\delta$ are inattentive or otherwise unable to contact the fund in period 1. All investors who have not yet redeemed can contact the fund to redeem in periods 2 and 3.

2.2 Discussion of key assumptions

Sequential service. Unlike the usual banking arrangement studied in Diamond and Dybvig (1983), Engineer (1989), and many others, the fund in our model does not need to serve redeeming investors one-at-a-time. Instead, it collects all redemption requests in a period before making any payments. The assumption matches the operation of an open-end mutual fund that pays redeeming investors at the end of each *pricing period*, which often corresponds

[5] As in Wallace (1988) and others, this isolation assumption implies investors are unable to trade shares in the fund or other claims with each other in periods 1 and 2. This approach is consistent with our focus on open-end mutual funds rather than exchange-traded funds or other market-based arrangements.

to a business day. In a standard two-period model, the optimal contract when there is no sequential service rules out bank runs under very general conditions (see Green and Lin, 2003, Section 3).[6] We show below that this result does not extend to the model with three consumption periods that we study here. A key feature of this model is that investors redeeming in period 1 must be served by the end of that period, before the fund observes what will happen in period 2. In other words, even though investors are not paid one-at-a-time, a form of sequential service *across periods* arises naturally in any longer-horizon model, and this fact potentially opens the door to a run on the fund.

Liquidation costs. Money market funds hold assets that are fairly liquid most of the time but may become less liquid in periods of financial stress. It is likely no coincidence that the runs observed in 2008 and 2020 both occurred during periods of significant stress and lower market liquidity. In the analysis below, we consider scenarios where the fund's investment is liquid ($r_1 \approx 1$) and where it is illiquid ($r_1 < 1$). In both cases, we show that investors' expectations about the future liquidation value r_2 play an important role in shaping redemption behavior. We allow r_2 to be random to capture situations where investors are concerned that market conditions may deteriorate. We assume there is only downside risk, with $r_2 \leq r_1$, to simplify the presentation, but this assumption is not necessary for our results. Throughout the analysis, we assume the liquidation values r_t are independent of the fund's liquidation choices. This assumption would hold, for example, if the fund is a relatively small player in the market for those securities.[7]

Log utility. Using the natural log utility function simplifies the presentation and some calculations. In addition, it facilitates interpretation of the model because (as we show below) the efficient allocation in the absence of runs pays redeeming investors in each period an amount that corresponds to the net asset value (NAV) of a share in that period. A version of the model with CRRA utility leads to broadly similar results regarding the effectiveness of redemption fees. When the coefficient of relative risk aversion is greater than one, however, the efficient allocation gives early-redeeming investors more than the NAV of a share to insure against liquidity risk (as in Diamond and Dybvig, 1983, and many others). Using log utility instead implies that the efficient operation of the fund in our model corresponds well with how money market funds operate in practice.

[6] See also the discussion in Andolfatto and Nosal (2024).

[7] It may be interesting to extend our analysis to settings where a fund's sale of assets drives down the market price as in Allen and Gale (2004), Acharya and Yorulmazer (2008), Parlatore (2016), Lewrick and Schanz (2017b), Izumi and Li (2025), and many others. We leave this issue for future work.

Partial runs. Our assumption that only a fraction δ of non-type-1 investors can redeem at $t = 1$ captures the idea that a run does not typically take place within a single day (or a single pricing period).[8] Instead, a run is typically spread over time, which makes identifying a run difficult in the early stages. If $\delta = 1$, a run will involve all of the fund's investors requesting redemption in the first period and, hence, is easily identified before the fund makes any payments to investors. When δ is below 1, in contrast, the fund may initially be unsure whether the observed redemption demand is fundamental (with a large realization of π_1) or instead reflects a run. We show below that such uncertainty is necessary for a run to arise in this setting. We begin by assuming δ is a known constant then characterize policies that are *robust* in the sense of being run-proof for all possible values of δ.

2.3 Contracts and feasibility

In each period, the fund collects all redemption requests and allocates consumption to the redeeming investors. Let m_t denote the fraction of investors who choose to redeem (i.e., send a redemption "message" to the fund) in period t. The operation of the fund is characterized by three functions. In period 1, the fund observes m_1 and pays an amount c_1 to each redeeming investor. In period 2, the fund observes m_2 and the realized liquidation value r_2, then pays c_2 to redeeming investors. Once an investor redeems her share in the fund, she consumes and exits the economy. Investors who remain in the fund in period 3 each receive c_3, which corresponds to a pro-rata share of the fund's matured assets. Formally, a *contract* is a collection of payment functions

$$\begin{aligned} &c_1 : M_1 \to \mathbb{R}_+ && \text{where} && M_1 \equiv [0, \pi + \delta(1-\pi)] \\ &c_2 : M_1 \times M_2 \times P \to \mathbb{R}_+ && && M_2 \equiv [0, 1 - m_1] \\ &c_3 : M_1 \times M_2 \times P \to \mathbb{R}_+ && && P \equiv [\underline{r}, \overline{r}] \,. \end{aligned}$$

A contract is *feasible* if the fund can generate the specified payments from its given asset portfolio for any profile of redemption requests. In period 1, the payment c_1 must satisfy

$$m_1 c_1 + e_1 = s + r_1 \ell_1 \quad \text{for all } m_1, \tag{1}$$

[8] Chen et al. (2010) and Zeng (2017) use similar assumptions to capture the possibility of some investors being temporarily inactive.

where $e_1 \in [0, s]$ is the amount of storage held until period 2 ("excess liquidity") and $\ell_1 \in [0, 1-s]$ is the amount of investment liquidated in period 1. In period 2, feasibility requires

$$m_2 c_2 + e_2 = e_1 + r_2 \ell_2 \quad \text{for all } (m_1, m_2, r_2), \tag{2}$$

where $e_2 \in [0, s - e_1]$ is the amount of storage held until period 3 and $\ell_2 \in [0, 1 - s - \ell_1]$ is the amount of investment liquidated in period 2. Finally, feasibility in period 3 requires

$$(1 - m_1 - m_2) c_3 = R\left(1 - s - \ell_1 - \ell_2\right) + e_2 \quad \text{for all } (m_1, m_2, r_2). \tag{3}$$

The payment functions $\{c_1, c_2, c_3\}$ are *feasible* if, for every (m_1, m_2, r_2), there exist portfolio management choices $\{e_1, \ell_1, e_2, \ell_2\}$ such that equations (1) - (3) are satisfied.

2.4 The first-best allocation

A planner's problem. Suppose, as a benchmark, that the fund were operated by a planner who could observe investors' preference types and choose when they redeem their shares. We also allow the planner to choose the fund's initial portfolio s and the payment functions $\{c_1, c_2, c_3\}$. The planner would clearly direct type t investors to redeem only in period t, which implies redemption requests will satisfy $m_1 = \pi_1$, $m_2 = \pi - \pi_1$ and $m_3 = 1 - \pi$. Because π is known, the planner will pay impatient investors in both periods 1 and 2 using goods in storage and will pay patient investors in period 3 using matured investment. In other words, the planner will set e_2, ℓ_1, and ℓ_2 to zero and, therefore, the efficient allocation will not depend on the liquidation values r_1 and r_2. We can then write the payments $\{c_1, c_2, c_3\}$ directly as functions of the state π_1. The planner would choose the fund's portfolio $(s, 1-s)$ and these payment functions to maximize investors' expected utility

$$\int_0^{\pi} \left[\pi_1 u\big(c_1(\pi_1)\big) + (\pi - \pi_1) u\big(c_2(\pi_1)\big) + (1-\pi) u\big(c_3(\pi_1)\big)\right] f(\pi_1) d\pi_1$$

subject to the feasibility constraints

$$\begin{aligned} \pi_1 c_1(\pi_1) + (\pi - \pi_1) c_2(\pi_1) &= s && \text{for all } \pi_1, \text{ and} \\ (1-\pi) c_3(\pi_1) &= R(1-s) && \text{for all } \pi_1. \end{aligned}$$

The first-order conditions for c_1 and c_2 imply

$$u'\big(c_1^*(\pi_1)\big) = u'\big(c_2^*(\pi_1)\big) \quad \Rightarrow \quad c_1^*(\pi_1) = c_2^*(\pi_1) \qquad \text{for all } \pi_1,$$

that is, the planner always gives the same consumption to type 1 and type 2 investors. The feasibility constraints associated with a given π_1 can then be combined and simplified to

$$\pi c_1^*(\pi_1) + (1-\pi)\frac{c_3^*(\pi_1)}{R} = 1. \tag{4}$$

The first-order condition for c_3 implies

$$u'\big(c_1^*(\pi_1)\big) = Ru'\big(c_3^*(\pi_1)\big). \tag{5}$$

With log utility, it is straightforward to show that the solution to this problem sets $c_1(\pi_1) = c_2(\pi_1) = 1$ and $c_3(\pi_1) = R$ for all π_1. If investment is valued at cost (i.e., 1) in the early periods, these payments correspond to the net asset value of a share in each period. To achieve this allocation, the planner sets the fraction of storage in its portfolio to $s = \pi$ and invests the remaining fraction $(1-\pi)$. The following result summarizes this discussion.

Proposition 1. *The first-best allocation gives* 1 *to each type 1 investor at* $t = 1$ *and to each type 2 investor at* $t = 2$*, and it gives* R *to each type 3 investor at* $t = 3$.

Equations (4) and (5) also characterize the efficient allocation in a standard Diamond-Dybvig model with only two consumption periods. In other words, having an extra time period and uncertainty about the timing of early consumption demand do not change the efficient allocation of resources. The planner wants all impatient investors to consume 1, regardless of whether they redeem in period 1 or 2, and wants all patient investors to consume R.

Decentralization. We assume the fund begins with the planner's chosen portfolio, $s = \pi$. It is straightforward to show that there exist contracts that implement the first-best allocation from Proposition 1 as a perfect Bayesian equilibrium of the redemption game played by investors. Any such contract must satisfy

$$c_1(m_1) = 1 \quad \text{for all } m_1 \leq \pi, \tag{6}$$

$$c_2(m_1, m_2, r_2) = 1 \quad \text{for all } m_1 + m_2 \leq \pi, \text{ and} \tag{7}$$

$$c_3(m_1, m_2, r_2) = R \quad \text{for all } m_1 + m_2 \leq \pi. \tag{8}$$

These payments are feasible when the fund sets

$$\begin{aligned} &e_1 = (\pi - m_1) && \text{and} && \ell_1 = 0 && \text{for } m_1 \leq \pi \\ &e_2 = 0 && \text{and} && \ell_2 = 0 && \text{for } m_1 + m_2 \leq \pi. \end{aligned}$$

In other words, the fund uses goods in storage to make payments exactly as the planner would whenever total redemption demand does not exceed π. In this equilibrium, only type 1 investors redeem in period 1, so $m_1 = \pi_1$, and only type 2 investors redeem in period 2, so $m_2 = \pi - \pi_1$. Because $R > 1$, non-type-1 investors have a strict incentive to wait in period 1 and type-3 investors have a strict incentive to wait in period 2 in this equilibrium. Note that unilateral deviations from equilibrium play do not change the fractions (m_1, m_2) because there is a continuum of investors. As a result, the payments in the contract associated with redemption demand greater than π have no effect on individual investors' incentives, and this equilibrium exists regardless of how those payments are specified.

2.5 Time consistency

Our interest is in studying whether the fund is *fragile* in the sense that another equilibrium exists in which investors run on the fund. The answer to this question depends crucially on the payments associated with levels of early redemption demand greater than π, which – as desribed above – potentially lie off the equilibrium path of play. To ensure our model delivers credible policy advice, we require that these payments be time consistent.

Diamond and Dybvig (1983) and others have shown how promising a "tough" response to high withdrawal demand can rule out bank runs in a standard two-period model. A similar result can be shown to hold here, although the form of the "tough" response is different. Consider a contract that, for all $m_1 > \pi$, sets $c_1 = 0$ and $e_1 = s$. In other words, when redemption demand in period 1 is large enough that the fund detects a run is underway, all investors who have redeemed their shares will receive nothing in return. If $m_1 \leq \pi$ but $m_1 + m_2 > \pi$, meaning the fund detects a run is underway only in period 2, the contract allocates the fund's remaining resources so that c_2 and c_3 are strictly positive and satisfy $c_2 < c_3$. If a non-type-1 investor expects some other investors to run, she knows that m_1 will be greater than π with positive probability. As long as consuming zero is sufficiently unattractive, she will strictly prefer to wait, so there cannot be an equilibrium where investors run in period 1. The fact that the contract sets $c_2 < c_3$ for all (m_1, m_2) implies there cannot be an equilibrium where investors run in period 2. Therefore, this type of contract implements the first-best allocation without introducing the possibility of a run.[9]

However, the threat to give investors nothing in exchange for their shares in period 1 would clearly not be time-consistent. Once the redemption requests have been submitted, the fund would have a strong incentive to change course and offer positive consumption to all redeeming investors.[10] Our goal is to provide policy advice for reforming MMFs, and we do not want this advice to rely on non-credible threats to punish investors in the event a run is detected. For this reason, we require that the contract offered by the fund be *time consistent* in the following sense: whenever total redemption demand exceeds π, clearly indicating that a run is underway, the fund must choose payments that maximize the sum of investors' utilities conditional on the given redemption demand, that is,

$$m_1 u(c_1) + \mathbb{E}\left[m_2 u(c_2) + (1 - m_1 - m_2)u(c_3) \mid m_1\right]. \tag{9}$$

In other words, the fund must always act in the best interests of its investors, even when a run is underway. Formally, we require the payment functions $\{c_1, c_2, c_3\}$ to satisfy the following two conditions.

(TC1) For $m_1 > \pi$, $\{c_1, c_2, c_3\}$ must maximize equation (9) subject to constraints in equations (1) - (3) with $m_2 = \pi + \delta(1 - \pi) - m_1$.

(TC2) For $m_1 \leq \pi$ and $m_1 + m_2 > \pi$, $\{c_2, c_3\}$ must maximize equation (9) subject to the constraints in equations (2) - (3) with $e_1 = s - m_1 c_1$, and $\ell_1 = 0$.

Run detected in period 1. Condition (TC1) applies when redemption demand in the first period is large enough to indicate a run is underway. Choosing the payment in period 1 to maximize investors' expected utilities requires forecasting redemption demand in period 2, which the fund does using the structure of the model. First, observing m_1 and knowing a run is underway allows the fund to infer π_1 using $m_1 = \pi_1 + \delta(1 - \pi_1)$. Second, it is straightforward to show that an investor who turns out to be type 3 will never have an incentive to redeem

[9] Note that the "tough" policy here is very different from suspending convertibility of shares, which Engineer (1989) showed is ineffective at preventing runs in a three-period model. When convertibility is suspended, investors who are unable to redeem in period 1 can try again in period 2. A rush to redeem in period 1 then creates a backlog of redemption demand, which implies period-2 requests will be large as well. The contract we describe here, in contrast, honors all redemption requests in period 1 but may set the redemption price to zero. This policy is effective because it punishes redemption requests more heavily during a run in period 1 while making future redemption more attractive.

[10] Alternatively, one could imagine investors who receive zero for their redeemed share might take legal action against the fund. The court system might then intervene to overrule the tough response to a run, as discussed by Ennis and Keister (2009a) in the banking context.

in period 2.[11] Therefore, only those investors who were inattentive in period 1 and turn out to be type 2 will redeem in period 2, that is, $m_2 = (1-\delta)(\pi - \pi_1)$. Time consistency requires that the fund act to maximize investors' utilities given the observed redemption demand m_1 and this forecast for m_2; the resulting payments are given in our next result.

Proposition 2. *When $m_1 > \pi$, condition (TC1) requires the fund to set $e_2 = \ell_2 = 0$ and*

$$c_1^{TC1} = c_2^{TC1} = \max\Bigg\{\underbrace{\frac{\pi}{\pi + \delta(1-\pi)}}_{\textit{no liquidation}}, \underbrace{r_1(1-\pi) + \pi}_{\textit{liquidation at } t=1}\Bigg\}$$

$$c_3^{TC1} = \min\Bigg\{\underbrace{\frac{R}{1-\delta}}_{\textit{no liquidation}}, \underbrace{R(1-\pi) + \frac{R}{r_1}\pi}_{\textit{liquidation at } t=1}\Bigg\}.$$

When δ is small, the fund divides the goods in storage among the investors who will be redeeming in the first two periods and keeps all investment until $t = 3$. These payments correspond to the first term in the min/max operators above. When δ is larger, the fund liquidates some investment to rebalance its portfolio and the second term on each line applies.[12] Because the fund accurately forecasts redemption demand at $t = 2$ and there is only downside risk ($r_2 \leq r_1$), the fund does all of this liquidation at $t = 1$.[13] In both cases, investors redeeming in the first two periods receive the same amount, while investors redeeming in the final period receive strictly more. A complete proof of this result is given in Online Appendix A.1.

Run detected in period 2. Condition (TC2) applies when it becomes apparent that a run is underway only in period 2. In this case, the m_1 investors who redeemed in period 1 have already each been paid c_1 out of goods held in storage. Time consistency requires that the remaining asset portfolio be used efficiently to make payments to the m_2 investors redeeming in the current period and the $1 - m_1 - m_2$ investors who will redeem in the final period. The following proposition characterizes these payments when the fee applied to the m_1 investors who have already redeemed was not too large.

[11] In period 2, the fund faces a two-period problem without sequential service. The best payments given any pattern of redemptions at this point will satisfy $c_2 \leq c_3$, similar to Green and Lin (2003, Section 3).

[12] This portfolio rebalancing is reminiscent of the results in Zeng (2017), but reflects a very different motive. In that paper, the fund may preemptively liquidate investment even though doing so is costly because it wants to maintain a particular ratio of liquid to illiquid assets. Here, in contrast, the fund is preemptively liquidating investment because it recognizes it will need more liquid assets next period and worries that liquidation costs may increase in the meantime.

[13] If we instead assumed $\overline{r} > r_1$, the fund might choose to liquidate investment at both $t = 1$ and $t = 2$, with the latter choice being made after r_2 is realized. Our main results below hold for this case as well, but the analysis and exposition are more complex.

Proposition 3. *When $m_1 \leq \pi$, $c_1(m_1)$ is sufficiently close to 1, and $m_1+m_2 > \pi$, condition (TC2) requires the fund to set*

$$c_2(m_1, m_2, r_2) = \max\Bigg\{\underbrace{\frac{\pi - m_1 c_1(m_1)}{m_2}}_{\textit{no liquidation}}, \underbrace{\frac{r_2(1-\pi) + \pi - m_1 c_1(m_1)}{1 - m_1}}_{\textit{liquidation at } t=2}\Bigg\}$$

$$c_3(m_1, m_2, r_2) = \min\Bigg\{\underbrace{\frac{R(1-\pi)}{1 - m_1 - m_2}}_{\textit{no liquidation}}, \underbrace{\frac{R(1-\pi) + \frac{R}{r_2}[\pi - m_1 c_1(m_1)]}{1 - m_1}}_{\textit{liquidation at } t=2}\Bigg\}.$$

Depending on redemption demand and the realization of r_2, the fund's choice of liquidation in period 2 may be zero (corresponding to the first term in the *max/min* operators) or positive (the second term). Either way, the time-consistent policy imposes a redemption fee in period 2 ($c_2 < 1$) unless there is no cost of liquidating investment ($r_2 = 1$). In Online Appendix A.2, we prove a generalized version of this result (labeled Proposition 3′) that applies for an arbitrary period-1 payment rule $c_1(m_1) \leq 1$. If the rule imposes large fees for some values of m_1, a third possibility arises: the fund may not pay out all of its liquid assets at $t = 2$ and instead hold some excess liquidity ($e_2 > 0$). As long as the liquidation values r_t are not too low, however, the best run-proof policy will not fall in this region and the simpler result in Proposition 3 above applies. We focus on this case in the main text and in the examples below.

Our focus is on preventing runs, that is, identifying contracts that generate a unique equilibrium in which investors only redeem when they need to consume. For such contracts, (TC1) and (TC2) only restrict off-equilibrium payments. The fund has full freedom to choose the payments that are made in each state along the equilibrium path.

In the next section, we ask whether the fund can implement the first-best allocation without introducing a run. We show the answer is often "no." The remaining sections then characterize second-best run-proof contracts.

3 Preemptive runs

In this section, we study investors' equilibrium redemption behavior when the fund aims to implement the first-best allocation described in Proposition 1. We show that an equilibrium often exists where non-type 1 investors run preemptively, that is, they redeem in period 1 because they worry a fee may be imposed in period 2.

3.1 The incentive to run

The requirement that the fund (i) follows the planner's allocation in equations (6) - (8) when redemption demand is below π and (ii) satisfies the time-consistency constraints (TC1) - (TC2) when redemption demand is above π together fully determine the payment functions $\{c_1, c_2, c_3\}$. In other words, there is a unique contract that both implements the first-best allocation as an equilibrium and satisfies time consistency. This contract pays redeeming investors 1 in periods 1 and 2 unless doing so would require some investment to be liquidated, in which case the time-consistent redemption fee is applied. In this section, we show that this contract often also admits a run equilibrium.

Consider a non-type-1 investor who has the option to redeem in period 1. Suppose she expects all other attentive non-type-1 investors to redeem in period 1, meaning redemption demand will be $m_1 = \pi_1 + \delta(1 - \pi_1)$. What will she receive if she joins the run? There are two possibilities. If the realization of π_1 is sufficiently small,

$$\pi_1 \leq \frac{\pi - \delta}{1 - \delta}, \tag{10}$$

then $m_1 \leq \pi$ will hold. In this case, no fee will be applied and she will receive 1. If this inequality is reversed, m_1 will be larger than π, in which case the fund will impose the redemption fee in Proposition 2. The investor's expected payoff from joining the run is then

$$\text{Redeem:} \quad \int_0^{\frac{\pi-\delta}{1-\delta}} u(1) f_n(\pi_1) d\pi_1 + \int_{\frac{\pi-\delta}{1-\delta}}^{\pi} u\big(c_1^{TC1}\big) f_n(\pi_1) d\pi_1, \tag{11}$$

where f_n denotes the density of π_1 conditional on the investor not being type 1.[14] If the investor instead chooses to wait in period 1, she will redeem in period 2 if she turns out to be type 2 and will otherwise wait until period 3. Let $p_{\pi_1} \equiv \frac{\pi - \pi_1}{1 - \pi_1}$ denote the probability of being type 2 conditional on not being type 1. Then her expected payoff from waiting is

$$\text{Wait:} \quad \int_0^{\frac{\pi-\delta}{1-\delta}} \Big\{ p_{\pi_1} \mathbb{E}\left[u\big(c_2(m_1, m_2, r_2)\big)\right] + (1 - p_{\pi_1}) \mathbb{E}\left[u\big(c_3(m_1, m_2, r_2)\big)\right] \Big\} f_n(\pi_1) d\pi_1 \tag{12}$$

$$+ \int_{\frac{\pi-\delta}{1-\delta}}^{\pi} \left[p_{\pi_1} u\big(c_2^{TC1}\big) + (1 - p_{\pi_1}) u\left(c_3^{TC1}\right)\right] f_n(\pi_1) d\pi_1,$$

where $\mathbb{E}$ represents the expectation with respect to r_2. As before, there are two situations.

[14] This conditional density function can be written as $f_n(x) = \frac{(1-x)f(x)}{\int_0^{\pi}(1-z)f(z)dz}$.

The first line contains the realizations of π_1 small enough that the inequality in equation (10) holds and no fee is applied in period 1. In these cases, the payments c_2 and c_3 come from Proposition 3 using the redemption demand corresponding to a run:

$$m_1 = \pi_1 + \delta(1 - \pi_1) \quad \text{and} \quad m_2 = (1 - \delta)(\pi - \pi_1). \tag{13}$$

The second line contains the realizations of π_1 large enough that a fee is applied in period 1. In these cases, payments in periods 2 and 3 come from Proposition 2.

It will often be convenient to focus on the net benefit of redeeming early during a run for a non-type-1 investor, which is the difference between equations (11) and (12). Regrouping terms, this net benefit can be written as

$$\begin{aligned}\mathcal{R} \equiv & \int_0^{\frac{\pi-\delta}{1-\delta}} \Big\{ u(1) - p_{\pi_1}\mathbb{E}\left[u\big(c_2(m_1, m_2, r_2)\big)\right] - (1 - p_{\pi_1})\mathbb{E}\left[u\big(c_3(m_1, m_2, r_2)\big)\right] \Big\} f_n(\pi_1) d\pi_1 \\ & - \underbrace{\int_{\frac{\pi-\delta}{1-\delta}}^{\pi} (1 - p_{\pi_1}) \left\{ u\left(c_3^{TC1}\right) - u\left(c_1^{TC1}\right) \right\} f_n(\pi_1) d\pi_1}_{\equiv T > 0}.\end{aligned} \tag{14}$$

The first line of equation (14) is the net benefit of redeeming early conditional on π_1 being small enough that no fee is applied in period 1. The second line is the net cost of redeeming early when π_1 is large enough that a fee is imposed. This cost is always strictly positive and proportional to the probability that the investor will be type 3. Note that this cost, which we denote $T > 0$, is determined entirely by the time-consistency condition (TC1) and, therefore, will remain unchanged when we introduce redemption fees for $m_1 \leq \pi$ below.

Overall, equation (14) highlights why an investor may have an incentive to run. If she knew π_1 would be large enough to trigger the time-consistent fee in period 1, she would have no incentive to join the run, exactly as in a two-period model without sequential service. The potential incentive to join a run comes from the possibility that π_1 is small enough that no fee is applied in period 1. Redemption demand in period 2 will reveal the run, and the fund will then impose a fee. If the investor turns out to be type 2, she will need to redeem in period 2 and pay this fee. Given this possibility, redeeming preemptively – before the fee is imposed – may be attractive.

3.2 A run equilibrium

Figure 1 presents an example that illustrates how a preemptive run can exist.[15] The red curve depicts equation (11): the expected value of redeeming in period 1 for a non-type 1 investor when all other attentive investors also redeem. If δ is small, few investors participate and the run is very unlikely to trigger a fee in period 1, which means the expected payoff is close to $\ln(1) = 0$. As δ increases, the probability a run will trigger a fee at $t = 1$ increases, which makes redeeming less attractive. The blue curve depicts equation (12): the expected value of waiting to redeem in this same situation. If δ is small, a run will cause little or no liquidation in period 2, which implies the investor will receive close to 1 if she redeems in period 2 and close to R if she redeems in period 3. In this situation, waiting is better than joining the run, as shown in the figure, and no run equilibrium exists. On the other side, when δ is greater than π, period-1 redemption demand in a run is always large enough that a fee is applied, which makes waiting to redeem attractive and again no run equilibrium exists.

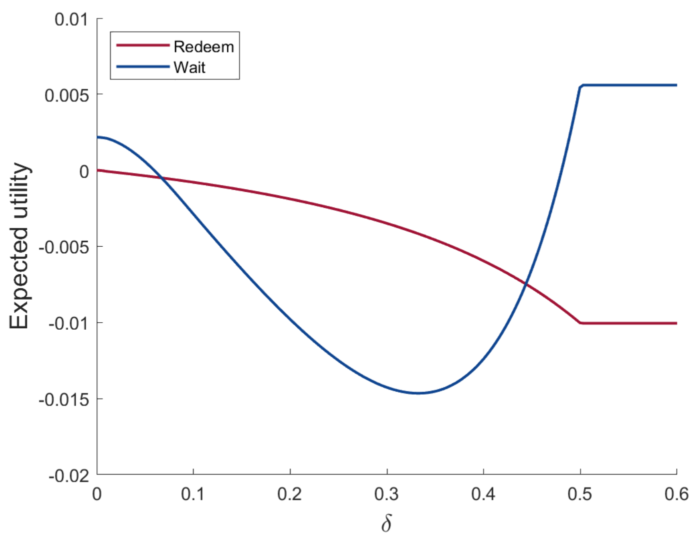

Figure 1: A run equilibrium exists for intermediate values of δ

For a range of values of δ in between these extremes, however, the figure shows that the investor has a strict incentive to join the run. In this region, the run is large enough that it will lead to a significant redemption fee in period 2 if the liquidation value of investment

[15] The parameter values for this example are $\pi = 0.5$, $R = (1.04)^{\frac{1}{12}}$, $r_1 = 0.98$, and $r_2 \in \{0.8, 0.98\}$ with probability 0.5 of each. The fraction π_1 of type 1 investors follows a uniform distribution on $[0, \pi]$. In Online Appendix B, we present a series of examples to show that a preemptive run exists for a wide range of parameter values and to explore how the fragility of the fund depends on parameter values.

turns out to be low. At the same time, however, the run is small enough that there is a significant chance it will not trigger a fee in period 1. This combination gives the investor an incentive to join the run and try to redeem before a fee is imposed. This pattern is general and appears throughout our analysis below: a run equilibrium is most likely to exist in this framework when the fraction δ of attentive investors lies in an intermediate range.

3.3 Discussion

We interpret the contract in this section as capturing some key features of the reforms to prime and tax-exempt MMFs that were adopted in the U.S. in October 2014. Under those rules, funds would redeem shares at NAV unless high redemption demand pushed the fund's liquid assets below a threshold level. Once this threshold was passed, a fund had the ability to impose a redemption fee of up to 2% and was directed to do so if it was deemed to be in the best interests of shareholders.[16] The threshold was set so that it would be hit only in extraordinary circumstances, not in normal times. We interpret this policy as attempting to rule out runs by only imposing redemption fees that lie off the path of play in the no-run equilibrium, as in our contract above. We interpret our time consistency constraints as capturing the spirit of the directive that the fund act in its shareholders' best interests in setting the fees.

Engineer (1989) and Cipriani et al. (2014) have shown that a policy of restricting withdrawals once demand is unusually high may be ineffective at preventing runs. Our environment with redemption fees and no sequential service gives the fund more flexibility to shape incentives in the redemption game. One might have hoped that such flexibility would allow the fund to implement the efficient allocation without also introducing a run equilibrium. Indeed, our reading of the 2014 reforms is that they were based on this type of reasoning. Policymakers emphasized, for example, that a redemption fee can correct the negative externality that arises when current redemptions leave the fund with a less liquid portfolio.

The example above (and the additional examples in Appendix B) show that, in many cases, this approach does not work. The problem arises when (i) investors believe a run is starting but will be small enough that a fee is unlikely to be applied in the current period, and (ii) liquidating investment in future periods may be costly. An attentive investor will then recognize that redemption fees are likely to be larger in the future, creating the incentive to

[16] The 2014 rules also allowed funds to impose redemption gates (that is, to suspend convertibility) once the threshold was passed. The approach in Engineer (1989) can be adapted to show that such a suspension policy is ineffective at preventing runs in our setting. For this reason, we focus our analysis on the more promising part of the 2014 rules: allowing redemption fees.

redeem preemptively. This type of incentive appears to have played an important role in the runs on money market funds in March 2020.[17] Put differently, this type of policy corrects the negative externality associated with early redemptions only if redemption demand is large enough to immediately indicate a run is underway, which is not always the case.

What should policymakers do? One common approach in the literature is to assign a "sunspot" probability to the run equilibrium and derive the best fee policy taking into account the possibility of a run. This policy will typically permit a run to occur if the sunspot probability is small enough and will eliminate the run equilibrium if the probability is larger (see, for example, Cooper and Ross, 1998, and Peck and Shell, 2003, or more recent work by Dávila and Goldstein, 2023). Determining this probability in practice is quite difficult, however. Moreover, it seems possible that the policy process may systematically underestimate the probability of a future run, leading to reforms that are insufficiently aggressive.[18,19] For these reasons, we focus on policies that eliminate the run equilibrium. In the next section, we derive the best such policy in the sense of making the fund as attractive as possible to investors without introducing a run equilibrium. We first characterize the best run-proof policy, both in general and within a simple class of policies, when the fraction δ of attentive investors and the distribution of the period-2 liquidation value r_2 are known. We then characterize the best policy that is robust to changes in δ and the distribution of r_2.

4 Preventing runs

In situations where the policy in Section 3 admits a run equilibrium, preventing runs requires the fund to impose a fee in at least some circumstances where redemption demand is in the normal range, that is, $m_1 \leq \pi$. It is always possible to prevent runs by using a large enough fee. For example, setting $c_1(m_1)$ to the lowest possible liquidation value of investment ($\underline{r}$) for all $m_1 \in [0, \pi]$ completely removes the strategic complementarity in investors' actions. Early redemptions by other investors would then *increase* the payments to investors who redeem in

[17] For example, the reform proposal in Securities and Exchange Commission (2022) states that "the possibility of an imposition of a fee ... appears to have contributed to incentives for investors to redeem." See Li et al. (2021) for more empirical evidence of this effect.

[18] Funds and investors that expect to be bailed out in a crisis have an incentive to downplay the possibility of a run, for example, and lobby for weaker reforms. The fact that the run in March 2020 came only a few years after the previous reform is consistent with the idea that those reforms were insufficiently aggressive.

[19] Another alternative is to try to endogenize the probability of a run using global games, as in Goldstein and Pauzner (2005) or private sunspots as in Mitkov (2025). However, these approaches require restricting the structure of the model and the types of contracts allowed, which would conflict with our goal of deriving the best policy within a broad class.

subsequent periods and there would be no incentive to join a run. While this policy prevents runs, it also gives low consumption to type-1 investors in the no-run equilibrium, sharply reducing the fund's attractiveness.

In this section, we study the least costly way for the fund to rule out preemptive runs. We first derive a general characterization of the best run-proof fee policy. We then characterize the best *simple* policy, in which the fee is zero up to a threshold level of redemption demand and constant above it. In both cases, we illustrate how the best run-proof policy depends on parameters, especially the fraction δ of attentive investors and the distribution of the future liquidation value r_2. Finally, we study simple policies that are *robust* in the sense of being run-proof for a range of values for these parameters. We show that the threshold and fee in the best robust policy are set to protect the fund against runs of different sizes.

4.1 General fee policies

In a *general* policy, the fund can choose any feasible payment function $c_1(m_1) \leq 1$ in period 1 when redemption demand is consistent with fundamentals, that is, when $m_1 \in [0, \pi]$. When a fee is imposed in period 1, the proceeds are divided efficiently among the remaining investors in periods 2 and 3. Other parts of the analysis, including the time-consistency constraints when redemption demand is greater than π, remain the same. It is straightforward to show that any such policy will again generate a no-run equilibrium in the redemption game. Our goal is to find the best policy that does not also admit a run equilibrium.

The run-proof constraint. To prevent a run, the fund must set the function $c_1(m_1)$ so that, in the event of a run, the net benefit of redeeming early for an attentive non-type-1 investor is non-positive. Introducing a redemption fee policy into the expression for this net benefit in equation (14) yields

$$\mathcal{R} \equiv \int_0^{\frac{\pi-\delta}{1-\delta}} \left\{ \begin{array}{c} u\big(c_1(m_1)\big) - p_{\pi_1}\mathbb{E}\left[u\big(c_2(m_1, m_2, r_2)\big)\right] \\ -(1-p_{\pi_1})\mathbb{E}\left[u\big(c_3(m_1, m_2, r_2)\big)\right] \end{array} \right\} f_n(\pi_1) d\pi_1 - T \leq 0, \tag{15}$$

where the payments c_2 and c_3 come from Proposition 3 when redemption demand (m_1, m_2) corresponds to a run as in equation (13). As before, the potential benefit of redeeming in period 1 comes when redemption demand π_1 is small, which corresponds to the integral term in $\mathcal{R}$. The redemption fee policy must decrease this benefit until it is no more than the cost T of redeeming early in states where high redemption demand triggers the time-consistent fee (see equation (14)).

Welfare. When the contract satisfies equation (15), no run will occur in equilibrium and investors' expected utility at $t = 0$ is

$$\mathcal{W} = \int_0^{\pi} \left[\pi_1 u\big(c_1(\pi_1)\big) + (\pi - \pi_1) u\big(c_2^N(\pi_1)\big) + (1-\pi) u\big(c_3^N(\pi_1)\big)\right] f(\pi_1) d\pi_1, \tag{16}$$

with the payments in periods 2 and 3 given by

$$c_2^N(\pi_1) = \min\Bigg\{\underbrace{\frac{\pi - \pi_1 c_1(\pi_1)}{\pi - \pi_1}}_{\textit{no excess liquidity}}, \underbrace{\frac{R(1-\pi) + \pi - \pi_1 c_1(\pi_1)}{1 - \pi_1}}_{\textit{excess liquidity at } t=2}\Bigg\} \tag{17}$$

and

$$c_3^N(\pi_1) = \max\Bigg\{\underbrace{R}_{\textit{no excess liquidity}}, \underbrace{\frac{R(1-\pi) + \pi - \pi_1 c_1(\pi_1)}{1 - \pi_1}}_{\textit{excess liquidity at } t=2}\Bigg\}, \tag{18}$$

where the N superscript indicates "no run". Depending on the size of the fee imposed at $t = 1$, these payments will take one of two forms: either all remaining liquid assets will be given to redeeming investors at $t = 2$, or only some liquid assets will be paid out at $t = 2$ and the remainder will be held until $t = 3$. In the latter case, we say the fund is holding *excess liquidity* at $t = 2$. If no excess liquidity is held, type 3 investors earn the return on investment R and the fee revenue is paid out to type 2 investors. If following this plan would give type 2 investors more than R, the fund will instead set $e_2 > 0$ so that type 2 and type 3 investors earn the same amount, as shown in the expressions above.

Optimal policy. The fund's problem is to choose the payment function $c_1(m_1)$ for $m_1 \in [0, \pi]$ to maximize equation (16) subject to the run-proof constraint (15). Let $c_1^*(m_1)$ denote the solution to this problem, which we call the best *general* run-proof policy. This problem is complicated by the fact that both the objective and the constraint are non-smooth functions of $c(m_1)$ due to the max and min operators above. Nevertheless, we can gain insight into the properties of the optimal policy by looking at the first-order condition in regions where the objective and constraint functions are locally smooth.[20] Suppose, for example, the optimal policy in a neighborhood of m_1 involves no excess liquidity when there is no run (see equations (17) –(18)) and liquidation in the event a run is detected at $t = 2$ (see Proposition 3). Then

[20] We provide a more general analysis of the solution and its properties in Online Appendix C.

the first-order necessary condition for the optimal choice of $c_1(m_1)$ is

$$m_1 \Big\{ \underbrace{u'\big(c_1(m_1)\big) - u'\big(c_2(m_1, \pi - m_1)\big)}_{\equiv F(c_1(m_1), m_1)} \Big\} f(m_1) \tag{19}$$
$$= \lambda \Big\{ \underbrace{u'\big(c_1(m_1)\big) + \frac{m_1}{1 - m_1} \mathbb{E}\left[u'\big(c_2(m_1, m_2, r_2)\big) \right]}_{\equiv G(c_1(m_1), m_1)} \Big\} f_n\left(\frac{m_1 - \delta}{1 - \delta} \right).$$

The left-hand side of this condition is the increase in expected utility associated with a marginal increase in c_1. The m_1 type 1 investors will have marginally more consumption, which they value at $u'(c_1)$. Type 2 investors will consume less; this cost is captured by the $u'(c_2)$ term. If $c_1(m_1) < 1$, the difference between these terms is strictly positive, since increasing c_1 toward 1 moves the allocation closer to the planner's solution. The expected benefit of this increase is proportional to the probability that this level of redemption demand will occur in equilibrium, $f(m_1)$.

The right-hand side of equation (19) is the cost of increasing c_1 in terms of tightening the run-proof constraint, which is proportional to the Lagrange multiplier λ. Increasing $c_1(m_1)$ raises the incentive to run both directly, by increasing the payoff of redeeming at $t = 1$, and indirectly, by decreasing the payoff of redeeming in later periods; both terms in the curly brackets are strictly positive. The overall cost is also proportional to the probability of m_1 arising in the event of a run, which is the probability (as seen by a non-type 1 investor) that the proportion of type 1 investors is equal to $(m_1 - \delta)/(1 - \delta)$.

Some general properties of the optimal policy can be seen from equation (19). First, for any $m_1 < \delta$, the f_n term on the right-hand side is zero; if a run occurs, redemption demand cannot be less than δ. Equation (19) then requires $c_1 = c_2$ on the left-hand side, which implies $c_1 = c_2 = 1$ and $c_3 = R$. In other words, when redemption demand is inconsistent with a run, there should be no fee and the optimal policy implements the planner's allocation. For any $m_1 \geq \delta$, the f_n term is strictly positive and equation (19) implies $c_1(m_1) < 1$ must hold. When redemption demand is consistent with a run, a fee should be applied. In this region, the optimal c_1 varies with redemption demand m_1.

Solving for the full optimal fee schedule requires taking into account the max and min operators in all of the payment functions that enter the objective and the run-proof constraint, which we do numerically. Figure 2 presents two examples that illustrate how the optimal fee varies with redemption demand. Panel (a) is based on the same parameter values as Figure 1 above, including a uniform distribution for π_1, while panel (b) uses a truncated normal distribution.[21] As expected, both panels show there is no fee when redemption demand m_1 is less than $\delta = 0.2$, and the time-consistent fee from Proposition 2 is applied when $m_1 > \pi = 0.5$. The interesting region is in the middle, where $m_1 \in [\delta, \pi]$. In panel (a), the payment c_1 is an *increasing* function of redemption demand m_1 in this region, meaning the fee is smaller when more investors redeem. In panel (b), the fee is non-monotone in this region. These patterns are counterintuitive at first: one might have expected larger redemptions to lead to a higher redemption fee. Why might the fee become *smaller* when more investors redeem?

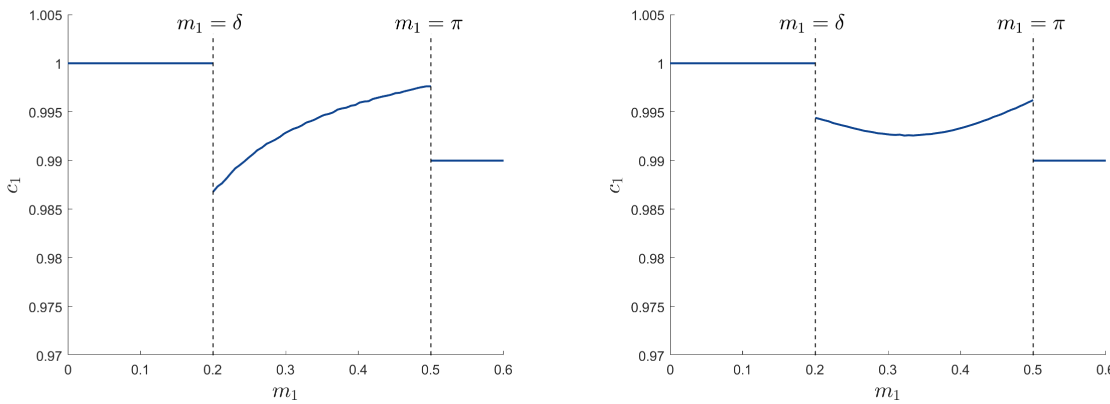

(a) Uniform distribution (b) Truncated normal distribution

Figure 2: The best run-proof general policy

To understand these patterns, rewrite the first-order condition in equation (19) as

$$L(m_1, c_1) \equiv \left(\frac{F\big(m_1, c_1\big)}{G\big(m_1, c_1\big)} \right) \frac{f(m_1) m_1}{f_n\left(\frac{m_1 - \delta}{1 - \delta}\right)} = \lambda. \tag{20}$$

This equation implicitly defines the best run-proof contract $c_1(m_1)$ depicted in the middle region of both panels. The first term in L is a marginal benefit-cost ratio: the numerator F is the increase in expected utility per type 1 investor when c_1 is increased, while the denominator G is the increase in the incentive for a non-type-1 investor to run. We show in Online Appendix C.2 (Lemma 4) that this ratio is decreasing in c_1: when c_1 is larger, the marginal benefit of increasing c_1 further is smaller relative to the marginal cost of increasing the incentive to run. The ratio can be either increasing or decreasing in m_1, depending on parameter values.

The second term in L is a *weighted likelihood ratio*: the numerator is proportional to the

[21] The distribution f in panel (b) is a truncated normal distribution with a mean of 0.3 and a standard deviation of 0.2. All other parameter values are the same as in panel (a).

probability of m_1 conditional on investors not running, while the denominator is the probability of m_1 if a run were to occur (as viewed by a non-type-1 investor). The numerator is weighted by the number of redeeming investors m_1 because the welfare cost of imposing a fee is proportional to the number of investors who pay the fee. The denominator, in contrast, reflects the incentive to run for an individual investor. For many distributions f, including the uniform distribution used in panel (a), the weighted likelihood ratio is increasing in m_1. If this effect is strong enough, the entire function L will be increasing in m_1.[22]

Because equation (20) shows that L must equal the constant λ for all values of m_1, and because L is decreasing in c_1 as discussed above, higher values of m_1 will be associated with larger values of c_1 if and only if L is increasing in m_1.[23] In these cases, $c_1(m_1)$ is increasing in the middle region as shown in panel (a). Intuitively, imposing a large fee when redemption demand is high tends to be costly because many investors will be hit by the fee in those states. Unless large redemption demand when there is no run is sufficiently unlikely, the best run-proof fee will be *smaller* when many investors redeem.

Panel (b) in Figure 2 shows that the optimal fee schedule is not always decreasing in this region, however. Depending on the distribution f and other parameter values, the best run-proof fee can be a complex and non-monotone function of redemption demand. Because such policies may be difficult to implement in practice, we now shift our focus to a simpler class of policies that take the same form as the 2023 MMF reform in the U.S.

4.2 Simple policies

In this section, we study a class of simpler policies in which the fee is zero for redemption demand up to a threshold and constant afterward, that is,

$$c_1(m_1) = \begin{cases} 1 & \text{if } m_1 \in [0, \bar{m}) \\ \bar{c} & \text{if } m_1 \in [\bar{m}, \pi] \end{cases}$$

for some pair $(\bar{m}, \bar{c})$ with $\bar{c} \leq 1$. If redemption demand m_1 were larger than π, the time-consistent redemption fee from Proposition 2 would be applied, as before.

We divide the run-proof condition for a simple policy $(\bar{m}, \bar{c})$ into two parts. If the threshold is relatively low, with $\bar{m} \leq \delta$, then the fee will apply for certain in the event of a run. For

[22] In Online Appendix C.2 (Lemma 5), we provide a sufficient condition on the weighted likelihood ratio for L to be increasing in m_1.

[23] See Proposition 7 in Online Appendix C.3 for a formal statement and proof of this result.

this case, the run-proof condition in equation (15) can be written as

$$\mathcal{R}(\bar{m},\bar{c}) \equiv \int_0^{\frac{\pi-\delta}{1-\delta}} \left\{ \begin{array}{l} u(\bar{c}) - p_{\pi_1}\mathbb{E}\left[u\big(c_2(m_1,m_2,r_2)\big)\right] \\ -(1-p_{\pi_1})\mathbb{E}\left[u\big(c_3(m_1,m_2,r_2)\big)\right] \end{array} \right\} f_n(\pi_1)d\pi_1 - T \leq 0. \tag{21}$$

The payments c_2 and c_3 above are determined by Proposition 3, but with $c_1(m_1)$ set to the constant $\bar{c}$ for all $m_1 \in [\bar{m}, \pi]$. If the threshold is higher, with $\bar{m} > \delta$, the fee will only apply in a run if the realization of π_1 is large enough that $m_1 = \pi_1 + \delta(1-\pi_1) \geq \bar{m}$. The run-proof condition in this case is

$$\begin{aligned} \mathcal{R}(\bar{m},\bar{c}) \equiv & \int_0^{\frac{\bar{m}-\delta}{1-\delta}} \Big\{ u(1) - p_{\pi_1}\mathbb{E}\left[u\big(c_2\big)\right] - (1-p_{\pi_1})\mathbb{E}\left[u\big(c_3\big)\right] \Big\} f_n(\pi_1)d\pi_1 \\ & + \int_{\frac{\bar{m}-\delta}{1-\delta}}^{\frac{\pi-\delta}{1-\delta}} \Big\{ u(\bar{c}) - p_{\pi_1}\mathbb{E}\left[u\big(c_2\big)\right] - (1-p_{\pi_1})\mathbb{E}\left[u\big(c_3\big)\right] \Big\} f_n(\pi_1)d\pi_1 - T \leq 0 \end{aligned} \tag{22}$$

where the indexes (m_1, m_2, r_2) have been omitted from c_2 and c_3 to save space.
Define the set of simple run-proof policies P as

$$P \equiv \big\{ (\bar{m},\bar{c}) \mid \mathcal{R}(\bar{m},\bar{c}) \leq 0 \big\}.$$

Figure 3 depicts this set in blue, using the same parameters as the example in Figure 1. For any $\bar{m} \leq \delta$, the incentive to run $\mathcal{R}(\bar{m},\bar{c})$ is independent of the threshold $\bar{m}$, as shown in equation (21). It is straightforward to show that, in this region, there is a unique $\bar{c}^*$ such that the incentive to run is positive for $\bar{c} > \bar{c}^*$ and negative for $\bar{c} < \bar{c}^*$. In other words, $\bar{c}^*$ is the highest payment that can be made at $t = 1$ without introducing a run equilibrium *if* the redemption fee would always be activated by a run. When $\bar{m} > \delta$, the fee may or may not be activated by a run, depending on the realized value of π_1. In this region, the incentive to run is strictly increasing in both $\bar{m}$ and $\bar{c}$, which implies the boundary of the run-proof set is strictly decreasing as shown in the figure. Intuitively, if the fee is less likely to be activated in a run, it must be larger to make running unattractive.

Optimal policy. Under a simple policy $(\bar{m},\bar{c})$, welfare in the no-run equilibrium from equation (16) can be written as

$$\begin{aligned} \mathcal{W}(\bar{m},\bar{c}) \equiv & \int_0^{\bar{m}} [\pi_1 u(1) + (\pi-\pi_1)u(1) + (1-\pi)u(R)] f(\pi_1)d\pi_1 \\ & + \int_{\bar{m}}^{\pi} \left[\pi_1 u(\bar{c}) + (\pi-\pi_1)u\big(c_2^N(\pi_1)\big) + (1-\pi)u\big(c_3^N(\pi_1)\big) \right] f(\pi_1)d\pi_1, \end{aligned} \tag{23}$$

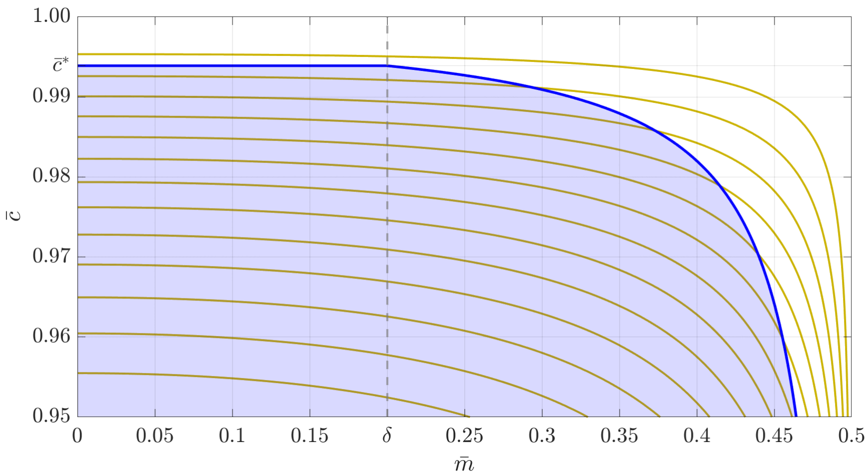

Figure 3: Run-proof set and the best simple policy

where the payments c_2^N and c_3^N are given by equations (17) and (18) with $c_1(\pi_1)$ equal to $\bar{c}$ for all $\pi_1 \in [\bar{m}, \pi]$. It is straightforward to show that welfare is strictly increasing in $\bar{m}$ for any $\bar{c} < 1$ and is strictly increasing in $\bar{c}$ for any $\bar{m} < \pi$. For a given fee, welfare is higher if the fee is imposed less often and, for a given threshold, welfare is higher if the fee imposed above that threshold is smaller. As a result, the indifference curves in the space of policies $(\bar{m}, \bar{c})$ are strictly decreasing, as depicted by the yellow curves in Figure 3. It follows immediately that the best run-proof policy cannot have $\bar{m} < \delta$. It may be at the kink point, where $\bar{m} = \delta$, or it may lie on the downward-sloping portion, with $\bar{m} > \delta$.

Our next result shows that if the best general policy $c_1^*(m_1)$ is increasing on $[\delta, \pi]$, as in panel (a) of Figure 2, the best simple policy is at the kink point in Figure 3. As discussed above, this condition is met if the distribution f is uniform and in many other cases as well.

Proposition 4. *If $\bar{c}^* < 1$ and the best general policy $c_1^*(m_1)$ is increasing on $[\delta, \pi]$, then the unique best simple run-proof policy sets $\bar{m} = \delta$ and $\bar{c} = \bar{c}^*$.*

Intuitively, a given redemption fee is costlier in states where π_1 is large because more investors are hit by the fee, as emphasized in Section 4.1. If the policy were to set the threshold higher than δ, the corresponding payment would be less than $\bar{c}^*$, at a point on the downward-sloping part of the frontier in Figure 3. Starting from this situation, decreasing the threshold and the fee together shifts the burden of the fee away from states where many investors pay the fee and into states where fewer investors pay the fee. When the best general policy is an increasing function, this shift unambiguously increases investors' expected utility. It bears emphasizing that this condition is sufficient, but not necessary, for the best simple policy to have $\bar{m} = \delta$. For a wide range of parameter configurations, the best policy imposes a fee whenever redemption demand at $t = 1$ is consistent with a run.[24]

Discussion. Unlike the (complex) general policy depicted in Figure 2, the policy in Proposition 4 is simple enough to use as a guide for regulation. Implementing the policy in practice would require calibrating the model parameters, of course. Some parameters would be fairly straightforward. The distribution f, for example, could be mapped to the distribution of daily outflows from the fund in normal times, and π could correspond to the upper bound of the support of this distribution. We provide such a calibration in Section 5 below; see Figure 6. However, the best run-proof policy also depends on parameters that are less easily observed and for which changes may be difficult to monitor.

The parameter δ, for example, could be calibrated to match the outflows in the first pricing period of past run episodes, but its value might be quite different in future episodes. The runs on Silicon Valley Bank and some other U.S. regional banks in early 2023 provide a useful lesson in this regard. The speed of the withdrawals from these banks was much faster than in past banking crises, highlighting that the value of δ in our model could depend on factors such as who a fund's investors are and how frequently/easily they communicate with each other.[25] Shifts in social media or other communication patterns could lead to changes in δ that would be difficult for a fund or regulator to discern in real time. Alternatively, δ could reflect the time of day that a run begins; a run that starts early in the morning would likely have more first-day redemptions than a run that starts later in the day. Past data on first-day redemptions would have limited predictive power in this case as well.[26] The best policy also depends on investors' belief about the future liquidation value r_2, which may be difficult to measure in a way that is verifiable enough to be used in the redemption fee calculation. We deal with this issue next by introducing a notion of *robust* policies.

4.3 Robust policies

We say a fee policy is *robust* if it is run-proof for *any* fraction δ of attentive investors and *any* distribution of the future liquidation value r_2 on $[\bar{r}, \underline{r}]$. In this section, we extend our results above to characterize the best robust simple policy. Designing a robust policy typically

[24] In Online Appendix D, we prove a generalized version of this result (labeled Proposition $4'$) that applies to any interval $[m^T, \pi]$ with $m^T \geq \delta$ and is used in Section 4.3 below.

[25] Cipriani and La Spada (2020) and Allaire et al. (2024) provide evidence that the composition of investors within a fund affects redemption behavior.

[26] One could make δ stochastic, but the best policy would then depend on the distribution of δ, which would be equally difficult to measure and monitor over time.

involves focusing on a worst-case scenario. For a given simple fee policy, it is straightforward to show that investors' incentive to run in equations (21) and (22) is strictly decreasing in r_2. This fact implies there is a clear worst-case scenario: r_2 takes its smallest possible value $\underline{r}$ with certainty, which we refer to as the "adverse liquidity scenario." A fee policy that is run-proof for this scenario is also run-proof for any distribution of r_2 on $[\underline{r}, \overline{r}]$.

Our analysis in Section 3, especially Figure 1, shows that identifying a worst-case scenario for the parameter δ is less straightforward. When there is no redemption fee, the incentive to run is highest when δ takes an intermediate value: large enough that the run will cause damage, but small enough that it may not trigger a fee at $t = 1$. We show below that when the fund uses a simple fee policy, there are two "worst-case" scenarios for δ, one where the fee is applied in all states in the event of a run and another where the fee is only applied in some states. These two scenarios determine the best run-proof fee and threshold, respectively.

The best robust policy maximizes investors' expected utility in equation (16) subject to the constraint that the policy lies in the run-proof set depicted in Figure 3 for all $\delta \in [0, 1]$ under the adverse liquidity scenario. Formally, the problem can be written as

$$\max_{\{\bar{m},\bar{c}\}} \quad \mathcal{W}(\bar{m}, \bar{c})$$
$$s.t. \quad \mathcal{R}(\bar{m}, \bar{c}; \delta, \underline{r}) \leq 0 \quad \text{for all } \delta.$$

The objective function is unchanged because there is no run in equilibrium and, therefore, δ and r_2 have no effect on equilibrium welfare or the indifference curves in Figure 3. The function $\mathcal{R}$ in the constraints of this problem is again given by equations (21) and (22); the notation now emphasizes that this function depends on δ and r_2 and is being evaluated at $r_2 = \underline{r}$ with certainty. The constraint set for this problem is, therefore, the intersection of the run-proof set in Figure 3 across all values of δ, with each set evaluated in the adverse liquidity scenario.

The set of robust policies. We characterize the constraint set in two steps.

Step 1. First, suppose the threshold $\bar{m}$ were set to zero, so the redemption fee would be applied for certain at $t = 1$ regardless of δ and π_1. Ask: what is the largest payment $\bar{c}$ that is run-proof for all δ in this case? Define

$$\bar{c}^*_R = \min_{\delta \in [0,1]} \left\{ \bar{c}^*(\delta) : R\left(0, \bar{c}^*(\delta); \delta, \underline{r}\right) = 0 \right\}.$$

Graphically, for each value of δ, find the payment $\bar{c}^*(\delta)$ that corresponds to the flat part of the boundary of the run-proof set in Figure 3.[27] Then $\bar{c}_R^*$ (the subscript indicates "robust") is the smallest of these payments. Because $\bar{c}^*(\delta)$ is continuous, such a minimum always exists. Let δ_H denote the value of δ where this minimum is reached. Then δ_H is the worst-case scenario, in the sense of requiring the highest fee to remove the incentive to run, if the fee will be applied in all states.

Step 2. Next, suppose the payment in the fee region is set to $\bar{c}_R^*$, but allow the threshold $\bar{m}$ to be positive. Ask: what is the highest threshold that is run-proof for all δ in this case? Define

$$\bar{m}_R^* = \min_{\delta \in [0,1]} \left\{ \bar{m}(\delta) : R\left(\bar{m}(\delta), \bar{c}_R^*; \delta, \underline{r}\right) = 0 \right\}.$$

Graphically, for each value of δ, find the threshold $\bar{m}(\delta)$ where the boundary of the run-proof set in Figure 3 crosses $\bar{c}_R$,[28] and let $\bar{m}_R^*$ denote the smallest of these thresholds. Let δ_L denote the value of δ for which this minimum attains. Then δ_L is the worst-case scenario, in the sense of requiring the lowest threshold to remove the incentive to run, if the fee is set to $\bar{c}_R^*$.

These two steps are illustrated in Figure 4, using the same parameters as the example in Figure 1. Panel (a) shows the boundary of the run proof set $P(\delta)$ for a collection of $\delta \geq \delta_H$. As δ decreases from 1, the set of run-proof policies becomes strictly smaller until δ_H is reached. Panel (b) shows the run-proof boundaries for a collection of $\delta \leq \delta_L$. In this panel, as δ increases from zero, the set of run-proof policies becomes strictly smaller until δ_L is reached. As panel (c) illustrates, δ_L is strictly less that δ_H, and the run-proof sets for these two values are overlapping. Because δ_H is associated with larger runs, a larger fee (i.e., a lower $\bar{c}$) is needed to make the policy run proof for much of the region, but the threshold can be set higher and still trigger a fee with high probability. Under δ_L, in contrast, the run is smaller and can be prevented with a smaller fee, but the threshold associated with any given fee must be lower so the fee will be triggered with sufficient probability in a run. Panel (c) also depicts the run-proof set for a value of δ strictly between δ_H and δ_L; this set lies strictly outside the robust run-proof set.

The best robust policy. As the figure illustrates, the boundary of the set of robust run-proof policies is flat at $\bar{c}_R^*$ up to the threshold $\bar{m}_R^*$ and then downward sloping. The indifference curves generated by investors' expected utility are the same as in Figure 3. Our next

[27] If $\mathcal{R}(0,1;\delta,\underline{r}) \leq 0$, i.e., making the efficient payment is already run-proof for the given δ and $\underline{r}$, set $\bar{c}^*(\delta) = 1$.

[28] If the boundary does not cross $\bar{c}_R$, for example because δ is small and all policies are run-proof, then set $\bar{m}(\delta) = \delta_H$.

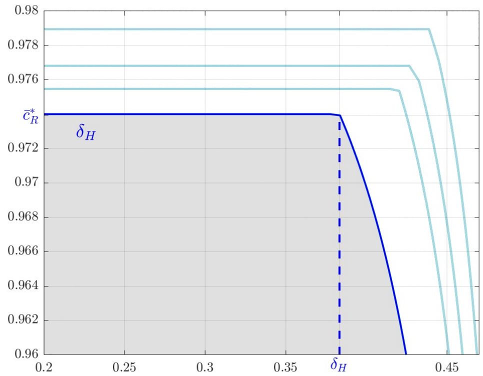

(a) $\delta \geq \delta_H$

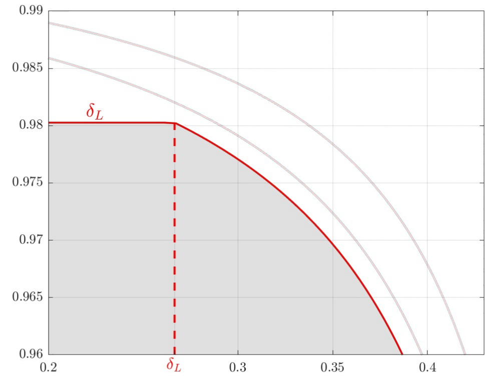

(b) $\delta \leq \delta_L$

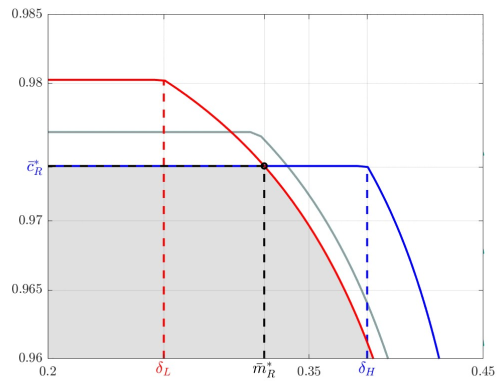

(c) $\delta_L \leq \delta \leq \delta_H$

Figure 4: Robust simple policies

result shows that, under the same conditions studied in the earlier sections, the best policy is again at the kink point of this set, $(\bar{m}_R^*, \bar{c}_R^*)$. A proof of this result follows directly from Proposition 4′ in Appendix D.

Proposition 5. *If $\bar{c}_R^* < 1$ and the the best general policy $c_1^*(m_1)$ based on δ_L is increasing on $(\bar{m}_R^*, \pi)$, then $(\bar{m}_R^*, \bar{c}_R^*)$ is the unique best robust simple policy.*

The two elements of the best robust simple policy have clear and distinct interpretations. The payment $\bar{c}_R^*$, and hence the redemption fee, is determined entirely by δ_H. In other words, the fee is set to guard against the worst-case *large* run that would trigger the fee for certain. If the threshold were be set to δ_H, however, the fund would be susceptible to a slightly smaller run (i.e., a lower value of δ) because then the fee would not be triggered at $t = 1$ in some states. To be robust run-proof, the policy must set the threshold to $\bar{m}_R^*$, which guards against the worst-case *small* run given $\bar{c}_R^*$.

Figure 5 depicts the expected payoff of redeeming and waiting when the fund adopts the best robust simple policy, using the same parameter values as Figure 1.[29] The figure shows that the policy makes waiting at least weakly better than redeeming for all values of δ. The blue curve has a kink at $\delta = \bar{m}_R^*$, reflecting the fact that redemption demand in a run will be large enough to trigger the fee for sure when $\delta > \bar{m}_R^*$, but only with some probability when $\delta < \bar{m}_R^*$. The figure illustrates how the run-proof condition binds at both δ_L and δ_H under the best robust policy.

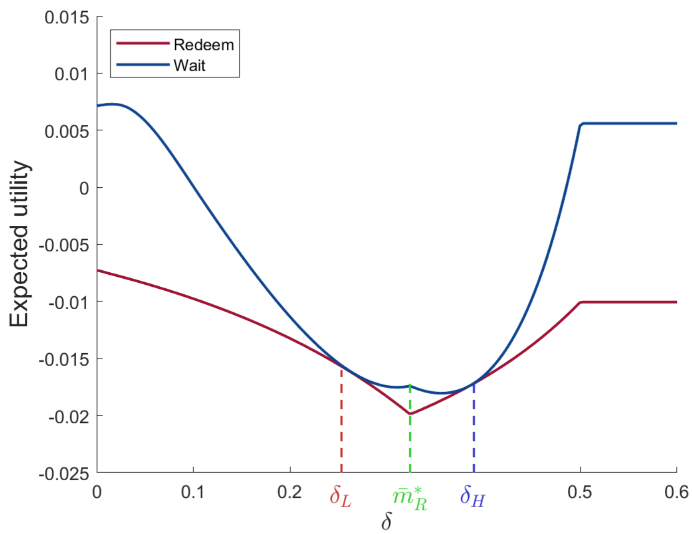

Figure 5: No preemptive run equilibrium under the best robust policy

We view the result in Proposition 5 as the main policy prescription of our model. In the next section, we compare this prescription to the 2023 MMF reform in the U.S. and discuss its broader implications.

5 Evaluating the 2023 Reform

In this section, we use our model to evaluate the MMF reform adopted in the U.S. in 2023. After describing how redemption fees are set in the new policy, we show that our analysis identifies two important weaknesses. First, when liquidity conditions may deteriorate, the policy can leave funds vulnerable to a preemptive run. Second, when liquidity conditions are stable, the fee imposed by the policy is inefficiently large, making funds less attractive to investors. Finally, we show that another feature of the 2023 reform, increased liquidity requirements, is less desirable when fees are set according to the best robust policy.

[29] The best robust policy for this example sets $\bar{m}_R^* = 0.324$ and $\bar{c}_R^* = 0.974$.

Vertical slice rule. Under the 2023 reform, institutional prime and tax-exempt MMFs are required to impose a redemption fee in a pricing period if net redemptions exceed 5% of the fund's assets. This threshold is low enough that it is expected to be met with some regularity in normal times. When the threshold is met, a fund is required to impose a redemption fee equal to the cost it would face if it were to sell a pro-rata share of each security in its portfolio. In other words, redeeming investors will receive the current liquidation value of a "vertical slice" of the fund's portfolio.

In our model, this new rule corresponds to a simple policy with the threshold $\bar{m}$ set to 5% and the payment

$$\bar{c}_V = r_1(1 - \pi) + \pi. \tag{24}$$

This payment corresponds to the time-consistent payment in our analysis when redemption demand in period 1 is greater than π and the fund chooses to liquidate some investment (see Proposition 2). It is important to note that $\bar{c}_V$ relies solely on the current liquidation value r_1, whereas the best robust fee $\bar{c}_R^*$ also depends on the future adverse liquidity scenario $\underline{r}$. This lack of forward-looking considerations creates the policy's first weakness.

Vulnerability. Our model identifies situations where preemptive runs may still occur under the new policy. Suppose, for example, markets are currently perfectly liquid ($r_1 = 1$), but investors believe conditions may worsen: the lower bound $\underline{r}$ of the future liquidation value r_2 is less than 1. The vertical-slice rule in equation (24) will then set $c_1 = 1$ for all m_1, meaning no redemption fee is applied. Investors recognize, however, that if a run were to occur, the fund may need to liquidate assets in period 2 and a redemption fee may apply then. This possibility can give non-type-1 investors an incentive to run preemptively.

A calibrated example. Figure 6 illustrates this result using a distribution for π_1 that is calibrated to match features of the data. The blue bars in panel (a) represent, for different levels of the threshold $\bar{m}$, the fraction of prime and tax-exempt MMFs that would have exceeded the threshold on a given day in the period December 2016 - October 2021.[30] We use a truncated normal distribution for π_1 and choose the upper bound π, mean μ, and standard deviation σ to approximate this data; the orange bars show the calibrated distribution.[31] Panel (b) presents the expected value of redeeming in period 1 and of waiting for a non-type 1 investor when all other attentive investors redeem. Because no fee is applied in period 1

[30] This data comes from Table 6 of SEC (2023).

[31] The calibrated parameters are $\pi = 0.1301, \mu = 0.0032$ and $\sigma = 0.0427$. The other parameter values used in panel (b) of Figure 6 are $R = (1.04)^{\frac{1}{180}}$, $r_1 = 1$, and $\underline{r} = 0.8$.

regardless of redemption demand, the value of redeeming is $\ln(1) = 0$ for all δ. The panel shows that run equilibrium exists for a wide range of values of δ in this case.

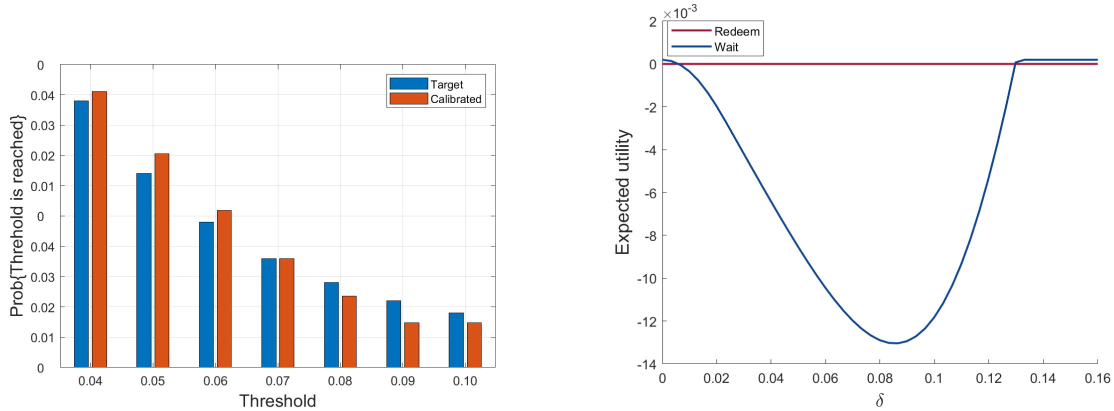

(a) Calibrated distribution of π_1 (b) Incentive to run

Figure 6: Calibrated example

This vulnerability is also illustrated in panel (a) of Figure 7, which uses the same parameter values. As in Figures 3 and 4 above, the shaded area represents the set of simple, robust run-proof contracts. In this example, the possibility of costly liquidation next period implies that a positive fee is needed to make the fund run-proof for some values of δ. The dashed line indicates the period 1 payout associated with the vertical slice rule, which has zero redemption fee in this case. The figure shows that no choice of threshold $\bar{m}$ would make the fund robust run-proof. Overall, this first example demonstrates that the current redemption fee must take future liquidation values into account to be effective. The 2023 reform does not have this feature.

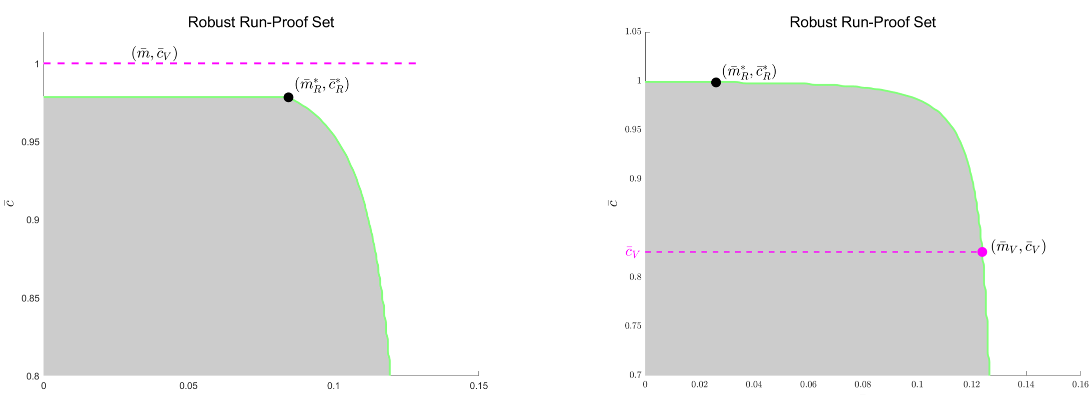

(a) Vulnerability (b) Inefficiency

Figure 7: Effectiveness of a vertical-slice fee rule

Inefficiency. The second weakness identified by our model arises when liquidity conditions will not deteriorate. Suppose, for example, r_1 is already at the lower bound $\underline{r} < 1$. In this case, a vertical-slice rule policy is robust run-proof as long as the threshold is chosen appropriately. However, the following result shows that the fee assigned by the rule in this case is inefficiently large.

Proposition 6. *Suppose $r_1 = \underline{r} < 1$ and the conditions of Proposition 5 are satisfied. Let $\bar{m}_V$ be the largest threshold such that the policy $(\bar{m}_V, \bar{c}_V)$ is robust run proof. Then the best robust policy has $\bar{c}_R^* > \bar{c}_V$ and $\bar{m}_R^* < \bar{m}_V$.*

Proposition 5 applies whenever the best general fee policy based on δ_L is strictly increasing on the interval $(\bar{m}_R^*, \pi)$, which is satisfied by our calibrated example above.[32] The result in Proposition 6 is illustrated by the example in panel (b) of Figure 7. The vertical-slice payment $\bar{c}_V$ in this case cuts across the robust run-proof set. The best vertical-slice policy sets the threshold to the highest value in the set, $\bar{m}_V$, to minimize the probability the fee is applied in equilibrium. As is clear from the figure, this policy lies on the downward-sloping part of the frontier. It follows from Proposition 5 that welfare is improved by moving along the frontier to the northwest: using a a smaller fee and a lower threshold. Intuitively, in this example, the vertical-slice rule imposes a large fee in states where the fraction of type 1 investors is high. In these states, the conditional probability of being type 1 – and having to pay the fee – is relatively large for each investor. Risk averse investors would prefer a smaller fee, even though that fee would be applied more often.

It is worth noting that there is a case in which the vertical-slice fee is equal to the best robust fee and, therefore, a vertical slice policy can be both effective and efficient. Given other parameter values, there is a unique value of r_1 such that the payment $\bar{c}_V$ from equation (24) exactly equals $\bar{c}_R^*$, which implies the dashed line in Figure 7 coincides with the flat part of the boundary of the robust run-proof set. This coincidence is a knife-edge case, however. If r_1 is slightly higher, the vertical-slice policy will be vulnerable to a run, as in panel (a), while if r_1 is slightly lower, it will have an inefficiently high fee as in panel (b).

Desirability of liquidity requirements. As a final step in our analysis of the 2023 reform, we examine the desirability of requiring funds to hold more liquid assets. Under the new rules, funds must maintain at least 25% of their total assets in daily liquid assets and at least 50% in weekly liquid assets, up from the previous thresholds of 10% and 30%, respectively. So far in our analysis, we have held the fraction of the fund's portfolio in the liquid asset

[32] Figure 13 in Online Appendix C.4 presents the best general policy for the calibrated example.

constant at π. We interpret the new rules as increasing this fraction to a higher level, which we denote $s > \pi$. We provide examples to show that increasing s can reduce fragility and increase investors' expected utility when the redemption fee policy follows the vertical slice rule. However, in these same examples, increasing s makes investors *worse off* under the best robust simple policy. These examples illustrate that liquidity requirements are less desirable as a policy tool when redemption fees are used more efficiently.

It is straightforward to extend our analysis above to allow the fraction of the funds' portfolio held in storage to be any $s \geq \pi$. The payment associated with the vertical-slice rule in equation (24) becomes

$$\bar{c}_V(s) = r_1(1 - s) + s. \tag{25}$$

Note that this payout is increasing in s whenever $r_1 < 1$. In situations where the vertical slice payment is inefficiently low, like panel (b) of Figure 7, this increase will move the policy closer to the best robust policy, at least in the fee dimension. Using these same parameter values, panel (b) of Figure 8 shows that investors' expected utility increases as the fraction of liquid assets s is raised above π. If an (inefficient) vertical-slice rule must be used, holding more liquid assets can potentially mitigate the inefficiency.

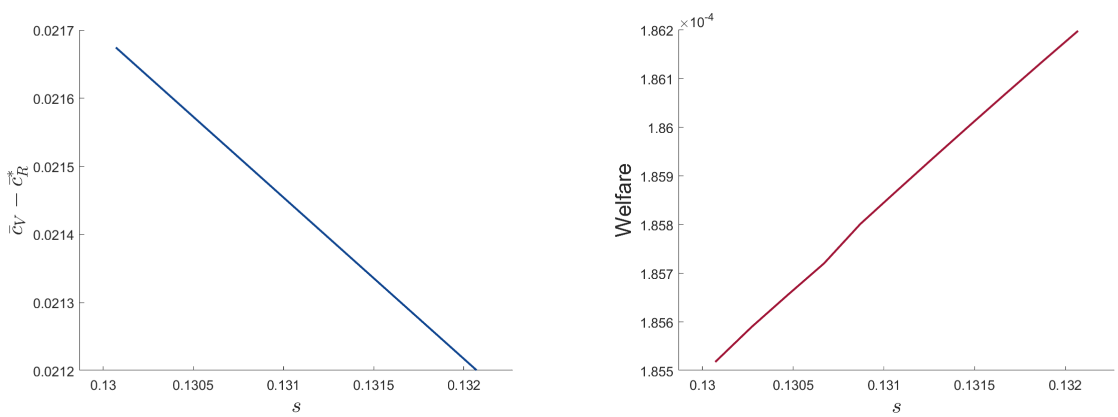

(a) Lower fragility

(b) Higher welfare

Figure 8: Effect of a higher liquidity requirement under the vertical slice rule

A higher liquidity requirement can also be useful in the situation depicted in panel (a) of Figure 7, where the fee imposed by the vertical slice rule is too small to rule out preemptive runs. When the fund holds more liquid assets ($s > \pi$), the robust run-proof set in the figure expands, and a run can be prevented with a smaller fee and/or higher threshold. In this case, even though holding more liquid assets increases the vertical-slice payout $\bar{c}_V(s)$, it can decrease the *difference* between $\bar{c}_V(s)$ and the best robust run-proof payout $\bar{c}^*_R(s)$. To illustrate this point, panel (a) of Figure 8 plots this difference and shows it to be decreasing in s. When liquidity conditions may worsen, holding more liquid assets can bring the vertical-slice fee policy closer to being robust run-proof.

Figure 9 studies this same change when the best robust fee policy is used in place of the vertical slice rule. In both examples, investors' expected utility is now strictly *decreasing* as s is raised above π. In other words, when an efficient redemption-fee policy is used, raising liquidity requirements is less attractive. To see the intuition for this result, suppose we start from the first-best allocation and compare a small redemption fee with a small increase in the fund's liquid assets. Both of these changes decrease the incentive for investors to run on the fund, but they have different effects of investors' expected utility. The redemption fee transfers consumption from type 1 to type 2 and 3 investors in some states. Starting from the first-best allocation, the utility loss from this transfer is second-order. Changing the fund's portfolio, in contrast, decreases the present value of investors' total consumption, which creates a first-order utility loss. If only a small change is needed to make the fund run-proof, therefore, it is generally better to use a redemption fee than a more liquid portfolio.[33] Figure 9 shows that this logic extends to the two examples presented here: with a well-designed fee policy in place, holding excess liquidity lowers investors' expected utility. These results suggest that future reforms should concentrate on improving the design of redemption fees and perhaps consider easing liquidity requirements.

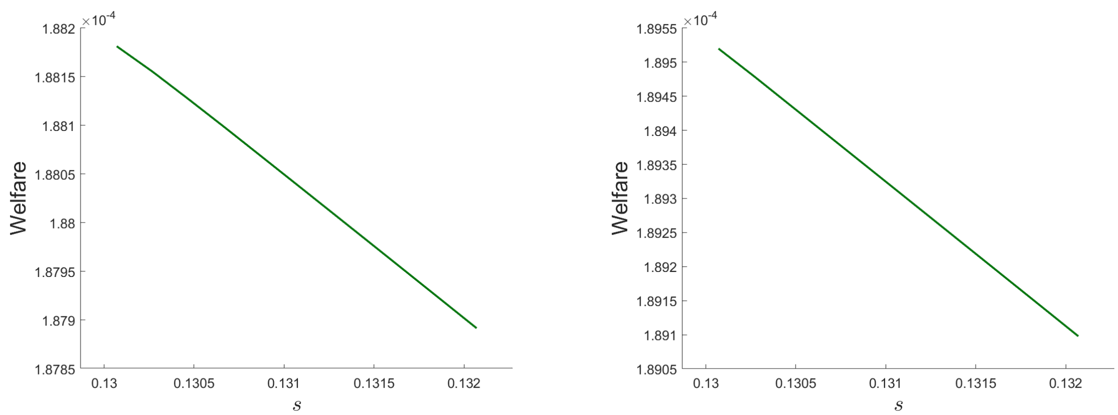

(a) Scenario 1: $r_1 = 1$ and $\underline{r} < r_1$

(b) Scenario 2: $r_1 = \underline{r} < 1$

Figure 9: Effect of a higher liquidity requirement under the best robust policy

[33] In a model with two consumption periods and sequential service, Ennis and Keister (2006) use this logic to show that the best run-proof contract has $s = \pi$ if the cost of liquidating investment is sufficiently small.

6 Concluding Remarks

In the years since the 2008 global financial crisis, there has been a wide-ranging debate about how to best prevent runs on banks and other financial intermediaries. Proposals have included requiring intermediaries to hold more capital and more liquid assets, placing limits on maturity transformation, expanding government guarantees, and more. A variety of policy reforms have been enacted for different types of intermediaries, including commercial banks, money market funds, and other shadow-banking arrangements. However, as the runs in the U.S. on prime MMFs in 2020 and on regional banks in 2023 demonstrate, significant vulnerabilities remain. Moreover, there is little consensus on the basic principles that should guide future policy reforms.

We contribute to this debate by focusing on one particular tool of financial stability policy: redemption fees. The idea that depositors or other investors should receive less from their intermediary when withdrawal demand is high has a long history in the banking theory literature. Wallace (1990) called this property a *partial suspension of convertibility* and argued that "a good banking system would have the partial suspension property" (p.19). The 2023 MMF reform in the U.S. can be viewed as an attempt to implement this observation in practice. MMFs offer a particularly useful setting for studying this type of reform, as their structure more closely resembles the banks in Diamond-Dybvig–style models than do commercial banks or many other intermediaries. The insights gained from studying reforms to MMFs can also help inform the design of future reforms and regulatory frameworks for other types of run-prone financial intermediaries.

Our analysis above shows how redemption fees can be effective at preventing runs. It also illustrates the principles that should guide the design of a fee policy. In general, the fee assigned to a given level of redemption demand should depend on three factors. All else equal, it should be higher for levels of redemption demand that are more likely to occur during a run, and it should be lower for levels of redemption demand that are more likely to occur in normal times. In addition, the fee should tend to be smaller in states where many investors need to redeem. Together, these forces imply the best general run-proof fee schedule is complex and may be a non-monotone function of redemption demand. If attention is restricted to a simpler class of policies, with a fixed fee that applies whenever redemption demand is above a threshold, our results show how these principles guide the optimal choice of the threshold and fee.

We have focused on environments like mutual funds where redeeming investors are all paid at the end of a period. Our analysis highlights the importance of a key parameter in these settings: the fraction of investors who would redeem on the first day when a run starts. We showed that the incentive for investors to run is non-monotone in this parameter; the danger comes from a run that is large enough to damage the fund but small enough that it might initially go undetected. Because this parameter is particularly difficult for funds and regulators to monitor, we advocate for policies that are robust in the sense of being immune to a run of any size. We derived the best simple, robust run-proof policy and showed how the general principles above determine the threshold and fee in this policy. The optimal fee protects against larger runs, which will trigger the fee for sure, while the optimal threshold protects against smaller runs, which are more likely to initially go undetected.

Our analysis also emphasizes that preventing runs requires the redemption fee to be forward-looking in the sense that it depends on possible future liquidation costs. The 2023 reform to MMFs in the U.S. does not have this feature, which leaves funds vulnerable to a preemptive run if investors believe liquidity conditions may worsen. Measuring the probability distribution of future liquidation costs in a way that can be used for regulatory purposes may be difficult. If so, an alternative is to specify an *adverse scenario* for future liquidation values and derive the redemption fee policy with r_2 set to this adverse scenario. It is important that this adverse scenario be constant over time, to avoid a situation where investors run preemptively because they expect the adverse scenario to be revised. This concern indicates that the adverse scenario must truly reflect a worst-case situation. However, using a very low liquidation value to calculate redemption fees may make funds unattractive to investors. One way to balance these concerns would be to require funds to have a contractual backstop that allows them to sell their investments at a specified minimum price. This minimum price would create a credible adverse scenario for use in the redemption fee policy.

As a final comment, our model abstracts from uncertainty about the fundamental value of funds' illiquid investment, which makes the analysis particularly well suited for money market mutual funds. However, many of the same considerations would arise in setting a redemption fee (or *swing pricing*) policy for other open-end mutual funds that hold illiquid assets. Extending our model to study corporate bond funds, for example, would require introducing uncertainty into the fundamental value R of matured investment. This extension may be a promising avenue for future research.

Appendix

This appendix provides additional proofs, examples, and generalized results that complement the material in the main text.

A Time consistent payments

In this appendix, we prove the two propositions in Section 2.5 that characterize time-consistent contracts. We first prove Proposition 2, which applies when a run is detected in the first period. We then prove a more general version of Proposition 3, which determines time-consistent payments when a run is detected in the second period.

A.1 Run detected in period 1

Proposition 2. *When $m_1 > \pi$, condition (TC1) requires the fund to set $e_2 = \ell_2 = 0$ and*

$$c_1^{TC1} = c_2^{TC1} = \max\left\{\frac{\pi}{\pi+\delta(1-\pi)}, r_1(1-\pi)+\pi\right\}$$
$$c_3^{TC1} = \min\left\{\frac{R}{1-\delta}, R(1-\pi)+\frac{R}{r_1}\pi\right\} \quad \textit{for all } (m_1, m_2, r_2).$$

We prove this result in the following steps. First, we present the optimization problem the fund will solve when $m_1 > \pi$ and the conditions that characertize its solution. We then present and prove three lemmas that establish properties of the solution. Finally, we use these lemmas to prove the proposition.

Under condition (TC1), the fund's problem at $t = 1$ when $m_1 > \pi$ is

$$\begin{aligned}
\max \quad & m_1 u(c_1) + [\pi+\delta(1-\pi)-m_1]\mathbb{E}\left[u\big(c_2(r_2)\big)\right] + (1-\pi)(1-\delta)\mathbb{E}\left[u\big(c_3(r_2)\big)\right] \qquad \text{[P1]}\\
s.t. \quad & m_1 c_1 + e_1 = \pi + r_1\ell_1\\
& [\pi+\delta(1-\pi)-m_1]\,c_2(r_2) + e_2(r_2) = e_1 + r_2\ell_2(r_2) \;\text{ for each } r_2\\
& (1-\pi)(1-\delta)c_3(r_2) = R\,[1-\pi-\ell_1-\ell_2(r_2)] + e_2(r_2) \;\text{ for each } r_2\\
& e_1 \geq 0, \ell_1 \geq 0\\
& e_2(r_2) \geq 0, \ell_2(r_2) \geq 0 \;\text{ for each } r_2.
\end{aligned}$$

Let μ_1, $\mu_2(r_2)$, $\mu_3(r_2)$, ν_1, w_1, $\nu_2(r_2)$, and $w_2(r_2)$ be the corresponding multipliers, which are all nonnegative. For simplicity, we consider the case when r_2 is a discrete random variable; similar arguments apply when r_2 is a continuous random variable. The first-order necessary conditions for a solution are then:

$$u'(c_1) = \mu_1 \quad [c_1]$$

$$q_{r_2} u'\left(c_2(r_2)\right) = \mu_2(r_2) \quad [c_2(r_2)]$$

$$q_{r_2} u'(c_3(r_2)) = \mu_3(r_2) \quad [c_3(r_2)]$$

$$\sum_{r_2} \mu_2(r_2) + \nu_1 = \mu_1 \quad [e_1]$$

$$\mu_3(r_2) + \nu_2(r_2) = \mu_2(r_2) \quad [e_2(r_2)]$$

$$r_1\mu_1 + w_1 = R\sum_{r_2} \mu_3(r_2) \quad [\ell_1]$$

$$r_2\mu_2(r_2) + w_2(r_2) = R\mu_3(r_2) \quad [\ell_2(r_2)],$$

where q_{r_2} denotes the probability of r_2, that is, $q_{r_2} = Prob\{\tilde{r}_2 = r_2\}$. The complementarity slackness [CS] conditions are:

$$\nu_1 e_1 = 0, \quad \nu_2(r_2)e_2(r_2) = 0, \quad w_1\ell_1 = 0, \quad \text{and} \quad w_2(r_2)\ell_2(r_2) = 0.$$

We establish three lemmas that characterize various portfolio management choices in the solution and then use these lemmas to prove the proposition. The first lemma shows that it is never optimal for the fund to hold excess liquidity in period 2.

Lemma 1. *The solution to [P1] has $e_2(r_2) = 0$ for all r_2.*

Proof. To begin, note that for any r_2, setting both $e_2(r_2) > 0$ and $\ell_2(r_2) > 0$ cannot be optimal. In other words, it is never optimal for the fund to simultaneously hold excess liquidity and liquidate investment. To see this, note that if $e_2(r_2) > 0$ and $\ell_2(r_2) > 0$ both held, the [CS] conditions and first order conditions for $e_2(r_2)$ and $\ell_2(r_2)$ would imply both $\mu_2(r_2) = \mu_3(r_2)$ and $r_2\mu_2(r_2) = R\mu_3(r_2)$, which contradict each other.

The proof is by contradiction in two parts. First, suppose $e_2(r_2) > 0$ held for some but not all r_2. Then there would exist r_2' and r_2'' such that $e_2(r_2') > 0$ and $e_2(r_2'') = 0$. The fact that $e_2(r_2') > 0$ would imply $\ell_2(r_2') = 0$, which through the budget constraints would imply both $c_2(r_2') < c_2(r_2'')$ and $c_3(r_2') > c_3(r_2'')$. However, because $e_2(r_2') > 0$ implies $v_2(r_2') = 0$, the first-order condition for $e_2(r_2')$ would imply $\mu_2(r_2') = \mu_3(r_2')$ and, therefore, $c_2(r_2') = c_3(r_2')$. It would follow that $c_2(r_2'') > c_3(r_2'')$, which would require $\mu_2(r_2'') < \mu_2(r_3'')$, and, therefore, $\nu_2(r_2'') < 0$, which is a contradiction.

Second, suppose $e_2(r_2) > 0$ held for all r_2, which would imply $\ell_2(r_2) = 0$ and $\nu_2(r_2) = 0$

for all r_2. We would then have $\mu_2(r_2) = \mu_3(r_2)$ and $c_2(r_2) = c_3(r_2)$ for all r_2. The budget constraints would also imply $e_1 > 0$, and the [CS] conditions and the first order conditions together would imply both

$$\mu_1 = \sum_{r_2} \mu_2(r_2) = \sum_{r_2} \mu_3(r_2) \quad \text{and}$$
$$r_1\mu_1 + w_1 = R \sum_{r_2} \mu_3(r_2).$$

These two equations would imply $r_1\mu_1 + w_1 = R\mu_1$, which in turn would imply $w_1 > 0$ and, therefore, $\ell_1 = 0$. However, using $m_1 > \pi$, the period 1 budget constraint (with $\ell_1 = 0$) would then require $e_1 < 0$, which is a contradiction.
Together, these twp steps show that $e_2(r_2)$ must equal zero for all values of r_2. □

The next lemma shows that the fund will always hold excess liquidity in period 1.

Lemma 2. *The solution to [P1] has* $e_1 > 0$*.*

Proof. The proof is again by contradiction. Suppose the solution had $e_1 = 0$. Then it must have $\ell_2(r_2) > 0$ for all r_2. Otherwise, the remaining type 2 investors would have zero consumption, which cannot be optimal. Since the fund would be liquidating investment in period 2 for all r_2, it would not hold excess liquidity, that is, $e_2(r_2) = 0$ for all r_2.
The [CS] conditions and first order conditions for ℓ_1 and $\ell_2(r_2)$, would then imply

$$r_1\mu_1 + w_1 = R \sum_{r_2} \mu_3(r_2) = \sum_{r_2} r_2\mu_2(r_2). \tag{26}$$

or

$$\mu_1 = \sum_{r_2} \frac{r_2}{r_1}\mu_2(r_2) - \frac{1}{r_1}w_1.$$

Combining this equation with the first-order condition for e_1 to eliminate the variable μ_1 would then yield

$$\nu_1 + \sum_{r_2} \mu_2(r_2) = \sum_{r_2} \frac{r_2}{r_1}\mu_2(r_2) - \frac{1}{r_1}w_1$$

or, multiplying both sides by r_1 and rearranging terms,

$$r_1\nu_1 + w_1 = \sum_{r_2}(r_2 - r_1)\mu_2(r_2).$$

Note that, for any $r_2' > r_2''$, the first-order conditions and budget constraints would then imply $c_2(r_2') > c_2(r_2'')$; in other words, consumption must be higher in the state where investment is worth more. Therefore, we would have

$$r_1\nu_1 + w_1 = \sum_{r_2}(r_2 - r_1)u'(c_2(r_2))q_{r_2} < \sum_{r_2}(r_2 - r_1)u'(c_2(\underline{r}))q_{r_2} = (E[r_2] - r_1)u'(c_2(\underline{r})) \leq 0,$$

where the last inequality follows from $r_1 \geq E[r_2]$. The strict inequality in the middle of this line contradicts the fact that $r_1\nu_1$ and w_1 are both non-negative. □

Lemmas 1 and 2 show that the fund will hold some excess liquidity exiting period 1, but will use all of its liquid assets to make payments in period 2. The third lemma uses these conditions to show that the fund will not liquidate investment in period 2.

Lemma 3. *The solution to [P1] has $\ell_2(r_2) = 0$ for all r_2.*

Proof. To begin, note that Lemma 2 implies $\nu_1 = 0$ and, therefore, the first-order condition for e_1 becomes

$$\mu_1 = \sum_{r_2} \mu_2(r_2)$$

The proof of this lemma is again by contradiction in two steps. First, suppose the fund sets $\ell_2(r_2) > 0$ for all r_2. If $\ell_1 > 0$ also held, the first-order conditions for ℓ_1 and $\ell_2(r_2)$ would imply

$$r_1\mu_1 = R\sum_{r_2}\mu_3(r_2) = \sum_{r_2} r_2\mu_2(r_2).$$

Combining these equations would then yield

$$\sum_{r_2} r_1\mu_2(r_2) = \sum_{r_2} r_2\mu_2(r_2),$$

which is a contraction because $r_1 \geq r_2$ for all r_2 with strict inequality in some states. If, instead, $\ell_1 = 0$ held, then for any $r_2' > r_2''$, it would follow from the budget constraints that

$c_2(r_2') > c_2(r_2'')$. Therefore, we would have

$$w_1 = \sum_{r_2}(r_2 - r_1)\mu_2(r_2) < (E[r_2] - r_1)u'(c_2(\underline{r})) \leq 0,$$

which contradicts the fact that w_1 is non-negative.

Second, suppose $\ell_2(r_2) > 0$ held for some but not all r_2. Note that it is never optimal for the fund to choose $\ell_2(r_2') = 0$ and $\ell_2(r_2'') > 0$ for any $r_2' > r_2''$. Intuitively, liquidation is less attractive when r_2 is lower. (To verify this statement, note that since $e_2(r_2') = e_2(r_2'') = 0$, this pattern of liquidation would imply $c_2(r_2') < c_2(r_2'')$ and $c_3(r_2') > c_3(r_2'')$. It would then follow from the first-order conditions for $\ell_2(r_2')$ and $\ell_2(r_2'')$ that

$$w_2(r_2') = Rq_{r_2'}u'(c_3(r_2')) - r_2'q_{r_2'}u'(c_2(r_2')) < Rq_{r_2'}u'(c_3(r_2'')) - r_2'q_{r_2'}u'(c_2(r_2'')) = 0,$$

which contradicts the fact that $w_2(r_2')$ is non-negative.) Therefore, it would have to be the case that there exists $\hat{r}_2$ such that $\ell_2(r_2) > 0$ for any $r_2 > \hat{r}_2$ and $\ell_2(r_2) = 0$ for any $r_2 \leq \hat{r}_2$. The [CS] conditions and first-order conditions for $e_1 > 0$ and $\ell_2(\bar{r}) > 0$ would then imply $\mu_1 = \sum_{r_2} \mu_2(r_2)$ and $\bar{r}\mu_2(\bar{r}) = R\mu_3(\bar{r})$. For any $r_2 \leq \hat{r}_2$, since we would have $\ell_2(\bar{r}) > 0$ and $\ell_2(r_2) = 0$, the budget constraints would then imply $c_2(\bar{r}) > c_2(r_2)$ and $c_3(\bar{r}) < c_3(r_2)$. As a result, we would have

$$Ru'(c_3(r_2)) < Ru'(c_3(\bar{r})) = \bar{r}u'(c_2(\bar{r})) < \bar{r}u'(c_2(r_2)).$$

It would then follow from the first-order condition for $\ell_2(r_2)$ that

$$w_2(r_2) = Rq_{r_2}u'(c_3(r_2)) - r_2q_{r_2}u'(c_2(r_2)) < \bar{r}q_{r_2}u'(c_2(r_2)) - r_2q_{r_2}u'(c_2(r_2)) = (\bar{r} - r_2)\mu_2(r_2).$$

But the first-order conditions for ℓ_1 and $r_1 \geq \bar{r}$ would then imply

$$\begin{aligned}
w_1 = R\sum_{r_2}\mu_3(r_2) - r_1\mu_1 &= \sum_{r_2} r_2\mu_2(r_2) + \sum_{r_2=\underline{r}}^{\hat{r}_2} w_2(r_2) - r_1\sum_{r_2}\mu_2(r_2) \\
&< \sum_{r_2}(r_2 - r_1)\mu_2(r_2) + \sum_{r_2=\underline{r}}^{\hat{r}_2}(\bar{r} - r_2)\mu_2(r_2) \\
&< \sum_{r_2=\underline{r}}^{\hat{r}_2}(\bar{r} - r_1)\mu_2(r_2) \\
&\leq 0,
\end{aligned}$$

where the last inequality follows from $\bar{r} \leq r_1$. The strict inequality here contradicts the fact that w_1 is non-negative. □

With these lemmas in hand, we are ready to prove the proposition.

Proof of Proposition 2. Lemmas 1-3 establish that the solution has $e_1 > 0$ and $e_2(r_2) = \ell_2(r_2) = 0$ for any r_2. In other words, the fund will hold excess liquidity exiting period 1, it will use all of its liquid assets in period 2, and it will not liquidate investment in period 2. The only remaining question is whether it will liquidate investment in period 1, that is, whether $\ell_1 = 0$ or $\ell_1 > 0$. The first-order condition for an interior choice of ℓ_1 combined with the budget constraints and the results of Lemmas 1-3 can be solved for

$$c_1^l = c_2^l(r_2) = r_1(1-\pi) + \pi \quad \text{and} \quad c_3^l(r_2) = R(1-\pi) + \frac{R}{r_1}\pi \quad \text{for any } r_2. \tag{27}$$

This choice will solve the problem if and only if sets $\ell_1 \geq 0$, which requires

$$c_1^l \geq \frac{\pi}{\pi + \delta(1-\pi)}. \tag{28}$$

Otherwise, the solution will set $\ell_1 = 0$, which generates

$$c_1^n = c_2^n(r_2) = \frac{\pi}{\pi + \delta(1-\pi)} \quad \text{and} \quad c_3^n(r_2) = \frac{R}{1-\delta} \quad \text{for any } r_2. \tag{29}$$

Together, equations (27) - (28) establish the result. □

A.2 Run detected in period 2

Next, we prove Proposition 3′, which determines the time-consistent payments when a run is detected in period 2. Proposition 3, as stated in in Section 2.5, assumes the redemption fee in period 1 is small enough that the fund would not choose to hold excess liquidity in period 2. The more general result we prove here does not impose this assumption, and Proposition 3 follows from it as a special case. We focus on this special case in the main text to simplify the exposition, but we use the generalized result in Appendix C to provide a more complete characterization of the best general run-proof policy.

Proposition 3′. *When $m_1 \leq \pi$ and $m_1 + m_2 > \pi$, condition (TC2) requires the fund to set*

$$c_2(m_1,m_2,r_2) = \min\Bigg\{\max\Bigg\{\underbrace{\frac{\pi - m_1c_1(m_1)}{m_2}}_{\textit{no liquidation}}, \underbrace{\frac{r_2(1-\pi)+\pi-m_1c_1(m_1)}{1-m_1}}_{\textit{liquidation at } t=2}\Bigg\}, \underbrace{\frac{R(1-\pi)+\pi-m_1c_1(m_1)}{1-m_1}}_{\textit{excess liquidity at } t=2}\Bigg\}$$

$$c_3(m_1,m_2,r_2) = \max\Bigg\{\min\Bigg\{\underbrace{\frac{R(1-\pi)}{1-m_1-m_2}}_{\textit{no liquidation}}, \underbrace{\frac{R(1-\pi)+\frac{R}{r_2}[\pi-m_1c_1(m_1)]}{1-m_1}}_{\textit{liquidation at } t=2}\Bigg\}, \underbrace{\frac{R(1-\pi)+\pi-m_1c_1(m_1)}{1-m_1}}_{\textit{excess liquidity at } t=2}\Bigg\}.$$

Proof. Under condition (TC2), when $m_1 \leq \pi$ and $m_1 + m_2 > \pi$, the fund's problem in period 2 given $c_1(m_1)$ and r_2 is

$$\begin{aligned} \max \quad & m_2u(c_2) + (1-m_1-m_2)u(c_3) \\ s.t. \quad & m_2c_2 + e_2 = \pi - m_1c_1(m_1) + \ell_2 r_2 \\ (1-m_1-m_2)c_3 &= R(1-\pi-\ell_2) + e_2 \\ e_2 &\geq 0, \ell_2 \geq 0. \end{aligned}$$

Letting μ_1, μ_2, ν_1, and w_1 be the corresponding multipliers, the first order conditions are

$$u'(c_2) = \mu_1 \qquad [c_2]$$

$$u'(c_3) = \mu_2 \qquad [c_3]$$

$$\mu_2 + \nu_1 = \mu_1 \qquad [e_2]$$

$$r_1\mu_1 + w_1 = R\mu_2 \qquad [\ell_2]$$

The complementarity slackness [CS] conditions are $\nu_1 e_2 = 0$ and $w_1\ell_2 = 0$. The same argument used in the proof of Lemma 1 shows that having both $e_2 > 0$ and $\ell_2 > 0$ cannot be optimal. That leaves three possilbe configurations for the solution: (i) $e_2 = \ell_2 = 0$, (ii) $e_2 > 0$ and $\ell_2 = 0$, and (iii) $e_2 = 0$ and $\ell_2 > 0$. We consider each of these cases in turn.

(i) If the solution has $e_2 = \ell_2 = 0$ (no excess liquidity and no liquidation), it is given by

$$c_2^n(m_1,m_2) = \frac{\pi - m_1}{m_2}, c_3^n(m_1,m_2) = \frac{R(1-\pi)}{1-m_1-m_2}.$$

(ii) If the solution has $e_2 > 0$ and $\ell_2 = 0$ (excess liquidity and no liquidation), the first order condition and [CS] condition for e_2 imply$\mu_1 = \mu_2$ and, therefore, $c_2 = c_3$. The budget constraints then imply

$$c_2^e(m_1, m_2, r_2) = c_3^e(m_1, m_2, r_2) = \frac{R(1-\pi) + \pi - m_1}{1 - m_1}.$$

Note that $e_2 = \pi - m_1 c_1(m_1) - m_2 c_2^e > 0$ if and only if $c_2^e(m_1, m_2, r_2) > \frac{\pi - m_1 c_1(m_1)}{m_2} = c_2^n(m_1, m_2)$.

(iii) If the solutions has $e_2 = 0$ and $\ell_2 > 0$ (liquidation and no excess liquidity), the first order condition and [CS] condition for ℓ_2 imply

$$r_2 u'(c_2) = R u'(c_3) \Rightarrow c_3 = \frac{R}{r_2} c_2.$$

Substitutiing this relationship into the budget constraints yields

$$c_2^l(m_1, m_2, r_2) = \frac{r_2(1-\pi) + \pi - m_1}{1 - m_1}, \quad c_3^l(m_1, m_2, r_2) = \frac{R(1-\pi) + \frac{R}{r_2}(\pi - m_1)}{1 - m_1}.$$

Note that $\ell_2 = \frac{1}{r}[m_2 c_2^l(m_1, m_2; r_2) - \pi + m_1] > 0$ if and only if $c_2^l(m_1, m_2; r_2) > \frac{\pi - m_1}{m_2} = c_2^n(m_1, m_2)$. Therefore, the optimal solution to the fund's problem at $t = 2$ when $m_1 \leq \pi$ and $m_1 + m_2 > \pi$ is the following:

$$c_2(m_1, m_2, r_2) = \min\{\max\{c_2^n(m_1, m_2), c_2^l(m_1, m_2, r_2)\}, c_2^e(m_1, m_2, r_2)\}$$
$$c_3(m_1, m_2, r_2) = \max\{\min\{c_3^n(m_1, m_2), c_3^l(m_1, m_2, r_2)\}, c_3^e(m_1, m_2, r_2)\},$$

which completes the proof. □

To see that Proposition 3 follows as a special case of this result, first note that for any m_1, m_2 and r_2, we have $c_2^l(m_1, m_2, r_2) < c_2^e(m_1, m_2, r_2)$. If $c_1(m_1) \geq \frac{1}{m_1}\left[\pi - \frac{R(1-\pi)m_2}{1-m_1-m_2}\right]$, then $c_2^n(m_1, m_2) \leq c_2^e(m_1, m_2, r_2)$, also holds for any m_1, m_2 and r_2. In other words, as long as $c_1(m_1)$ is not too small, c_2 will necessarily equal the larger of c_2^n and c_2^ℓ. The budget constraints then imply that c_3 will necessarily equal the smaller of c_3^n and c_3^ℓ, and we have

$$c_2(m_1, m_2, r_2) = \max\{c_2^n(m_1, m_2), c_2^l(m_1, m_2, r_2)\}$$
$$c_3(m_1, m_2, r_2) = \min\{c_3^n(m_1, m_2), c_3^l(m_1, m_2, r_2)\},$$

as stated in Proposition 3.

B Additional examples of preemptive runs

In this appendix, we expand on the example presented in Section 3.2 showing that a preemptive run equilibrium can exist when the fund aims to implement the first-best allocation. We present additional examples that show this equilibrium exists for a wide range of parameter values and that provide insight into how the fragility of the fund depends on parameters.

The first set of examples demonstrates that a preemptive run equilibrium can exist for distributions of π_1 besides the uniform distribution used in Figure 1.

Example 1: Let $\pi = 0.5$, $R = 1.04^{\frac{1}{12}}$, $r_1 = \bar{r} = 0.98$, $\underline{r} = 0.8$, and $q = 0.5$. We consider two probability distributions for $\pi_1 \in [0, \pi]$:

(a) A truncated normal distribution with mean $\mu = 0$ and standard deviation $\sigma = 0.3$;

(b) A truncated normal distribution with mean $\mu = 0.5$ and standard deviation $\sigma = 0.3$;

The density function of distribution (a) is strictly decreasing on $[0, \pi]$, while that of distribution (b) is strictly increasing. In this sense, the two distributions represent opposite departures from the uniform distribution used for the example in Figure 1. Figure 10 shows the expected values of redeeming in period 1 (in red) and waiting (in blue) for a non-type 1 investor when all other attentive investors redeem. We refer to the range of values for δ for which the run equilibrium exists as the *fragile set*. The figure shows that the fragile set is significant in both cases.

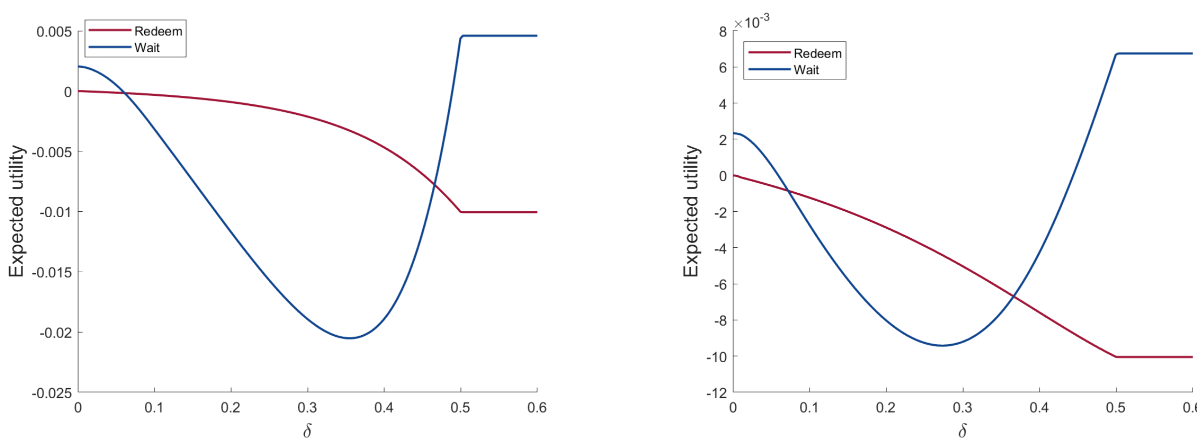

(a) Truncated normal with $\mu = 0$ and $\sigma = 0.3$ (b) Truncated normal with $\mu = 0.5$ and $\sigma = 0.3$

Figure 10: Fragility region of δ with different distribution of π_1

Note that distribution (b) assigns higher probability to larger values of π_1, which implies that a run is more likely to be detected in the first period. Comparing the two panels of the figure shows that the fragile set becomes smaller in this case. If we increase the mean of the distribution further, to $\mu = 0.9$, the blue line will lie everywhere above the red line and a run equilibrium will not exist for any value of δ (not pictured here). These patterns illustrate that a run equilibrium is more likely to exist when fundamental redemption demand in period 1 is smaller. As discussed in the main text, the incentive for an investor to run in this model comes from the possibility that the run will not be detected until period 2, which gives the investor an opportunity to withdraw before a redemption fee is imposed. If π_1 is likely to be very close to π, a run is likely to be detected in period 1, in which case investors are unable to redeem before the fee is applied and have no incentive to join a run.

The second set of examples explores how the fragile set changes as we vary other parameter values, holding the distribution of π_1 fixed.

Example 2: Let $r_1 = \bar{r} = 0.9$ and assume π_1 follows a truncated normal distribution with mean $\mu = 0.2$ and standard deviation $\sigma = 0.3$. We consider the following four cases:

(a) $\pi = 0.5, R = 1.04^{\frac{1}{12}}, \underline{r} = 0.8, q = 0.5;$

(b) $\pi = 0.5, R = 1.04^{\frac{1}{12}}, \underline{r} = 0.8, q = 0.2;$

(c) $\pi = 0.6, R = 1.04^{\frac{1}{12}}, \underline{r} = 0.8, q = 0.5;$

(d) $\pi = 0.5, R = 1.02^{\frac{1}{12}}, \underline{r} = 0.8, q = 0.5;$

Figure 11 presents the fragility diagram for each of these four cases. Take panel (a) as a baseline. Notice that, relative to the example in Figure 10, we have decreased the liquidation value r_1 from 0.98 to 0.9. This lower liquidation value leads to higher time-consistent redemption fees, which cause the fragile set to shrink. This result is counterintuitive: in standard models, a bank/fund becomes more fragile when the liquidation value of its assets decreases because these assets will be depleted more quickly during a run. Recall, however, that investors in our model do not run because they fear the fund will exhaust its resources in period 1; the time-consistent redemption fee policy prevents that outcome. Instead, investors run if they think the redemption fee in period 1 is will be too low on average and the redemption fee in period 2 will likely be larger. A decrease in the current liquidation value of assets r_1, holding the distribution of r_2 fixed, makes this outcome less likely and, therefore, reduces the incentive to run.

Moving from panel (a) to panel (b) in Figure 11, there is more downside risk in the future liquidation value, that is, the lower value $\underline{r}$ is more likely to occur. This change cases the fragile set to expand. In other words, a preemptive run equilibrium is more likely to exist when the market liquidity conditions are more likely to deteriorate. This result is intuitive because, as emphasized above, the incentive to run in our model arises from the possibility that the redemption fee will be larger in the future. A lower realization of r_2 would make the time-consistent redemption fee in period 2 larger, increasing this concern.

Moving from panel (a) to panel (c), the total fraction of impatient investors increases. Holding the distribution of π_1 fixed, this change implies that a run is less likely to be detected in the first period. As a result, the fragile set increases in size, as shown in the figure. Finally, moving from panel (a) to panel (d), the return on matured investment decreases. The figure shows that this change also increases the size of the fragile set (slightly), at least in this example, because the benefit of waiting to withdraw is smaller.

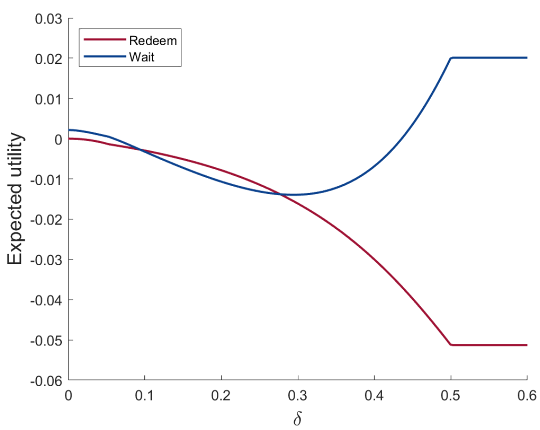

(a) $\pi = 0.5, R = 1.04^{\frac{1}{12}}, \underline{r} = 0.8, q = 0.5$

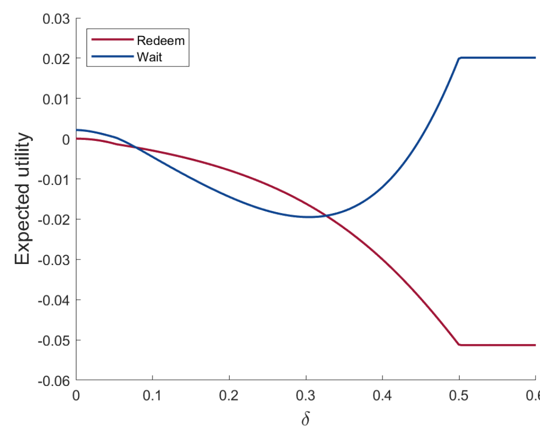

(b) $\pi = 0.5, R = 1.04^{\frac{1}{12}}, \underline{r} = 0.8, q = 0.2$

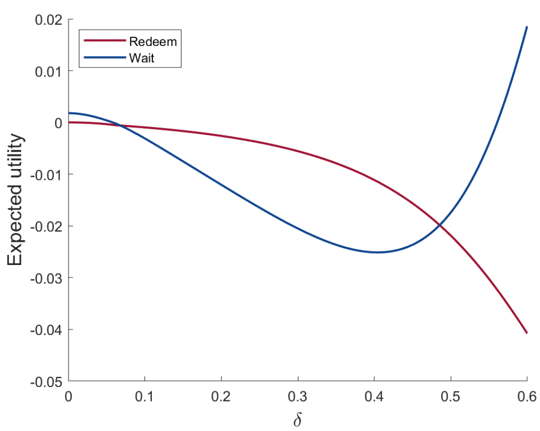

(c) $\pi = 0.6, R = 1.04^{\frac{1}{12}}, \underline{r} = 0.8, q = 0.5$

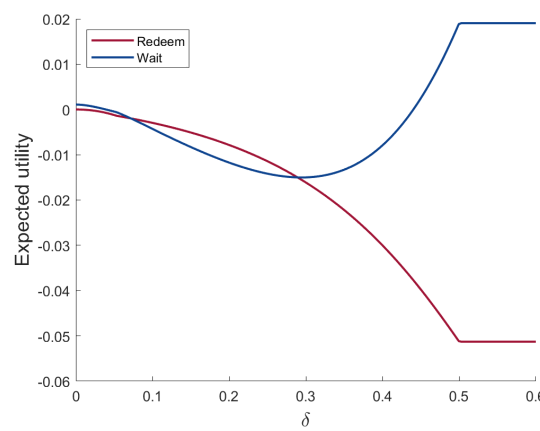

(d) $\pi = 0.5, R = 1.02^{\frac{1}{12}}, \underline{r} = 0.8, q = 0.5$

Figure 11: Fragility region of δ with different parameters

C Properties of the best general run-proof policy

In this appendix, we provide a detailed characterization of the best general run-proof policy studied in Section 4.1. As part of the analysis, we prove several results that are referred to in the discussion in the main text.

C.1 Setup and first-order conditions

Let $\mathcal{C}$ denote the set of all continuous payment functions $c_1(m_1) \leq 1$ when $m_1 \in [\delta, \pi]$. For any $c_1(m_1) \in \mathcal{C}$, by change of variables ($m_1 = \pi_1 + \delta(1 - \pi_1)$), we can rewrite the run-proof constraint (15) as

$$\mathcal{R}\big(c_1(m_1)\big) = \int_{\delta}^{\pi} \left\{ \begin{array}{c} \ln\big(c_1(m_1)\big) - p_{\frac{m_1-\delta}{1-\delta}} \mathbb{E}_{r_2}\left[\ln\big(c_2(m_1, m_2, r_2)\big)\right] \\ -\left(1 - p_{\frac{m_1-\delta}{1-\delta}}\right) \mathbb{E}_{r_2}\left[\ln\big(c_3(m_1, m_2, r_2)\big)\right] \end{array} \right\} f_n\left(\frac{m_1 - \delta}{1 - \delta}\right) dm_1 - (1 - \delta)T \leq 0,$$

where $m_2 = \pi + \delta(1 - \pi) - m_1$, the payments c_2 and c_3 are given by Proposition $3'$, and T is defined in equation (14). For any contract $c_1(m_1) \in \mathcal{C}$, welfare in the no-run equilibrium (16) is given by

$$\mathcal{W}\big(c_1(m_1)\big) = \int_{\delta}^{\pi} \left[m_1 \ln\big(c_1(m_1)\big) + (\pi - m_1)\ln\big(c_2^N(m_1)\big) + (1 - \pi)\ln\big(c_3^N(m_1)\big)\right] f(m_1)dm_1 + C,$$

where $C = \int_0^{\delta}(1-\pi)\ln(R)f(m_1)dm_1$ and the payments c_2^N and c_3^N are as defined in equations (17) - (18) with $\pi_1 = m_1$. The best general run-proof policy then solves the following problem:

$$\max_{c_1(m_1) \in \mathcal{C}} \quad \mathcal{W}\big(c_1(m_1)\big) \qquad [P^G]$$
$$\text{s.t.} \quad \mathcal{R}\big(c_1(m_1)\big) \leq 0.$$

As shown in Proposition $3'$, there are three different regions for the payments c_2 and c_3 in the run-proof constraint, depending on the choice of $c_1(m_1)$, m_1 and the realization of r_2. For simplicity, we focus on the case where c_2 and c_3 are in the liquidation region, with

$$c_2(m_1, m_2, r_2) = \frac{r_2(1 - \pi) + \pi - m_1 c_1(m_1)}{1 - m_1}, \quad c_3(m_1, m_2, r_2) = \frac{R}{r_2} c_2(m_1, m_2, r_2).$$

All of the examples we study satisfy this condition under the optimal schedule $c_1^*(m_1)$. The analysis can be extended to the case where the payments c_2 and c_3 in the run-proof constraint lie in the other regions identified in Proposition 3′ at the cost of additional notation and complexity.

For the payments c_2^N and c_3^N in the objective function, we retain the full expressions defined in equations (17) and (18), which include the max and min operators, for two reasons. First, in some examples we study, c_2^N and c_3^N move between regions depending on the value of m_1. Second, maintaining the max and min operators in the objective function illustrates the complication involved in solving problem $[P^G]$.

Let $W\big(m_1, c_1(m_1)\big)$ denote the integrand of the welfare measure $\mathcal{W}\big(c_1(m_1)\big)$ If

$$c_1(m_1) \geq \bar{c}(m_1) = R - (R-1)\frac{\pi}{m_1} \Leftrightarrow \frac{\pi - m_1 c_1(m_1)}{\pi - m_1} \leq \frac{R(1-\pi) + \pi - m_1 c_1(m_1)}{1 - m_1}$$

then we can write this integrand as

$$W(m_1, c_1) = \left[m_1 \ln(c_1) + (\pi - m_1) \ln\left(\frac{\pi - m_1 c_1}{\pi - m_1}\right) + (1-\pi)\ln(R) \right] f(m_1).$$

Taking its partial derivative with respect to c_1 yields

$$W_{c_1}(m_1, c_1) = m_1 \left(\frac{1}{c_1} - \frac{\pi - m_1}{\pi - m_1 c_1}\right) f(m_1) = m_1 \frac{1 - c_1}{c_1} \frac{\pi}{\pi - m_1 c_1} f(m_1).$$

If, instead, we have $c_1(m_1) < \bar{c}(m_1)$, then the integrand is

$$W(m_1, c_1) = \left[m_1 \ln(c_1) + (1 - m_1) \ln\left(\frac{R(1-\pi) + \pi - m_1 c_1}{1 - m_1}\right) \right] f(m_1),$$

and its partial derivative with respect to c_1 is

$$W_{c_1}(m_1, c_1) = m_1 \left(\frac{1}{c_1} - \frac{1 - m_1}{A_R - m_1 c_1}\right) f(m_1) = m_1 \frac{A_R - c_1}{c_1(A_R - m_1 c_1)} f(m_1),$$

where $A_R = R(1-\pi) + \pi$. Note that $W_{c_1^-}(m_1, \bar{c}) = W_{c_1^+}(m_1, \bar{c})$ but $W_{c_1 c_1^-}(m_1, \bar{c}) \neq W_{c_1 c_1^+}(m_1, \bar{c})$. Therefore, $W(m_1, c_1)$ is a non-smooth function of c_1 given any m_1.

Similarly, let $R\big(m_1, c_1(m_1)\big)$ denote the integrand of the incentive to run $\mathcal{R}\big(c_1(m_1)\big)$:

$$R(m_1, c_1) = \left\{ \begin{array}{c} \ln(c_1) - p_{\frac{m_1-\delta}{1-\delta}} \mathbb{E}_{r_2}\left[\ln\left(\frac{r_2(1-\pi)+\pi-m_1c_1}{1-m_1}\right)\right] \\ -\left(1 - p_{\frac{m_1-\delta}{1-\delta}}\right) \mathbb{E}_{r_2}\left[\ln\left(\frac{R}{r_2}\frac{r_2(1-\pi)+\pi-m_1c_1}{1-m_1}\right)\right] \end{array} \right\} f_n\left(\frac{m_1-\delta}{1-\delta}\right).$$

Taking its partial derivative with respect to c_1 yields

$$R_{c_1}(m_1, c_1) = \left(\frac{1}{c_1} + \frac{m_1}{1-m_1}\mathbb{E}_{r_2}\left[\frac{1-m_1}{A_{r_2} - m_1 c_1}\right]\right) f_n\left(\frac{m_1-\delta}{1-\delta}\right)$$
$$= \mathbb{E}_{r_2}\left[\frac{A_{r_2}}{c_1(A_{r_2} - m_1 c_1)}\right] f_n\left(\frac{m_1-\delta}{1-\delta}\right),$$

where $A_{r_2} = r_2(1-\pi) + \pi$. Following equation (20), define

$$L(m_1, c_1) \equiv \frac{W_{c_1}(m_1, c_1)}{R_{c_1}(m_1, c_1)} = \frac{F(m_1, c_1)}{G(m_1, c_1)} H(m_1), \tag{30}$$

where $\frac{F}{G}$ is the marginal benefit-cost ratio and is given by

$$\frac{F(m_1, c_1)}{G(m_1, c_1)} = \begin{cases} \left(\mathbb{E}_{r_2}\left[\frac{A_{r_2}(\pi - m_1 c_1)}{\pi(1-c_1)(A_{r_2} - m_1 c_1)}\right]\right)^{-1} & \text{if } c_1 \geq \bar{c}(m_1) \\ \left(\mathbb{E}_{r_2}\left[\frac{A_{r_2}(A_R - m_1 c_1)}{(A_R - c_1)(A_{r_2} - m_1 c_1)}\right]\right)^{-1} & \text{if } c_1 < \bar{c}(m_1), \end{cases}$$

and

$$H(m_1) = \frac{m_1 f(m_1)}{f_n\left(\frac{m_1-\delta}{1-\delta}\right)} \tag{31}$$

is the weighted likelihood ratio. As in the main text, the best policy $c_1^*(m_1)$ is implicitly defined by the first-order condition

$$L(m_1, c_1) = \lambda \quad \text{for all } m_1 \in [\delta, \pi].$$

The properties of the best policy are thus determined by the properties of the L function, which we study in the next section.

C.2 Properties of the L function

In this section, we prove two lemmas that establish properties of the L function defined in equation (30). The first result shows that L is decreasing in c_1, while the second result provides a sufficient condition for L to be increasing in m_1.

Lemma 4. *$L(m_1, c_1)$ is decreasing in $c_1 \in [0, 1]$ for any $m_1 \in [0, \pi]$.*

Proof. First note that the weighted likelihood ratio $H(m_1)$ does not depend on c_1, so the we only need to show that the the marginal benefit-cost ratio $\frac{F}{G}$ is decreasing in c_1. Choose any

$m_1 \in [0, \pi]$. First, we show that $\frac{F(m_1,c_1)}{G(m_1,c_1)}$ is decreasing in $c_1 \in [\bar{c}(m_1), 1]$, which is equivalent to showing that $\frac{\partial}{\partial c_1}\left(\frac{(1-c_1)(A_{r_2}-m_1c_1)}{\pi - m_1c_1}\right) \le 0$ for any r_2. Note that

$$\frac{\partial}{\partial c_1}\left(\frac{(1-c_1)(A_{r_2}-m_1c_1)}{\pi - m_1c_1}\right) = \frac{1}{(\pi - m_1c_1)^2}[(A_{r_2}-\pi)(m_1-\pi) - (m_1c_1-\pi)^2] \le 0.$$

The inequality follows from the facts that $A_{r_2} > \pi$, $m_1 \le \pi$ and $c_1 \le 1$.
Next, we show that $\frac{F(m_1,c_1)}{G(m_1,c_1)}$ is strictly decreasing in $c_1 \in \big[0, \bar{c}(m_1)\big)$, which is equivalent to showing that $\frac{\partial}{\partial c_1}\left(\frac{(A_R-c_1)(A_{r_2}-m_1c_1)}{A_R - m_1c_1}\right) < 0$ for any r_2. Note that

$$\frac{\partial}{\partial c_1}\left(\frac{(A_R-c_1)(A_{r_2}-m_1c_1)}{A_R - m_1c_1}\right) = -\frac{A_R(1-m_1)(A_{r_2}-m_1c_1) + m_1(A_R-c_1)(A_R-m_1c_1)}{(A_R-m_1c_1)^2} < 0,$$

where the strict inequality follows from $m_1 < 1$, $R > 1$ and $m_1c_1 < A_{r_2} < A_R$. Since $\frac{F(m_1,c_1)}{G(m_1,c_1)}$ is continuous at $c_1 = \bar{c}(m_1)$, it is decreasing in $c_1 \in [0,1]$. □

Lemma 5. *If $H'(m_1) \ge 0$ and $\frac{H'(m_1)}{H(m_1)} \ge \frac{1}{\underline{r}(1-\pi)+\pi-m_1} - \frac{1}{R(1-\pi)+\pi-m_1}$, then $L_{m_1}(m_1,c_1) \ge 0$ for any $m_1 \in [0,\pi]$ and $c_1 \in [0,1]$.*

Proof. For any (m_1, c_1) such that $c_1 \ge \bar{c}(m_1)$, we have

$$L(m_1,c_1) = \left(\mathbb{E}_{r_2}\left[\frac{A_{r_2}(\pi - m_1c_1)}{\pi(1-c_1)(A_{r_2}-m_1c_1)}\right]\right)^{-1} H(m_1).$$

Note that $\frac{\pi - m_1c_1}{A_{r_2}-m_1c_1}$ is strictly decreasing in m_1. Therefore, it follows from $H'(m_1) \ge 0$ that $L_{c_1}(m_1,c_1) \ge 0$. For any (m_1,c_1) such that $c_1 < \bar{c}(m_1)$, we have

$$L(m_1,c_1) = \left(\mathbb{E}_{r_2}\left[\frac{A_{r_2}(A_R - m_1c_1)}{(A_R-c_1)(A_{r_2}-m_1c_1)}\right]\right)^{-1} H(m_1).$$

Let $g(m_1,c_1;r_2) = \frac{A_R-m_1c_1}{A_{r_2}-m_1c_1}$, which is a strictly increasing function in m_1 since $A_{r_2} < A_R$. Note that

$$\left|\frac{g_{m_1}(m_1,c_1;r_2)}{g(m_1,c_1;r_2)}\right| = \frac{c_1(A_R - A_{r_2})}{(A_R-m_1c_1)(A_{r_2}-m_1c_1)} \le \frac{A_R - A_{r_2}}{(A_R-m_1)(A_{r_2}-m_1)} \le \frac{1}{A_{\underline{r}}-m_1} - \frac{1}{A_R - m_1}.$$

Here, the first inequality follows from $c \le 1$, and the second inequality follows from $r_2 \ge \underline{r}$. Therefore, we have

$$\left|\frac{H'(m_1)}{H(m_1)}\right| \ge \frac{1}{A_{\underline{r}}-m_1} - \frac{1}{A_R-m_1} \ge \left|\frac{g_{m_1}(m_1,c_1;r_2)}{g(m_1,c_1;r_2)}\right|,$$

which implies that $\frac{g(m_1,c_1;r_2)}{H(m_1)}$ is a decreasing function for any r_2. As a result, we have $L_{m_1}(m_1, c_1) \geq 0$. □

Lemma 5 shows that the function L is increasing in m_1 if the weighted likelihood ratio is increasing at a relatively high rate. It is straightforward to see that if $\pi_1 \sim U[0, \pi]$, the weighted likelihood ratio $H(m_1)$ is increasing and satisfies the relative growth rate condition in the lemma. In addition, the conditions are satisfied by a variety of probability distributions for π_1. For example, consider a truncated normal distribution with mean $\mu = 0$ and standard deviation $\sigma = 0.2$, while keeping other parameters the same as the example in Section 3. As shown in Figure 12, the weighted likelihood ratio is increasing and satisfies the relative growth rate condition for all $m_1 \in [\delta, \pi]$.

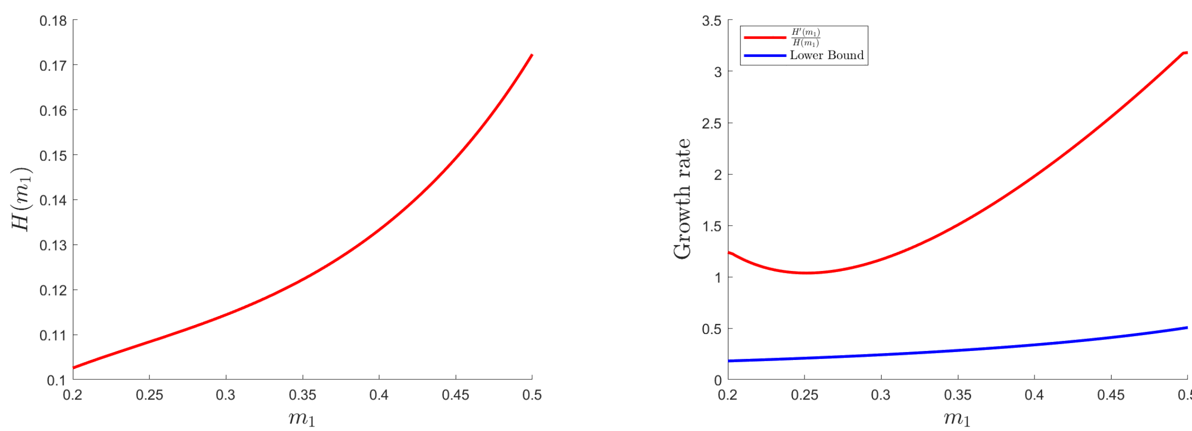

(a) Weighted likelihood ratio $H(m_1)$ (b) Growth rate vs Lower bound

Figure 12: Sufficient conditions on weighted likelihood ratio in Lemma 5

C.3 When is the fee decreasing in redemption demand?

The next proposition relates the monotonicity of the best general run-proof policy $c_1^*(m_1)$ to the function L. In particular, it shows that $c_1^*(m_1)$ is an increasing function over some interval if and only if $L(m_1, c_1)$ is increasing in m_1 on that interval.

Proposition 7. *The best general run-proof policy $c_1^*(m_1)$ is increasing on $\mathcal{I} \subseteq [\delta, \pi]$ if and only if $L_{m_1}(m_1, c_1^*(m_1)) \geq 0$ for $m_1 \in \mathcal{I}$.*

Proof. We first prove the "if" part by contradiction. Pick any $\mathcal{I} \subseteq [\delta, \pi]$. Suppose that $L_{m_1}(m_1, c_1^*(m_1)) \geq 0$ for $m_1 \in \mathcal{I}$. Suppose that there exists an interval $(a, b) \subseteq \mathcal{I}$ such that $c_1^*(m_1)$ is a strictly decreasing function. Without loss of generality, pick $x_1, x_2 \in (a, b)$

such that $x_1 < x_2$, $c_1^*(x_1) \geq \bar{c}(x_1)$ and $c_1^*(x_2) \geq \bar{c}(x_2)$. We want to construct a $c_1'(m_1) \in \mathcal{C}$ that is run-proof but attains higher welfare in the no-run equilibrium than $c_1^*(m_1)$. Pick any $\Delta \in (0, \frac{x_2 - x_1}{2})$. Consider the following $c_1'(m_1)$:

$$c_1'(m_1) = \begin{cases} c_1^*(m_1) - \alpha \left[\frac{\eta_1(m_1)}{R_{c_1}\big(m_1, c_1^*(m_1)\big)} + K\eta_1(m_1) \right] & \text{if } m_1 \in [x_1, x_1 + \Delta] \\ c_1^*(m_1) + \alpha \frac{\eta_2(m_1)}{R_{c_1}\big(m_1, c_1^*(m_1)\big)} & \text{if } m_1 \in [x_2 - \Delta, x_2] \\ c_1^*(m_1) & \text{Otherwise,} \end{cases}$$

where $\alpha > 0$ and $K > 0$ can be made arbitrarily small, and

$$\eta_1(m_1) = -\left(m_1 - \frac{2x_1 + \Delta}{2}\right)^2 + \frac{\Delta^2}{4};$$
$$\eta_2(m_1) = -\left(m_1 - \frac{2x_2 - \Delta}{2}\right)^2 + \frac{\Delta^2}{4}.$$

First, note that $\eta_1(x_1) = \eta_1(x_1 + \Delta) = 0$ and $\eta_2(x_2 - \Delta) = \eta_2(x_2) = 0$. Therefore, $c_1'(m_1)$ is a continuous function, i.e., $c_1'(m_1) \in \mathcal{C}$. Next, it is straightforward to check that $\eta_1(m_1) = \eta_2(m_1')$ whenever $|m_1 - x_1| = |m_1' - x_2 + \Delta|$ for any $m_1 \in [x_1, x_1 + \Delta]$ and $m_1' \in [x_2 - \Delta, x_2]$. Furthermore, define

$$M_1 = \int_{x_1}^{x_1+\Delta} R_{c_1}\big(m_1, c_1^*(m_1)\big)\eta_1(m_1)dm_1;$$
$$M_2 = \int_{x_2-\Delta}^{x_2} W_{c_1}\big(m_1, c_1^*(m_1)\big)\eta_1(m_1)dm_1;$$
$$M_3 = \int_{x_2-\Delta}^{x_2} L\big(m_1, c_1^*(m_1)\big)\eta_2(m_1)dm_1 - \int_{x_1}^{x_1+\Delta} L\big(m_1, c_1^*(m_1)\big)\eta_1(m_1)dm_1.$$

Since $W_{c_1} > 0$ and $R_{c_1} > 0$, it follows from the definition of $\eta_1(m_1)$ that $M_1 > 0$ and $M_2 > 0$. Also, since $x_1 < x_2$ and $c_1^*(x_1) > c_1^*(x_2)$, it follows from the proof of Lemma 4 and $L_{m_1}(m_1, c_1^*(m_1)) \geq 0$ that $M_3 > 0$. We choose K such that $K < \frac{M_3}{M_1 + M_2}$.
To see if $c_1'(m_1)$ is run-proof, we first define the following two functionals:

$$J_{x_1}(c_1) = \int_{x_1}^{x_1+\Delta} R\big(m_1, c_1(m_1)\big)dm_1;$$
$$J_{x_2}(c_1) = \int_{x_2-\Delta}^{x_2} R\big(m_1, c_1(m_1)\big)dm_1.$$

By the definition of the first variation of a functional, or the Gateau derivative of a functional, we have

$$J_{x_1}\left(c_1^* - \alpha\left(\frac{1}{R_{c_1}} + K\right)\eta_1\right) = J_{x_1}(c_1^*) - \delta J_{x_1}|_{c_1^*}\left(\left(\frac{1}{R_{c_1}} + K\right)\eta_1\right)\alpha + o(\alpha),$$

where $\lim_{\alpha\to 0}\frac{o(\alpha)}{\alpha} = 0$, and

$$\begin{aligned}\delta J_{x_1}|_{c_1^*}\left(\left(\frac{1}{R_{c_1}} + K\right)\eta_1\right) &= \int_{x_1}^{x_1+\Delta} R_{c_1}\big(m_1, c_1^*(m_1)\big)\left[\frac{1}{R_{c_1}(m_1, c_1^*(m_1))} + K\right]\eta_1(m_1)dm_1 \\ &= \int_{x_1}^{x_1+\Delta}\eta_1(m_1)dm_1 + K\int_{x_1}^{x_1+\Delta} R_{c_1}\big(m_1, c^*(m_1)\big)\eta_1(m_1)dm_1.\end{aligned}$$

Therefore, we have

$$\begin{aligned}&J_{x_1}(c_1^*) - J_{x_1}\left(c_1^* - \alpha\left(\frac{1}{R_{c_1}} + K\right)\eta_1\right) \\ &= \alpha\left[\int_{x_1}^{x_1+\Delta}\eta_1(m_1)dm_1 + K\int_{x_1}^{x_1+\Delta} R_{c_1}\big(m_1, c_1^*(m_1)\big)\eta_1(m_1)dm_1\right] - o(\alpha).\end{aligned} \tag{32}$$

Similarly, we have

$$J_{x_2}(c_1^*) - J_{x_2}\left(c_1^* + \alpha\frac{\eta_2}{R_{c_1}}\right) = -\alpha\int_{x_2-\Delta}^{x_2}\eta_2(m_1)dm_1 - o(\alpha). \tag{33}$$

It follows from equations (32) and (33) that

$$\begin{aligned}&\int_\delta^\pi R\big(m_1, c_1^*(m_1)\big)dm_1 - \int_\delta^\pi R\big(m_1, c_1'(m_1)\big)dm_1 \\ &= J_{x_1}(c_1^*) - J_{x_1}\left(c_1^* - \alpha\left(\frac{1}{R_{c_1}} + K\right)\eta_1\right) + J_{x_2}(c_1^*) - J_{x_2}\left(c_1^* + \alpha\frac{\eta_2}{R_{c_1}}\right) \\ &= \alpha\left[\int_{x_1}^{x_1+\Delta}\eta_1(m_1)dm_1 - \int_{x_2-\Delta}^{x_2}\eta_2(m_1)dm_1\right] + \alpha K\int_{x_1}^{x_1+\Delta} R_{c_1}\big(m_1, c_1^*(m_1)\big)\eta_1(m_1)dm_1 - 2o(\alpha) \\ &= \alpha K\int_{x_1}^{x_1+\Delta} R_{c_1}\big(m_1, c_1^*(m_1)\big)\eta_1(m_1)dm_1 - 2o(\alpha),\end{aligned}$$

where the last equality follows from the way we construct $\eta_1(m_1)$ and $\eta_2(m_1)$. By the definition of $o(\alpha)$, we can find a small enough $\alpha > 0$ such that $|o(\alpha)| < \frac{1}{2}\alpha K M_1$. Therefore, we have

$$\int_\delta^\pi R\big(m_1, c_1^*(m_1)\big)dm_1 - \int_\delta^\pi R\big(m_1, c_1'(m_1)\big)dm_1 = \alpha K M_1 - 2o(\alpha)$$

$$\geq \alpha K M_1 - 2|o(\alpha)|$$
$$> 0.$$

Since $\int_\delta^\pi R\left(m_1, c_1^*(m_1)\right) dm_1 \leq (1-\delta)T$, it follows that $c_1'(m_1)$ is run-proof. Next, we show that $c_1'(m_1)$ attains a higher welfare in the no-run equilibrium than $c_1^*(m_1)$. Following the same analysis above yields

$$\begin{aligned}
&\int_\delta^\pi W\big(m_1, c_1^*(m_1)\big)dm_1 - \int_\delta^\pi W\big(m_1, c_1'(m_1)\big)dm_1 \\
&= \alpha \begin{bmatrix} \int_{x_1}^{x_1+\Delta} \frac{W_{c_1}\left(m_1, c_1^*(m_1)\right)}{R_{c_1}\left(m_1, c_1^*(m_1)\right)} \eta_1(m_1) dm_1 \\ - \int_{x_2-\Delta}^{x_2} \frac{W_{c_1}\left(m_1, c_1^*(m_1)\right)}{R_{c_1}\left(m_1, c_1^*(m_1)\right)} \eta_2(m_1) dm_1 \end{bmatrix} + \alpha K M_2 - 2o(\alpha) \\
&= \alpha \begin{bmatrix} \int_{x_1}^{x_1+\Delta} L\left(m_1, c_1^*(m_1)\right) \eta_1(m_1) dm_1 \\ - \int_{x_2-\Delta}^{x_2} L\left(m_1, c_1^*(m_1)\right) \eta_2(m_1) dm_1 \end{bmatrix} + \alpha K M_2 - 2o(\alpha) \\
&\leq -\alpha M_3 + \alpha K M_2 + 2|o(\alpha)| \\
&< \alpha(K M_2 - M_3) + \alpha K M_1 \\
&< 0,
\end{aligned}$$

where the last strict inequality follows from $K < \frac{M_3}{M_1+M_2}$. Therefore, there exists a general policy $c_1'(m_1) \in \mathcal{C}$ that is run-proof and achieves a higher welfare in the no-run equilibrium than $c_1^*(m_1)$, which contradicts the fact that $c_1^*(m_1)$ is the best general run-proof policy.
Lastly, we prove the "only if" part by contradiction as well. Pick any $\mathcal{I} \subseteq [\delta, \pi]$. Suppose that the best general run-proof policy $c_1^*(m_1)$ is increasing in $\mathcal{I}$. Suppose that there exists an interval $(a, b) \subseteq \mathcal{I}$ such that $L_{m_1}\left(m_1, c_1^*(m_1)\right) < 0$ for $m_1 \in (a, b)$. By following a similar argument as above, we can construct a general policy that is run-proof and offers a higher welfare in the no-run equilibrium than $c_1^*(m_1)$, leading to a contradiction. □

Together with Lemma 5, Proposition 7 provides a sufficient condition for the redemption fee to be a decreasing function of redemption demand over any sub-interval of $[\delta, \pi]$. However, this condition is only sufficient; many examples that do not satisfy the condition nevertheless have a strictly decreasing fee policy. The example in the next section is one such case.

C.4 The calibrated example

In Section 5, we calibrate the distribution of π_1 to match data on daily redemption demand from SEC (2023). (See Figure 6.) Our calibrated distribution is a truncated normal with upper bound $\pi = 0.1301$, mean $\mu = 0.0032$, and standard deviation $\sigma = 0.0427$. The other parameter values we use in the example are $R = 1.04^{\frac{1}{180}}$, $r_1 = \bar{r} = 1$, $\underline{r} = 0.8$, $q = 0$. First, given those parameter values, the best robust policy $\bar{c}_R^* = 0.978$, $\bar{m}_R^* = 0.084$ and $\delta_L = 0.039$. Panel (a) of Figure 13 plots the weighted likelihood ratio $H(m_1)$ in $[\bar{m}_R^*, \pi]$ for this example. The figure shows that the ratio is decreasing in this case, meaning it does not satisfy the sufficient condition identified in Lemma 5. Nevertheless, panel (b) shows that the best general run-proof policy $c_1^*(m_1)$ based on δ_L is strictly increasing on the interval $[\bar{m}_R^*, \pi]$. This example demonstrates that the sufficient condition is not tight; the increasing property holds in many other cases as well. We use the fact that this property holds for the calibrated distributon in our policy discussion in Section 5.

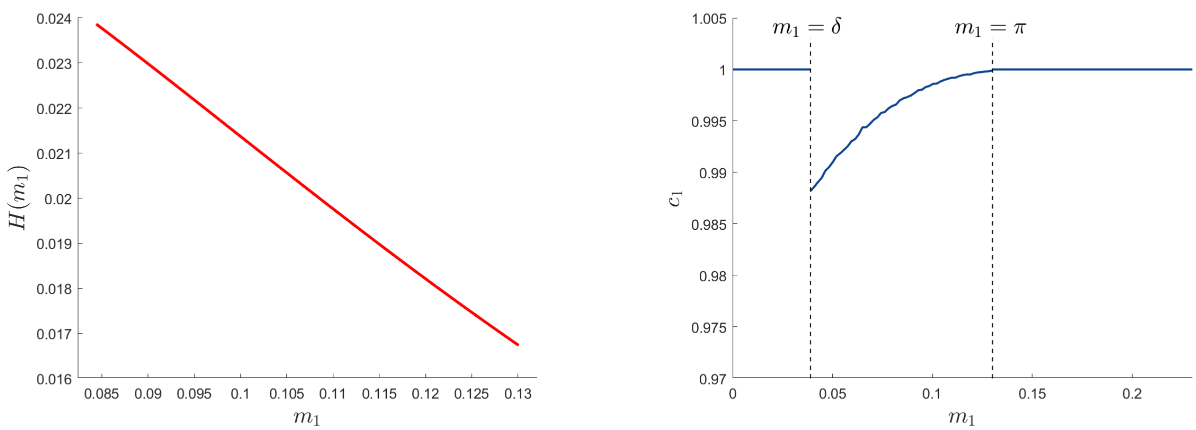

(a) Weighted likelihood ratio $H(m_1)$

(b) Best general run-proof policy

Figure 13: Weighted likelihood ratio and best general policy for the calibrated example

D The best simple run-proof policy

In this appendix, we prove Proposition 4′ which characterizes the best simple run-proof policy when the parameter δ is known. Proposition 4, as stated in Section 4.2, requires the best general run-proof policy to be increasing on the interval $[\delta, \pi]$. The more general result we prove here only requires it to be increasing on a subset of this interval, $[m^T, \pi]$ for some $m^T \geq \delta$, and determines the best simple policy with a threshold in this subset. Proposition 4 follows immediately as a special case with m^T set to δ. The more general result is used to establish Proposition 5 at the end of Section 4.3, which characterizes the best robust simple policy.

Consider the following "truncated" problem of solving for the best simple run-proof policy with a threshold in $m^T \in [0, \pi]$:

$$\max_{(\bar{m},\bar{c})\in P} \quad \mathcal{W}\big(\bar{m}, \bar{c}\big) \qquad [P^S(m^T)]$$
$$\text{s.t.} \quad \bar{m} \geq m^T.$$

Note that setting $m^T = 0$ yields the optimal policy problem discussed in Section 4.2. The following proposition characterizes the solution to the more general problem $[P^S(m^T)]$.

Proposition 4′. *Given any $m^T \geq \delta$, suppose that the best general policy $c_1^*(m_1)$ based on δ is increasing on $[m^T, \pi]$. Suppose that there exists a unique $c^T < 1$ such that $\mathcal{R}(m^T, c^T) = 0$. Then (m^T, c^T) is the unique optimal solution to problem $[P^S(m^T)]$.*

Proof. Choose any $m^T \geq \delta$. For $\bar{m} \geq m^T$, by change of variables $\big(m_1 = \pi_1 + \delta(1-\pi_1)\big)$, we can rewrite (22) as

$$\mathcal{R}(\bar{m}, \bar{c}; \delta) = \int_{\delta}^{\bar{m}} \mathbb{E}_{r_2}\left[\ln\left(\frac{1-m_1}{A_{r_2}-m_1}\right)\right] f_n\left(\frac{m_1-\delta}{1-\delta}\right) dm_1 + \int_{\bar{m}}^{\pi} \mathbb{E}_{r_2}\left[\ln\left(\frac{\bar{c}(1-m_1)}{A_{r_2}-m_1\bar{c}}\right)\right] f_n\left(\frac{m_1-\delta}{1-\delta}\right) dm_1 - (1-\delta)T,$$

where $A_{r_2} = r_2(1-\pi) + \pi$. Taking the partial derivatives of $\mathcal{R}(\bar{m}, \bar{c})$ gives us

$$\mathcal{R}_{\bar{m}}(\bar{m}, \bar{c}; \delta) = \mathbb{E}_{r_2}\left[\ln\left(\frac{A_{r_2}-\bar{m}\bar{c}}{\bar{c}(A_{r_2}-\bar{m})}\right)\right] f_n\left(\frac{\bar{m}-\delta}{1-\delta}\right) > 0,$$
$$\mathcal{R}_{\bar{c}}(\bar{m}, \bar{c}; \delta) = \int_{\bar{m}}^{\pi} \mathbb{E}_{r_2}\left[\frac{A_{r_2}}{\bar{c}(A_{r_2}-m_1\bar{c})}\right] f_n\left(\frac{m_1-\delta}{1-\delta}\right) dm_1 > 0.$$

Next, let $\hat{m}(\bar{c}) = \frac{(R-1)\pi}{R-\bar{c}}$. Note that $\frac{\pi-\hat{m}(\bar{c})\bar{c}}{\pi-\hat{m}(\bar{c})} = \frac{A_R-\hat{m}(\bar{c})\bar{c}}{1-\hat{m}(\bar{c})}$. If $\bar{m} \leq \hat{m}(\bar{c})$, we have

$$c_2^N(m_1, \bar{c}) = \begin{cases} \frac{\pi - m_1\bar{c}}{\pi - m_1} & \text{if } m_1 \in [\bar{m}, \hat{m}(\bar{c})] \\ \frac{A_R - m_1\bar{c}}{1-m_1} & \text{if } m_1 \in (\hat{m}(\bar{c}), \pi). \end{cases}$$

Therefore, the corresponding welfare function in the no-run equilibrium is

$$\begin{aligned} \mathcal{W}^A(\bar{m}, \bar{c}) = & \int_{\bar{m}}^{\hat{m}(\bar{c})} \left[m_1 \ln(\bar{c}) + (\pi - m_1) \ln\left(\frac{\pi - m_1\bar{c}}{\pi - m_1}\right) + (1-\pi)\ln(R) \right] f(m_1) dm_1 \\ & + \int_{\hat{m}(\bar{c})}^{\pi} \left[m_1 \ln(\bar{c}) + (1 - m_1) \ln\left(\frac{A_R - m_1\bar{c}}{1 - m_1}\right) \right] f(m_1) dm_1 + \int_0^{\bar{m}} (1-\pi)\ln(R) f(m_1) dm_1. \end{aligned}$$

Taking the partial derivatives of $\mathcal{W}^A(\bar{m}, \bar{c})$ gives us

$$\begin{aligned} \mathcal{W}^A_{\bar{m}}(\bar{m}, \bar{c}) &= -[\bar{m}\ln(\bar{c}) + (\pi - \bar{m})\ln(\frac{\pi - \bar{m}\bar{c}}{\pi - \bar{m}})] f(\bar{m}), \\ \mathcal{W}^A_{\bar{c}}(\bar{m}, \bar{c}) &= \int_{\bar{m}}^{\hat{m}(\bar{c})} \frac{\pi(1-\bar{c})}{\bar{c}(\pi - m_1\bar{c})} m_1 f(m_1) dm_1 + \int_{\hat{m}(\bar{c})}^{\pi} \frac{A_R - \bar{c}}{\bar{c}(A_R - m_1\bar{c})} m_1 f(m_1) dm_1. \end{aligned}$$

If $\bar{m} > \hat{m}(\bar{c})$, we have $c_2^N(m_1, \bar{c}) = \frac{A_R - m_1\bar{c}}{1-m_1}$ for $m_1 \geq \bar{m}$. The corresponding welfare function in the no-run equilibrium is

$$\mathcal{W}^B(\bar{m}, \bar{c}) = \int_0^{\bar{m}} (1-\pi)\ln(R) f(m_1) dm_1 + \int_{\bar{m}}^{\pi} \left[m_1 \ln(\bar{c}) + (1 - m_1)\ln\left(\frac{A_R - m_1\bar{c}}{1 - m_1}\right) \right] f(m_1) dm_1.$$

Taking the partial derivatives of $\mathcal{W}^B(\bar{m}, \bar{c})$ gives us

$$\begin{aligned} \mathcal{W}^B_{\bar{m}}(\bar{m}, \bar{c}) &= (1-\pi)\ln(R) f(\bar{m}) - [\bar{m}\ln(\bar{c}) + (1-\bar{m})\ln(\frac{A_R - \bar{m}\bar{c}}{1-\bar{m}})] f(\bar{m}), \\ \mathcal{W}^B_{\bar{c}}(\bar{m}, \bar{c}) &= \int_{\bar{m}}^{\pi} \frac{A_R - \bar{c}}{\bar{c}(A_R - m_1\bar{c})} m_1 f(m_1) dm_1. \end{aligned}$$

Note that, at the point $(\bar{m}, \bar{c})$ such that $\bar{m} = \hat{x}(\bar{c})$, we have $\mathcal{W}^A_{\bar{m}}(\bar{m}, \bar{c}) = \mathcal{W}^B_{\bar{m}}(\bar{m}, \bar{c})$ and $\mathcal{W}^A_{\bar{c}}(\bar{m}, \bar{c}) = \mathcal{W}^B_{\bar{c}}(\bar{m}, \bar{c})$. Therefore, the indifference curve of $\mathcal{W}(\bar{m}, \bar{c})$ is differentiable everywhere. To complete the proof, we will show that the indifference curve of $\mathcal{W}(\bar{m}, \bar{c})$ is strictly flatter than the frontier of the run-proof set P, i.e., $\frac{\mathcal{W}^A_{\bar{m}}(\bar{m},\bar{c})}{\mathcal{W}^A_{\bar{c}}(\bar{m},\bar{c})} < \frac{\mathcal{R}_{\bar{m}}(\bar{m},\bar{c};\delta)}{\mathcal{R}_{\bar{c}}(\bar{m},\bar{c};\delta)}$ and $\frac{\mathcal{W}^B_{\bar{m}}(\bar{m},\bar{c})}{\mathcal{W}^B_{\bar{c}}(\bar{m},\bar{c})} < \frac{\mathcal{R}_{\bar{m}}(\bar{m},\bar{c};\delta)}{\mathcal{R}_{\bar{c}}(\bar{m},\bar{c};\delta)}$, at any $(\bar{m}, \bar{c}) \in P$ with $\bar{m} \geq m^T$ and $\bar{c} \leq c^T < 1$.
We begin this final step by presenting a lemma whose proof is straightforward and, therefore, omitted.

Lemma 6. *Suppose the functions $f_1(x)$, $f_2(x)$, and $g(x)$ satisfy the following conditions: (i) $g(x) \geq 0$ for $x \in [a,b]$, (ii) $\frac{f_1(x)}{g(x)}$ is increasing in $[a,c]$, (iii) $\frac{f_2(x)}{g(x)}$ is increasing in $[c,b]$, and (iv) $\frac{f_1(c)}{g(c)} = \frac{f_2(c)}{g(c)}$. Then*

$$\frac{\int_a^c f_1(x)dx + \int_c^b f_2(x)dx}{\int_a^b g(x)dx} \geq \frac{f_1(a)}{g(a)}.$$

To adapt the result in Lemma 6 to our problem here, let $f_1(m_1) = \frac{(1-\bar{c})\pi}{\bar{c}(\pi - m_1\bar{c})} m_1 f(m_1)$, $f_2(m_1) = \frac{A_R - \bar{c}}{\bar{c}(A_R - m_1\bar{c})} m_1 f(m_1)$, and

$$g(m_1) = \mathbb{E}_{r_2}\left[\frac{A_{r_2}}{\bar{c}(A_{r_2} - m_1\bar{c})}\right] f_n\left(\frac{m_1 - \delta}{1-\delta}\right).$$

Since the best general policy $c_1^*(m_1)$ based on δ is increasing in $[\delta, \pi]$, it follows from Proposition 7 that $L_{m_1}(m_1, \bar{c}) \geq 0$ for $m_1 \in [\bar{m}, \pi]$. Therefore, $\frac{f_1(m_1)}{g(m_1)}$ is increasing in $[\bar{m}, \hat{m}(\bar{c})]$, $\frac{f_2(m_1)}{g(m_1)}$ is increasing in $[\hat{m}(\bar{c}), \pi]$, and $\frac{f_1(\hat{m}(\bar{c}))}{g(\hat{m}(\bar{c}))} = \frac{f_2(\hat{m}(\bar{c}))}{g(\hat{m}(\bar{c}))}$. Using Lemma 6, we then have

$$\frac{\mathcal{W}^A_{\bar{c}}(\bar{m},\bar{c})}{\mathcal{R}_{\bar{c}}(\bar{m},\bar{c};\delta)} = \frac{\int_{\bar{m}}^{\hat{m}(\bar{c})} f_1(m_1)dm_1 + \int_{\hat{m}(\bar{c})}^{\pi} f_2(m_1)dm_1}{\int_{\bar{m}}^{\pi} g(m_1)dm_1} \geq \frac{f_1(\bar{m})}{g(\bar{m})} = \underbrace{\frac{\pi\bar{m}(1-\bar{c})}{\mathbb{E}_{r_2}\left[\frac{A_{r_2}(\pi - \bar{m}\bar{c})}{A_{r_2} - \bar{m}\bar{c}}\right]}}_{\equiv T_1^A(\bar{m},\bar{c})} \frac{f(\bar{m})}{f_n\left(\frac{\bar{m}-\delta}{1-\delta}\right)},$$

and

$$\frac{\mathcal{W}^B_{\bar{c}}(\bar{m},\bar{c})}{\mathcal{R}_{\bar{c}}(\bar{m},\bar{c};\delta)} = \frac{\int_{\bar{m}}^{\pi} f_2(m_1)dm_1}{\int_{\bar{m}}^{\pi} g(m_1)dm_1} \geq \underbrace{\frac{\bar{m}(A_R - \bar{c})}{\mathbb{E}_{r_2}\left[\frac{A_{r_2}(A_R - \bar{m}\bar{c})}{A_{r_2} - \bar{m}\bar{c}}\right]}}_{\equiv T_1^B(\bar{m},\bar{c})} \frac{f(\bar{m})}{f_n\left(\frac{\bar{m}-\delta}{1-\delta}\right)}.$$

Note that

$$\frac{\mathcal{W}^A_{\bar{m}}(\bar{m},\bar{c})}{\mathcal{R}_{\bar{m}}(\bar{m},\bar{c};\delta)} = \underbrace{\frac{-\left[\bar{m}\ln(\bar{c}) + (\pi - \bar{m})\ln\left(\frac{\pi - \bar{m}\bar{c}}{\pi - \bar{m}}\right)\right]}{\mathbb{E}_{r_2}\left[\ln\left(\frac{A_{r_2} - \bar{m}\bar{c}}{\bar{c}(A_{r_2} - \bar{m})}\right)\right]}}_{\equiv T_2^A(\bar{m},\bar{c})} \frac{f(\bar{m})}{f_n\left(\frac{\bar{m}-\delta}{1-\delta}\right)},$$

and

$$\frac{\mathcal{W}^B_{\bar{m}}(\bar{m},\bar{c})}{\mathcal{R}_{\bar{m}}(\bar{m},\bar{c};\delta)} = \underbrace{\frac{(1-\pi)\ln(R)f(\bar{m}) - \left[\bar{m}\ln(\bar{c}) + (1-\bar{m})\ln\left(\frac{A_R - \bar{m}\bar{c}}{1-\bar{m}}\right)\right]}{\mathbb{E}_{r_2}\left[\ln\left(\frac{A_{r_2} - \bar{m}\bar{c}}{\bar{c}(A_{r_2} - \bar{m})}\right)\right]}}_{\equiv T_2^B(\bar{m},\bar{c})} \frac{f(\bar{m})}{f_n\left(\frac{\bar{m}-\delta}{1-\delta}\right)}.$$

Note that, for any $\bar{m} \in [0, \pi]$, $T_1^A(\bar{m}, 1) = T_2^A(\bar{m}, 1)$ and $T_1^B(\bar{m}, 1) \geq T_2^B(\bar{m}, 1)$. It follows from $\bar{c} < 1$ that

$$T_1^A(\bar{m}, \bar{c}) - T_2^A(\bar{m}, \bar{c}) > T_1^A(\bar{m}, 1) - T_2^A(\bar{m}, 1) = 0$$
$$T_1^B(\bar{m}, \bar{c}) - T_2^B(\bar{m}, \bar{c}) > T_1^B(\bar{m}, 1) - T_2^B(\bar{m}, 1) \geq 0,$$

which completes the proof. □

Note that Proposition 4 in Section 4.2 follows directly from this result by setting $m^T = \delta$, and Proposition 5 in Section 4.3 follows by setting $m^T = \bar{m}_R^*$.

References

Acharya, V. V. and Yorulmazer, T. (2008). Cash-in-the-market pricing and optimal resolution of bank failures. *Review of Financial Studies*, 21(6):2705–2742.

Allaire, N., Breckenfelder, J., and Hoerova, M. (2024). Fund fragility: The role of fund ownership base. SSRN Working Paper.

Allen, F. and Gale, D. (2004). Financial intermediaries and markets. *Econometrica*, 72(4):1023–1061.

Andolfatto, D. and Nosal, E. (2024). Efficient run-proof and run-prone banking. SSRN Working Paper.

Andolfatto, D., Nosal, E., and Sultanum, B. (2017). Preventing bank runs. *Theoretical Economics*, 12(3):1003–1028.

Capponi, A., Glasserman, P., and Weber, M. (2020). Swing pricing for mutual funds: Breaking the feedback loop between fire sales and fund redemptions. *Management Science*, 66(8):3581–3602.

Capponi, A., Glasserman, P., and Weber, M. (2023). Swing pricing: Theory and evidence. *Annual Review of Financial Economics*, 15(1):617–640.

Cavalcanti, R. d. O. and Monteiro, P. K. (2016). Enriching information to prevent bank runs. *Economic Theory*, 62(3):477–494.

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

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

Cipriani, M., Martin, A., McCabe, P. E., and Parigi, B. M. (2014). Gates, fees, and preemptive runs. *FRB of New York Staff Report 670*.

Cooper, R. and Ross, T. W. (1998). Bank runs: Liquidity costs and investment distortions. *Journal of Monetary Economics*, 41(1):27–38.

Dávila, E. and Goldstein, I. (2023). Optimal deposit insurance. *Journal of Political Economy*, 131(7):1676–1730.

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

Engineer, M. (1989). Bank runs and the suspension of deposit convertibility. *Journal of Monetary Economics*, 24(3):443–454.

Ennis, H. M. (2013). Some theoretical considerations regarding net asset values for money market funds. *FRB Economic Quarterly*, 98(4):231–254.

Ennis, H. M. and Keister, T. (2006). Bank runs and investment decisions revisited. *Journal of Monetary Economics*, 53(2):217–232.

Ennis, H. M. and Keister, T. (2009a). Bank runs and institutions: The perils of intervention. *American Economic Review*, 99(4):1588–1607.

Ennis, H. M. and Keister, T. (2009b). Run equilibria in the Green-Lin model of financial intermediation. *Journal of Economic Theory*, 144(5):1996–2020.

Ennis, H. M., Lacker, J. M., and Weinberg, J. A. (2023). Money market fund reform: Dealing with the fundamental problem. *Journal of Risk and Financial Management*, 16(1):42.

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

Government Accountability Office (2023). Money market mutual funds: Pandemic revealed unresolved vulnerabilities. https://www.gao.gov/assets/gao-23-105535.pdf.

Green, E. J. and Lin, P. (2003). Implementing efficient allocations in a model of financial intermediation. *Journal of Economic Theory*, 109(1):1–23.

Huang, X. (2024). Sophisticated banking contracts and fragility when withdrawal information is public. *Theoretical Economics*, 19(1):285–324.

Izumi, R. and Li, Y. (2025). Financial stability with fire sale externalities. *Journal of Money, Credit and Banking*, forthcoming.

Jin, D., Kacperczyk, M., Kahraman, B., and Suntheim, F. (2022). Swing pricing and fragility in open-end mutual funds. *Review of Financial Studies*, 35(1):1–50.

Lewrick, U. and Schanz, J. F. (2017a). Is the price right? Swing pricing and investor redemptions. *BIS Working Paper 644*.

Lewrick, U. and Schanz, J. F. (2017b). Liquidity risk in markets with trading frictions: What can swing pricing achieve? *BIS Working Paper 643*.

Li, L., Li, Y., Macchiavelli, M., and Zhou, X. (2021). Liquidity restrictions, runs, and central bank interventions: Evidence from money market funds. *Review of Financial Studies*, 34(11):5402–5437.

Ma, Y., Xiao, K., and Zeng, Y. (2025). Bank debt, mutual fund equity, and swing pricing in liquidity provision. *Review of Financial Studies*, forthcoming.

McCabe, P. E., Cipriani, M., Holscher, M., and Martin, A. (2013). The minimum balance at risk: A proposal to mitigate the systemic risks posed by money market funds. *Brookings Papers on Economic Activity*, 2013(1):211–278.

Mitkov, Y. (2025). Private sunspots in games of coordinated attack. *Theoretical Economics*, 20(2):427–451.

Parlatore, C. (2016). Fragility in money market funds: Sponsor support and regulation. *Journal of Financial Economics*, 121(3):595–623.

Peck, J. and Shell, K. (2003). Equilibrium bank runs. *Journal of Political Economy*, 111(1):103–123.

President's Working Group on Financial Markets (2020). Overview of recent events and potential reform options for money market funds. https://home.treasury.gov/system/files/136/PWG-MMF-report-final-Dec-2020.pdf.

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

Securities and Exchange Commission (2014). Money market fund reform; Amendments to Form PF. Federal Register Vol. 79, No. 157 (47736), August 14. https://www.federalregister.gov/d/2014-17747.

Securities and Exchange Commission (2022). Money market fund reforms. Federal Register Vol. 87, No. 26 (7248), February 8. https://www.federalregister.gov/d/2021-27532.

Securities and Exchange Commission (2023). Money market fund reforms and related amendments. Federal Register Vol. 88, No. 148 (51404), August 3. https://www.federalregister.gov/d/2023-15124.

Sultanum, B. (2014). Optimal Diamond-Dybvig mechanism in large economies with aggregate uncertainty. *Journal of Economic Dynamics and Control*, 40:95–102.

Voellmy, L. (2021). Preventing runs with fees and gates. *Journal of Banking & Finance*, 125:106065.

Wallace, N. (1988). Another attempt to explain an illiquid banking system: The Diamond and Dybvig model with sequential service taken seriously. *Federal Reserve Bank of Minneapolis Quarterly Review*, 12(Fall):3–16.

Wallace, N. (1990). A banking model in which partial suspension is best. *Federal Reserve Bank of Minneapolis Quarterly Review*, 14(Fall):11–23.

Zeng, Y. (2017). A dynamic theory of mutual fund runs and liquidity management. *Working Paper No. 42, European Systemic Risk Board, April.*