

Federal Reserve Bank of New York
Staff Reports

The Minimum Balance at Risk: A Proposal to Mitigate the Systemic Risks Posed by Money Market Funds

Patrick E. McCabe
Marco Cipriani
Michael Holscher
Antoine Martin

Staff Report No. 564
July 2012



This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the authors and are not necessarily reflective of views at the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

The Minimum Balance at Risk: A Proposal to Mitigate the Systemic Risks Posed by Money Market Funds

Patrick E. McCabe, Marco Cipriani, Michael Holscher, and Antoine Martin

Federal Reserve Bank of New York Staff Reports, no. 564

July 2012

JEL classification: G01, G20, G28

Abstract

This paper introduces a proposal for money market fund (MMF) reform that could mitigate systemic risks arising from these funds by protecting shareholders, such as retail investors, who do not redeem quickly from distressed funds. Our proposal would require that a small fraction of each MMF investor's recent balances, called the "minimum balance at risk" (MBR), be demarcated to absorb losses if the fund is liquidated. Most regular transactions in the fund would be unaffected, but redemptions of the MBR would be delayed for thirty days. A key feature of the proposal is that large redemptions would subordinate a portion of an investor's MBR, creating a disincentive to redeem if the fund is likely to have losses. In normal times, when the risk of MMF losses is remote, subordination would have little effect on incentives. We use empirical evidence, including new data on MMF losses from the U.S. Treasury and the Securities and Exchange Commission, to calibrate an MBR rule that would reduce the vulnerability of MMFs to runs and protect investors who do not redeem quickly in crises.

Key words: money market funds, runs, redemption restrictions, systemic risk

McCabe: Board of Governors of the Federal Reserve System (e-mail: patrick.e.mccabe@frb.gov). Cipriani, Holscher, Martin: Federal Reserve Bank of New York (e-mail: marco.cipriani@ny.frb.org, michael.holscher@ny.frb.org, antoine.martin@ny.frb.org). The authors are grateful to Ken Anadu, Steffanie Brady, Darrell Duffie, Matt Eichner, Josh Gallin, Jamie McAndrews, Susan McLaughlin, Eric Rosengren, and seminar participants at the Federal Reserve Bank of New York, the Federal Reserve Bank of Boston, and the Board of Governors of the Federal Reserve System for comments on earlier versions of this paper, and they thank Greg Nini, Henry Shilling, and Sarah ten Siethoff for help providing and interpreting data. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Board of Governors, the Federal Reserve Bank of New York, or the Federal Reserve System.

1. Introduction

By many measures, U.S. money market funds (MMFs or “money funds”) are a very popular financial product. With \$2.7 trillion in assets under management at the end of 2011, MMFs represented nearly a quarter of all U.S. mutual fund assets (ICI, 2012). MMFs are key intermediaries of short-term funding, particularly for financial institutions: These funds owned over 40 percent of U.S. dollar-denominated financial commercial paper outstanding at the end of 2011 and about one-third of dollar-denominated negotiable certificates of deposit.¹ In part because of their record in maintaining a stable, \$1.00 net asset value (NAV), MMFs serve as an important cash-management tool for individuals, firms, institutions, and governments.

However, MMFs are vulnerable to runs and, given the size of the money fund industry and its importance in allocating short-term funding to financial institutions, this vulnerability poses considerable systemic risk to the U.S. financial system. The potentially dire consequences of a run on MMFs were evident in September 2008, when the Lehman bankruptcy caused one fund to “break the buck” (its NAV fell below \$1) and triggered massive and widespread redemptions from other MMFs. These outflows were a key factor in the freezing of short-term funding markets and a broader curtailment of credit supply (see, for example, Federal Open Market Committee, 2008; Board of Governors of the Federal Reserve System, 2009; U.S. Securities and Exchange Commission (SEC), 2009; President’s Working Group on Financial Markets (PWG), 2010). The severity of the damage to financial stability caused by the run in 2008 led to unprecedented government interventions to support MMFs in order to halt the run. Less than two years after that crisis, the heavy exposures of MMFs to European financial institutions put the funds at risk of transmitting strains from Europe very rapidly to U.S. short-term funding markets (Financial Stability Oversight Council (FSOC), 2011; Chernenko and Sunderam, 2012).

In light of the systemic risk stemming from MMFs’ susceptibility to runs, calls for reform have come from government agencies (e.g., Schapiro, 2010, 2011, 2012; PWG, 2010; FSOC, 2011), academics (e.g., Squam Lake Group, 2011), and industry (e.g., Mendelson and Hoerner, 2011; Goebel, Dwyer, and Messman, 2011). Some have proposed eliminating MMFs’ stable \$1 NAVs by requiring that the funds have “floating” NAVs that reflect the mark-to-market values of their shares more precisely (e.g., Volcker, 2011; Squam Lake Group, 2011; Lacker, 2011). Others have argued for mechanisms to bolster the stable NAV, particularly capital buffers that could absorb portfolio losses. Although either a floating NAV or a capital buffer could provide additional stability to MMFs, it is worth noting that investors in an MMF with a floating NAV would still face strong incentives to redeem shares quickly at the first sign of trouble—before other redemptions deplete the fund’s most-liquid assets—and similar incentives to run would face shareholders in a fund that is perceived to be at risk of losses that exceed its capital buffer.

¹ Sources: MMFs’ filings of Securities and Exchange Commission (SEC) form N-MFP, certificates of deposit data from DTCC, and commercial paper data from DTCC and the Federal Reserve Board.

This paper proposes another approach to mitigating the vulnerability of MMFs to runs by introducing a “minimum balance at risk” (MBR) that could provide a *disincentive* to run from a troubled money fund. The MBR would be a small fraction (for example, 5 percent) of each shareholder’s recent balances that could be redeemed only with a delay. The delay would ensure that redeeming investors remain partially invested in the fund long enough (we suggest 30 days) to share in any imminent portfolio losses or costs of their redemptions. However, as long as an investor’s balance exceeds her MBR, the rule would have no effect on her transactions, and no portion of any redemption would be delayed if her remaining shares exceed her minimum balance.

The motivation for an MBR is to diminish the benefits of redeeming MMF shares quickly when a fund is in trouble and to reduce the potential costs that others’ redemptions impose on non-redeeming shareholders. Thus, the MBR would be an effective deterrent to runs because, in the event that an MMF breaks the buck (and only in such an event), the MBR would ensure a *fairer allocation* of losses among investors.

Importantly, an MBR rule also could be structured to create a *disincentive* for shareholders to redeem shares in a troubled MMF, and we show that such a disincentive is necessary for an MBR rule to be effective in slowing or stopping runs. In particular, we suggest a rule that would subordinate a portion of a redeeming shareholders’ MBR, so that the redeemer’s MBR absorbs losses before those of non-redeemers. Because the risk of losses in an MMF is usually remote, such a mechanism would have very little impact on redemption incentives in normal circumstances. However, if losses became more likely, the expected cost of redemptions would increase. Investors would still have the option to redeem, but they would face a choice between redeeming to preserve liquidity and staying in the fund to protect principal. Creating a disincentive for redemptions when a fund is under strain is critical in protecting MMFs from runs, since shareholders otherwise face powerful incentives to redeem in order to simultaneously preserve liquidity and avoid losses.

One concern about some disincentives for redemptions is that they might dampen market discipline that results from investors’ flexibility to redeem. On the contrary, the MBR should *strengthen* incentives for *early* market discipline for MMFs by clarifying that investors cannot quickly redeem all shares from a fund *during* a crisis. Since investors would have strong incentives to identify potential problems well before any losses are realized, the market discipline encouraged by the MBR would more likely be based on investors’ assessments of the riskiness of a fund’s strategy or operations, rather than on headlines that trigger runs. Thus, redemptions that result from incentives created by the MBR would probably be diffuse, rather than destabilizing.

An MBR could be introduced in a manner that preserves the basic features of MMFs. Funds could, for example, maintain stable \$1 NAVs and honor most transactions without any delay or restrictions. The loss-allocation rules that are central to the MBR concept would *only* affect shareholders’ actual balances in the event that a fund breaks the buck and closes. Implementation could be fairly straightforward, as a fund would only have to track two additional variables for each investor—her minimum balance at risk and any portion of her

MBR that she has requested to redeem. Importantly, the MBR concept could be introduced together with other reforms, such as a capital-buffer requirement. We discuss some of the advantages and disadvantages of an MBR rule, including one that accompanies a small capital buffer.

An appropriately designed MBR would have several benefits. By discouraging investors from redeeming from a troubled MMF, the MBR would help the fund avoid the need for fire sales of assets to raise cash. This not only benefits the fund and its investors, but it reduces the contagion risk that one fund's strains can propagate throughout the financial system. Moreover, the MBR rules that we introduce here would have important benefits for investors who are not prone to redeem shares quickly at the first sign of trouble for an MMF. In particular, retail investors, who historically have been less quick to run from distressed funds, would enjoy additional protections, particularly because the MBR would result in a more equitable allocation of any losses between redeeming and non-redeeming shareholders in the event that a fund breaks the buck.

The paper analyzes the effects of an MBR rule on investor incentives by examining how such a rule would shift the allocation of any fund losses among redeeming and non-redeeming investors. We focus on options that create disincentives for redemptions because investors facing an MBR rule without such disincentives would always benefit by redeeming from a troubled fund. We illustrate how investors' choices to redeem shares or not might vary with preferences for liquidity and with expected losses in a troubled MMF. Hence, an MBR intended to discourage redemptions from a troubled MMF should reflect these factors. We discuss a number of possible adjustments to MBR rules to highlight the flexibility of the concept and its ability to accommodate normative concerns; for example, we describe an option to exempt small accounts from some elements of the MBR. We employ empirical evidence, including new data from the U.S. Treasury and the SEC on MMF losses during September and October 2008, to calibrate an MBR rule that would reduce the vulnerability of MMFs to runs and protect investors who do not redeem quickly in crises. We conclude that an MBR rule that creates a *disincentive* to redeem, particularly in tandem with a capital buffer, would be a practical way to provide meaningful mitigation of the systemic risks associated with MMFs.

The remainder of the paper begins with a discussion in section 2 of the structural vulnerability of MMFs to runs and recent proposals for reform. Section 3 outlines the MBR concept and uses an example to illustrate its workings and benefits. Section 4 analytically describes investors' incentives to redeem from distressed MMFs using a simple model. Section 5 provides a more detailed analysis of how an MBR might work in practice to reduce or counter investors' incentives to redeem shares in a troubled MMF. Section 6 reviews empirical evidence—including new data on MMF losses in 2008—on the principal and liquidity losses that investors might expect when MMFs encounter serious strains, and shows how MBR rules would affect redeeming and non-redeeming investors' expected losses. Section 7 examines some key elements of an MBR rule, including the appropriate size of the MBR. Section 8 provides a policy-oriented discussion of some potential concerns about an MBR rule and a comparison to some other approaches to MMF reform, and section 9 concludes the paper.

2. Background, policy context, and literature

The vulnerability of MMFs to runs can, in large measure, be traced back to their stable \$1 NAVs, the investors who are attracted to stable-value funds, and the methods used to maintain price stability. Like other mutual funds, MMFs provide maturity and liquidity transformation by holding a combination of highly liquid assets and less-liquid, longer-dated securities while allowing shareholders to redeem shares on demand. But unlike other mutual funds, MMFs redeem shares at a price that almost never varies from \$1 per share.² Indeed, the historical success of the funds in maintaining principal stability (only two money funds have “broken the buck” since 1983, when the SEC adopted rule 2a-7 to govern MMFs) has attracted a large, highly-risk averse shareholder base that includes many sophisticated institutional investors. These shareholders reportedly value principal stability above all other features of MMFs (ICI, 2009; Stevens, 2011). As such, they are prone to pulling away from money funds quickly at any sign of trouble.

Although the stable NAV is critical for many MMF investors, no capital buffer or insurance guarantees a money fund’s \$1 share value. Instead, MMFs have relied on a combination of strict SEC rules on portfolio composition, an ability to round their NAVs to the nearest cent, and, when all else fails, voluntary financial support from their sponsors (investment management firms and their affiliates) when they have the wherewithal to provide it (Moody’s, 2010; Rosengren, 2012; Schapiro, 2012). However, if shareholders have reason to question the ability of these mechanisms to prevent losses, they have strong incentives to redeem shares before others do. Institutional shareholders appear to be particularly attuned to these incentives (PWG, 2010; McCabe, 2010; Chernenko and Sunderam, 2012). The imperative to be the first to exit means that the first sign of serious strains for an MMF—or concerns that other shareholders perceive such a problem—may be enough to trigger a run.

The incentive to redeem before others do arises because investors who redeem shares from a troubled MMF may benefit by imposing costs on other shareholders. The uneven allocation of risks and losses between redeeming and non-redeeming shares is largely a consequence of money fund practices and features that help to maintain the stable NAV. Most importantly, because MMFs round their NAVs to the nearest cent, an investor who redeems shares from a fund that has incurred a loss of less than 0.5 percent may still be able to obtain \$1 per share. In effect, the fund transfers a redeeming shareholder’s pro rata share of the loss to the fund’s non-redeeming shareholders, as the loss is concentrated over a shrinking number of shares.

In addition, MMFs meet redemptions by disposing of their highly liquid assets, rather than selling a cross-section of all of their holdings, which typically include less-liquid securities such as term commercial paper and term repo. By managing liquidity in this manner, the funds can avoid realizing losses from sales of less-liquid investments and, as long as all

² A few MMFs maintain stable \$10 share prices.

investors do not redeem at once, this strategy gives investors (in any collective investment vehicle) more liquidity than they would have individually. But during periods of market strain, when liquidity is at a premium, redeemers who receive \$1 per share bear none of the liquidity costs of their redemptions. Instead, non-redeeming investors are left with claims on a less-liquid portfolio. Hence, MMFs' liquidity management practices, as well as NAV rounding, contribute to principal stability but also allow shareholders who redeem from troubled money funds to leave behind risks and costs that must be borne by those who remain invested in the fund.

A third feature of MMFs that reflects their imperative to maintain a stable NAV (and comply with rule 2a-7) contributes to the risk that investors' redemptions may harm shareholders *in other MMFs*. Money funds generally can hold only assets with the highest short-term ratings. In practice, given the relatively small number of private firms with such ratings, MMFs that hold private debt securities and other obligations of private firms (such as repo) have exposures to similar sets of issuers and counterparties—mostly large financial institutions. Hence, redemptions that force one MMF to sell less-liquid assets at a loss may put downward pressure on asset prices, place other MMFs at risk of suffering losses, and prompt shareholders in those funds to redeem shares preemptively.

The severity of the run on MMFs in September 2008 and its broader consequences prompted calls for reforms to mitigate the systemic risks arising from MMFs' structural vulnerability to runs. Although the SEC adopted amendments to rule 2a-7 to make MMFs more resilient to market disruptions, SEC Chairman Mary Schapiro recognized a need "to pursue more fundamental changes to the structure of money market funds to further protect them from the risk of runs" (Schapiro, 2010). The President's Working Group on Financial Markets "agree[d] with the SEC that more should be done to address MMFs' susceptibility to runs" (PWG, 2010). Subsequent concerns that the funds might transmit strains from the European debt crisis to U.S. short-term funding markets provided further motivation for reform (FSOC, 2011; Schapiro, 2011; Rosengren, 2012).

Proposals for MMF reform largely have focused on two possible strategies for mitigating systemic risks. First, the problems that arise from MMFs' stable NAVs have led to many calls for MMFs to switch to floating NAVs. Advocates for such a change include, for example, the Group of Thirty (2009), Lacker (2011), Wall Street Journal (2011), and Volcker (2011). A floating NAV would eliminate some of the destabilizing consequences of NAV rounding and remove the discontinuity in MMF pricing—and hence, some of the resulting disruptions—when share values fall below the 99.5 cent threshold. Importantly, a floating NAV might also make fluctuations in share prices a regular occurrence, improve MMF investors' understanding of the risks inherent in their funds, cause the most risk-averse investors to seek more appropriate investments, and reduce the importance of discretionary sponsor support for MMFs.³ But MMFs with floating NAVs probably would

³ However, not all proposals for requiring floating NAVs would actually result in regular fluctuations in share prices. In particular, requiring that MMFs maintain \$10 NAVs, while continuing to allow them to round to the

still be at some risk of runs, in part because the funds' liquidity management practices would continue to give an advantage to early redeemers in a crisis.

In addition, because a floating NAV requirement would eliminate what appears to be a key attraction for many MMF investors, such a change might lead to a precipitous decline in MMF assets and in these funds' capacity to provide short-term funding. While a shrinking of the MMF industry might reduce the systemic risks arising from the money funds themselves, stable-value investment vehicles would continue to pose systemic risks if assets migrate to other, less regulated, less transparent stable-NAV products (such as offshore MMFs and some private liquidity funds). Alternatively, if institutional investors move cash to banks, the banking system might experience a large increase in uninsured, "hot money" deposits.⁴

A second strategy for mitigating the systemic risks associated with MMFs would bolster, rather than eliminate, their stable NAVs. Two broad approaches have been suggested. First, to improve the ability of MMFs to maintain principal stability, funds might be required to hold a capital buffer that would absorb portfolio losses. A buffer also could potentially replace the principal-stabilizing function of NAV rounding and perhaps eliminate the need for sponsor support for troubled funds, so—depending on its size and other features—a capital buffer might curtail the problems that arise from these practices. Proposals for capital buffers for MMFs include Squam Lake Group (2011), Mendelson and Hoerner (2011), Goebel, Dwyer, and Messman (2011), and McCabe (2011). Some advocates for reform, notably the Squam Lake Group, as well as government officials (SEC Chairman Schapiro, 2011; FSOC, 2011) have suggested that MMF reforms should either require a floating NAV or a capital buffer.

Capital buffers have drawbacks, however. A small buffer *on its own* would do little to mitigate systemic risks; investors would likely flee from MMFs in any crisis out of fear that losses would exceed the size of the buffer. Even so, capital might blunt MMF portfolio managers' incentives for prudent risk management and investors' incentives to monitor risks in their funds, since capital could absorb losses associated with small mistakes. Of course, a very large buffer could diminish these concerns, but raising sufficient capital to absorb the losses that might be associated with systemic events would be challenging, particularly in light of the very low yields that MMFs earn when short-term rates are low. Proponents of capital buffers have argued that capital might be raised directly from MMF shareholders by retaining income that would normally be distributed to them (e.g., Goebel, Dwyer, and Messman, 2011), from third-party investors in capital markets (e.g., McCabe,

nearest cent, would probably preserve investors' expectation that NAV changes are rare. Such a "reform" would have very little utility.

⁴ See PWG (2010) for a more thorough discussion of the advantages and disadvantages of requiring MMFs to have floating NAVs.

2011 and BlackRock, 2011), or from MMF sponsors (e.g., BlackRock 2010). Each of these potential sources has its own complications.⁵

A second approach to bolstering MMFs' stable NAVs would impose restrictions or fees on redemptions to address the risks arising from investors' ability to redeem shares on demand. For example, redemptions might be delayed for several days to reduce their immediate drain on a fund's liquidity and give it time to respond to sudden outflows.⁶ Alternatively, fees might be levied to discourage redemptions and ensure that redeeming investors bear the liquidity costs of their actions. One benefit of such restrictions or fees is that they might reinforce investors' understanding that MMF shares are not perfect substitutes for insured deposits with guarantees of principal and liquidity.

A disadvantage of redemption restrictions and fees is that they generally would only be feasible if imposed conditionally. A redemption fee that is charged in all circumstances would negate the principal stability that is critical for many MMF investors. Similarly, an unconditional delay of *every* redemption would undermine the liquidity of shares that is established in the Investment Company Act of 1940 for all mutual funds (not just MMFs). Either change, if applied at all times, would likely have impacts similar to the consequences of a floating NAV.

Clearly, restrictions or fees on redemptions could be imposed conditionally. But applying them only in circumstances when a fund (or the entire MMF industry) is under strain, or when a spike in demand for liquidity boosts the cost of redemptions, would raise the risk of preemptive runs by investors who anticipate that restrictions or fees might soon be put in place. For example, news that one fund has imposed restrictions or fees could cause shareholders in *other* funds to redeem shares on concerns that their funds might do the same. Arguably, conditional fees or restrictions might *increase* the vulnerability of MMFs to runs.

This discussion merely touches on the range of advantages and disadvantages of reform proposals that have been offered elsewhere. We return to a discussion of how the minimum balance at risk might compare to other reform proposals in sections 8.5, 8.6, and 8.7, although a full discussion of the relative merits of other reform options is beyond the scope of this paper. However, a quick review suggests that none of the other options is without shortcomings and that the potential benefits of an MBR rule should be explored in detail.

⁵ For example, a capital buffer obtained from a fund's retained income would take many years to become large enough to provide substantial protection from losses, while capital raised from third parties would require creation of a market for a new, untested type of security. And capital provided by MMF sponsors would potentially lead to further consolidation of the MMF industry among affiliates of large, systemically important financial institutions (and shift the risks associated with money funds toward these institutions).

⁶ A delay in redemptions alone might provide some additional valuable time for a response to a crisis but probably would not add much stability to MMFs, since investors most likely would still react to strains at a fund by submitting large redemption requests.

3. The minimum balance at risk

The minimum balance at risk that we propose in this paper is a new approach to MMF reform that would mitigate the risk of runs by limiting the ability of redeeming investors to benefit by imposing risks, costs, and losses on non-redeemers. The MBR would be a small fraction (for example, 5 percent) of each investor's recent balances that could be redeemed only with a delay. That delay would be sufficient (we suggest 30 days) to ensure that redeeming investors remain partially invested in the fund long enough to share in any imminent portfolio losses or costs of their redemptions. The MBR would have no effect on an investor's transactions in the fund as long as her remaining shares exceeded the MBR. Moreover, the MBR rules that we describe would impose no losses on *any* investor unless that investor's MMF breaks the buck.

Conceptually, the MBR straddles the two primary approaches to shoring up a stable-NAV MMF: capital and redemption restrictions. By identifying a minimum portion of each investor's balance that would be at risk for absorbing losses—whether or not he redeems shares—the MBR essentially serves a function similar to capital. As a form of redemption restriction, the MBR avoids the conditionality problem that limits the usefulness of most such restrictions for reducing systemic risk. That is, the MBR rule would always be in place, so investors would not be able to redeem preemptively to avoid the restriction. Yet, the MBR would have no effect on most transactions in a fund, particularly during normal times. Only when an MMF appears to be at risk of losses would the MBR materially affect investors' incentives to redeem shares.

Properly designed, the MBR would diminish both the benefits to investors of redeeming MMF shares quickly and the potential costs of their redemptions to other (non-redeeming) shareholders. Importantly, an MBR rule could also create an important *disincentive* for shareholders to redeem shares in a troubled MMF: Redemptions would put a portion of a redeeming shareholder's MBR into a subordinated position relative to those of other investors, so that redeemers' MBRs absorb losses before non-redeemers' MBRs. Because the chance of loss in an MMF is almost always remote, such a disincentive normally would be immaterial. However, the disincentive to redeem a fund's shares would strengthen if the fund were to experience an increased risk of losses.

We focus on rules that would subordinate redeemers' MBRs, for several reasons. First, creating a *disincentive* to redeem is essential to reducing the risk of runs on MMFs. An MBR without subordination of redeemers' balances *merely reduces* the incentive to run, and given the intensity of incentives to redeem from a troubled MMF, investors without a disincentive to do so should still be expected to run from such funds. Second, the automatic strengthening of the disincentive when an MMF encounters trouble would allow the subordination rule to be in effect at all times without imposing an undue burden on a fund and its investors. Only in the event that a fund experienced problems would the disincentive become large enough to offset powerful incentives to redeem. Notably, the strengthening of the disincentive would be endogenous and would not require intervention by a fund's board or government officials. Third, a disincentive based on subordination is

fair: Losses *only* would be incurred by investors if the MMF itself were to suffer a loss (in excess of any capital buffer), and losses would be allocated first to those whose actions in redeeming put the most burden on the fund.

3.1 The minimum balance at risk: basic features. The purpose of the minimum balance at risk is to absorb losses in an MMF if (and only if) the fund “breaks the buck” and liquidates. That is, in the event of a fund’s closure, the MBR would be used to allocate losses in a manner that would be different from the current loss-allocation rules, which divide losses based solely on each investor’s total shares in the fund at the time of closure. Importantly, an MBR would allow investors who otherwise would have redeemed all of their shares to shoulder some of a fund’s losses.

The MBR would be determined individually for each investor based on his recent account activity. Specifically, the MBR would be based on the shareholder’s “reference amount,” R , which might be, for example, his maximum balance (“high water mark”) over the previous 30 days. Alternatively, the reference amount might be the investor’s 30-day-average balance or some other function of his recent balances.⁷

The minimum balance at risk for investor i is a fraction, m , of the investor’s reference amount. That is,

$$MBR_i = mR_i$$

For illustration, we discuss a rule with $m = 5$ percent; in any case, m should be large enough to give investors confidence that, *in a crisis*, an MMF’s losses would not exceed its aggregate MBR, that is, the sum of all of its shareholders’ MBRs. Section 7.3 discusses the size of an MBR that would be necessary to provide investors with confidence that they are generally better off *not* redeeming from a distressed MMF.

An investor with a total balance B would have an available balance equal to $B - MBR$. As long as the investor’s redemption requests do not bring his balance below his MBR, his transactions would be unaffected by the MBR rule. When the investor places a redemption request that would bring his balance below the MBR, however, any shortfall relative to the MBR would be disbursed with a delay of D days. We propose a delay of 30 days.

More generally, the redemption delay, D , and the period over which R is calculated should be long enough to minimize the opportunity for preemptive runs. If the delay is short, informed shareholders may continue to see redemptions at the first sign of a problem as an effective way to obtain full payment for all their shares before any losses are realized. If so, the MBR with a short delay might even be *destabilizing*, since such investors might redeem all shares as fast as possible in the hope that the delay period might pass before any

⁷ In either case, the reference amount probably should be calculated based on balances *excluding* any shares that have been held back (subject to redemption with a delay). Section 7.2 below discusses different approaches for setting the reference amount.

material losses are realized. The 30-day delay that we propose should be long enough to minimize the effectiveness of such a strategy for avoiding imminent losses.⁸

3.2 Creating a *disincentive* for shareholders to redeem. To be effective in braking redemptions from a troubled MMF, an MBR rule must create a *disincentive* for redemptions that is strong enough to counter investors' powerful reasons to redeem when an MMF is under strain. Since the MBR works by affecting the allocation of losses only in the event that a fund breaks the buck, a sensible way to create a disincentive is by stipulating that redeeming investors absorb losses in such a fund before other investors.

Specifically, the MBR rules that we propose would cause some or all of a redeemer's MBR to be subordinated relative to non-redeemers' MBRs, so that if a fund breaks the buck, redeeming shareholders absorb losses before non-redeemers do. Thus, as long as losses do not exceed a fund's aggregate MBR, redeeming shareholders would shoulder a larger share of the losses than non-redeemers. We illustrate these rules in detail in sections 3.4 and 5 below.

3.3 An MBR rule with an NAV-stabilizing buffer. The MBR rules that we propose would work particularly well in tandem with a capital buffer. The disincentives for redemptions from a troubled MMF that are created by a well-calibrated MBR rule would make liquidity-related losses less likely for MMFs and hence make a buffer of any size more effective in providing stability. Clearly, a capital buffer that is large enough to provide substantial loss-absorption capacity also would complement an MBR in mitigating investors' incentives to redeem when a fund suffers a loss.

To illustrate how an MBR might work together with an NAV buffer, much of this paper's analyses focus on an MMF with both an MBR and a small, "NAV-stabilizing" buffer. Our intent is not to suggest that such a small buffer would provide meaningful loss-absorption capacity, but instead to illustrate how a buffer of any size interacts with the MBR. That said, even the 50 basis point buffer that we analyze might be very helpful if only in replacing the problematic role of NAV rounding in providing share-price stability for MMFs. For example, the NAV-stabilizing buffer might be used to absorb routine price fluctuations in MMFs, such as those resulting from increases in interest rates, which under current rules might push a fund's market-value (or "shadow") NAV slightly below \$1. A buffer that absorbs such fluctuations could provide the price stability that investors seek without allowing funds to redeem shares at \$1 when their shadow NAVs are less—a practice that dilutes the value of shares held by non-redeeming shareholders.

⁸ Section 7.1 below provides further discussion of the appropriate delay period for disbursing investors' MBRs. In addition, the period over which R is computed probably should be as long as the delay period, D , since a shorter period for calculating R would allow investors to circumvent the delay period for all but a tiny portion of their shares by redeeming all available balances and waiting a few days for the reference amount to decline.

In the event that an MMF with an MBR rule suffers losses, we assume that losses would be allocated in the following order:

- First, to the NAV-stabilizing buffer (if any). If losses exceed the size of the buffer, the fund is liquidated, and remaining losses are distributed according to the MBR rule.
- Second, any losses in excess of the buffer would be absorbed on a pro rata basis by shareholders' subordinated MBRs (if any).
- Third, additional losses would be absorbed on a pro rata basis by the remainder of shareholders' MBRs.
- Fourth, any remaining losses would be divided on a pro rata basis over all other shares in the fund.

3.4 An illustrative example. A simple example can illustrate the MBR concept and the effectiveness of different MBR rules in discouraging runs on MMFs. Consider a money market fund with a 5 percent MBR ($m = 5$ percent) and two investors, Alice and Bob. On Monday, each has 100 shares in the fund (for simplicity, assume that each investor's reference amount, R , is also \$100), so Alice and Bob can both redeem as much as \$95 immediately, but must leave \$5 in the fund for 30 days.

Although the fund has net assets of \$200, assume that a distressed security constitutes 2 percent of its value. On Tuesday, recognizing this potential problem, Alice redeems *all* of her available balance. Her \$95 redemption represents almost half of the fund's assets and forces the fund to sell some of its less liquid assets at a small discount to pay Alice. We assume a realized loss of 20 cents. Bob does not redeem any shares. The fund finishes the day on Tuesday with 105 shares and \$104.80 in assets. Its shadow NAV is \$0.9981, but rounding allows it to maintain its stable \$1 NAV for purchases and redemptions.

On Wednesday, the distressed security loses 95 percent of its value, so the fund suffers a \$3.80 loss. This forces the fund to break the buck; with \$101 in assets and 105 shares, its NAV is only \$0.9619. Note that cumulative losses (\$4) are 2 percent of the fund's assets on Monday, but Alice's redemptions have concentrated the loss on a smaller number of shares, and those remaining shares have lost 3.81 percent of their value.

The MBR is primarily a mechanism for allocating losses in an MMF in the unlikely event that it breaks the buck, and the different types of MBR rules we describe here vary only in how they affect that loss allocation. In our illustrative example, a key question is how the fund's \$4 loss is divided between Alice and Bob. However, because investors in an MMF that closes may lose access to their cash for a prolonged period, we also examine the costs associated with the impairment of liquidity for Alice and Bob.

No MBR (*the status quo*). Before we examine different MBR rules, a review of the allocation of losses between Alice and Bob under current MMF rules is instructive. Without an MBR, Alice would have been able to redeem *all* 100 of her shares on Tuesday. Hence, as shown in the first line in table 1 below, Alice's loss would have been \$0 (column 1), and Bob

would lose \$4 (column 2).⁹ Importantly, some of the losses that Bob suffers are those specifically associated with Alice’s redemptions; the redeeming shareholder not only transfers her share of the fund’s loss to the non-redeemer, but redemptions may have additional costs that the non-redeemer must bear.

Table 1. Allocation of MMF losses

Type of MBR rule	Rule for allocating losses	No liquidity cost for investor		With liquidity cost for investor	
		Alice’s loss (1)	Bob’s loss (2)	Alice’s loss (3)	Bob’s loss (4)
1. None (status quo)	Proportional to current shares	\$0.00	\$4.00	\$0.00	\$4.48
2. Weak	Proportional to current shares	\$0.19	\$3.81	\$0.21	\$4.29
3. Simple	Proportional to MBR	\$2.00	\$2.00	\$2.02	\$2.49
4. Strong	Redemptions can subordinate the MBR so that it absorbs losses first; additional losses proportional to remaining MBR	\$4.00	\$0.00	\$4.01	\$0.50
5. Effective	Redemptions can subordinate a <i>portion</i> of the MBR so that it absorbs losses first; additional losses proportional to remaining MBR	\$3.29	\$0.71	\$3.29	\$1.21

In addition to the principal losses that Bob must shoulder, he suffers a loss of liquidity because his investment in the closed MMF may be inaccessible through a long liquidation process. To illustrate the effect of this liquidity cost for a shareholder during a crisis when liquidity is in short supply, we assume that Alice and Bob would be willing to pay a 50 basis point fee for obtaining all of their assets immediately, rather than having them locked up during a liquidation phase of indefinite duration.¹⁰

Columns 3 and 4 of line 1 show the effect of including liquidity costs to Alice and Bob under current rules. Since Alice has redeemed everything, she suffers no loss of liquidity; her

⁹ Presumably, the additional \$5 that Alice would have been able to redeem on Tuesday without the MBR might have increased the fund’s total losses somewhat (by increasing the cost of liquidating securities).

¹⁰ In section 6.2, we provide empirical evidence to suggest that 50 basis points is a reasonable estimate of the magnitude of the liquidity cost to shareholders in a closed fund. Note that the liquidity costs (described here) to investors whose MMF shares are unavailable for an indefinite period after a fund’s closure are distinct from the liquidation costs (discussed above) that a still-open fund may pay as it sells assets to meet shareholders’ redemptions.

losses are still zero. In contrast, Bob has 100 shares locked up in the fund; his loss of liquidity costs the equivalent of another 48 cents (50 basis points of the value of his shares, which has fallen to \$96), and his total loss is \$4.48.

An MBR with a “weak rule” for allocating losses. To be effective in mitigating the risk of runs, an MBR rule *cannot* simply allocate an MMF’s losses in proportion to the number of shares each shareholder owns at the time a fund experiences losses (the manner in which mutual fund losses would normally be apportioned). To illustrate the ineffectiveness of this approach, consider a “weak” MBR that would divide the \$4 loss between Alice and Bob based on each investor’s current shares.

With a 5 percent MBR, Alice can only redeem 95 shares on Tuesday, so when the fund breaks the buck, she has 5 shares and Bob has 100. The first two columns of line 2 show losses excluding liquidity costs with a weak rule: Alice loses \$0.19 ($\$4 \times 5/105$), while Bob’s loss is \$3.81. When liquidity costs to the investors are included (columns 3 and 4), Alice’s total loss rises by 2.4 basis points (50 basis points of \$4.81), and Bob’s loss increases by 48 basis points. Although the weak rule would prevent the redeeming investor from avoiding *all* losses, redemptions still allow her to avoid the vast majority of losses that a non-redeemer must shoulder. Hence, the weak rule would do little to improve the inequitable allocation of losses between redeemers and non-redeemers and would have little impact on the incentive to redeem.

An MBR with a “simple rule” for allocating losses. A more sensible MBR rule might divide losses among investors in proportion to their MBRs, rather than their total shareholdings. Based on this “simple rule,” since Alice and Bob both have an MBR of \$5 on Wednesday, they would share the loss equally: Each loses \$2 (line 3, columns 1 and 2). Moreover, as long as the fund’s total loss does not exceed the combined MBRs of Alice and Bob (\$10), losses would be shared equally. Bob, however, would bear alone any losses in excess of \$10.

The simple rule represents a notable improvement on the status quo; this rule ameliorates the inequitable allocation of losses between redeeming and non-redeeming investors and reduces incentives to redeem. Alice bears losses in proportion to her holdings on Monday, despite her redemptions on Tuesday. As such, her incentive to redeem to avoid a loss is reduced (as is Bob’s incentive to redeem to avoid losses resulting from Alice’s departure).

However, the simple rule has some very serious limitations. Most importantly, although the simple rule *reduces* the incentive to redeem shares in a troubled MMF, it does not eliminate that incentive. Indeed, Alice still benefits by redeeming her shares on Tuesday, for two reasons. First, she limits her potential losses to the size of her MBR; only non-redeeming shareholders can shoulder any losses in excess of a fund’s aggregate MBR.

Second, even if the fund’s total loss is smaller than its aggregate MBR and Alice and Bob share the principal loss equally, Alice’s redemptions give her immediate access to 95 percent of the value of her original investment in cash. At a time when liquidity is likely to be particularly scarce, Bob’s money may be locked up in a prolonged liquidation of the fund.

When we take these liquidity costs into account, Alice's loss increases to \$2.02 while Bob, who has much more cash tied up in the fund, loses a total of \$2.49.

Thus, investors who redeem under the simple MBR rule incur no additional cost but preserve their own liquidity and cap their losses. Hence, the simple MBR rule does little to mitigate the risk of a run and only partially protects unsophisticated investors from the costs of those runs.

In addition, the simple rule falls short of equitably allocating losses. Alice's early redemptions were costly for the fund, but she does not bear any additional loss that reflects the liquidation costs that she forced upon the fund by redeeming all of her available shares.

An MBR with a "strong rule" for allocating losses. As the discussion of the simple rule makes clear, to be effective in countering shareholders' incentives to redeem from a troubled MMF, an MBR rule must create a *disincentive*. A sensible way to do so is by specifying that redeeming investors should absorb losses before other investors in a break-the-buck scenario. The strong rule creates such a disincentive by stipulating that redemptions cause a redeemer's MBR to be subordinated relative to non-redeemers' MBRs.

More specifically, this rule (described in more detail in section 5.3) would subordinate a fraction of Alice's MBR equal to the fraction of her available shares that she redeems. Since Alice has redeemed *all* of her available shares, she would subordinate all of her MBR. This subordinated buffer absorbs the fund's full \$4 loss, and Bob would lose nothing (line 4 of table 1). Bob does lose liquidity, however, so his loss is 50 cents when this effect is taken into account, while Alice's total loss is \$4.01. With the outcome described here, Alice's redemptions clearly have made her worse off. Importantly, however, even with a strong MBR rule, Alice's choice on Tuesday still might have been to redeem if, for example, she had been concerned that the fund's losses might exceed its aggregate MBR or if she had placed high value on the preservation of her own liquidity.

Clearly, the MBR rule makes redemptions potentially costly, particularly if a fund is in danger of breaking the buck. Hence, the strong rule can help to offset the incentives to redeem that arise from investors' preference for liquidity and the concern that losses might exceed a fund's aggregate MBR. The cost of redemptions may also offset the liquidation costs that redeemers like Alice impose on other investors. In addition, under this rule, redemptions *protect* non-redeeming investors by providing a subordinated buffer to absorb losses. In sum, the strong rule may brake and even reverse the incentive to redeem from a troubled fund before other investors do, so this rule may help to stabilize a distressed MMF, allow it to weather a difficult period, and prevent contagion that might result if redemptions caused further strains.

Even so, the strong rule may fail to stop redemptions if shareholders expect large losses in the fund. If its losses exceed its aggregate MBR buffer, redeeming investors' losses will be limited to their MBRs, but other investors' losses will extend beyond their MBRs.

An MBR with an “effective rule” for allocating losses. One means of attenuating the strong rule’s effect on the allocation of losses would be to limit the fraction of a redeeming shareholder’s MBR that can be subordinated. For example, with an MBR of 5 percent and a subordination limit of 60 percent, only 3 percent of an investor’s recent balances could be subordinated.

Under such an “effective rule,” an investor like Alice, who redeems *all* of her available shares, would have 60 percent of her MBR (\$3.00) subordinated. Any remaining losses (in our example, \$1) are allocated pro rata among non-subordinated MBR balances. Bob’s MBR is \$5, and the remainder of Alice’s MBR is \$2, so Alice’s loss (excluding liquidity costs) is \$3.29, and Bob loses \$0.71 (line 5 of table 1).¹¹ Losses including liquidity costs are \$3.29 and \$1.21 for Alice and Bob, respectively.

One advantage of both the effective and the strong MBR rules is that they offer investors a *choice* between redeeming to preserve liquidity and remaining invested in the fund to preserve principal. In contrast, under current rules for MMFs, investors who redeem quickly *preserve both* liquidity and principal, because they gain immediate access to their cash and avoid any potential losses. Meanwhile, under current rules, investors who do not redeem quickly may *lose both* liquidity and principal, as their money may be tied up in a lengthy liquidation process and they may ultimately suffer a disproportionate share of the fund’s losses. These consequences of the current MMF structure are not only destabilizing, but they are also more likely to hurt retail investors and others who are unlikely to redeem quickly from a distressed fund.

4. A simple model of an MMF investor’s choice to redeem

In this section, we illustrate the incentives of a shareholder in an MMF that has an MBR of fraction m and which will break the buck and close with probability x . Our model sets the context for the analysis of investors’ incentives under different types of MBR rules. Importantly, however, the model in this section does not consider the important effects that other shareholders’ redemptions can have on an investor’s realized losses in an MMF.

Suppose that Dana, an MMF shareholder, must choose whether or not to redeem her shares. If the likelihood of fund closure is nil ($x = 0$), Dana’s choice only reflects whether the \$1 per share that she can receive by redeeming exceeds the value to her of continuing to hold MMF shares. Let r be the net benefit of staying invested in the MMF for another day rather than using the cash for something else. For example, r might be positive because of the net yield advantage of MMF shares over other cash-like investments, while a negative value might reflect the opportunity cost to Dana of leaving cash in the fund, rather than deploying it to a more productive use (suggesting that she is about to redeem her shares for reasons

¹¹ The subordinated portion of Alice’s MBR absorbs the first \$3 in losses. The fund’s non-subordinated MBRs total \$7 (Bob’s \$5 and Alice’s remaining \$2). Since Alice owns 29 percent of the non-subordinated total, she absorbs 29 cents of the remaining \$1 loss. Her total losses are thus \$3.29. Bob absorbs 71 percent of the remaining \$1 loss, or \$0.71.

other than concerns about MMF losses). To summarize, the value that Dana privately assigns to each of her MMF shares is $\$1 + r$.

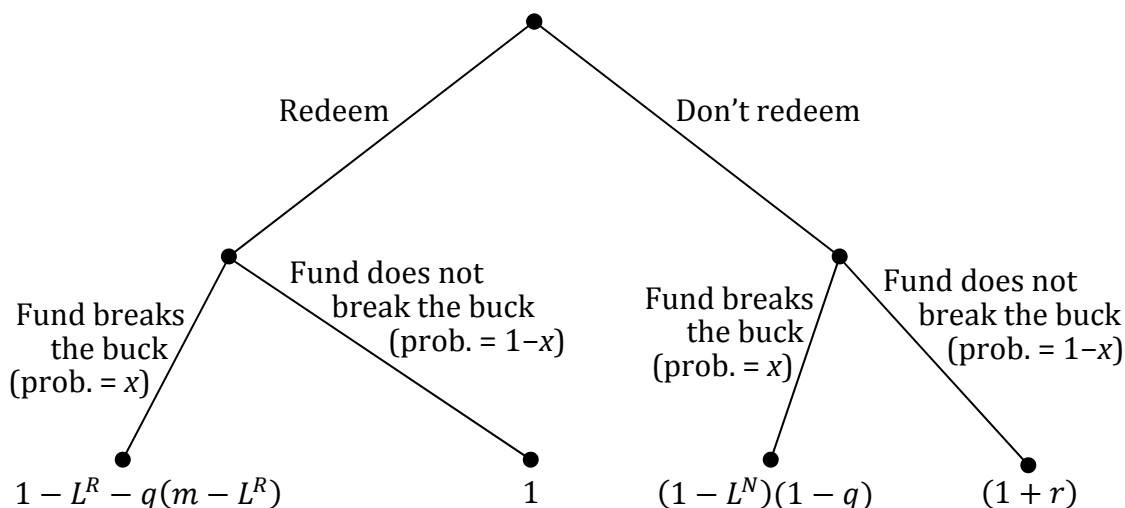
Suppose that Dana knows that if her MMF breaks the buck, she will lose principal of L^R per share if she redeems and L^N per share if she does not. In addition, in the event that her MMF does break the buck, we assume that it closes and liquidates, and Dana thus experiences a liquidity loss of fraction $q \in (0,1)$ of the remaining value of her shares. So, if Dana does not redeem and her fund breaks the buck, her per-share payoff is $\$1$ less principal losses and liquidity costs, that is: $1 - L^N - q(1 - L^N)$. If she does redeem, she loses access only to her MBR, net of any principal losses. As long as her pro rata share of the fund's losses is smaller than her MBR, her remaining stake in the fund is $m - L^R$, her liquidity cost is $q(m - L^R)$, and her per-share payoff is $1 - L^R - q(m - L^R)$. But if her pro rata share is at least as large as her MBR, her payoff is just $1 - m$ per share.

Importantly, even when r is positive, we assume that it is very small relative to Dana's possible losses in the event that her MMF breaks the buck. For example, $r \ll L^N$, $r \ll L^R$, and $r \ll q(1 - L^N)$.

The figure below illustrates Dana's choice and her payoffs for possible outcomes. In terms of expected payoffs, Dana is better off redeeming than not redeeming if:

$$x[1 - L^R - q(m - L^R)] + (1 - x)1 > x[(1 - L^N)(1 - q)] + (1 - x)(1 + r)$$

$$x(L^N - L^R)(1 - q) > (1 - x)r - x(1 - m)q$$



If breaking the buck is certain ($x = 1$), then Dana should redeem if $(L^N - L^R)(1 - q) > -(1 - m)q$, or equivalently, if

$$L^N + \left(\frac{q}{1 - q}\right)(1 - m) > L^R.$$

The two terms on the left-hand side of the inequality are principal losses that Dana would suffer if she does not redeem and the incremental liquidity costs—which are always

positive—that she would suffer by not redeeming. She should redeem if these combined losses exceed her principal losses after redeeming.¹²

Dana’s redemption incentives can be summarized by defining B as the *net* benefit of redeeming shares when breaking the buck is certain. That is,

$$B = L^N + \left(\frac{q}{1-q}\right)(1-m) - L^R$$

We can write the conditions under which Dana should redeem in terms of B :

$$\text{Redeem if } \begin{cases} r < 0, & \text{for } x = 0 \\ r < \left(\frac{x}{1-x}\right)B, & \text{for } 0 < x < 1 \\ B > 0, & \text{for } x = 1 \end{cases}$$

This simple model illustrates several features of Dana’s choice. First, for *very small* values of x , all that matters is Dana’s net preference for holding MMF shares, r . If r is positive, she holds her shares; if it is negative, she redeems. But as the likelihood of MMF losses, x , increases, Dana’s choice hinges on the sign of B .

Second, once B becomes the dominant factor in Dana’s decision, she will always benefit from redeeming unless L^R exceeds L^N by a wide enough margin to cover the incremental liquidity costs of not redeeming. Under current rules, $L^R = 0$, so *nothing* offsets her incentive to redeem. Third, if L^R is large enough that $B < 0$, then the resulting disincentive to redeem shares strengthens as the risk of MMF losses increases.

Dana generally would have to make her decision about whether to redeem shares well before the magnitude of her fund’s eventual realized losses can be observed. As such, Dana would only know an *expected* value of B in advance. Adding to Dana’s uncertainty would be the possible effects of other shareholders’ redemptions on Dana’s losses. Indeed, others’ redemptions may increase or reduce Dana’s expected principal losses, $E(L^R)$ and $E(L^N)$, and Dana’s decision would thus reflect her expectations about others’ behavior. For example, under current rules, $E(L^R) = L^R = 0$, but $E(L^N)$ increases as others redeem, so Dana’s incentive to redeem strengthens when she expects others to do so. We address these issues in the next sections.

¹² If Dana redeems, her principal losses cannot exceed m . When $L^R = m$, Dana should redeem if $L^N + q(1 - L^N) > m$, that is, if the principal losses and liquidity costs she incurs by not redeeming exceed her MBR.

5. The minimum balance at risk: a closer look

In this section, we examine the effects of MBR rules on the incentives of MMF shareholders to redeem when they believe that the likelihood of losses is no longer *de minimis*. As illustrated in section 4, once an MMF is in distress (x is no longer very small), investors' decisions to redeem depend on whether resulting losses exceed those that would be realized by staying invested in the fund *plus* any net liquidity cost associated with losing access to shares for a prolonged period. We examine how MBR rules affect incentives to redeem by describing the relationship between an MMF's total losses and losses for shareholders who redeem immediately before the loss is realized and losses for shareholders who remain invested in the fund. We also illustrate how MBR rules affect the linkages between one investor's losses and *other* investors' redemptions.

We consider an MMF with a 5 percent MBR and focus on outcomes for small investors who each own 0.1 percent of the fund's assets. To illustrate how an MBR might complement a small NAV-stabilizing buffer, we assume that the fund has a 50 basis point buffer, so that in normal times it maintains a mark-to-market ("shadow") NAV of \$1.005. Importantly, we assume that the fund must be liquidated once the NAV-stabilizing buffer is exhausted, so if losses exceed 50 basis points and NAV falls below \$1.00, the fund is closed and the MBR rules are employed to allocate losses among shareholders. This assumption rules out dilutive redemptions (that is, investors receiving \$1 for shares that are worth less than \$1), which are costly for non-redeeming shareholders.¹³

We assume that losses are absorbed first by the NAV-stabilizing buffer. So, for losses smaller than 50 basis points, although the NAV of the fund is eroded and the fund may be required to take action to bolster its buffer, it can remain open and no losses are imposed on shareholders. Any losses exceeding the buffer force the shutdown of the fund, however, and are allocated first to subordinated MBRs, then to non-subordinated MBRs, and finally to other MMF shares, as delineated in section 3.3.

We begin by briefly reviewing allocations of losses in a fund without an MBR rule. Next, we examine the simple, strong, and effective MBR rules. Finally, we discuss some refinements that might be made to these rules.

5.1. No MBR (status quo with a 50 basis point buffer). Chart 1 shows how individual shareholders' losses (the vertical axis) would vary with the losses in an MMF with a 50 basis point buffer and no MBR (the horizontal axis). We consider three investors, one who redeems all shares just before the losses are realized (the thick red line), one who redeems 25 percent of her shares (the thinner, yellow line), and one who redeems no shares (the very thin blue line). Here, we assume that other redemptions from the fund are trivial.

¹³ Relaxing this assumption would generally increase the advantages that are enjoyed by redeeming shareholders, so MMFs that allow dilutive redemptions would need stronger MBR rules to reduce the threat of runs than the funds that we discuss.

Because our aim is to study the choices of shareholders to redeem or not from a troubled MMF, the chart plots investors' losses as a fraction of their pre-redemption assets.¹⁴

As long as the fund's losses do not exceed 50 basis points, the fund stays open and none of the investors incurs any loss. However, once losses exceed 50 basis points, stark differences emerge. The redeeming shareholder bears no loss under any circumstance. The investor who redeems nothing suffers losses in proportion to the fund's losses in excess of 0.5 percent. For example, as the dashed black lines show, if the fund loses 5.5 percent, the non-redeeming investor loses 5 percent. The investor who redeems 25 percent of her available balance just before the fund's loss experiences only 75 percent of the losses that the non-redeeming shareholder suffers.

Chart 1 illustrates only part of investors' strong incentive to run under current rules. Those who redeem immediately before a loss is reflected in a fund's share price are able to shield themselves from any loss of principal. However, there are several additional reasons to redeem from an MMF in trouble. First, other shareholders' redemptions at \$1 per share concentrate losses on remaining shareholders.

We illustrate that point in Chart 2, which shows the losses incurred by a redeeming investor (thick red lines) and a non-redeeming investor (thin blue lines) under different assumptions about the behavior of other investors. As in chart 1, the solid lines show how each investor's losses vary with the fund's losses under the assumption that other redemptions from the fund are trivial. In contrast, the dashed lines in chart 2 illustrate losses when other redemptions total 25 percent of the fund's shares. The dotted lines represent losses when the fund experiences redemptions of 50 percent of assets.

The thin blue dashed and dotted lines show that the losses borne by a non-redeeming investor grow with the fraction of shares that others redeem. That is, when redeeming investors receive \$1 per share, losses are concentrated on remaining shareholders. In contrast, other shareholders' redemptions have no effect on the fully redeeming shareholder: The thick red dashed and dotted lines lie on top of the thick red solid line. By pulling out of the fund, the redeemer has shielded herself not only from unrealized losses in the fund but also from the costly effects of other shareholder's redemptions. The stark difference in outcomes for redeeming and non-redeeming investors highlights the problem that, under current rules, when a fund is perceived to be in trouble, investors have a strong incentive to rush for the exits before other investors do.

Two additional factors may motivate investors to redeem from an MMF in trouble. First, redemptions can force the fund to liquidate assets, creating losses that non-redeeming

¹⁴ In contrast, a comparison of investors' losses as a percentage of the shares they hold at the time of a fund's closure would fail to take into account the cash that redeeming investors have removed from a fund. Thus, it would overstate redeeming investors' total losses as a share of their holdings at the time they make the decision to redeem or remain invested in an MMF.

investors must absorb or depleting the fund's liquidity and leaving it vulnerable to further redemptions. Second, an investor who has redeemed shares preserves her own liquidity if the fund closes. In contrast, an investor who remains in the fund may lose both principal and liquidity if the fund breaks the buck and is liquidated. We examine the incentive to preserve liquidity in section 5.2.

5.2. Simple MBR rule. Under the simple rule, losses in excess of a fund's capital buffer are allocated in proportion to shareholders' MBRs. Hence, principal losses are equally shared, as long as the fund's losses do not exceed its capital plus its aggregate MBRs, although losses of *liquidity* are not equally shared, as illustrated below. Importantly, however, a shareholder can cap his losses at his MBR by redeeming all of his available shares, so if the fund's losses exceed its capital plus its aggregate MBR, non-redeeming shareholders bear all of the additional losses.

Chart 3 illustrates that principal losses that do not exceed the fund's capital and its aggregate MBR are shared equally. The very thin blue, thin yellow, and thick red lines depict losses of investors who redeem none, 25 percent, and all of their *available* shares, which constitute 95 percent of their balances.¹⁵ The lines lie atop one another for investor losses of less than 5 percent. But the fully redeeming investor's losses never exceed his MBR; only investors that have not redeemed all available balances (the very thin blue and thin yellow lines) can incur additional losses.

The incentive to redeem is reduced but not eliminated. Chart 3 makes the point that redeeming everything is clearly beneficial if losses are large, and it does not make an investor worse off if losses are small.

In addition, even with the MBR, shareholders have an incentive to redeem to safeguard their own liquidity. Chart 4 illustrates this point using the simple assumption that shareholders would be indifferent between having their investments locked up in a closed MMF during a liquidation phase of indefinite duration and a 50 basis point fee for obtaining all of their assets immediately.¹⁶

The solid lines in chart 4, which reproduce those in chart 3, exclude liquidity costs. The dotted lines show losses for investors that include the opportunity cost of lost liquidity. These costs depend on the value of the investor's remaining shares in the fund and whether the fund must close (if losses erode the value of an investor's shares in a closed fund, the value of his lost liquidity declines proportionally).

¹⁵ In contrast, in the previous charts showing losses without the MBR, *all* of each investor's shares are available.

¹⁶ In section 6.2, we provide empirical evidence that suggests that this assumption about the cost of lost access to MMF shares is not unreasonable. But the qualitative points we make here are important as long as the liquidity cost is material to investors.

For the non-redeeming investor, liquidity-related losses can be substantial, since her entire balance is locked up for the duration of the fund's liquidation process. The loss of liquidity upon fund closure causes her loss function inclusive of liquidity cost (the very thin, dotted blue line) to jump markedly at the point when the fund is forced to close. Importantly, with a simple MBR rule, when liquidity costs are taken into account, any loss large enough to cause fund closure causes greater losses for non-redeemers than for redeeming investors.¹⁷

Even for a fully redeeming investor, there is a small liquidity cost associated with fund closure (if her shares are still worth something), since that investor still has 5 percent of her assets in the fund. But her liquidity loss is at most 2.5 basis points (50 basis points of 5 percent of assets), and it declines as her MBR is absorbed. Hence, the thick red dotted line lies just a bit above the thick red solid line for MMF losses of less than 5.5 percent.

Chart 5 shows how investors' losses with a simple MBR depend upon other investors' behavior. Here, for simplicity, we ignore preferences for liquidity. With a simple MBR rule, redeemers bear losses *pari passu* with other investors only until redeemers' MBRs are exhausted, and redeemers' losses are not affected by others' redemptions. In contrast, not only can non-redeemers suffer losses in excess of their MBRs, but if losses exceed the fund's aggregate MBR, others' redemptions concentrate losses on non-redeemers. This effect reinforces the incentive to redeem shares when a fund is perceived to be in trouble.

In summary, the simple MBR rule would offer some important improvements over current rules. This rule would reduce redeemers' first-mover advantage and allocate any losses that a fund might suffer more equitably. Hence, it would have some benefit, for example, in protecting retail investors, who have proven to be much less likely to run than institutional investors.

However, the simple MBR rule probably would not provide much additional stability to MMFs in a crisis, particularly for funds with large institutional investors that may be highly motivated to preserve liquidity. The rule falls short because it leaves investors with a substantial incentive to run from distressed MMFs. A redeeming investor still can limit her losses to the size of her MBR, preserve her liquidity, and impose some liquidity costs on others. That is, the simple MBR rule does not eliminate the inequitable allocation of risks and costs between redeeming and non-redeeming shareholders. The former still can shift some costs and losses to the latter, so non-redeemers—particularly retail investors and others who do not redeem quickly at the first sign of trouble—would still be at greatest risk of suffering losses. Thus, even with a simple MBR rule, MMFs still would be vulnerable to runs and constitute a source of systemic risk.

5.3. Strong MBR rule. The strong MBR rule addresses the shortcomings of the simple rule by creating a disincentive for shareholders to redeem. It does so by subordinating the

¹⁷ Importantly, this result does not depend on our assumption here that preservation of liquidity is worth 50 basis points to investors; *any* non-trivial price for liquidity makes redeemers relatively better off.

MBRs of redeeming investors relative to those of non-redeemers, so that redeemers may experience greater losses than non-redeemers if the fund suffers a loss that exceeds its capital buffer.

Under the strong MBR rule, the fraction of a redeeming investor's MBR that is subordinated would be the fraction of his available shares that he has redeemed. That is, his subordinated balance—the amount in a first-loss position—is given by the following formula:¹⁸

$$\text{Subordinated balance} = \text{MBR} \times (\text{cumulative net redemptions} / \text{potential redemptions}).$$

Cumulative net redemptions in this formula are the investor's reference amount less his current balance (cumulative redemptions are zero if his balance exceeds his reference amount).¹⁹ For example, if the reference amount is defined to be his maximum balance over the past 30 days, his cumulative net redemptions would be his maximum balance less his current balance. Potential redemptions are his reference amount less his MBR. So, the formula can be written

$$\text{Subordinated balance} = \text{MBR} \left(\frac{\text{reference amount} - \text{current balance}}{\text{reference amount} - \text{MBR}} \right)$$

as long as the reference amount exceeds the current balance (otherwise, the subordinated balance is zero). All that is needed to compute the subordinated balance is the investor's reference amount, his MBR (which is just a fixed proportion of his reference amount), and his current balance. For example, if an investor redeems all his available shares, his current balance is his MBR, the ratio of cumulative net redemptions to potential redemptions is 1, and the investor's entire MBR would be subordinated. If his net redemptions total only half of his potential redemptions, then half of his MBR would be subordinated.

Chart 6 illustrates how the strong MBR rule might apply to three investors in an MMF: Alice, Bob, and Charlie. We assume that the reference amount for the MBR is each investor's 30-day maximum balance (high-water mark), and all three investors' maximum balances over the past 30 days were \$100. However, Alice has redeemed all of her available shares—her current balance is \$5—so her entire MBR is subordinated. Bob has redeemed nothing, and none of his MBR is subordinated. And Charlie has a current balance of \$62, so his redemptions are 40 percent of his potential redemptions and he has a subordinated balance of \$2 (40 percent of \$5).

Chart 7 aggregates all investors' positions in an MMF and illustrates the order in which losses in the fund would be absorbed. Losses that are smaller than 50 basis points are absorbed by the small NAV-stabilizing buffer and do not cause the fund to close, so MMF

¹⁸ The subordinated balance is in a first-loss position relative to other shareholder assets, but the capital buffer would absorb any losses *before* subordinated balances are affected.

¹⁹ If the investor's reference amount is his high-water mark, his balance cannot exceed his reference amount.

shareholders bear no principal loss and MBR rules for allocating losses are immaterial. Losses exceeding 50 basis points cause the fund to close, so MBR rules govern the seniority of investors' claims and thus how any losses would be allocated. Any MBRs that are in a first-loss position would absorb losses first, pro rata to each investor's share of the fund's aggregate subordinated MBRs. Next, losses exceeding aggregate subordinated MBRs are allocated over remaining (non-subordinated) MBRs, pro rata to each investor's share of the fund's aggregate non-subordinated MBRs. Finally, if losses exceed the fund's aggregate MBR, then all remaining shares bear any additional loss on a pro rata basis.

Chart 8 shows the important effects of the strong MBR rule on the allocation of a fund's losses among investors with different recent redemptions. Here, we assume that all other redemptions from the fund are trivial, and we ignore the liquidity costs associated with lost access to shares in a closed fund. The fully redeeming investor is no longer at least as well off as the non-redeeming investor; for small losses in the fund, the fully redeeming investor (the thick red line) experiences much larger losses than the investor who redeems less (the thinner yellow line) or not at all (the very thin blue line). In this stark example, the fully redeeming investor loses her entire MBR for losses just slightly larger than 50 basis points because other investors do not redeem anything. By redeeming her entire available balance, she put her full MBR in a subordinated position, but without other redemptions, her MBR is the fund's total subordinated MBR.²⁰ (Below, we discuss some reasons to believe that this type of outcome would be extremely unlikely, but we also propose some adjustments that would make such stark allocations less feasible, including one that would eliminate subordination for most retail investors.)

Chart 8 illustrates that the strong MBR rule can provide a *disincentive* to redeem if the fund is expected to incur losses that exceed the NAV-stabilizing buffer but not the aggregate MBR buffer. In addition to providing stability to a fund by making redemptions (and consequent losses) less likely in a crisis, subordination of redeeming investors' MBRs ensures that they pay something for any liquidation costs that their redemptions impose upon the fund. Notwithstanding these benefits, if losses exceed a fund's aggregate MBR, then investors still would be better off redeeming.

Chart 9 adds consideration of the potential liquidity costs to investors who may have their money locked up in a closed MMF through a prolonged liquidation process. The dotted lines represent losses including these liquidity costs. As long as the investor's principal losses plus the liquidity costs of losing access to her shares do not exceed her MBR, she is better off if she does not redeem shares.

Chart 10 shows how an investors' losses depend on the behavior of other investors (for simplicity, we exclude the costs of lost liquidity). The thick red dashed and dotted lines

²⁰ In this example, because *no* other investor redeems, the redeeming investor is effectively leveraged by the ratio of the fund's assets to her assets when the fund's losses exceed its capital buffer, up until the point at which her subordinated MBR is absorbed.

show that a redeeming investor's losses are less extreme, for small losses to the fund, when other investors redeem as well. That is, when other investors redeem shares, they also contribute to the fund's subordinated MBR, so each investor owns a smaller share of the aggregate subordinated MBR.²¹ The thin blue dashed and dotted lines show that a non-redeeming investor's losses are also smaller when other investors redeem, because the subordinated MBRs of redeeming investors provide remaining investors some protection from losses.²²

The benefits of an MBR rule that subordinates redeeming investors' MBRs are evident in a comparison of charts 2 and 10. Chart 2 shows that under current rules, an investor's losses depend on other investors' behavior in a *destabilizing* way. The more other investors are expected to redeem, the greater the expected losses for non-redeeming investors. This dynamic provides a strong incentive to redeem at the first sign of trouble, before others do. In contrast, chart 10 shows that with a strong MBR rule, an investor's losses depend on other investors' behavior in a *stabilizing* way. Redeeming investors contribute to the fund's subordinated MBR and provide more protection for non-redeeming investors. This dynamic stabilizes the fund and benefits retail investors and others who may not redeem quickly when MMFs encounter problems. Even so, it is important to recognize that if expected losses are large enough (or if the MBR is too small), investors may be better off if they redeem shares.

5.4. Effective MBR rule. One possible concern with the strong MBR rule is that the size of the MBR itself is also the maximum amount of redemption-triggered subordination. This introduces a tension in determining the appropriate size for an MBR. A larger MBR provides more protection for investors by reducing the likelihood that a fund's losses would exceed the MBR, and thus lowers the chance that shareholders can impose losses on others by redeeming. But under the strong rule, a larger MBR could potentially cause redeeming investors to experience more disproportionate losses if a fund breaks the buck. Of course, this effect of subordination is a key part of providing incentives not to redeem when the fund is experiencing difficulties. Nevertheless, it may be possible to obtain most of the beneficial incentives of a larger MBR with a rule that would not have the potential to lead to more extreme outcomes.

²¹ That is, when many investors redeem, each individual redeeming investor's MBR is less highly leveraged than would be the MBR of a single redeeming investor.

²² Chart 10 can also illustrate the effect of relaxing our assumption that all investors are small. Consider, for example, the case in which no other investor redeems. Then, the loss function of an investor who owns 25 percent of all outstanding shares in the fund and who redeems all of her shares is shown by the thick red dashed line. Because she owns a greater portion of the fund, the subordination creates less effective leverage for her MBR than it would for the MBR of a smaller investor (in the limit, an investor who owns *all* of the shares in a fund is unaffected by subordination). The loss function for the same investor, if she does *not* redeem, is the thin solid blue line.

The effective MBR rule accomplishes these aims by limiting the amount of an investor's MBR that can be subordinated. Specifically, her subordinated balance would be determined as follows:

$$\text{Subordinated balance} = s \times \text{MBR} \times (\text{cumulative net redemptions} / \text{potential redemptions}).$$

Here, s is the fraction of an investor's MBR that can be subordinated (the strong MBR rule sets $s = 1$). If $s = 1/2$, for example, then a *fully* redeeming investor would have half of his MBR subordinated.

Chart 11 shows how investors' losses vary with the losses of an MMF with an effective MBR rule when other redemptions are trivial. Here, we set $s = 0.6$, so with an MBR of 5 percent, a maximum of 3 percent of an investor's reference amount can be subordinated. Comparing the thick red lines in charts 8 and 11 shows that a redeeming investor's losses increase more gradually under the effective MBR rule than with the strong rule. Hence, redeeming investors absorb a smaller share of the fund's losses under the effective rule.

Chart 12 illustrates how an investor's losses would depend on other investors' behavior. As was the case with the strong MBR rule, losses for each shareholder—the one who redeems everything (thick red lines) and the one who redeems nothing (thin blue lines)—are diminished as others' redemptions contribute to the fund's aggregate subordinated MBR. This dynamic provides stability to the fund; shareholders are better off if they do not redeem and others do. The disincentive to redeem, as indicated by the distance between the thick red and thin blue lines, is somewhat weaker with the effective rule than with the strong rule. But the effective rule results in a less uneven allocation of losses among investors.

5.5. Further refinements. The stark allocation of losses with the strong MBR rule when only one investor redeems shares from an MMF that subsequently breaks that buck (illustrated, for example, in chart 8) is admittedly an extreme example. Data from the Investment Company Institute on MMF redemptions show that monthly redemptions in MMFs from 2009 to 2011 averaged 45 percent of the funds' assets (ICI, 2012), so the likelihood that a 30-day period might pass in which only one investor (or a small number of investors) redeems shares seems remote.²³

Still, it is worthwhile to consider adjustments to the MBR rule that would protect investors who make incidental redemptions from triggering subordination of their MBRs. We consider a couple of potential adjustments.

The first adjustment would exempt the first \$50,000 of an investor's redemptions from subordination. Importantly, an investor who redeems less than \$50,000 would still be

²³ The likelihood of such an outcome could be estimated precisely only with account-level time-series data for individual shareholders. We know of no such data that are publicly available.

subject to the minimum balance at risk, but none of her MBR would be subordinated.²⁴ Specifically, that rule would be:

$$\text{Subordinated balance} = s \times \text{MBR} \times \left(\frac{\text{cumulative net redemptions} - \$50,000}{\text{potential redemptions}} \right).$$

for cumulative net redemptions in excess of \$50,000, and zero otherwise.²⁵

One consequence of this adjustment is that any investor with a balance of under \$50,000 would *never* have any portion of her MBR subordinated, so the adjustment would exempt many retail MMF accounts from subordination. This may be appropriate, given that retail investors are generally much less prone to run from distressed MMFs than institutional investors.²⁶ Of course, a \$50,000 exemption for redemptions would reduce protections for non-redeeming investors, especially in retail MMFs. However, investors would continue to be protected by the MBR itself, which would ensure that redeeming investors share proportionally in any losses of principal as long as losses do not exceed the fund's aggregate MBR.

A second possible adjustment to the MBR subordination rules would put a small fraction of *every* investor's MBR in a first-loss position, regardless of his redemption behavior. Redemptions would increase the share of an investor's MBR that is subordinated. The purpose of such an adjustment would be to reduce the effective leverage of investors who redeem. For example, with a 5 percent MBR, if 10 percent of every investor's MBR (0.5 percent of her reference amount) is automatically subordinated, then even if no one else redeems, a fully redeeming investor would be effectively leveraged 10 to 1 (recall that

²⁴ As described throughout this paper, although an MBR without subordination might do little to brake runs, particularly by institutional investors, the MBR itself provides important protections for retail investors and others who are less likely to redeem shares in a crisis. Allowing an exemption from the *MBR itself* probably would lead some MMFs to attempt to avoid implementing *any* MBR by limiting all balances to \$50,000 or less. Ironically, this outcome might lead to small (primarily retail) investors enjoying less protections than larger institutional investors.

²⁵ We are grateful to Eric Rosengren for suggesting, during discussions of the MBR concept with staff at the Federal Reserve Bank of Boston, the possibility of a de minimis threshold to which subordination would not apply. Importantly, as formulated here, this rule would not allow an investor to avoid subordination by making a series of redemptions, each of which is just under \$50,000, because the investor's cumulative net redemptions are her reference amount less her current balance. How she arrives at her current balance is immaterial.

²⁶ Retail MMFs were much less likely than their institutional counterparts to suffer very large outflows during the run in 2008 (PWG, 2010; McCabe, 2010) and retail funds experienced relatively small net outflows in the summer of 2011, when concerns about MMFs' European exposures and the possible consequences of a breakdown in debt-ceiling negotiations triggered large redemptions from institutional MMFs (see, for example, Chernenko and Sunderam, 2012). Indeed, differences in the riskiness of retail and institutional funds led the SEC to consider different liquidity requirements for retail and institutional funds in its proposed amendments to rule 2a-7 (SEC, 2009). Importantly, however, many MMFs have both retail and institutional investors—this was one factor that led the SEC not to adopt a rule that distinguished retail and institutional funds (SEC, 2010).

without such an adjustment, his effective leverage would be the ratio of the fund's assets to his assets, which is 1000 to 1 in our example).

Chart 13 depicts investors' losses (assuming that other redemptions are trivial) from a fund with an effective MBR rule combined with a minimum subordination rule. This combination implies that the subordinated portion of each investor's MBR ranges from 10 percent to 60 percent (based on $s = 0.6$) of his total MBR. A comparison of the thick red line in this chart with that in chart 8 (the strong MBR rule) shows that the adjustments shown in chart 13 cause the fully redeeming investor's losses to grow more slowly as the fund's losses increase than when the fund employs the strong MBR rule. Hence, the adjustments weaken the disincentive to redeem. Still, as chart 13 shows, the rules would preserve the primary benefits of the MBR.

6. Investors' expected losses

Our analyses in section 5 of the allocations of an MMF's losses among its investors are informative about incentives and how investors fare for a range of realized MMF losses. We show, for example, that investors are unambiguously better off redeeming from a distressed fund that has no MBR or one that has only a simple MBR rule. However, for the strong and effective MBR rules, investors' incentives to redeem from a distressed fund depend critically on the likelihood of possible outcomes, since the net benefits of redeeming change with the scale of a fund's realized losses. When a fund's losses are smaller than its aggregate MBR, non-redeeming investors generally are better off than redeemers, but relative outcomes reverse when losses exceed the aggregate MBR. Moreover, the effects of other shareholders' redemptions on an investor's losses depend on whether the fund's losses exceed its aggregate MBR.

In this section, we use empirical evidence on historical realized losses in MMFs—including a new data set that researchers have not previously examined—and the opportunity costs of lost liquidity during crises to examine investors' *expected* losses and thus their incentives.²⁷ Our objectives are conceptual as well as empirical. By showing how an investor's expected losses may vary with his own redemptions and those of others, we provide additional insights on investors' incentives under different MBR rules. Moreover, the empirical nature of our analysis is informative about the appropriate size of an MBR for mitigating the risk of runs, which we discuss in section 7.3.

We begin with a discussion of the empirical evidence that is available to help calibrate expected losses in a distressed MMF, including new data on funds' real-time losses in September and October 2008. We also examine evidence on the possible magnitude of the opportunity costs of lost liquidity for investors who lose access to shares in a closed MMF. Next, we use the empirical evidence on fund losses and liquidity costs as a starting point for assessing how an investor's expected losses might vary with his own redemptions and with

²⁷ Our focus on expected losses would be consistent with an assumption that investors are risk neutral. Explicitly considering risk aversion would complicate the analysis without modifying the key insights.

those of others under various MBR rules. While these *quantitative* results depend on our empirical evidence and assumptions about the distribution of fund losses in a crisis, most of the *qualitative* properties of investor incentives are not sensitive to assumptions about a fund's expected losses or the liquidity cost to investors in a closed MMF.

6.1. Expected losses in a distressed MMF. To be effective in curbing runs on MMFs, an MBR rule must brake redemptions from distressed funds at precisely the times when investors would otherwise have the strongest incentives to redeem. As the model in section 4 illustrates, when an MMF encounters strains and the likelihood of losses increases, investors' decisions to redeem or not depend on how the choice affects their expected losses (including the value of lost liquidity) *in the event that the fund breaks the buck and closes*. That is, for a fund in distress, conditional expectations are critical, and investors will redeem when doing so reduces their expected principal and liquidity losses in a break-the-buck scenario (when $E(B) > 0$).

Our analysis in section 5 shows how the losses of redeeming and non-redeeming MMF investors vary with a fund's realized losses under various MBR rules. But evaluation of investors' *expected* losses, conditional on a fund's closure, requires information about the (conditional) distribution of MMF losses when funds break the buck.

6.1.1. Previously available data on MMF losses. Although hundreds of MMFs have suffered material losses over the last 30 years, information on the magnitude of those losses is not readily available. The historical record of losses in MMF portfolios has been obscured by fund sponsors' longstanding practice of providing discretionary financial support to funds that would otherwise have been in danger of breaking the buck. The record of sponsor support for MMFs is quite extensive. For example, SEC Chairman Schapiro has reported that sponsors have intervened more than 300 times to support MMFs since the funds were introduced in the 1970s (Schapiro, 2012). Moody's found 144 cases in which U.S. MMFs received support from sponsors from 1989 to 2003 (Moody's, 2010). The SEC reported that 44 MMFs received support because of investments in distressed asset-backed commercial paper in the wake of the ABCP crisis that began in mid 2007, and the Commission noted that 100 MMFs received support in September and October 2008 (SEC, 2009). MMF sponsors reportedly intervened as recently as November 2011 to support their MMFs (Henriquez, 2012).

The MMF industry's reliance on sponsor support poses considerable systemic risk, however, because implicit support opens channels for transmission of strains between the funds and their sponsors that can be destabilizing in a crisis (McCabe, 2010; PWG, 2010). If an MBR rule is intended to provide stability to MMFs that is *not* predicated on discretionary support from sponsors, the rule should be robust to the types of MMF losses that would have occurred in the absence of such support.

However, little data are available on MMF losses that might have occurred in the absence of sponsor support. Prior to the SEC's 2010 amendments to MMF rules, there was no requirement that sponsors disclose support actions publicly. Some incomplete information can be garnered from published "no-action letters" that the SEC issued to permit specific

sponsor interventions, as well as from the financial statements of MMFs and their publicly-traded sponsors. But these sources usually do not provide enough information to estimate the magnitude of NAV decline that would have occurred without intervention.²⁸ The SEC's 2010 rule amendments for MMFs enhanced requirements for disclosure of support, and MMFs must now report at each month end what their mark-to-market NAVs would be without support.

The two U.S. money market funds that have broken the buck since the adoption of rule 2a-7 in 1983 are one source of information about MMF losses in the absence of support. The Community Bankers U.S. Government Money Market Fund lost 3.9 percent of its value in 1994 (SEC, 1999), and the Reserve Primary Fund lost approximately 1.6 percent in 2008 (SEC 2009c).²⁹ Research by staff at the Federal Reserve Bank of Boston found that in addition to guarantees from sponsors, 47 MMFs between 2007 and 2010 benefitted either from a direct cash contribution from a sponsor or a sponsor's purchase of distressed assets from the fund. These included nine cases in which estimated support exceeded one percent of the funds' NAVs: "four cases between 1 and 2 percent of assets under management, two cases between 2 and 3 percent, and three cases over 3 percent" (Rosengren, 2012).³⁰ Analysis by Moody's of MMFs' exposures to Lehman at the time of its bankruptcy showed 15 MMFs held Lehman obligations that ranged from 0.25 percent to 5.6 percent of the funds' assets and averaged 1.9 percent of assets—so, Reserve's position (1.2 percent) was less than average among funds that held Lehman (Moody's 2012).

²⁸ The SEC has posted many of the recent no-action letters for MMFs at: <http://www.sec.gov/divisions/investment/im-noaction.shtml>. Most of the no-action letters and financial statements provide, at best, information about the total amounts of distressed assets purchased from funds or the notional size of guarantees provided for distressed assets. Another hurdle to estimating the per-share value of support actions is the fact that sponsor financial statements often do not specify which of the sponsor's funds received support.

²⁹ The Primary Fund's losses were caused largely by its \$785 million in holdings of Lehman debt obligations (1.2 percent of the fund's assets) at the time of Lehman's bankruptcy. RCMI, the adviser to the Reserve Primary Fund, announced at about 4 pm on Tuesday, September 16, 2008, that the NAV of the fund's shares had dropped 3 percent to 97 cents (SEC, 2009a), presumably because large redemptions had further eroded the NAV. However, a June 2009 SEC order regarding distribution of the Primary Fund's assets indicated that the value of those assets ultimately would allow shareholders receive 98.4 cents per share (SEC, 2009c), although this amount includes income earned on the assets after the Primary Fund was closed. Some Reserve statements (for example, "Including this seventh distribution, \$50.7 billion, or approximately 99.04% of Fund assets as of the close of business on September 15, 2008, will have been returned to investors" (Anderson, 2010)) have referred to the portion of the fund's assets—some of which lost significant value—that have been returned to investors as part of the fund's liquidation process, rather than to investors' principal losses as claimants upon those assets.

³⁰ The Federal Reserve Bank of Boston results reported here are preliminary. In addition, Moody's (2010) found that between 2007 and 2010, 21 MMF sponsors spent at least \$12.1 billion to maintain funds' stable NAVs, although the figure includes support for some European funds and some non-registered investment pools. However, Moody's totals were not stated as shares of supported funds' NAVs.

6.1.2. New data on MMF losses. We analyze a new data set that provides additional evidence on the scale of losses that have occurred in MMFs. The data come from the Department of the Treasury and the SEC, which collected information about certain MMFs that participated in the Treasury's Temporary Guarantee Program for Money Market Funds. Specifically, MMFs with shadow NAVs below \$0.9975 were required to report information about their portfolios, including what their NAVs would have been without sponsor-provided guarantees, such as capital support agreements (U.S. Treasury, 2008). One shortcoming in the data, however, is that the effects of some forms of sponsor support, such as direct cash infusions to a fund and outright purchases of securities from a fund at above-market prices, are *not* excluded from reported shadow NAVs. In addition, the data do not reflect portfolio losses that might have occurred in the absence of extensive government support in 2008 for MMFs, short-term funding markets, and some financial institutions. As such, the NAV data do not reflect the full extent of losses that might have occurred without sponsor and government interventions.

We obtained data from the U.S. Treasury and the SEC on the shadow NAVs of all reporting MMFs (but not the funds' identities) for the days from September 5 to October 17, 2008.³¹ Over this period, the number of daily reports ranges from four to 63. Because of a requirement that funds with NAVs below \$0.9975 report at least weekly, the number of observations jumps on each Friday, and the number of funds reporting shadow NAVs on Fridays from September 5 to October 17 ranges from 19 to 63 (see line 1 of table 2).

As shown on line 3 of the table, over this period, as many as 11 MMFs on any particular day reported that they would have broken the buck without guarantees from sponsors (including as many as five funds that reported shadow NAVs of less than \$0.995 *before* the Lehman bankruptcy).³² A total of 29 funds reported shadow NAVs that would have fallen below \$0.995 at some point during this episode.³³ Average shadow NAVs, excluding the effects of guarantees, for *all* reporting funds dropped to as low as \$0.993 on October 3 and October 17 (line 4). Perhaps more important for an investor who is concerned about expected losses in a break-the-buck event, average shadow NAVs of funds that would have broken the buck without guarantees from sponsors fell as low as \$0.975 on October 3 (line 6). Even excluding two outliers (funds that reported shadow NAVs of \$0.903 and \$0.929, as shown on line 8), the average shadow NAV of funds that would have broken the buck was as low as \$0.977, on October 17 (line 7).

³¹ The Treasury's Temporary Guarantee Program was only announced on September 19, but some funds provided data for shadow NAVs on dates before the program's inception.

³² A number of MMFs still had embedded losses due to holdings of ABCP that had lost value beginning in August 2007.

³³ This total could exclude MMFs that would have broken the buck without sponsor support in the form of direct cash infusions to the fund or outright purchases of securities from the fund at above-market prices.

Table 2. MMFs' reported shadow NAVs (excluding effects of guarantees from sponsors), September 5 to October 17, 2008*

Date	Sept. 5	Sept. 12	Sept. 19	Sept. 26	Oct. 3	Oct. 10	Oct. 17
1. Number of reporting funds	19	31	46	46	50	63	28
2. With shadow NAV < \$1.00	13	31	46	46	50	62	28
3. With shadow NAV < \$0.995	3	5	8	11	9	11	6
Average shadow NAVs (dollars, excluding effects of guarantees from sponsors)							
4. All reporting MMFs	0.9967	0.9963	0.9950	0.9937	0.9929	0.9936	0.9927
5. Shadow NAV < \$1.00	0.9952	0.9963	0.9950	0.9937	0.9929	0.9935	0.9927
6. Shadow NAV < \$0.995	0.9867	0.9894	0.9850	0.9844	0.9747	0.9775	0.9774
7. \$0.95 < shadow NAV < \$0.995	0.9867	0.9894	0.9850	0.9844	0.9836	0.9824	0.9774
8. Minimum shadow NAV	0.9800	0.9801	0.9788	0.9750	0.9034	0.9285	0.9637

*Reported shadow NAVs may include some forms of sponsor support, such as direct cash infusions to the fund or outright purchases of securities from the fund at above-market prices.

Source: U.S. Department of the Treasury and SEC.

6.1.3. What might investors expect? Ultimately, the effectiveness of an MBR rule in stemming runs during a crisis would depend, in part, on investors' expectations of losses in distressed funds. The data we have reviewed suggest that expected (average) losses of about 2 percent for MMFs that break the buck would not be unreasonable.

Clearly, there are multiple challenges in interpreting the data on past losses in MMFs and using them to forecast investors' conditional expectations for losses in the future. One drawback of the SEC-Treasury shadow NAV data, as noted above, is that they do not reflect all portfolio losses (for example, after a sponsor's outright purchase of securities from a fund's portfolio at above-market prices). At the same time, MMFs' historical losses have reflected not only declines in the value of portfolio assets but also the effects of heavy redemptions, which may have forced funds to suffer additional losses through fire sales of securities and may have concentrated losses on smaller numbers of shares. To the extent that an MBR makes heavy redemptions less likely, these components of historical losses might be less relevant for future losses with a robust MBR rule.

In addition, the losses that we have documented here mostly occurred during a period of very serious strains both for MMFs and for the entire financial system. During that episode—particularly in September and October 2008—market quotes were unreliable, MMFs' shadow NAVs were difficult to compute, and the SEC provided pricing relief that may have boosted shadow NAVs (SEC, 2008). Unfortunately, we have even less information about the losses that MMFs suffered in other episodes, such as the ABCP crisis that began in 2007. While one might hope that the types of losses that MMFs experienced in 2008 will never be seen again, we believe that those losses may serve to inform investors' expectations about future losses.

6.2. Liquidity costs of losing access to shares in a closed MMF. When an MMF breaks the buck and closes, its shareholders suddenly lose access to their cash for what may be a prolonged liquidation process. Indeed, the Reserve Primary Fund, which broke the buck on

September 16, 2008, still had not completed the distribution of all of its assets to shareholders as of this writing over three years later (Schapiro, 2012). Clearly, a fund's closure is costly to shareholders who lose liquidity.

There are many challenges in estimating this cost and thus in assessing the strength of shareholders' incentives to redeem shares in a distressed fund to avoid lost liquidity. The opportunity cost to shareholders of an unplanned loss of access to cash could depend on their individual circumstances. For investors whose cash needs can be met by holdings outside the closed fund, the costs may be small. For those with immediate cash needs, the opportunity cost may be the cost of liquidating other assets or the cost of borrowing funds (see, for example, McCoy, 2008a and 2008b), although such investors might not need immediate access to their entire MMF balances. In addition, the costs of losing access to cash would depend on broader financial conditions; liquidity may be inexpensive during "normal" times but very costly during periods of financial stress, when MMFs are most likely to be under strain.

An additional factor in assessing the cost of lost liquidity in a closed MMF is the length of time over which investors are likely to lose access to their cash. That interval could extend for years, as the Reserve Primary Fund example shows. More generally, if the process of disposing securities proceeds smoothly, shareholders might expect to receive cash as the securities held by their fund mature, although an MMF with a strong or effective MBR rule would not be able to distribute cash to shareholders with subordinated claims until all other claims have been paid.³⁴ Among prime MMFs, the weighted average life (WAL) of portfolio holdings averaged 76 days in 2011, while the WAL for securities outside of funds' buffers of weekly liquid assets averaged 114 days, and the maximum maturity of any security (excluding government securities) held by each individual fund averaged 327 days.³⁵

We use a couple of approaches to estimate the value of shareholders' lost liquidity. First, we examine the costs that an MMF investor might face in raising cash from outside sources to replace money that is temporarily stuck in a closed fund. Second, we review evidence on liquidity premiums for financial instruments and the potential costs of liquidating securities, particularly during periods of financial strain. We assume that investors lose access to their cash for six months.

Our first measure is the average net cost to an investor who must borrow from a bank to replace cash in a closed MMF. We assume that the investor takes out a business loan and that the net cost is the rate on that loan less the yield on the investor's MMF shares before they are liquidated. Hence, we take the average rate on commercial and industrial loans

³⁴ If subordinated claims are paid before all other claims are made whole and one of the fund's portfolio assets subsequently loses value, shareholders with subordinated MBRs ultimately could receive a greater portion of the fund's assets than they are entitled to under the MBR rule.

³⁵ Data are from MMFs' filings of SEC form N-MFP.

with 31 to 365 days to maturity from the Federal Reserve’s Survey of Terms of Business Lending and subtract the asset-weighted net yield of prime institutional funds. This net-cost measure is reported, at an annual rate, on line 1 of table 3. The first column shows that, on average from 2005 to 2011, the net cost was 2.25 percent. The second and third columns show that this measure of net cost was actually a bit smaller during the financial crisis in 2007 and 2008.

Table 3. Some indicators of the liquidity costs to investors in closed MMFs

Indicator	(1) Average, 2005-2011 (unless noted)	(2) Early- crisis period	(3) Severe- crisis period
<u>A. Costs (percent, annual rate unless noted)</u>			
1. Net cost of bank loan	2.25	1.66	2.08
2. Liquidity component of Libor-OIS spread	0.31	0.59	2.25
3. OIS-Treasury bill spread	0.28	0.62	0.89
4. Treasury note-bond premium (percent)	0.79	0.86	1.46
<u>B. Time intervals</u>			
5. Net cost of bank loan		July 2007- Sept. 2008	Oct. 2008- Dec. 2008
6. Liquidity component of Libor-OIS spread	July 2005- June 2010	Aug. 8, 2007- Sept. 15, 2008	Sept. 16, 2008- Oct. 17, 2008
7. OIS-Treasury bill spread		Aug. 1, 2007- Sept. 12, 2008	Sept. 15, 2008- Oct. 17, 2008
8. Treasury note-bond premium	Feb. 2005- June 2010	Aug. 1, 2007- Sept. 12, 2008	Sept. 15, 2008- Oct. 17, 2008

Sources. Net cost of bank loan: Federal Reserve Survey of Terms of Bank Lending, iMoneyNet, and authors’ calculations; liquidity component of Libor-OIS spread: Poskitt, 2011; OIS-Treasury bill spread: Bloomberg and U.S. Treasury; Treasury note-bond premium: Musto, Nini, and Schwarz, 2011 and additional data provided by Greg Nini.

A second approach to estimating liquidity costs focuses on liquidity premiums for financial instruments, particularly during periods of financial strain that would probably characterize an episode when investors were considering redeeming MMF shares to avoid losses. A drawback of this approach is that during these episodes, credit and liquidity risks become particularly difficult to disentangle, and researchers generally must focus on the least risky instruments and markets to isolate liquidity premiums.

One such approach aims to extract the liquidity component of the spread between the three-month London Interbank Offered Rate (Libor) and overnight indexed swaps (OIS).³⁶

³⁶ Spreads at six-month maturities would more appropriate as measures of liquidity costs to investors who lose access to cash for six months, but six-month interbank lending is reportedly nonexistent during periods of even moderate stress.

The liquidity-premium portion of this spread shows the degree to which liquidity costs drive three-month interbank funding costs above the expected average of overnight rates (Bank of England, 2007; Poskitt, 2011). Line 2 reports averages of the estimated liquidity premium for three-month maturities from Poskitt (2011) for different episodes. The average from mid-2005 to mid-2010 was 31 basis points. But the premium rose to as much as 2.25 percent during the severe phase of the crisis in 2008, when the disastrous run on MMFs occurred.

A related approach examines the premium that investors are willing to pay (in terms of lower yields) to own very highly liquid assets, such as Treasury bills, rather than receive the expected short-term interbank rate, as measured by OIS. Line 3 of the table reports that this spread averaged 28 basis points from 2005 to 2011, but widened in periods of financial stress as Treasury bill rates fall well below OIS rates. Over the month following Lehman's bankruptcy and the Reserve Primary Fund's announcement that it had broken the buck (column 3), the spread averaged 89 basis points.

Still another approach examines the *price* differential between relatively more liquid Treasury notes and less liquid Treasury bonds with the same maturities (Musto, Nini, and Schwarz, 2011). That measure, which might be understood as the instantaneous cost of switching from highly liquid instruments to very highly liquid instruments, averaged 79 basis points between 2005 and 2010, but rose to about 1.5 percent during the crisis period that included the run on MMFs.³⁷

Many of the estimated annual liquidity costs reported in table 3 probably would understate opportunity costs to most MMF investors. One exception is net borrowing costs, shown on line 1, which might be lower for the subset of large corporations with top-tier short-term ratings. (For retail investors, who cannot obtain C&I loans, net costs could be much higher since borrowing costs likely exceed C&I rates, and net yields on retail MMFs are generally lower than those of institutional funds.) The liquidity-premium estimates shown on line 2 are for market participants that generally have access to government-provided liquidity (such as the Federal Reserve's discount window), and hence would understate liquidity costs for most investors. And the Treasury note-bond premium on line 4 captures the liquidity premium only for very highly liquid assets (Treasury notes) compared to highly liquid assets (Treasury bonds). For investors who must sell less-liquid assets to raise cash, liquidity costs presumably would be much larger.

In summary, the opportunity costs of lost liquidity to investors in a closed MMF are probably material and significant, although they are difficult to measure. Based on the indicators we have reviewed here, we believe that an MBR should be designed to withstand redemption pressures that incorporate investors' incentives to avoid lost-liquidity costs of

³⁷ Musto, Nini, and Schwarz found that this cost continued to rise after October 17, 2008, and reached as high as 5 percent later that year.

at least 50 basis points. That would correspond, for example, to an investor paying a net rate of 2 percent on a loan for half of her MMF balance for a six-month period, or paying a premium of 100 basis points to liquidate assets to replace half of that balance immediately to raise cash. Alternatively, 50 basis points would be the cost of paying a 1 percent liquidity premium for six months. However, given the uncertainties in the indicators of liquidity costs, our sensitivity analysis considers opportunity costs as low as 25 basis points and as high as 1 percent.

6.3. Modeling investors' expected losses in funds with MBR rules. To illustrate how MBR rules affect investors' expected losses, we assume that when a fund gets into trouble, losses are distributed exponentially with an expected value of 1.5 percent. This assumption implies that losses for funds that break the buck average 2 percent, which is roughly consistent with the empirical evidence we reviewed in section 6.1.³⁸ Importantly, this also implies that investors recognize that losses could possibly exceed the size of the MBRs that we analyze. In addition, as outlined in section 6.2, we assume that the liquidity costs to investors who lose access to shares in a closed fund total 50 basis points of the value of their shares (although not all the examples below include that cost). We note below how these assumptions affect our illustrations. Also, as in section 5, we assume that the fund has an NAV-stabilizing buffer of 50 basis points and that the investors we study each hold 0.1 percent of the MMF's shares.

Chart 14 shows expected losses for an investor in an MMF with *no* MBR but with a 50 basis point buffer. The chart plots the relationship between the investor's expected losses and her own redemptions, under different assumptions about other investors' redemptions. The solid blue line shows her expected losses if other investors do not redeem shares, the dashed line shows expected losses if others' redemptions are 25 percent of the fund's assets, and the dotted line shows the effect of others redeeming 50 percent. For simplicity, we exclude from these expected-loss functions the costs of lost liquidity for investors in closed funds.

The chart demonstrates the strong incentives that investors have to run from a distressed MMF under current rules, even with the small capital buffer. The more an investor redeems, the smaller her expected losses—and losses are zero if she redeems everything. Moreover, incentives to redeem are magnified by the costly effect of others' redemptions. As those redemptions increase, the shareholder's losses increase, *unless she redeems all of her shares*.

These basic properties of investors' incentives are *not* sensitive to our assumptions about expected fund losses, as long as there is some material possibility that the fund's losses might exceed its capital buffer (if any). But our assumptions do allow us to illustrate these properties intuitively. For instance, the solid blue line indicates that if the shareholder

³⁸ That is, we assume that conditional on loss exceeding 0.5 percent, the expected loss is 2.0 percent.

redeems nothing, her expected loss is only 1.1 percent when other investors do the same.³⁹ But the dotted line shows that her expected loss jumps to 2.1 percent when others redeem half of the fund's assets and concentrate its losses upon her.

The effect of introducing an MBR is illustrated in chart 15, which plots investors' expected losses for a fund with an MBR of 5 percent under the simple rule described in section 5.2. The thin blue lines exclude investors' liquidity costs due to fund closures. (The dashed vertical line at 95 percent redemptions indicates that investors cannot redeem more than this amount.) Introducing the MBR has several important effects on investors' expected losses. Notably, these effects qualitatively are not dependent on our basic assumptions about expected losses, although those assumptions do affect the quantities shown on the charts.

First, the benefits of the investor's own redemptions are much reduced. Although both of the thin blue lines in chart 15 are still downward sloping, they are much flatter than their counterparts in chart 14, and the investor is unable to reduce expected losses to zero with any amount of redemptions. Thus, the allocation of losses is much more equitable between redeeming and non-redeeming shareholders. Second, the MBR diminishes the effects of others' redemptions: The thin blue lines in chart 15 are much closer together than those in chart 14. If other investors increase their redemptions from nothing to 50 percent, the expected loss for a non-redeeming investor only increases from 1.07 percent to 1.11 percent. That is, the simple MBR rule internalizes much of the externality imposed by redeeming investors on those who do not redeem.

Nevertheless, with the simple MBR rule, all expected-loss functions are still downward sloping. This property depends only on there being *some* possibility of losses in excess of the sum of a fund's capital buffer and its aggregate MBR. Because the likelihood of that outcome is quite small with a 50 basis point capital buffer, a 5 percent MBR, and MMF expected losses of 1.5 percent, the slopes of the expected-loss functions are slight. Still, investors have an incentive to run from the fund; doing so reduces their expected losses, regardless of what other investors do. And although others' redemptions are far less damaging than without an MBR, those redemptions are still costly. (The concentration of losses caused by others' redemptions in this example only occurs when the fund's losses exceed 5.5 percent, which is unlikely, so the effect of others' redemptions on *expected* losses is small.) As a result, with a simple MBR rule, all investors have an incentive to run if they expect that the fund might suffer losses or that other investors' redemptions might lead to losses.

The motive to preserve liquidity by redeeming from a fund before it closes contributes significantly to investors' incentives to redeem, even with the MBR rule. The thin blue lines in chart 15 omit the effects of shareholders' lost liquidity when a MMF closes. The thick

³⁹ Recall that expected losses for the fund are 1.5 percent, and the capital buffer absorbs the first 50 basis points of any loss for the fund. Expected losses for the investor when *no one* redeems are $E(\max(\text{fund loss} - 50 \text{ bps}, 0))$.

red lines factor in the value of lost liquidity, as discussed in section 6.2. Including liquidity costs causes the expected-loss curves to tilt more steeply downward, indicating that liquidity preservation strengthens the incentive to run considerably.

Although the degree of steepening of the expected-loss curves depends on our estimate of the liquidity cost due to fund closure, *any* liquidity cost will steepen these curves. Importantly, even if there is no possibility of a loss that exceeds the combination of the fund's capital buffer and its aggregate MBR, a small liquidity cost would still leave investors with downward-sloping expected-loss curves. That is, with a simple MBR rule, fear of lost liquidity may drive investors to redeem, even if losses are not expected to be large.

The strong and the effective MBR rule have a dramatic effect on investors' incentives to redeem. The thin red and thick yellow lines in chart 16 show expected losses for an MMF with a strong rule and an effective rule, respectively. Under the effective rule considered here, subordination is limited to 60 percent of the shareholder's MBR (as in section 5.4). All of the expected losses shown here incorporate the liquidity costs to investors in closed funds.

Under both rules, the investor's expected losses *increase* as her redemptions increase, regardless of the assumptions here about other investors' redemptions.⁴⁰ That is, whether other investors redeem nothing at all or 50 percent of their shares, the investor minimizes her own expected losses by staying in the fund. Whereas investors in a distressed fund without an MBR or with a simple MBR rule have strong incentives to run, *not* running is advantageous for an investor in a fund with a strong or effective MBR rule. Indeed, these rules make *not* running a new equilibrium strategy for investors.

Moreover, other investors' redemptions *reduce* the expected losses for the investor. This again contrasts with what happens without an MBR or under the simple rule, as the vertical ordering of the curves that condition on other investors' redemptions is inverted in chart 16 compared with that in charts 14 and 15. That is, because redeeming shareholders' MBRs become subordinated, they provide some loss-absorbing protection for non-redeeming shareholders that compensates for the cost that redemptions impose on the fund. As a result, the more other investors redeem, the smaller the losses for an investor who does not join the run.

The effective rule preserves the qualitative features of the strong rule: The investor's expected losses increase with his own redemptions (so, the best course of action is not to

⁴⁰ Depending on the size of the MBR and assumptions about a fund's expected losses and investors' liquidity costs, very large redemptions by other investors can cause an investor's expected-loss curve to tilt downward, even with a strong or effective MBR. That is, very large redemptions by other investors can cause an investor to have an incentive to run. We discuss this possibility below. Still, a critical outcome here is that the strong and effective MBR rules can create a new equilibrium in which no investor has an incentive to redeem. Importantly, very extreme scenarios for others' redemptions probably are not feasible, since a distressed fund that faced redemption requests of virtually all of its shares most likely would be forced to close rather than honor those redemptions.

redeem), and expected losses decrease as other investors' redemptions grow. However, both effects are weakened. For example, when other investors do not redeem, an investor who redeems all of his available balance has an expected loss of 3.6 percent under the strong rule (substantially more than the fund's expected loss net of what is absorbed by the NAV-stabilizing buffer), but that same investor would expect to lose only 2.6 percent under the effective rule. That is, the effective rule reduces the expected cost of redeeming shares from a distressed fund, although—given our assumptions about expected fund losses and liquidity costs, as well as the parameters of the MBR rule and capital buffer analyzed here—the effective rule does preserve the incentive not to run.

The basic properties of the strong and effective rules that are illustrated by the comparison of chart 16 to charts 14 and 15 are not sensitive to our assumptions about the fund's expected losses. The strong rule always increases the slope of the investor's expected-loss curve, and the rule generally makes that slope positive, unless expected losses are very large or the MBR is small. Similarly, under the strong rule, other investors' redemptions reduce an investor's expected loss unless the fund's losses are expected to be large relative to its capital buffer and MBR. Since the effective rule lies somewhere between the strong and simple rule, its impact on investor incentives relative to the strong or simple rules will depend on the fraction of the investor's MBR that can be subordinated.

Chart 17 illustrates a tradeoff in the design of an MBR rule: Similar disincentives for redemptions can be achieved with different combinations of MBR size (m) and maximum potential subordination (s). The thick red lines in the chart show an investor's expected losses (including the liquidity costs of losing access to shares in a closed fund), as a function of her own redemptions from an MMF with a strong MBR rule and an MBR of just 3 percent. The thick red solid line plots expected losses when other shareholders do not redeem, while the thick red dotted line shows expected losses when others redeem 50 percent of their shares. Both lines slope up, indicating that—under our assumptions about expected losses in the fund, the costs of losing access to shares in a closed fund, and the 50 basis point capital buffer—this rule would still provide a disincentive for redemptions.

The thin blue lines in chart 17 show expected losses for an MBR of 5 percent with maximum subordination of 2 percentage points of the MBR ($s = 0.4$). The thin blue solid and dotted lines show expected losses when others' redemptions total zero and 50 percent of assets, respectively. Again, both lines slope up. Indeed, the similarity of the thin blue and thick red lines indicates that the strong and effective rules illustrated here provide similar incentives to shareholders. Hence, either combination—a strong rule with a smaller MBR (the thick red lines) or an effective rule with a larger MBR (the thin blue lines)—could be utilized to provide roughly the same disincentive to redeem.

In chart 18, we explore the effect of varying the size of an MBR in a fund with a strong MBR rule, given our baseline assumptions that the fund has a 50 basis point capital buffer, the MMF's expected losses are 1.5 percent, and liquidity costs to investors in a closed fund are 50 basis points. Here, we assume that other investors do not redeem. The slopes of the expected-loss curves steepen (become more positive) as the MBR is scaled up. For MBRs of 1 or 2 percent, expected-loss curves slope down, indicating that *even with a strong MBR*,

investors may have an incentive to run from a distressed MMF if the MBR is too small. When the MBR is 3 percent (and the fund has a 50 basis point capital buffer), the expected-loss function slopes slightly upward. MBRs of 4 or 5 percent provide more convincing disincentives to redeem.

Chart 19 repeats the exercise depicted in chart 18 but shows expected losses under the assumption that other investors' redemptions total 50 percent of assets. The basic results are similar to those shown on chart 18: When the MBR is 2 percent or less, an investor's redemptions can *reduce* her expected losses, so she has an incentive to run. Under our assumptions about the capital buffer, the fund's expected losses, and shareholders' potential liquidity costs, the disincentive to redeem appears with an MBR of 3 percent or more.

A comparison of charts 18 and 19 also shows that when the MBR is too small, other investors' redemptions increase an investor's expected losses. For example, with an MBR of 1 percent, expected losses depicted in chart 19 are generally larger than those in chart 18. An interesting feature of chart 19 is that the expected-loss curves all cross at redemptions of 50 percent: At this point, *every* investor is redeeming 50 percent, so all experience expected losses of 1.2 percent.

Charts 20 and 21 show the effects of varying our assumptions about expected MMF losses and liquidity costs to investors in a closed fund. Both charts depict investors' expected-loss functions for various sizes of strong MBRs (and a 50 basis point capital buffer), under the assumption that other investors' redemptions total 50 percent of fund assets. Chart 20 illustrates the effects of assuming that the fund's expected losses and investors' liquidity costs are *smaller* than our baseline assumptions, and chart 21 shows results when expected losses and costs are *larger* than baseline.

Specifically, in chart 20, we plot investors' expected losses when MMF losses are distributed exponentially with an expected value of 1 percent (so, conditional on a fund experiencing losses in excess of 50 basis points, expected losses are 1.5 percent) and the liquidity cost to investors of losing access to shares in a closed MMF is 25 basis points. Under these assumptions, for a fund with a strong MBR rule and a 50 basis point capital buffer, as long as the MBR is at least 2 percent, investors face a disincentive to redeem.

Chart 21 shows investors' expected losses when MMF losses are distributed exponentially with an expected value of 2 percent and the liquidity cost of losing access to shares in a closed MMF is 100 basis points. Here, the MBR must be at least 4 percent to provide a (slight) disincentive to run.

Finally, chart 22 returns to our baseline assumptions about the fund's expected losses and investors' liquidity costs, but shows investors' expected losses when others' redemptions total 80 percent. Such redemptions would be quite extreme and may not even be feasible, since a distressed fund that faced redemption requests of 80 percent of its assets likely

would be forced to close rather than honor those redemptions.⁴¹ Still, a comparison of charts 19 and 22 indicates that very heavy redemptions by other investors can undo the disincentive to redeem created by the strong MBR rule. For example, when others' redemptions total 50 percent of the fund's assets, a 3 percent MBR can create a disincentive to redeem (chart 19). But when others' redemptions are 80 percent of assets, the investor's expected loss function for a 3 percent MBR is downward sloping (chart 22). That is, she is better off redeeming.

7. Basic elements of an MBR rule

In this section, we discuss three key elements of the design of an MBR: the length of the delay (D) for redemptions of the MBR, the reference amount (R) used to compute the MBR, and the combination of MBR size (m) and maximum potential subordination (s).

7.1. Length of the delay, D , for redemption of the MBR. A key element of an MBR rule is the length of the delay for redemption of shareholders' MBRs. The delay period should be, at a minimum, long enough to inhibit "preemptive" runs, but not so long as to unnecessarily inconvenience shareholders or impede market discipline for MMFs.

Very short delays potentially would leave MMFs vulnerable to runs if investors believe that they can redeem quickly enough at the first sign of a problem to receive their MBRs before the problem finally causes a fund to break the buck. Because investors who hope to preserve principal and liquidity would have to exit a distressed fund earlier than they would under current rules, a short delay might make MMFs *more* vulnerable to runs. Longer delays would allow more time for a problem to be resolved before an investor's MBR is returned and thus provide more protection from preemptive runs.

At the same time, longer delays would be more inconvenient for MMF investors. Moreover, very long delays might impede market discipline for MMFs by discouraging investors from pulling away from a fund with poor management, even long before the emergence of any specific strains in its portfolio.

The possibility of preemptive runs suggests that a delay of a week or less likely would be inadequate for protecting non-redeeming investors and making MMFs more robust. MMFs maintain substantial liquidity buffers in the form of assets that mature (or can be readily disposed of) within a week. For example, the "weekly liquid assets," as defined in rule 2a-7, made up an average of 40 percent of prime MMF assets in 2011.⁴² However, outside this

⁴¹ iMoneyNet data indicate that two prime institutional MMFs suffered net redemptions of more than 70 percent in the weeks following the Lehman bankruptcy in 2008, but none had redemptions exceeding 80 percent.

⁴² Weekly liquid assets include cash, Treasury securities, government agency and government-sponsored enterprise discount notes with remaining maturities of 60 days or less, and securities that will mature within five business days or which have a put option that an MMF can exercise to receive payment within five business days. Data are from MMFs' filings of SEC form N-MFP.

pool of weekly liquid assets, the weighted average life (WAL) of securities held by prime MMFs averaged 114 days. An MMF under heavy redemption pressure would typically first tap its weekly liquidity buffer to avoid losses, so one concern with an MBR delay of just one week is that this would not provide enough time for realization of the immediate repercussions of investors' redemptions. Fast-acting investors might receive their full MBRs despite leaving non-redeeming investors with claims on a fund that has been largely stripped of its liquidity buffer.

The events of September 2008—before MMFs were required to maintain weekly liquidity buffers—also suggest that a delay of only a few days or a week would be inadequate to protect MMFs and their non-redeeming investors. Even amid the severe run triggered by Lehman's bankruptcy early on Monday, September 15, every MMF except the Reserve Primary Fund managed to survive until the end of the week without breaking the buck (though in some cases, only because of considerable sponsor support). Thus, with an MBR delay of just a few days, an MMF shareholder who redeemed from a fund *after* Lehman failed might have succeeded in exiting without a loss and shifting potential risks and losses to non-redeemers. While it is impossible to say what might have happened if the Treasury and Federal Reserve had not intervened on Friday, September 19, the run in 2008 suggests that redemptions of MBRs within a week might allow investors who take part in a run to exit MMFs without sharing in any of the resulting losses.⁴³

One consideration in setting the delay period is the maturity profile of MMF portfolio holdings. Because any problems in a fund's portfolio at the time of an investor's redemption should be resolved by the maturity date for its longest-lived assets, portfolio holdings might suggest a delay period of as long as 397 days, the limit on maturity for privately-issued assets that MMFs can hold. On average, the longest-dated security held in a prime MMF in 2011 (excluding government securities) matured in 327 days.⁴⁴ Alternatively, delay periods might be set to match the funds' maximum weighted average life of assets (120 days). Although the 30-day delay that we suggest is relatively short in comparison with the maturity of the typical MMF's assets, about half of prime MMFs' assets (on average) mature within 30 days, so a delay of that length should provide a reasonable amount of time to ensure that problems in a fund's portfolio when an investor requests a redemption are resolved prior to her receipt of final payment for her MBR.

In addition, because MMFs must report detailed portfolio-holdings data on a monthly basis on their websites, a 30-day delay ensures that an MMF usually would publish updated data at some point between an investor's request for redemption of her MBR and her receipt of

⁴³ On the morning of Friday, September 19, 2008, the Treasury announced its Temporary Guarantee Program for Money Market Funds and the Federal Reserve announced its Asset-Backed Commercial Paper Money Market Mutual Fund Liquidity Facility (AMLF). Redemptions from prime MMFs, which had averaged about \$100 billion *per day* on September 17 and 18, dropped to an average of \$13 billion per day over the following week (based on data from iMoneyNet, with adjustments described in McCabe, 2010).

⁴⁴ Data are from MMFs' filings of SEC form N-MFP.

payment. (A slightly longer delay, perhaps 35 days, would guarantee that updated data would become available between redemption requests and payments.) As such, a 30-day delay would help to ensure that investors who request redemptions based upon information about a problem that has already occurred in an MMF's portfolio would not be paid before updated information about the portfolio is publicly released.

7.2. Reference amount, *R*. The MBR would be a fraction of an investor's reference amount, so the method for calculating that amount from investors' recent balances is important. Two elements of this calculation are at issue: the "reference period," that is, the time period over which recent balances are analyzed, and the formula for computing the reference amount over that period. We believe that the most transparent and straightforward form of MBR rule would use the past 30 days as the reference period and set the reference amount equal to the investor's maximum balance, excluding any held-back shares, over the reference period. More generally, if payments of the MBR are delayed for D days, the reference period also should be D days.

To illustrate this approach, consider an investor, Alice, who has maintained a balance in an MMF for a year. The fund has an MBR of 5 percent and a delay period, D , of 30 days. Alice's balance has been fairly steady; it has only grown as interest has accrued, and on day T she has \$100 in the fund and requests to redeem all of her shares. On day T , her maximum balance over the reference period—the past 30 days—is \$100, so Alice's reference amount is \$100. She immediately receives \$95, while \$5 is held back for 30 days.

Similarly, over the next 30 days, Alice's reference amount remains at \$100, since that is her maximum balance over the reference period. But at the close of day $T+30$, at the end of the 30-day delay period, her maximum reference-period balance (excluding her \$5 in held-back shares) drops to 0, her MBR drops to zero, and she receives the delayed portion of her redemption. Setting the lengths of the reference period and delay period equal means that the redemption of her MBR occurs on the day when her MBR falls to zero.⁴⁵

⁴⁵ In theory, the reference period could be longer or shorter than the delay period. However, a reference period that is longer than the stated delay period would effectively extend that period: At the end of the delay period, an investor would still have a non-zero MBR, and full repayment of her shares would be delayed for the length of the reference period. For example, if Alice's MMF had a reference period of 60 days, she would still have an MBR of \$5 for 30 additional days after the delay period nominally had ended.

A reference period that is shorter than the delay period would diminish the benefits of the MBR by reducing the effective delay period (see section 7.1). A short reference period also would complicate investors' choices. For example, if the reference period were just 15 days, rather than 30, Alice would be better off not redeeming her full balance on day T . Instead, she should request to redeem only her full *available* balance (\$95) on day T . Then, at the end of day $T+15$, her reference amount would fall to \$5, and her MBR would decline to 25 cents. On that day, she would again face a choice between redeeming all \$5 of her remaining shares and receiving her 25 cent MBR in 30 days, or redeeming just \$4.75, which would allow her to receive 23.75 cents in 15 days, and so on.

The formula used to compute an investor's reference amount might take many forms. Two examples—the investor's *maximum* balance over the reference period and the investor's *average* balance—can illustrate some of the tradeoffs of using different formulas. A reference amount based on investors' maximum balances would be relatively straightforward and would offer MMFs more protection, for any given MBR fraction m , than other reasonable formulas. Alternatively, although an average-balance formula would lead to smaller reference amounts than a maximum-balance rule, the MBR fraction m could be adjusted to offset the difference. Investors with relatively volatile balances would benefit from an average-balance rule, which would make reference amounts less sensitive to short-lived spikes in shareholdings.

The average-balance formula would also result in shorter-lived MBRs, even for investors who have steady balances over the reference period, and such a formula could complicate redemption of held-back shares. Returning to the example above, after Alice redeems her entire \$100 balance, her MBR with a maximum-balance rule remains at \$5 for 30 days. In contrast, if Alice's reference amount is based on her average balance, her MBR would decline by 16.7 cents each day, and after 15 days, her MBR would be just \$2.50.⁴⁶ An average-balance rule in this case halves the effective protection afforded by Alice's MBR after her redemption. Moreover, with Alice's MBR falling each day, she would be entitled to payments of 16.7 cents per day for 30 days. With a maximum-balance rule, she would receive a single \$5 payment after 30 days.

Finally, an investor's held-back shares—that is, portions of an MBR that she has redeemed but which have not yet been disbursed—probably should be excluded from any formula for computing her reference amount. In the example above, if Alice's \$5 in held-back shares are included in calculating her reference amount, then 30 days after her redemption her reference amount would fall not to zero but to \$5. Including held-back shares in an investor's reference amount would unnecessarily complicate redemptions for investors and for MMFs.

7.3. Size of the MBR, m , and amount of potential subordination, s . The size of an MBR and the amount of potential subordination should be calibrated to diminish the vulnerability of MMFs to runs. Although the magnitudes necessary to achieve this aim will depend on investors' expectations in future crises and hence cannot be known with precision, the analysis in section 6 provides some useful guidelines. An MBR and potential subordination should be large enough to convince investors that running from MMFs during a crisis is not in their economic interest, but not so large as to impose unnecessary opportunity costs upon investors who could see increased portions of their balances subject to delayed redemption. Importantly, as the analysis in section 6.3 shows, an MBR rule can create disincentives for investors to run, even if they believe that there is some possibility that a fund's losses may exceed the aggregate MBR.

⁴⁶ After 15 days, her average balance would be \$50 over the 30-day reference period: \$100 for 15 days and \$0 for 15 days (note that her \$5 in held-back shares are not included in computing her reference amount).

One outcome to avoid is a combination of an MBR and potential subordination that are too small, which would impose costs on MMFs and their investors while leaving them vulnerable to runs. If so, the MBR would not achieve its objectives and MMFs would continue to pose risks to the financial system.

As we describe in section 6, an analysis of minimum MBR size is hampered by the paucity of available data on MMF losses, in part because sponsor support for MMFs has generally obscured the magnitude of the funds' portfolio losses. Nonetheless, our analysis suggests that if investors' expectations are based on past experiences of MMF losses, a strong MBR rule and an MBR of at least 3 to 4 percent of investors' balances, in combination with a 50 basis point capital buffer, probably would be sufficient to brake redemptions in a crisis. But given the uncertainties of estimating expected losses in a distressed MMF and the challenges of pinpointing liquidity costs to investors who lose access to shares in a closed fund, section 6 also includes analysis of investor incentives when expected losses and costs are smaller and larger than our baseline assumptions. In the scenario in which expected losses and costs are smaller than our baseline, a strong MBR rule and an MBR of just 2 percent might be adequate to brake redemptions in an MMF with a 50 basis point buffer. If expected losses and costs are at the higher end of the range we consider, a strong MBR rule and an MBR of 4 percent might be needed to provide disincentives to redeem. And the possibility of very heavy redemptions by other investors strengthens the case for a larger MBR.

As we discuss in section 6.3, other combinations of MBR size and maximum subordination (and a capital buffer) might create the necessary disincentives to redeem from troubled MMFs. For example, an effective MBR rule and an MBR of 5 percent, with subordination limited to 2 percentage points of the MBR ($s = 0.4$), in combination with a 50 basis point buffer, would achieve a degree of stability similar to that provided by a strong MBR rule and an MBR of 3 percent.

The size of an MBR necessary to provide stability to MMFs also would depend on other elements of the design of the MBR rule besides MBR size and the maximum amount of subordination. A larger capital buffer and a longer delay period (D) would also tend to make an MBR more robust. Also important is the formula used to compute the reference amount (R), with the maximum-balance (high-water mark) formula providing greater protection than other formulas.

8. Policy considerations

In this section, we discuss some policy considerations that may be important in assessing the utility of an MBR for mitigating the systemic risks posed by MMFs and protecting their investors. We begin by reviewing some possible criticisms of an MBR, notably arguments that (i) an MBR rule might cause preemptive runs, (ii) an MBR rule would dampen market discipline for MMFs by limiting investors' ability to redeem, and (iii) an MBR would diminish investor demand for MMFs. We then compare the MBR to some other proposals for MMF reform, including requiring MMFs to have floating NAVs, capital buffers, and other

forms of restrictions on redemptions. We also discuss how an MBR might complement other requirements in a broader reform package.

8.1. Could an MBR rule lead to preemptive runs? A strong or an effective MBR rule is unlikely to lead to harmful preemptive redemptions, as long as the delay period for redemption of the MBR is sufficiently long (see section 7.1). But such a rule may motivate investors to leave MMFs that manage risks poorly, well in advance of any specific problems materializing in a fund's portfolio, and hence encourage market discipline.

Concerns about preemptive runs are particularly acute for redemption restrictions and fees that are imposed conditionally. That is, shareholders may run *earlier* in a crisis, when they perceive that conditions that trigger restrictions may soon be met. Indeed, investors who exit immediately before conditional restrictions are imposed—like those who, under current rules, get out just before a fund breaks the buck—can avoid all principal and liquidity losses, even if their redemptions contribute to strains that eventually cause losses for others. As such, conditional restrictions likely would do little to change the incentives of investors in a distressed MMF other than moving forward the deadline for redemptions to avoid losses.

The MBR rules that we propose are fundamentally different from conditional restrictions, since the delay for disbursements of the MBR would *always* be in place. Hence, investors in an MMF with an MBR would not have incentives to run in advance of triggering events that might restrict or penalize redemptions.

Still, because an investor who redeems far enough ahead of a problem may have his entire MBR returned before the problem results in any losses, preemptive redemptions are, in principle, possible with an MBR rule. As discussed in section 7.1, preemptive runs may be a serious risk if the delay for payment of investors' MBRs is very short (for example, a week or less), but the risk of preemptive runs diminishes as the delay increases. But even with the 30-day delay in disbursement that we propose, investors who believe that an MMF is at risk might possibly redeem earlier than they otherwise would.

However, the dynamics of early redemptions from a fund with a strong or effective MBR rule clearly would be different from those in a preemptive run on a fund with conditional redemption restrictions, for several reasons. First, although the MBR does create a new 30-day-ahead deadline for emerging unscathed from a distressed MMF, this time interval is long enough that investors who do successfully exit on concerns about a fund probably would have to do so based on their own analyses of a fund's exposures, rather than because of headlines or other events that likely would cause heavy simultaneous redemptions by many investors. Hence, the redemptions motivated by an MBR rule with a 30-day delay would most likely be diffuse and manageable, rather than precipitous and widespread. Moreover, as we argue below, diffuse redemptions by investors who carefully monitor MMFs well in advance of any specific problems are *beneficial* insofar as they provide market discipline that discourages poor risk management in the MMF industry.

Second, with an effective or strong rule, as long as the MBR and potential subordination are large enough, incentives to redeem *diminish* as the fund's distress becomes more apparent. Because redemptions cause a portion of a shareholder's MBR to be subordinated, the implicit cost of redemptions rises as losses appear more likely.⁴⁷ This dynamic is automatic and inherent in the rule; no action is required by the fund (or the government) to ensure that disincentives to redeem strengthen as the fund's condition deteriorates.

In contrast, under current rules or with conditional restrictions, investors' incentives to redeem strengthen as a fund's distress becomes more obvious because those who redeem even just ahead of a triggering event are treated very differently than those who redeem—or attempt to redeem—immediately thereafter. Whether the triggering event is a fund breaking the buck (under current rules) or the imposition of conditional redemption restrictions, redeeming just beforehand allows an investor to avoid all losses, while those who wait risk principal and liquidity. Because other investors face the same choices, and their redemptions are likely to push a fund closer to a trigger, incentives to redeem strengthen as the fund's condition worsens.

Finally, it is possible that, notwithstanding shareholders' incentives *not* to redeem in a crisis from an MMF with an MBR, investors' irrational fears may cause them to do so anyway. Some may argue that misunderstandings of or apprehension about MBR rules would worsen such fears. To be sure, accurate predictions of irrational behaviors in a crisis are difficult. Notwithstanding this uncertainty, an MBR with subordination clearly would diminish or reverse pressures on *rational* investors to exit MMFs during crises. In addition, in the event that irrational investors run regardless of incentives, the rules that we propose would ensure that they bear some of the costs of their actions and provide extra protections for *other* investors, so that staying invested in the fund can continue to be a rational decision.

8.2. Would an MBR rule erode mechanisms for investor protection and market discipline for MMFs? No, a strong or effective rule with a sufficient delay period for redemptions of the MBR would encourage scrutiny of MMFs well in advance of any specific problems arising in fund portfolios. That is, the MBR would provide an incentive for industry discipline when it is most likely to be helpful, rather than during a crisis when discipline provides little net benefit. Investors could continue to protect themselves by redeeming from poorly run funds, but not at the expense of other investors.

The Investment Company Act of 1940 specifically protects the ability of mutual fund investors to redeem shares, in part because this option gives investors a degree of protection from poorly run funds and provides market discipline for the fund industry.

⁴⁷ A redeeming investor essentially writes a put option to non-redeeming investors (with a strike price equal to the fund's stable-value NAV), and as losses look more likely, the put comes closer to being in the money. In terms of the model in section 4, if $B < 0$, the disincentive to redeem strengthens as the probability of fund closure, x , increases.

Thus, one concern about redemption restrictions is that they might limit investors' ability to protect themselves and blunt market discipline in the MMF industry.

A well-designed MBR rule would not interfere with redemptions that help protect their investors *in aggregate* and provide useful discipline. But such a rule would discourage redemptions from a fund that is already in distress, since these redemptions generally only benefit redeeming investors *at the expense of non-redeeming shareholders* and the fund itself. Moreover, if losses have already occurred or are imminent, such redemptions provide little useful discipline for a fund's management.

Consider for example, two types of information that might cause an investor to determine that an MMF no longer suits his needs and hence to redeem shares. First, he might observe that a fund has changed its investment strategy. For example, in the summer of 2007, over a year before the Reserve Primary Fund broke the buck, it dramatically shifted its investment policies, began purchasing riskier assets, and boosted its yields (ICI, 2009; McCabe, 2010). Second, the investor might learn that a security held by his fund is in distress or has just defaulted.

Investors who redeem based on the first type of information may provide important discipline for MMFs. For example, if investors had responded in 2007 to Reserve's new investment strategy by redeeming shares, the MMF industry might have averted a catastrophe. (Instead, the Reserve Primary Fund's higher yields attracted very large inflows, and the fund's assets tripled over the following year.) Redemptions *before* a fund has encountered a problem are likely to be based on investors' individual assessments of funds' strategies and therefore probably would be diffuse, manageable, and unlikely to harm other shareholders. Indeed, to the extent that such redemptions punish risky investment strategies, they may reduce the likelihood of future losses.

In contrast, news that one or more of a fund's holdings is in serious distress is likely to trigger sudden, sharp redemptions, as shareholders seek to avoid losses by getting out of a fund while shares are still redeemable at \$1. Instead of reducing the likelihood of future losses, such redemptions at best only protect redeeming shareholders at the expense of non-redeemers. The resulting scramble may even destroy value: Large redemptions may themselves cause further losses for a fund that must dispose of assets quickly to raise cash, and such sales may have additional (unpriced) external effects on other investors and borrowers in short-term funding markets.

A well-designed MBR would encourage investors to redeem early from poorly run funds but discourage redemptions during a crisis. Ideally, the delay for disbursements of the MBR should be set so that investors who determine that a fund is poorly managed can redeem well before problems emerge, but investors who redeem when losses are imminent are not rewarded at the expense of other shareholders.

8.3. Would an MBR reduce demand for MMFs? The MBR would enhance the overall principal stability of MMFs by making damaging, crisis-driven redemptions less likely, and it would provide additional principal protection for investors who are unlikely to redeem

rapidly in a crisis. At the same time, the MBR rule would make less feasible the strategy of avoiding losses by redeeming MMF shares quickly at the first sign of a problem, and the rule would modestly reduce investors' liquidity and make investing in MMFs somewhat more complicated. The net effect of these changes on investor demand are difficult to predict, and would depend on other elements of a reform package. On balance, although reforms may reduce demand somewhat, in light of the limitations of other cash-management opportunities available for MMF investors, a sudden shift to alternative products appears unlikely.

The net demand effect of reforms that mitigate the systemic risks posed by MMFs would depend crucially on the broader elements of a reform package. Although the focus in this paper is on the potential benefits of an MBR for MMFs, a flexible reform package that allows for different types of MMFs with different risk-mitigation strategies would allow investors to choose the MMFs that best suit their preferences and minimize reductions in demand. For example, reforms might allow for funds with MBRs, as well as funds that are exempt from MBR rules because they have floating NAVs or hold only Treasury securities. Investors who need full access to their entire balances at all times, but are willing to tolerate minor fluctuations in principal, might choose a floating NAV fund. Investors who want principal stability but can live with a minor restriction on liquidity might choose an MMF with an MBR rule. And investors who desire both principal stability and fully unrestricted liquidity could sacrifice some yield by choosing Treasury-only MMFs.

Even with such flexibility, reforms that mitigate systemic risks arising from MMFs undoubtedly would be unpopular with some investors, and some net dampening of demand for MMFs is possible. This is to some degree inevitable, since the objective of reform is to internalize the costs that MMFs impose on the financial system. Costs associated with systemic risks are externalities: Runs on MMFs may have serious and costly consequences that extend well beyond the funds and their investors. For example, those consequences may adversely affect firms that rely on short-term funding, the financial system, and the economy. Indeed, the severity of the damage to financial stability caused by the run in 2008 was the impetus for unprecedented government interventions that supported MMFs in order to stop the run. Furthermore, even though MMFs suffered no material losses in the summer of 2011, sizeable redemptions from the funds motivated by concerns about their exposures to European banks caused reductions in the availability of short-term funding for U.S. nonfinancial firms (Chernenko and Sunderam, 2012).

Clearly, investors who redeem MMF shares for \$1 apiece in a crisis do not shoulder all the costs associated with the systemic risks that MMFs can create. So, rules that limit MMF investors' ability to impose costs on others may dampen demand for MMFs, even if reforms improve efficiency and also make the funds safer.⁴⁸ Perhaps not surprisingly, investors

⁴⁸ Other tools to mitigate systemic risks arising from financial activities also impose costs on market participants. For example, deposit insurance and bank capital requirements, which reduce the susceptibility of banks to runs, also impose some costs on investors.

who respond to surveys about MMF reform options generally show little enthusiasm for *any* option.⁴⁹

The effects of reforms on demand for MMFs probably would depend on how changes affect the features that make the funds attractive to investors: stability of principal, liquidity, and market-based yields. Below, we discuss how the impact of an MBR on different features of MMFs might affect retail and institutional investors' demand for the funds.

8.3.1. Enhanced principal stability. Some elements of our MBR proposal might be perceived quite positively by investors. One important advantage of the MBR is that it would preserve *and enhance* MMFs' principal stability, which is their "hallmark" feature and one that is key to their attraction for investors (ICI, 2009). By discouraging redemptions from funds that are under stress, the MBR rules would make MMFs less vulnerable to heavy redemptions that may force fire-sale related losses. Hence, an MBR should diminish the risk that a money fund would break the buck.⁵⁰ The net stabilizing effect of the MBR thus may make MMFs *more* attractive to investors who are particularly concerned about safety.

In addition, because the MBR rule would shift the allocation of losses in the event that an MMF does break the buck, this rule would provide additional safety to investors who are unlikely to redeem shares rapidly in a crisis. An MBR also would make less viable the strategy of holding MMF shares with the expectation that they can be redeemed—and any imminent losses shifted to other investors—at the first sign of trouble. While the option to employ that strategy might make MMFs attractive to some sophisticated investors, even under current rules it cannot be employed successfully by all shareholders. Arguably, forcing such shareholders to revise their expectations, or encouraging them to invest in other products, might be beneficial to other shareholders, MMFs, and the stability of short-term funding markets.

⁴⁹ See, for example, Fidelity (2012) and Treasury Strategies (2012). Notably, each of these surveys included a "holdback" option that would delay disbursement of 3 percent of *every* redemption by every investor for 30 days. Survey participants understandably responded negatively to this option, which clearly would impair the use of MMFs for cash management and would be extraordinarily complex for investors who purchase and redeem shares frequently. More generally, however, surveys that implicitly offer investors a choice between current rules and reforms designed to internalize costs arising from the systemic risks are bound to show little evidence of support for reform. An alternative perspective on the utility of reform might be provided by a survey of taxpayers who could be asked about their willingness to provide support for MMFs in the event of another widespread run on the funds.

⁵⁰ Because of the similarity of the portfolios that money funds hold, reducing one MMF's susceptibility to runs would benefit investors in other funds. For example, the reduced risk of one fund's forced sales, which could depreciate asset prices, diminishes the likelihood of losses in other MMFs that hold the same securities, and therefore benefits their investors. In addition, because of the contagion effects of runs, a reduction in the risk of losses and runs for one MMF makes runs on others less probable.

One concern about MBR rules might be that, in the event that an MMF breaks the buck, the rules for dividing losses based on recent redemptions would heighten uncertainty about outcomes for individual investors. To be sure, MBR rules would be unfamiliar and the possibility that their own redemptions might heighten losses may deter some investors from using MMFs. But investors' losses in an MMF that breaks the buck under *current rules* are already uncertain and depend—in a destabilizing manner—on each investor's own and others' recent redemptions. The MBR would shift the advantages from redeeming investors to non-redeemers, so that investors' incentives are stabilizing and lessen the likelihood of break-the-buck events.

8.3.2. Impact on liquidity. The enhanced principal stability provided by the MBR rule would come at some cost in terms of reduced liquidity for investors. We discuss the effects on demand from retail and institutional investors below. As noted above, the net effects of an MBR rule on demand for MMFs would depend on other elements of a reform package, including whether it allows for funds that would cater to investors with extreme liquidity needs.

An MBR rule would impose the largest liquidity costs on investors whose activity results in the highest liquidity costs for the fund. Investors who tend to hold relatively steady account balances would see very little impact of the new rules, except when they close accounts, but investors who regularly toggle between large and zero balances would see a higher liquidity cost. Such investors, it should be noted, also impose the highest liquidity costs on the MMFs (see, for example, Lyon, 1984).⁵¹ At the extreme, an investor who wishes to shift \$100 in and out of an MMF with a 5 percent MBR and avoid any delay would have to maintain an extra balance of \$5.26.⁵²

8.3.3. Effect on net yield. One advantage of an MBR rule is that it probably would have only a very minor effect on the market-based net yields that investors receive. Unlike a capital buffer, which would have ongoing funding costs, the costs associated with an MBR itself would likely be limited to one-time changes in the way that MMFs and financial services firms track investor's balances. For example, the transfer agents that normally handle MMF transactions, or fund distributors that handle MMF accounts, would have to develop systems to track each investor's MBR and ensure that redemptions of investors' MBRs are subject to the appropriate delay. Once those systems are in place, however, ongoing costs presumably would be very small.

8.3.4. Effects on retail demand for MMFs. The net effect on demand for MMFs would depend largely on how MBR rules affect the attractiveness of these funds relative to other

⁵¹ Flow volatility also can diminish significantly the performance of long-term mutual funds (Edelen, 1999).

⁵² That is, the investor would have to hold enough such that when she redeems \$100, she still has 5 percent of her original balance left. Before the redemption, she needs \$105.26 in her account; with a 5 percent MBR (\$5.26), she can withdraw \$100. In general, an investor who wants to regularly purchase and redeem \$100 worth of shares in an MMF with an MBR fraction m needs to maintain a balance of at least $\$100/(1-m)$.

investment alternatives. For retail investors who primarily seek principal stability, alternatives to MMFs generally would be insured deposits, although the impact of an MBR on the relative demand for retail MMFs and bank deposits is unclear. Since money funds with MBRs would be even less likely to break the buck than today's MMFs, and because our proposal would exempt the first \$50,000 of redemptions from triggering any subordination, an MBR rule probably would not cause retail investors to move money to banks because of concerns about MMF safety.

Another reason to believe that an MBR rule would likely have only modest effects on retail demand for MMFs is that one of the primary drivers of such demand has historically been the spread between retail MMF yields and those on bank deposits (see, for example, ICI, 2012). As noted above, the effects of an MBR rule itself on MMF yields would likely be small, so the funds presumably would continue to pay yields that typically would exceed rates paid on deposit accounts.

However, the additional complexity and reduced liquidity of MMFs with minimum balances might cause some retail investors to stop using MMFs—or to switch to MMFs without MBRs, if such funds are part of a broader reform package. For example, investors who temporarily park large sums in MMFs between transactions or between transfers of investments (for example, from one mutual fund to another within a single fund family), might find MMFs with MBRs less convenient. The transient nature of such investments suggests that a reduction in demand for them is unlikely to put a significant dent in aggregate retail demand for MMFs.

8.3.5. Effects on institutional demand for MMFs. Institutional investors who seek principal stability have a broader range of potential alternatives to MMFs than do retail investors. For example, institutional investors might turn to deposits, invest in “enhanced cash funds” or other short-term vehicles that are not registered with the SEC, put money into “offshore MMFs” that are beyond the reach of U.S. regulation, or invest directly in money market instruments. Since an MBR rule generally would enhance MMF principal stability but introduce some limits on liquidity of fund shares, the responses of institutional investors to an MBR rule probably would depend on the investors’ relative preferences for safety and liquidity.

Investors whose primary aim is to preserve principal probably would not be motivated to shift money away from MMFs with MBRs, since other cash-management options have safety drawbacks that limit their potential as substitutes for MMFs.⁵³ Lightly regulated or unregulated alternatives to MMFs, such as enhanced cash funds and offshore MMFs, would not offer principal-stability advantages relative to an SEC-registered MMF with an MBR.

⁵³ An exception might be investors who expect to be able to employ a strategy of avoiding losses by withdrawing money rapidly from funds that encounter trouble. As noted above, an important benefit of an MBR rule is its clarification that such a strategy, which has never been feasible for *all* investors, would no longer be viable for anyone.

Even bank deposits have safety disadvantages for large institutional investors whose cash holdings typically exceed by orders of magnitude the caps on deposit insurance coverage; for these investors, deposits are effectively large, unsecured exposures to a bank.⁵⁴ MMF shares—which represent claims on diversified, transparent, tightly regulated portfolios—would continue to offer important safety advantages relative to bank deposits. Hence, the most risk-averse and run-prone institutional investors likely would remain invested in MMFs, where an MBR could provide protection from runs.

However, because an MBR would put some restrictions on the liquidity of MMF shares, investors with heavy liquidity needs might turn to other options, including bank deposits. Investors conceivably also might turn to enhanced cash and offshore funds, but these alternatives have greater latitude than MMFs to suspend redemptions and employ other constraints on investors' liquidity, particularly during crises when preservation of liquidity is especially important.⁵⁵ Offshore MMFs organized in Europe, for example, have latitude under the European statute governing “undertakings for collective investment in transferable securities” (UCITS) to temporarily suspend redemptions “where suspension is justified having regard to the interests of the [shareholders]” (European Parliament, 2009).

Another possibility is that some institutional investors might elect to purchase money market instruments directly, rather than outsourcing their cash-management to MMFs. On net, such a development might provide further stabilization for short-term funding markets, because direct investments do not share some of the features of MMFs that make them vulnerable to runs. In particular, unlike MMF shareholders, direct investors who choose to sell assets in a crisis bear the liquidity costs of their own actions and have no ability to transfer risks and losses directly to those who do not sell assets. Hence, such investors have much smaller incentives to exit their positions abruptly than do MMF shareholders, who may have the option to redeem shares without bearing any of the associated liquidity costs of their actions—and even to receive \$1 for shares that are worth less.

8.4. Could migration of assets away from MMFs undermine financial stability? One concern about most options for MMF reform, including an MBR rule, is that new restrictions might motivate investors to shift their cash-management balances to other types of vehicles that would pose even greater financial stability risks than MMFs. As noted in section 8.3, the net effect of an MBR on investor demand for MMFs is difficult to predict, but the likelihood that shifts in investors' money would undermine financial stability seems small. The potential substitutes that would pose the most risk—such as enhanced cash

⁵⁴ Deposit insurance coverage is usually limited to the first \$250,000 of a depositor's balance, although the Dodd-Frank Wall Street Reform and Consumer Protection Act mandated the extension of unlimited insurance for noninterest-bearing transactions accounts in 2011 and 2012.

⁵⁵ In fact, several enhanced cash funds suspended redemptions in late 2007 and early 2008, and the problems of these funds prompted a very large migration of institutional cash *toward* MMFs in that period (Moody's, 2008; Baba et al., 2009).

funds and offshore MMFs—probably have limited attraction to investors who would migrate from MMFs out of concerns about reduced liquidity under an MBR.

Another substitution scenario of concern would be a shift by institutional investors into bank deposits. The size of institutional accounts would put them outside the scope of normal deposit insurance limits and potentially make banks more susceptible to runs (PWG, 2010).⁵⁶ Although institutional investors generally avoid the large, unsecured exposures to individual banks that uninsured deposits represent, an MBR rule for MMFs might cause some of these investors to seek the unrestricted free liquidity provided by deposits.

One factor that complicates the assessment of this risk is that temporary unlimited deposit-insurance coverage for noninterest-bearing transactions accounts appears already to have attracted large inflows of institutional cash to banks, probably at least in part from MMFs. Unlimited insurance is scheduled to expire at the end of 2012, so some of the institutional cash might be expected to return to MMFs, but the introduction of an MBR rule might dampen such return flows.⁵⁷ To some extent, current law mitigates concerns about large, uninsured deposits by providing greater capacity and flexibility to deal with liquidity strains in the banking sector than elsewhere; for example, banks have discount-window access that reduces their vulnerability to deposit outflows in a crisis.

8.5. How would an MBR rule compare to requiring MMFs to have floating NAVs? As noted in section 2, a requirement that all MMFs have floating NAVs has considerable appeal among policymakers, academics, and others for its simplicity, elimination of NAV rounding, and improvements in investors’ understanding of MMF risks. Ideally, a floating NAV would not only reduce the risk of runs on MMFs, but also clarify that MMF shareholders bear the risks of the securities in a fund’s portfolio. Moreover, if fluctuations in MMF net asset values become a regular, daily occurrence, a floating NAV might curtail the systemically risky reliance of MMFs on sponsor support.⁵⁸

At the same time, the floating NAV option as a *standalone* fix for the vulnerability of MMFs to runs has some important drawbacks, most notably the possibility that elimination of the “hallmark” feature of the funds would be tantamount to “eviscerating” them (ICI, 2009). Opponents of a floating NAV have cited a broad range of concerns about its potential

⁵⁶ As noted above, in 2011 and 2012, there is no limit on insurance coverage for noninterest-bearing transactions accounts.

⁵⁷ See note 54. Noninterest-bearing transactions accounts balances that exceed the normal \$250,000 limit on deposit insurance increased \$460 billion (54 percent) between the end of 2010 and the end of March 2012 (Federal Deposit Insurance Corporation (FDIC), 2012). During this period, ICI data indicate that institutional MMF assets declined \$173 billion (9 percent).

⁵⁸ However, as indicated in note 3, not all proposals for requiring floating NAVs would actually result in regular fluctuations in share prices. For example, allowing MMFs to maintain \$10 NAVs and round to the nearest cent might make NAV changes relatively rare and undermine any potential benefits of a floating NAV.

impacts on MMF investors, including tax and accounting complications that might substantially diminish the appeal of the funds. If so, a floating NAV might lead to a steep decline in investor demand for MMF shares and a migration of assets to less regulated vehicles that continue to offer stable NAVs. Moreover, even if MMFs with floating NAVs remain sizable, they might continue to be vulnerable to runs, since investors in distressed funds still would have strong incentives to redeem.

In contrast, while an MBR rule would reduce the liquidity of a small portion of each investor's MMF shares, it would *enhance* their principal stability and thus preserve the defining feature of MMFs. Indeed, an important impetus for development and analysis of the MBR proposal is the need for careful consideration of reform options for MMFs that would preserve their principal stability. Hence, the overall impact of the MBR rule on demand for MMFs likely would be far less stark than the effect of a floating NAV requirement for all MMFs. In addition, unlike the stable NAV, a strong or effective MBR rule would provide *disincentives* to run and protections for investors who do not redeem in a crisis. These advantages for an MBR would have to be balanced against the conceptual simplicity offered by a floating NAV requirement.

One option worth exploring would be to create at least two types of MMFs: floating NAV funds alongside others with stable NAVs and MBRs. This approach, which would be a form of "two-tier" system for MMFs, as described in the PWG Report (PWG, 2010), might have advantages over reforms that create only one type of fund.⁵⁹ Allowing both options might minimize the effects of reform on demand for MMFs; investors could choose between the principal stability of funds with MBRs and the unrestricted liquidity afforded by floating NAV funds without MBRs. The most risk-averse, run-prone investors probably would select the stable NAV funds with the protections of an MBR. Floating NAV funds would lack such protection, but presumably would have a less run-prone investor base that could not redeem specifically to exploit NAV rounding during a crisis.

8.6. How would an MBR rule compare to requiring MMFs to have capital buffers?

Both capital and the MBR rule effectively create buffers to absorb potential MMF losses, but there are important differences in how these buffers would provide stability. A *substantial* capital buffer—that is, one of magnitude comparable to the size of the MBRs that we suggest—would give MMFs some capacity to absorb losses *without* breaking the buck and could allow the funds to operate without disruption even after suffering material credit-related losses. An MBR, in contrast, would not prevent funds that suffer losses from breaking the buck and closing, nor would it shield investors (in aggregate) from losses. Instead, it would create from investors' own shares a buffer that would allow an allocation of losses that stabilizes MMFs by protecting non-redeeming investors. It is worth noting that because MBR rules and capital buffers have different strengths and weaknesses, the two options could complement one another well.

⁵⁹ The PWG report on MMF reform options provides a detailed description of some of the advantages of a two-tiered system.

A capital buffer and an MBR rule both would clarify the allocation of MMF risks and potential losses, but in different ways. A substantial pre-funded capital buffer would make clear that risks are borne by the owners of that capital, rather than by redeeming or non-redeeming shareholders or by MMF sponsors—that is, unless sponsors provide the capital. (However, a capital buffer that is too small to absorb reasonably foreseeable losses would have a much more ambiguous effect on the allocation of risks.) In contrast, an MBR primarily would clarify the allocation of risks *among* investors by ensuring that redeeming investors cannot shift risks and losses to non-redeemers. But an MBR rule would leave open the possibility of continued discretionary sponsor support for funds that are at risk of losses—a practice that would be less likely or seemingly necessary with a substantial capital buffer.

Although the allocation of risks resulting from capital and the MBR would differ, either approach would be an important improvement over the uncertain allocation arising from current practices. That uncertainty allows *all* parties to assume that they bear negligible risk.⁶⁰ Broadly speaking, the costs associated with taking on MMF risks should be similar, regardless of how losses are ultimately absorbed. Nonetheless, because current practices allow risks to go unacknowledged, reforms to ensure that risks are clearly allocated may *appear* relatively costly.

Since a capital buffer that is large enough to absorb most foreseeable losses would be the most ambitious option for reallocating risks, its *apparent* costs would likely exceed those of other options, including the MBR. Based on the size of the MMF industry at the end of May 2012, each percentage point of capital as a share of MMF assets would require up to \$29 billion in capital, depending on the scope of the capital requirement.⁶¹ Raising such sums introduces a variety of complications, and the advantages and disadvantages of a capital buffer would depend critically on how it is financed.⁶² Normally, the costs of raising a

⁶⁰ Investors may believe that they can redeem early enough in a crisis to avoid any MMF losses, and in the event that they do not redeem fast enough and losses do occur, investors appear to expect sponsors to provide support. Some sponsors were apparently so unprepared to acknowledge the possibility of losses that their MMFs—including the Reserve Primary Fund—had no systems for processing transactions at share prices that deviated from the customary \$1 stable NAV (SEC, 2009). Alternatively, sponsors can point to MMF prospectuses to argue that any risks are borne solely by investors, since MMF prospectuses and advertisements must warn that “it is possible to lose money by investing in the Fund” (SEC, 1998, 2003). And sponsors and investors alike may expect government support during a crisis.

⁶¹ This figure might be reduced if, for example, Treasury-only MMFs were exempt from a capital requirement.

⁶² At least three possible sources of capital have been suggested, and we sketch some of their pros and cons here only to provide a sense of the complex issues that may arise around financing a buffer. First, capital might be raised directly from MMF shareholders by retaining income that normally would be distributed to them (e.g., Goebel, Dwyer, and Messman, 2011). This approach would take many years to build a substantial buffer, unless net yields are reduced very sharply, and the ownership of such a buffer and the incentives to protect it would be unclear. Second, a buffer might be raised from third-party investors in capital markets (e.g., McCabe, 2011; BlackRock, 2011). This approach could potentially provide loss-absorption capacity more cost-effectively than other options, although efficiency gains might be offset somewhat by transactions

substantial capital buffer largely would be passed along to shareholders in the form of a reduction in net yields. But such costs could be a significant hurdle to the introduction of a meaningful capital buffer in the current environment of very low short-term interest rates, since MMFs have little room to increase the expenses they charge to pay for capital. Net yields of many MMFs are already close to zero, and *negative* yields for MMFs, which would erode their NAVs, would be potentially destabilizing under current rules.⁶³

In contrast, an MBR rule itself likely would have only minor effects on MMF expenses and yields, so introduction of an MBR probably would be less challenging than a similarly sized capital requirement, particularly in the current interest-rate environment. The costs of the MBR rule would primarily be reflected in the opportunity cost to shareholders of not having immediate access to a small portion of their balances—that is, in a reduction in the liquidity of MMF shares. In addition, the MBR could be viewed as costly to investors who otherwise had hoped to avoid losses on MMF shares by redeeming before other shareholders do, since the allocation of risks introduced by the MBR would make that strategy impossible.

As noted above, capital buffers and MBR rules also would have different effects on the linkages between MMFs and their sponsors. The MBR would not directly resolve the uncertain allocation of risks between MMF investors and sponsors; indeed, an MBR rule itself may preserve or even strengthen expectations of sponsor support for troubled MMFs.⁶⁴ In contrast, a pre-funded capital buffer that is large enough to absorb most foreseeable MMF losses also would help to sever the MMF-sponsor link; additional sponsor support would not be needed to keep a troubled MMF afloat. That said, investors and ratings agencies still might expect an MMF sponsor to “true up” a depleted capital buffer.

costs, such as underwriting costs, of selling claims on the buffer. Investors in the buffer also would have clear incentives to monitor MMF risks. A drawback is that buffers financed in capital markets would require the creation of a market for a new, untested type of security. Third, a buffer could be financed directly by MMF sponsors (e.g., BlackRock 2010). Although this approach would be simpler than obtaining funding in capital markets, it potentially would lead to further consolidation of the MMF industry among affiliates of large, systemically important financial institutions and shift the risks associated with money funds toward these institutions.

⁶³ The *net* yields that MMFs pay shareholders are the *gross* yields that the funds earn on their portfolio assets less the expenses the funds charge. In theory, the expenses charged by an MMF could exceed its gross yield and leave investors with a negative net yield, but an MMF that did so would erode its own NAV. As of June 2012, iMoneyNet data show that the asset-weighted average net yield of prime institutional MMFs was 14 basis points, while the average net yield of prime retail funds was just 2 basis points.

⁶⁴ Expectations of sponsor support create channels for transmission of systemic risks between MMFs and their sponsors (McCabe, 2010; PWG, 2010). Although an MBR rule probably would not diminish expectations for such support, the MBR rule would dampen one channel by which the sponsor-fund link can transmit strains during a crisis. Under current MMF rules, a loss of confidence in the financial strength of a sponsor may trigger or accelerate runs on its MMFs. The protections provided by an MBR make that outcome far less likely. Still, the MBR does not address the possibility that problems at an MMF that lead to expectations of support from its sponsor could contribute to the sponsor’s own financial distress.

Moreover, because large, already systemically important firms likely would face considerably lower hurdles to financing sizable capital buffers than would independent asset managers, a disadvantage of a capital requirement that allows sponsors to provide the capital is that it may lead to further consolidation of the MMF industry.

Capital and an MBR rule also would have different implications for investor incentives. By shielding MMF investors from losses, a meaningful capital buffer would shift the incentives for ensuring that MMF risks are well managed from MMF shareholders to the owners of the capital buffer.⁶⁵ In contrast, an MBR would be a buffer provided by MMF shareholders themselves, so it should clarify that they are at risk because they cannot quickly redeem from a troubled MMF. Hence, an MBR rule should strengthen shareholders' incentives for early monitoring of the funds, *before* problems materialize, while market discipline for funds with capital buffers would depend on monitoring by the owners of the buffer.

One clear advantage of an MBR over a capital buffer is the disincentive to run that would result from the strong or effective MBR rule. Indeed, with an MBR, investors can be better off *not* redeeming, even if they recognize that a fund's losses could possibly exceed the size of its aggregate MBR. In contrast, a capital buffer on its own can *reduce* the incentive to run from a fund that is at risk, but investors would still have strong incentives to exit a fund if there is any danger that losses might exceed its buffer. As such, a strong or effective MBR of fraction m of investors' balances may do more to mitigate the vulnerability of MMFs to runs than a standalone capital buffer of the same size.

Notwithstanding the relative strengths of either an MBR rule or capital on its own for mitigating the systemic risks posed by MMFs, these two approaches would complement one another well. For example, an MBR rule would bolster a capital buffer of any size by reducing the chance that it might be eroded by redemption-driven fire sales. For funds with MBRs, *substantial* pre-funded capital buffers might help to sever the systemically risky linkage to sponsors that historically have supported their off-balance-sheet MMFs whenever losses threatened the stable NAV. Without a pre-funded buffer, even with an MBR, expectations of sponsor support likely would remain important.

A capital buffer also would allow MMFs to eliminate the practice of redeeming shares at \$1 apiece when the funds' shadow NAVs are less. That is, with a capital buffer requirement that MMFs maintain shadow values *above* \$1, the funds could be prohibited from redeeming shares at \$1 if shadow NAVs fall below \$1. Such a rule thus would help to ensure that investors' redemptions do not mechanically erode funds' NAVs.

⁶⁵ See note 62. Under some formulations, the owners of that capital might be a new type of MMF investors. For example, if capital is raised by creating a subordinated class of shares that would absorb losses before the stable-value shares, the investors who purchase the new, subordinated class would face strengthened incentives to monitor MMF risks (McCabe, 2011; BlackRock, 2011). In contrast, under formulations in which a small buffer is raised by retaining a portion of the fund's own income over many years (e.g., Goebel, Dwyer, and Messman, 2011), the ownership of the buffer and the effect on incentives for monitoring and avoiding risks are less clear.

A buffer also would preclude the possibility that an MMF might operate for a prolonged period with a shadow NAV below \$1, which would create substantial uncertainty for investors and partially offset the benefits of an MBR.⁶⁶ Shareholders who redeem shares for \$1 from such a fund would concentrate losses over a smaller number of shares and push the shadow NAV lower. Although the effect of one investor's redemptions might be small, investors might be discouraged from redeeming out of concern that their actions might make breaking the buck more likely.⁶⁷ Such a scenario would be impossible if MMFs were required to maintain a buffer that kept shadow NAVs at or above \$1. Indeed, the possibility that a buffer might at least eliminate the practice of rounding NAVs up to \$1 is part of our motivation for analyzing an MBR rule that is accompanied by a small capital buffer.

8.7. How would an MBR rule compare to other forms of redemption fees and restrictions? Redemption fees and restrictions might be imposed at all times or only conditionally, such as when an MMF reaches a specified stress threshold. However, most form of restrictions or fees that are always in place would impair the utility of MMFs without offering significant protection from runs, while conditional restrictions could motivate damaging preemptive runs and thus might do more harm than good. In contrast to these other approaches, an MBR rule can mitigate the risk of runs while preserving investors' ability to use MMFs as a stable-NAV, cash-management product.

Redemption restrictions that are in place at all times might *slow* runs. For example, a delay in disbursement of all redemptions would provide some extra time for an MMF to raise cash to meet heavy redemptions in a crisis. But it would not reduce the *incentive* to redeem, and such a delay probably would substantially reduce the utility of MMFs for cash management. Similarly, a fee on all redemptions would undermine the principal stability of MMFs but, unless the fee was quite large, it might not meaningfully brake redemptions during a crisis. Applying such a fee only to redemptions in excess of some large amount might make it more palatable in normal times, but it would be even less effective in a crisis.

Conditional fees or restrictions would allow MMFs to function much as they currently do in normal times, but would penalize or impede redemptions when MMFs are under strain. For example, a fund might be required to charge a fee on any redemptions if its shadow NAV falls below 99.75 cents or if its "daily liquid assets" are depleted (HSBC, 2011).⁶⁸ But such a rule would give investors—particularly large, sophisticated investors with the most substantial MMF shareholdings and resources—the incentive to monitor MMFs daily and

⁶⁶ In fact, several MMFs that suffered losses during the financial crisis of 2007 and 2008 operated for more than a year with substantially impaired shadow NAVs (Condon, 2010).

⁶⁷ Thus, the combination of an MBR and an impaired NAV might create "zombie" funds that investors fear to leave. That said, when an investor does exit, as long as the fund's NAV does not cross the 99.5 cent threshold for 30 days thereafter, she would avoid losses by leaving them with non-redeeming shareholders.

⁶⁸ Rule 2a-7 requires that all taxable MMFs maintain at least 10 percent of their assets in daily liquid assets, which include cash, securities that mature overnight, and U.S. Treasury securities.

redeem quickly before conditions are met to activate restrictions or fees. As noted in section 8.1, conditional fees or restrictions would preserve investors' incentives to redeem quickly from troubled MMFs but would move up the deadline for redeeming shares at \$1 apiece to avoid all losses. Thus, one serious concern about conditional restrictions is that they could cause preemptive runs that otherwise never would have occurred.

9. Conclusion

This paper describes a new proposal, the minimum balance at risk, to mitigate the risk of runs in MMFs. The MBR would diminish or eliminate the advantages enjoyed by shareholders who redeem quickly—before others do—when an MMF is in distress, by ensuring that redeeming investors are not able to leave risks and losses behind to be absorbed by those who remain invested in a fund. By braking redemptions during crises, the MBR also would help MMFs avoid the need for fire sales of assets to raise cash. Thus, the MBR would not only benefit the funds and their investors, but it also would reduce the contagion risk that one MMF's strains could propagate through the financial system.

A key element of our proposal is that a portion of redeeming investors' MBRs would be subordinated to provide a deterrent to running from an MMF to avoid imminent losses. The disincentive to redeem would be negligible in normal times but would become salient on the rare occasions when investors grow concerned about the risk of losses. When the risk of losses is high, the MBR would create a tradeoff for investors, who could redeem to maintain their own liquidity or stay invested in the fund to safeguard principal. In contrast, under current rules, investors who redeem from a troubled MMF preserve *both* liquidity and principal, while those who remain behind are put at greater risk of losing both.

The MBR offers some important advantages over other proposals for reducing the vulnerability of MMFs to runs. Importantly, the MBR could allow MMFs to maintain features that are central to their attractiveness to investors, particularly their stable \$1 NAVs, their market-based yields, and the immediate liquidity of the vast majorities of investors' balances. The MBR would not require raising the large sums that would be needed to create a meaningful capital buffer, so the MBR likely would be more feasible than a standalone capital option (although a capital buffer could complement an MBR rule well). Moreover, the MBR rule creates a deterrent to redeeming in times of stress that cannot be provided by a floating NAV or a capital buffer. And unlike some proposals for conditional restrictions or fees on redemptions, an MBR rule would not set up incentives for preemptive runs. Indeed, the MBR likely would improve market discipline for MMFs by strengthening investors' incentives to monitor and respond to MMF risks when they first arise, rather than waiting to redeem until serious problems are imminent.

We illustrate how an MBR rule would alter the allocations of losses for those who do and do not redeem and show that, even with subordination, investors may still have incentives to run from distressed funds if expected losses are large or the MBR is too small. We then review historical data on MMF losses and introduce new data from the Treasury and the SEC on losses suffered by MMFs in 2008 to gauge the size of an MBR that would be needed

to protect MMFs from runs. Our analysis also incorporates evidence about the value of preserving liquidity in a crisis, which strengthens investors' incentive to redeem.

Based on the evidence we review, we find that, for an MMF with a 50 basis point capital buffer, a strong MBR rule with an MBR of at least 3 to 4 percent would probably be adequate to create disincentives for redemptions. However, given the uncertainties of estimating expected losses in a distressed MMF and the challenges of pinpointing liquidity costs to investors who lose access to shares in a closed fund, we test the sensitivity of our results to our assumptions. Under optimistic assumptions about expected losses and costs, a strong MBR rule with an MBR of just 2 percent might be adequate to brake redemptions. But under more pessimistic assumptions about expected costs and losses, an MBR might need to be at least 4 percent to brake redemptions. The possibility of very heavy redemptions by other investors strengthens the case for a larger MBR. We also show that the same degree of protection afforded by a strong MBR rule also could be created by an effective rule, which limits the amount of a shareholder's MBR that can be subordinated, if it is paired with a larger MBR.

MMFs present a number of challenges to policymakers. In normal times, they are not only popular cash-management vehicles for retail and institutional investors alike, but they are also important intermediaries for short-term funding. At the same time, MMFs are vulnerable to runs that can disproportionately harm their least flighty investors, interrupt the supply of credit to firms and institutions that rely on short-term funding, and threaten broader financial stability. In this paper, we offer evidence that an MBR rule could significantly reduce the risks that MMFs pose during crises while preserving their important economic roles in other times.

The run on MMFs in September 2008 provided clear evidence that the structural vulnerability of the funds to runs can have potentially deleterious consequences for the entire financial system. In addition, the 2008 run may have made the funds' institutional investors more skittish and thus more run-prone. Heavy redemptions by institutional investors during the summer of 2011, partly in response to increasing concerns about the funds' European holdings, indicate that institutional investors are more responsive to the risks in MMFs than they were before 2008.⁶⁹ But policymakers have *fewer* tools available now than in 2008 to address MMF runs; in particular, the Treasury's Temporary Guarantee Program for Money Market Funds, which effectively halted the run in 2008, would no longer be possible under current law.⁷⁰

⁶⁹ The outflows in 2011 came despite the fact that the European exposures caused no losses for MMFs (see Chernenko and Sunderam, 2012). In contrast, in August 2007, as the ABCP crisis unfolded and many MMFs suffered substantial losses due to holdings of distressed ABCP, institutional investors generally did *not* redeem MMF shares (McCabe, 2010).

⁷⁰ The Emergency Economic Stabilization Act of 2008 specifically prohibited the Treasury from using again the mechanism it had employed to finance the 2008 guarantee for MMFs.

In this environment, the need for MMF reform is particularly salient. The additional stability that an MBR rule could provide to MMFs also would enhance financial stability. Moreover, an MBR would help to ensure that MMFs can continue to play an important role in the financial system, even if another financial crisis should occur.

References

- Anderson, Neil**, "Less Than One Percent of Primary's Assets Are Still in Reserve," *The Wall Street Journal*, 16 July 2010. <http://www.mfwire.com/article.asp?storyID=32827&template=article&bhcp=1>.
- Bank of England**, "An indicative decomposition of Libor spreads," *Quarterly Bulletin*, Q4 2007, 47 (4), 498–499.
- BlackRock**, "ViewPoint: Money Market Funds. A Proposal for a Capitalized Special Purpose Entity," 7 February 2010. https://www2.blackrock.com/webcore/litService/search/getDocument.seam?venue=PUB_IND&source=GLOBAL&contentId=1111124986.
- , "Money Market Funds: Potential Capital Solutions," August 2011. https://www2.blackrock.com/webcore/litService/search/getDocument.seam?venue=PUB_IND&source=GLOBAL&contentId=1111146889.
- Board of Governors of the Federal Reserve System**, "Monetary Policy Report to the Congress," 24 February 2009. http://www.federalreserve.gov/monetarypolicy/mpr_20090224_part2.htm.
- Chernenko, Sergey and Adi Sunderam**, "The Quiet Run of 2011: Money Market Funds and the European Debt Crisis," mimeo, 12 May 2012.
- Condon, Christopher**, "Money Funds Get 'Micro Bailouts' as Disclosure Rules Take Effect," *Bloomberg.com*, 9 November 2010. <http://www.bloomberg.com/news/2010-11-09/money-funds-resorting-to-micro-bailouts-as-disclosure-rules-take-effect.html>.
- Edelen, Roger M.**, "Investor Flows and the Assessed Performance of Open-End Mutual Funds," *Journal of Financial Economics*, September 1999, 53 (3), 439–466.
- European Parliament**, "Directive of the European Parliament and of the Council on the coordination of laws, regulations and administrative provisions relating to undertakings for collective investment in transferable securities (UCITS)," register.consilium.europa.eu/pdf/en/09/st03/st03605-re01.en09.pdf, 19 June 2009.
- Federal Deposit Insurance Corporation**, "Insurance Fund Indicators," *Quarterly Banking Profile*, Q1 2012, 6 (2), 15–16.
- Federal Open Market Committee (FOMC)**, "Minutes of the Federal Open Market Committee," 28-29 October 2008. www.federalreserve.gov/monetarypolicy/files/fomcminutes20081029.pdf.

Financial Stability Oversight Council (FSOC), *Annual Report*, 2011.

Goebel, Scott C., Carrie E. Dwyer, and C. David Messman, Fidelity, Charles Schwab, and Wells Fargo letter to the Securities and Exchange Commission re: File No. 4-619; Release No. IC-29497 President's Working Group Report on Money Market Fund Reform, available at <http://www.sec.gov/comments/4-619/4619-97.pdf>, 3 May 2011.

Group of Thirty, *Financial Reform: A Framework for Financial Stability*, 15 January 2009.

HSBC Global Asset Management, "Liquidity Fees: A proposal to reform money market funds," 3 November 2011.

Investment Company Institute, "Report of the Money Market Working Group," available at http://www.ici.org/pdf/ppr_09_mmwg.pdf, 17 March 2009.

—, "Investment Company Fact Book," available at http://www.ici.org/pdf/2012_factbook.pdf, 2012.

Kacperczyk, Marcin and Philipp Schnabl, "How Safe are Money Market Funds?," mimeo, April 2012.

Lacker, Jeffrey M., Letter to the Securities and Exchange Commission re: File No. 4-619; Release No. IC-29497 President's Working Group Report on Money Market Fund Reform, available at <http://www.sec.gov/comments/4-619/4619-54.pdf>, 10 January 2011.

Lyon, Andrew B., "Money Market Funds and Shareholder Dilution," *Journal of Finance*, September 1984, 39 (5), 1011–1020.

McCabe, Patrick E., "The Cross Section of Money Market Fund Risks and Financial Crises," Working Paper 2010-51, Federal Reserve Board Finance and Economics Discussion Series, September 2010. Available at <http://www.federalreserve.gov/pubs/feds/2010/201051/201051pap.pdf>.

—, "An A-/B-share Capital Buffer Proposal for Money Market Funds," mimeo, 4 June 2011.

McCoy, Kevin, "Primary Fund shareholders put in a bind," *USA Today*, 11 November 2008. http://www.usatoday.com/money/perfi/funds/2008-11-11-market-fund-side_N.htm.

—, "Reserve fund's woes undercut mantra of safety, liquidity," *USA Today*, 14 November 2008. http://www.usatoday.com/money/perfi/funds/2008-11-10-reserve-primary-money-market-fund_N.htm.

Mendelson, Simon and Richard Hoerner, BlackRock letter to the Securities and Exchange Commission re: Release No. IC-29497; File No. 4-619, available at <http://www.sec.gov/comments/4-619/4619-22.pdf>, 10 January 2011.

Moody's Investors Service, "Sponsor Support Key to Money Market Funds," 8 August 2010.

– , "Lehman Support in Prime Money Market Funds," mimeo, 30 April 2012.

Musto, David, Greg Nini, and Krista Schwarz, "Notes on Bonds: Liquidity at All Costs in the Great Recession," mimeo, 11 October 2011.

Plaze, Robert, Letter to Karrie McMillan, General Counsel, Investment Company Institute, <http://www.sec.gov/divisions/investment/noaction/2008/ici101008.htm>, 10 October 2008.

Poskitt, Russell, "Do liquidity or credit effects explain the behavior of the LIBOR-OIS spread?," mimeo, 25 May 2011.

President's Working Group on Financial Markets (PWG), "Report of the President's Working Group on Financial Markets: Money Market Fund Reform Options," Available at <http://www.treasury.gov/press-center/press-releases/Documents/10.21%20PWG%20Report%20Final.pdf>, October 2010.

Rosengren, Eric S., "Money Market Mutual Funds and Financial Stability: Remarks at the Federal Reserve Bank of Atlanta's 2012 Financial Markets Conference," <http://www.bos.frb.org/news/speeches/rosengren/2012/041112/041112.pdf>, 11 April 2012.

Schapiro, Mary, "Statement on Money Market Funds Before the Open Commission Meeting," <http://www.sec.gov/news/speech/2010/spch012710mls-mmf.htm>, 27 January 2010.

– , "Remarks at SIFMA's 2011 Annual Meeting," <http://www.sec.gov/news/speech/2011/spch110711mls.htm>, 7 November 2011.

– , "Testimony on 'Perspectives on Money Market Mutual Fund Reforms' before U.S. Senate Committee on Banking, Housing, & Urban Affairs," http://banking.senate.gov/public/index.cfm?FuseAction=Hearings.Hearing&Hearing_ID=bba4146c-6b7f-47d0-93bc-ebc73189c9c0, 21 June 2012.

Squam Lake Group, "Reforming Money Market Funds," Letter to the Securities and Exchange Commission re: File No. 4-619; Release No. IC-29497 President's Working Group Report on Money Market Fund Reform, available at <http://www.sec.gov/comments/4-619/4619-57.pdf>, 14 January 2011.

Stevens, Paul Schott, Investment Company Institute letter to the Securities and Exchange Commission re: President's Working Group Report on Money Market Fund Reform Options (File No. 4-619), available at <http://www.sec.gov/comments/4-619/4619-49.pdf>, 10 January 2011.

U.S. Department of the Treasury, "Guarantee Agreement (Stable Value Single Fund)," www.treasury.gov/about/organizational-structure/offices/Domestic-Finance/Documents/Guarantee_Agreement_Stable-Value_Single-Fund.pdf, September 2008.

U.S. Securities and Exchange Commission, "Final Rule: Registration Form Used by Open-End Management Investment Companies," available at www.sec.gov/rules/final/33-7512r.htm, 18 March 1998. File No. S7-10-97.

—, "In the Matter of Craig S. Vanucci and Brian K. Andrew, Respondents: Order Instituting Public Administrative and Cease-and-Desist Proceedings," 11 January 1999. Administrative Proceeding File No. 3-9804.

—, "Final Rule: Amendments to Investment Company Advertising Rules," available at www.sec.gov/rules/final/33-8294.htm, 1 October 2003. File No. S7-17-02.

—, "Money Market Fund Reform: Proposed Rule," available at www.sec.gov/rules/proposed/2009/ic-28807.pdf, 30 June 2009. Release no. IC-28807.

—, "Securities and Exchange Commission v. Reserve Management Company, Inc. et al," 5 May 2009. Civil Action No. 09-CV-4346.

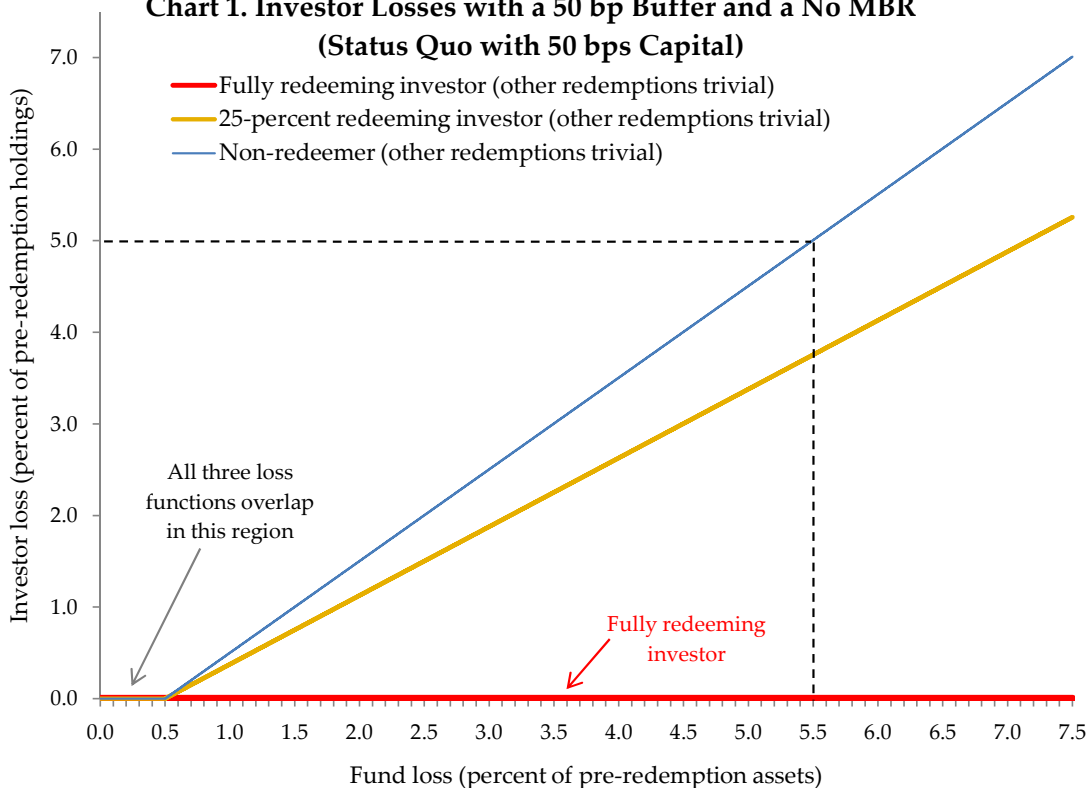
—, "Securities and Exchange Commission v. Reserve Management Company, Inc. et al," 8 June 2009. Case 1:09-cv-04346-PGG.

—, "Money Market Fund Reform: Final Rule," available at www.sec.gov/rules/final/2010/ic-29132.pdf, 23 February 2010. Release no. IC-29132.

Volcker, Paul A., Letter to the Securities and Exchange Commission re: File No. 4-619; Release No. IC-29497 President's Working Group Report on Money Market Fund Reform, available at <http://www.sec.gov/comments/4-619/4619-79.pdf>, 11 February 2011.

Wall Street Journal, "Taxpayers and Money Market Funds [Editorial]," 9 May 2011.

**Chart 1. Investor Losses with a 50 bp Buffer and a No MBR
(Status Quo with 50 bps Capital)**



**Chart 2. Investor Losses with a 50 bp Buffer and No MBR
(Status Quo with 50 bps Capital)**

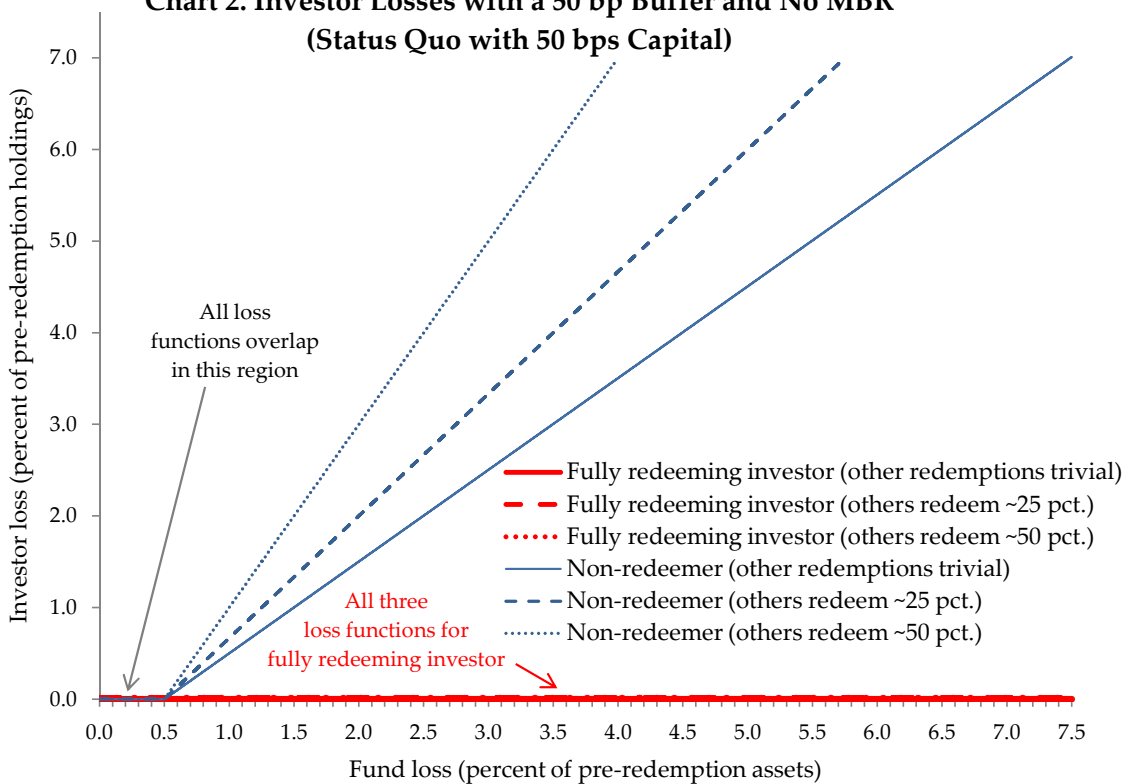


Chart 3. Investor Losses with a 50 bp Buffer and a 5 Percent MBR ("Simple" MBR Rule: No Subordination)

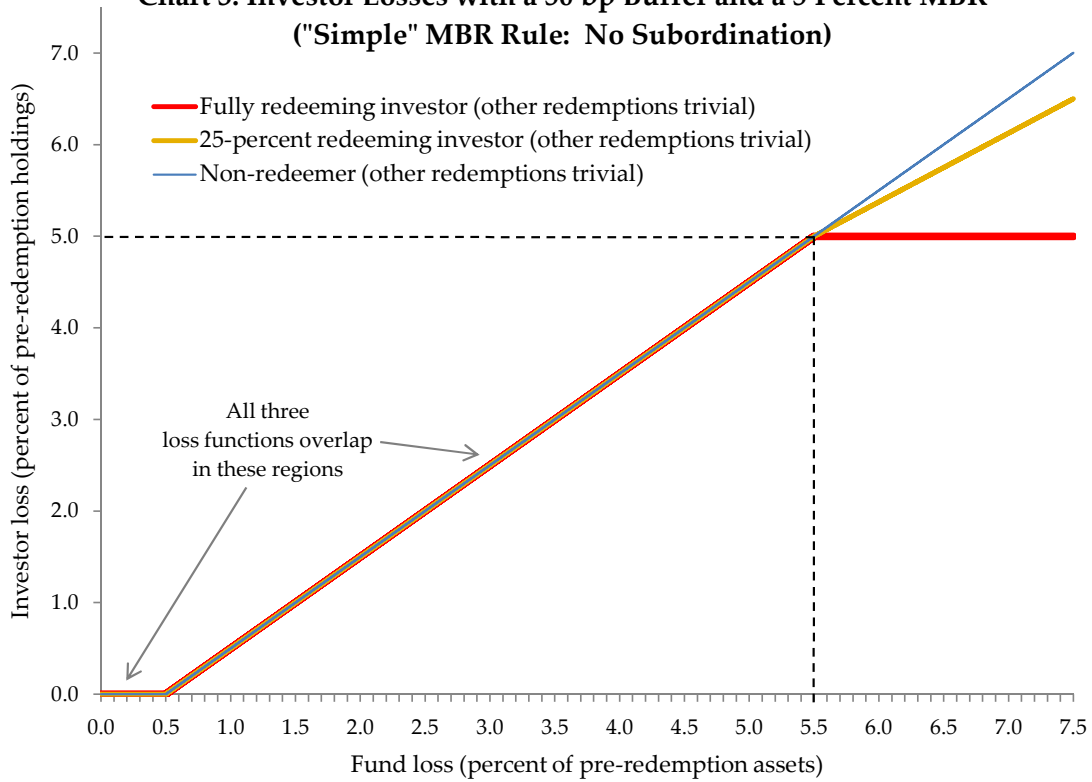


Chart 4. Investor Losses with a 50 bp Buffer and a 5 Percent MBR ("Simple" MBR Rule: No Subordination)

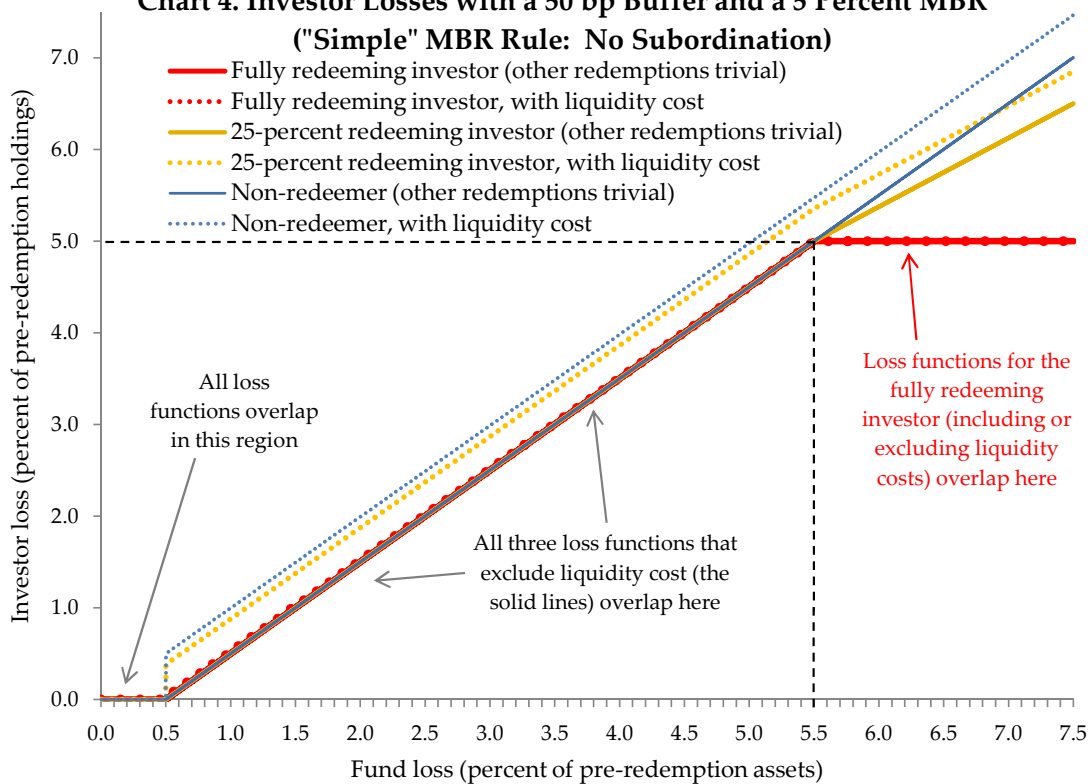


Chart 5. Investor Losses with a 50 bp Buffer and a 5 Percent MBR ("Simple" MBR Rule: No Subordination)

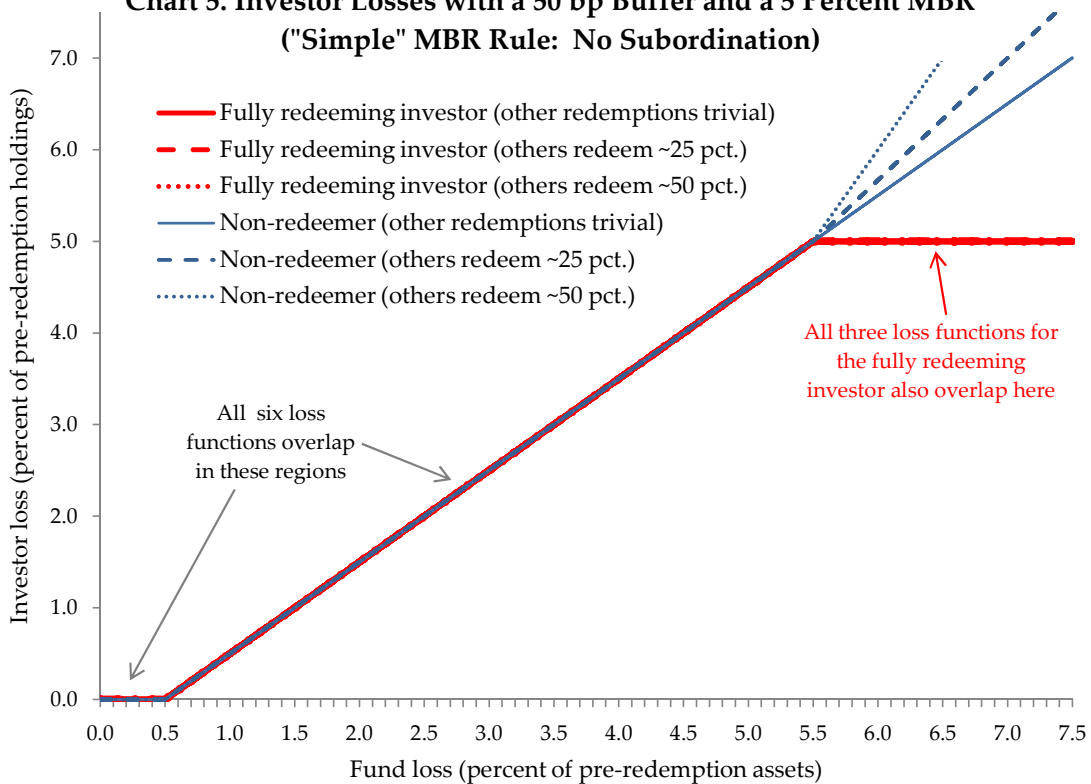


Chart 6. Investor Minimum Balances at Risk and Subordinated Positions under a Strong MBR Rule

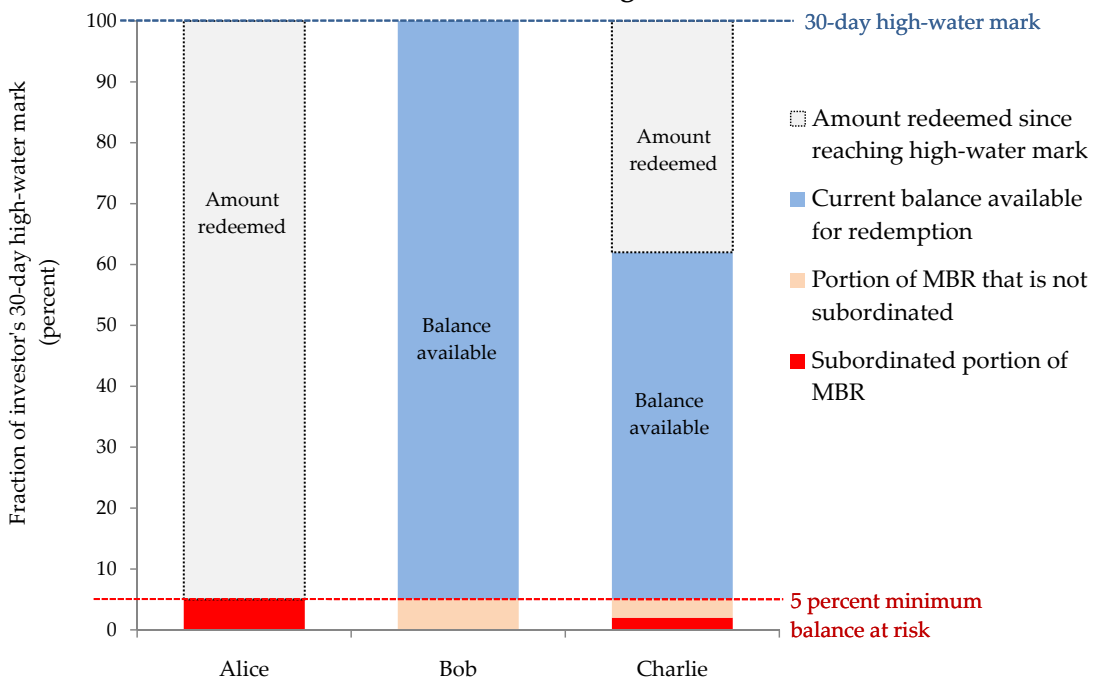


Chart 7. Composition of MMF \$1 NAV: Ranking of Claims

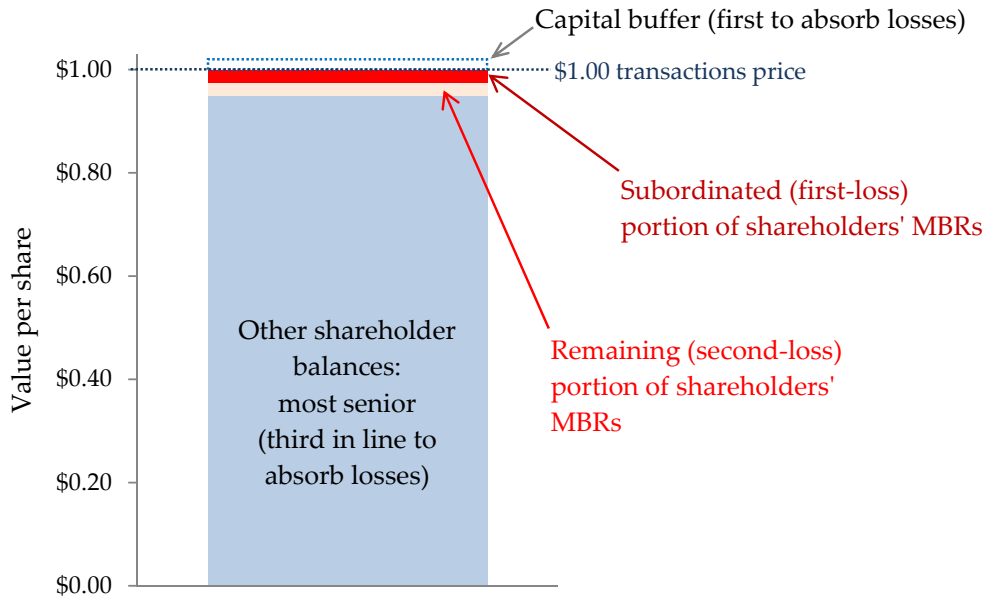


Chart 8. Investor Losses with a 50 bp Buffer and a 5 Percent MBR ("Strong" MBR Rule for Subordination)

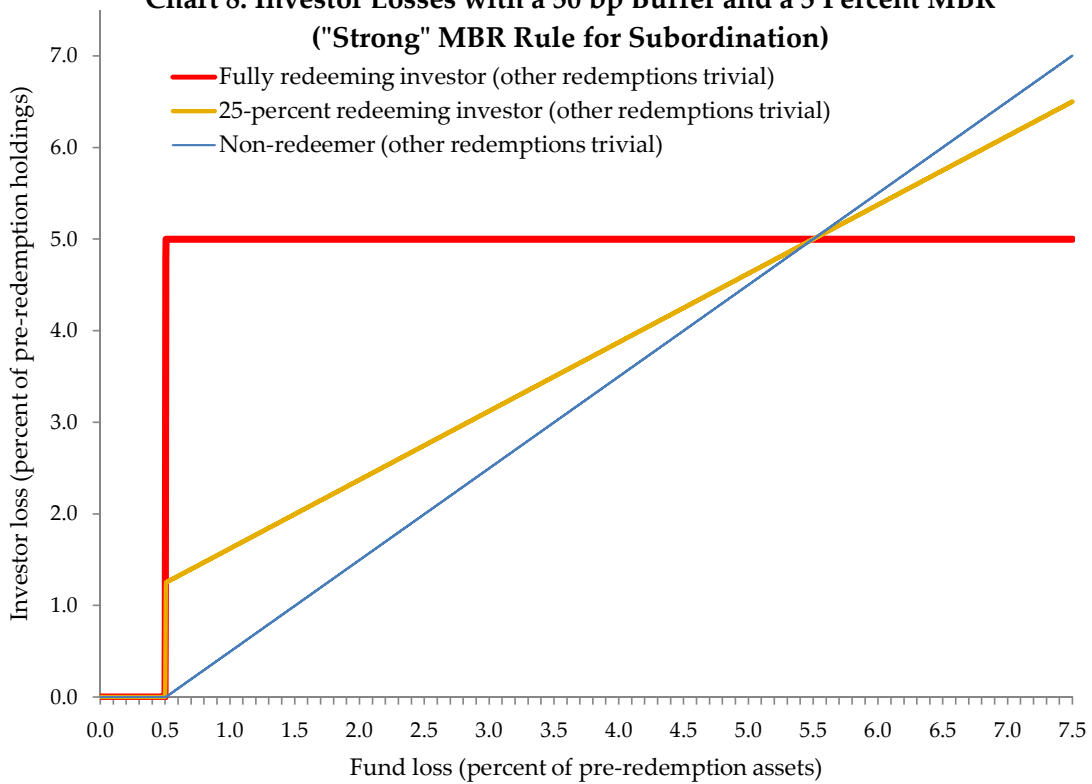


Chart 9. Investor Losses with a 50 bp Buffer and a 5 Percent MBR ("Strong" MBR Rule for Subordination)

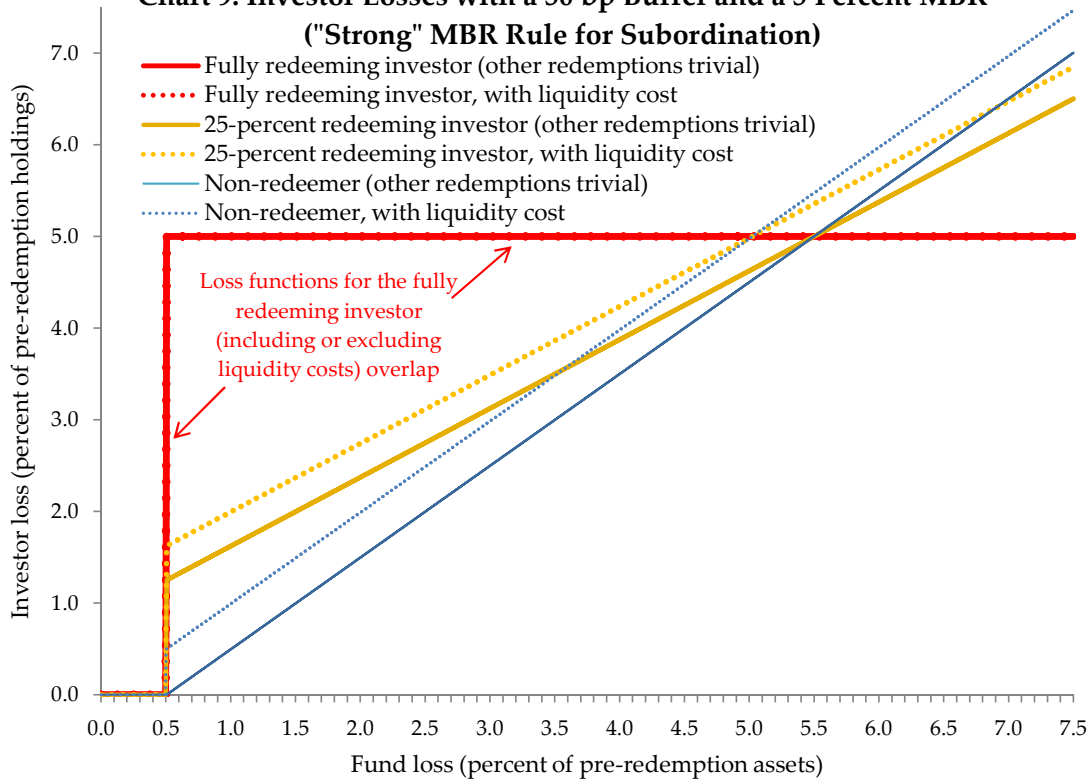


Chart 10. Investor Losses with a 50 bp Buffer and a 5 Percent MBR ("Strong" MBR Rule for Subordination)

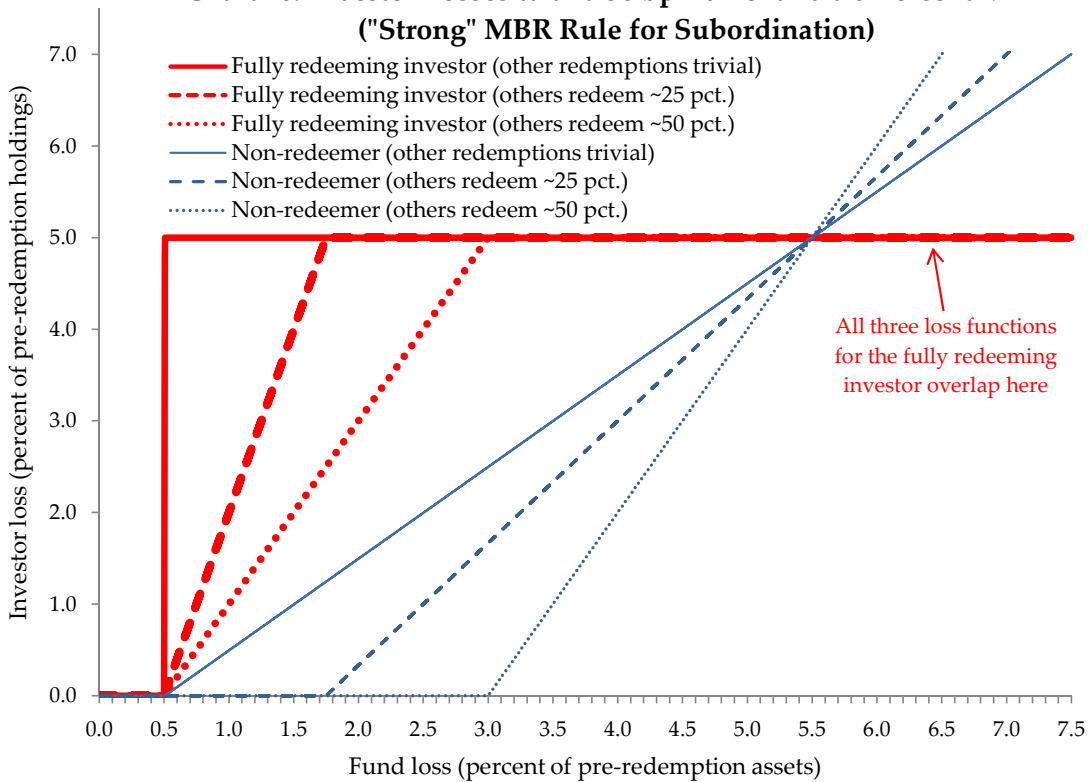


Chart 11. Investor Losses with a 50 bp Buffer and a 5 Percent MBR
 ("Effective" MBR Rule: 3 Pctg. Pts. of MBR Can Be Subordinated)

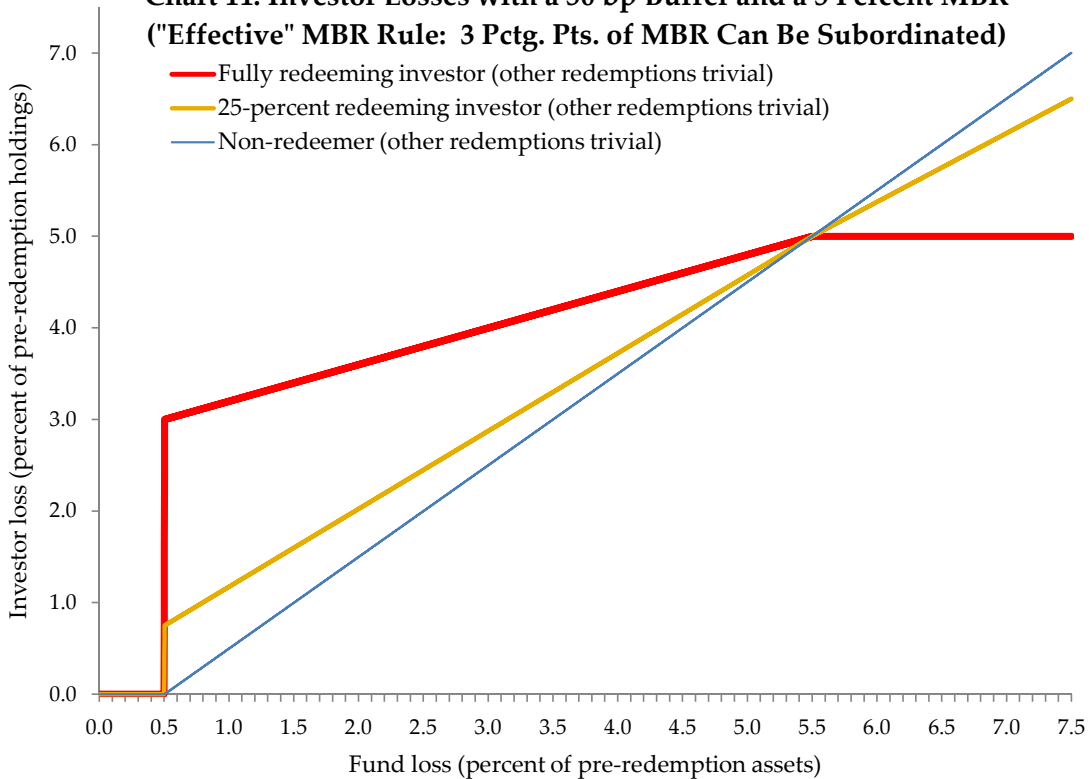


Chart 12. Investor Losses with a 50 bp Buffer and a 5 Percent MBR
 ("Effective" MBR Rule: 3 Pctg. Pts. of MBR Can Be Subordinated)

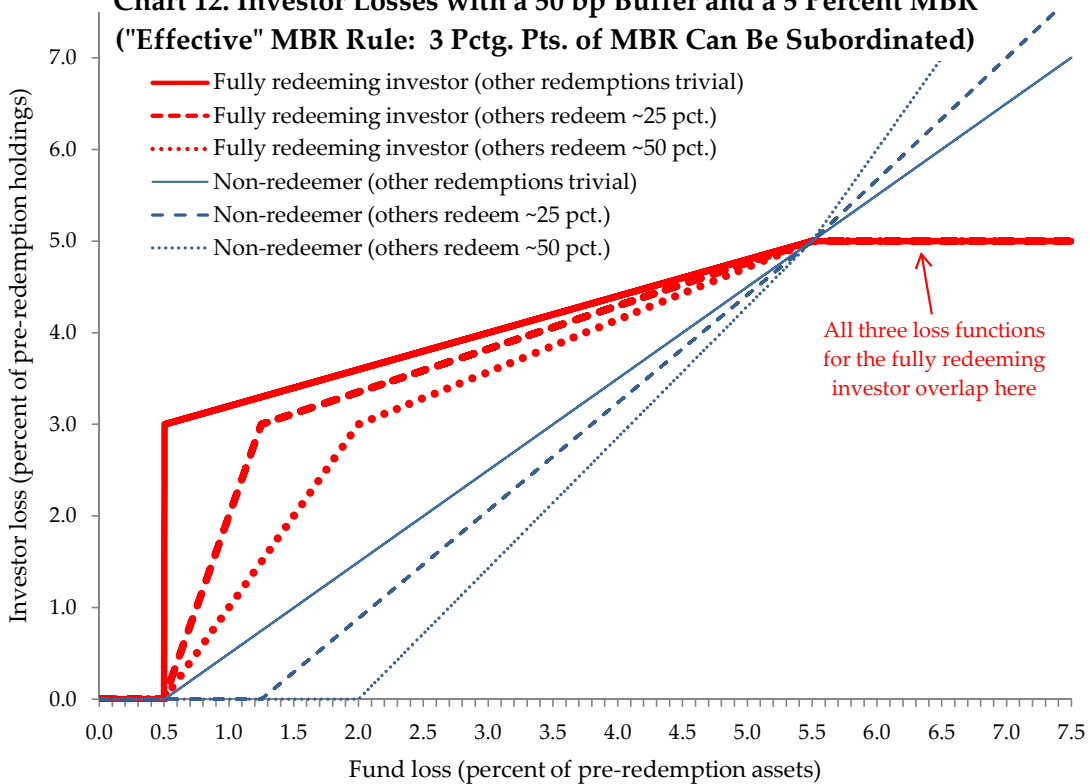
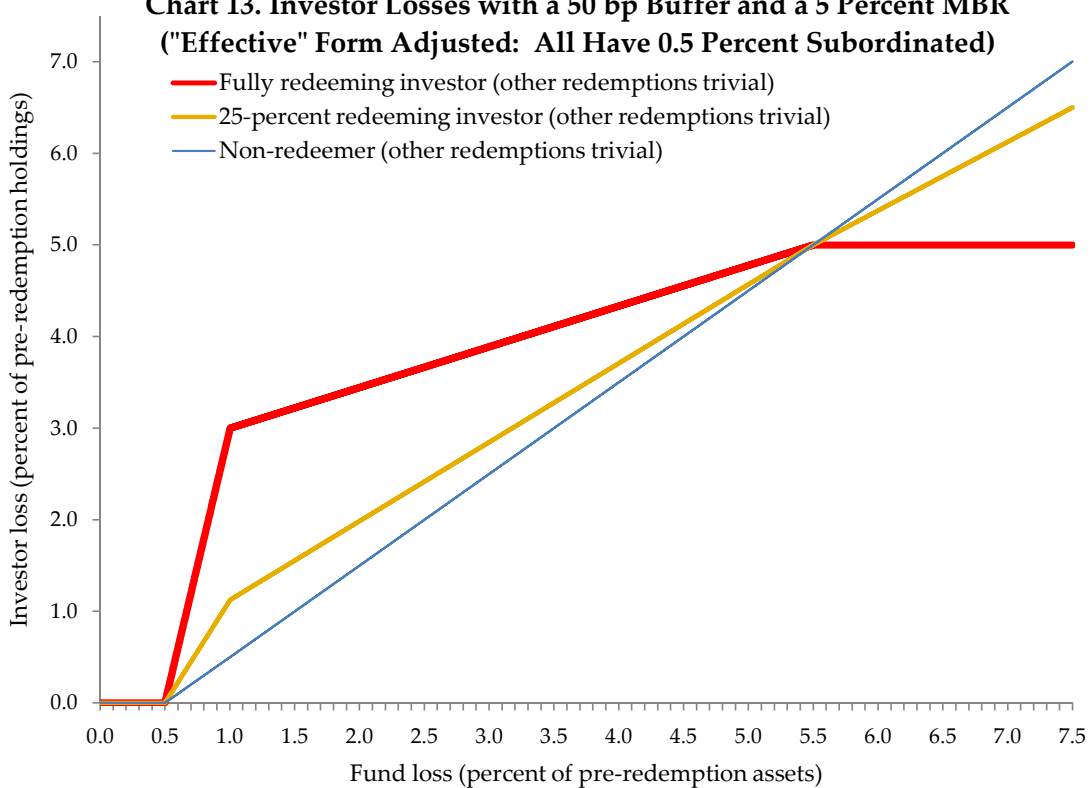


Chart 13. Investor Losses with a 50 bp Buffer and a 5 Percent MBR
("Effective" Form Adjusted: All Have 0.5 Percent Subordinated)



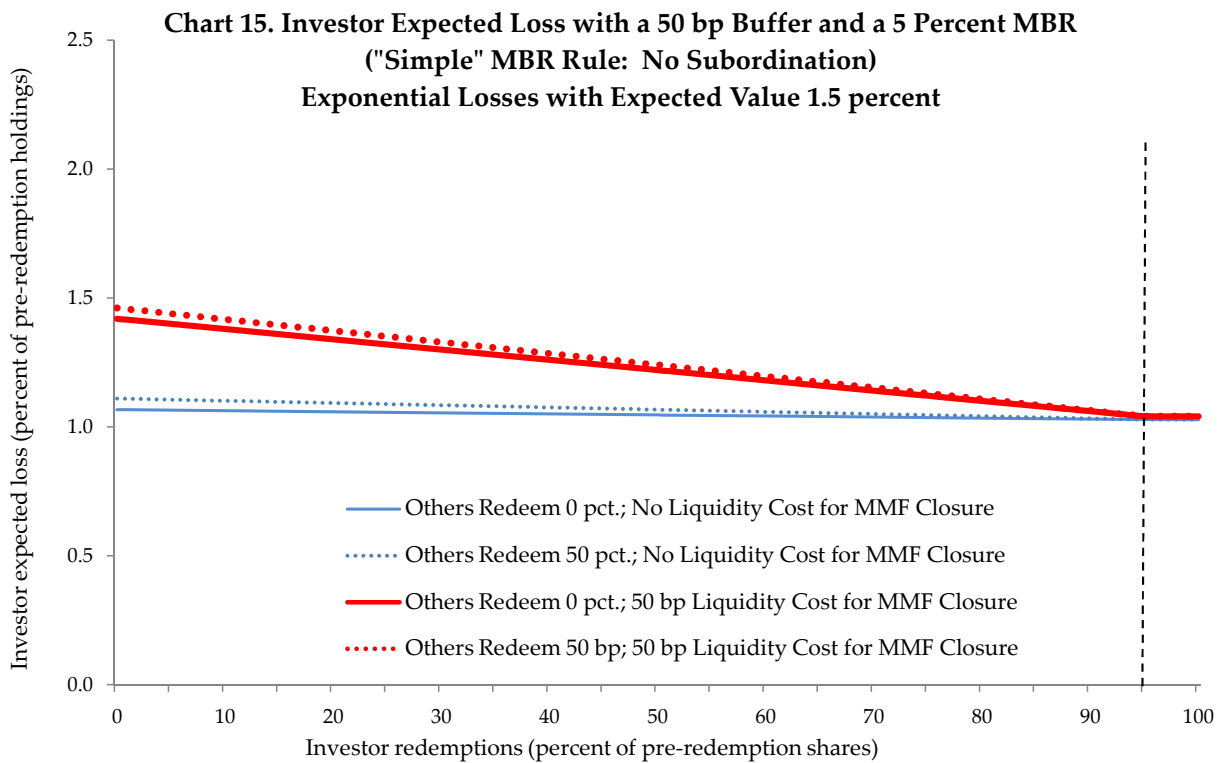
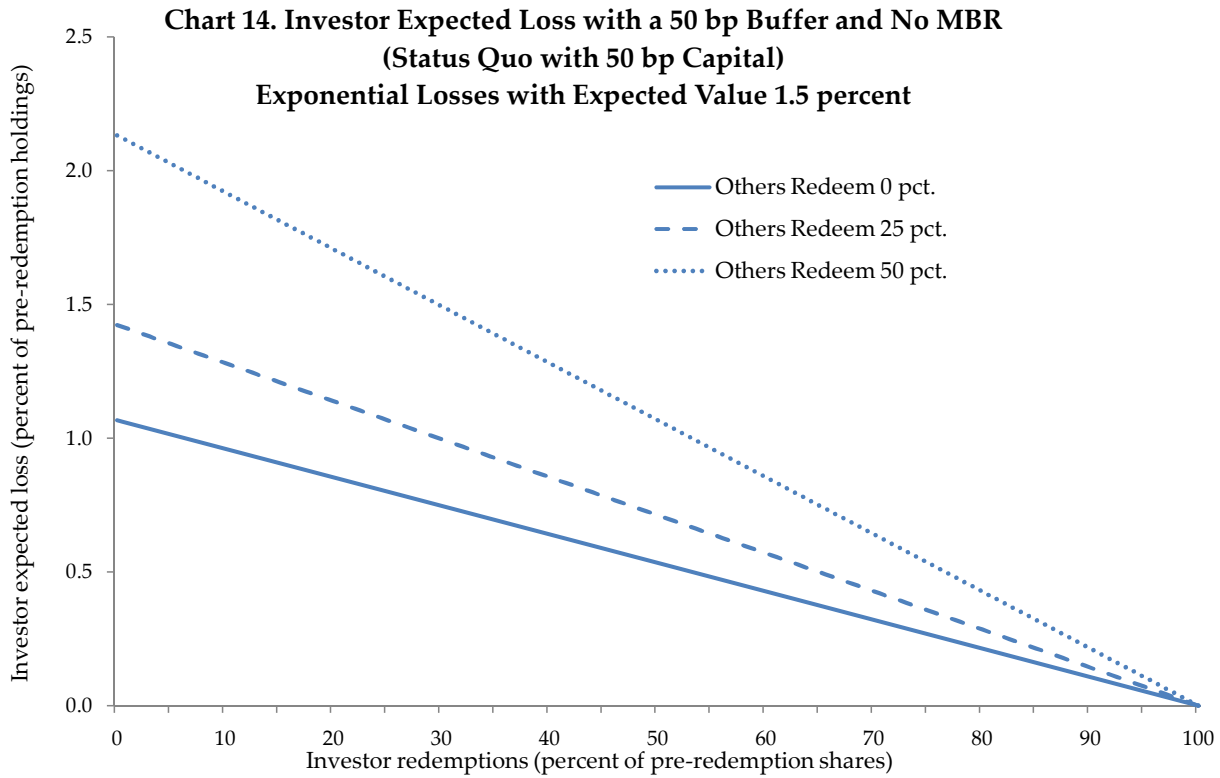


Chart 16. Investor Expected Loss with a 50 bp Buffer and a 5 Percent MBR
Exponential Losses with Expected Value 1.5 percent
("Effective" MBR Rule: 3 Pctg. Pts. of MBR can be Subordinated)
50 bp Liquidity Cost for MMF Closure

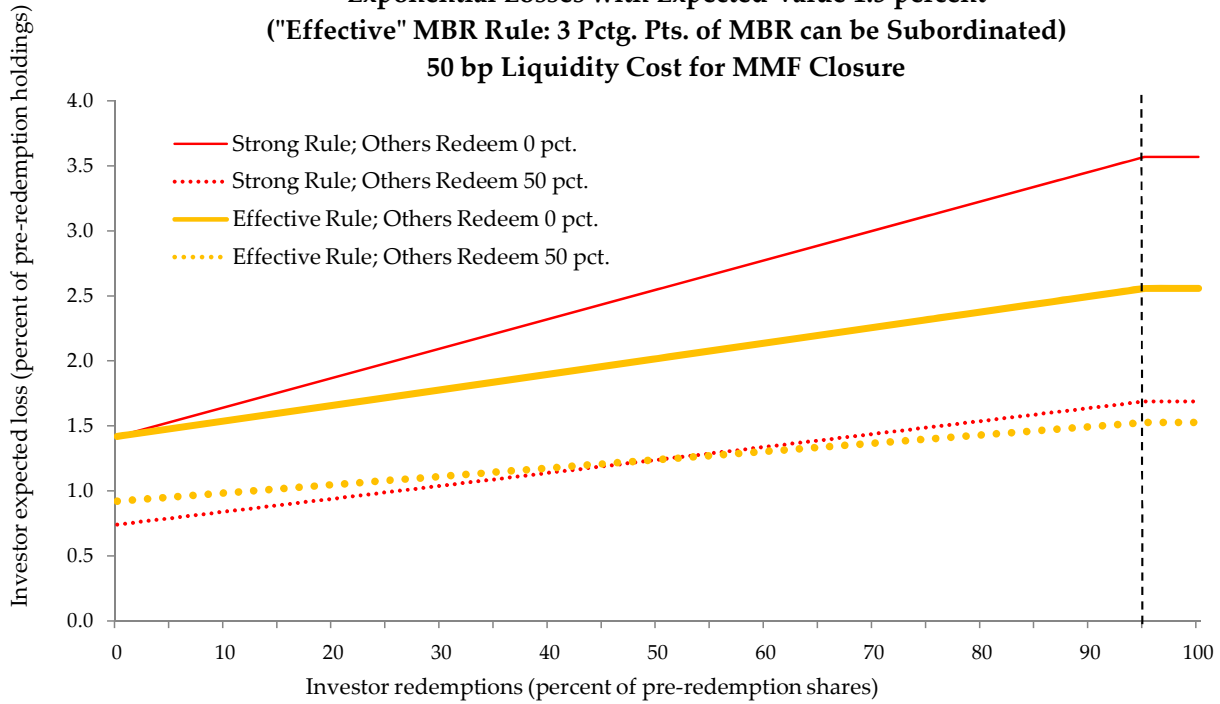
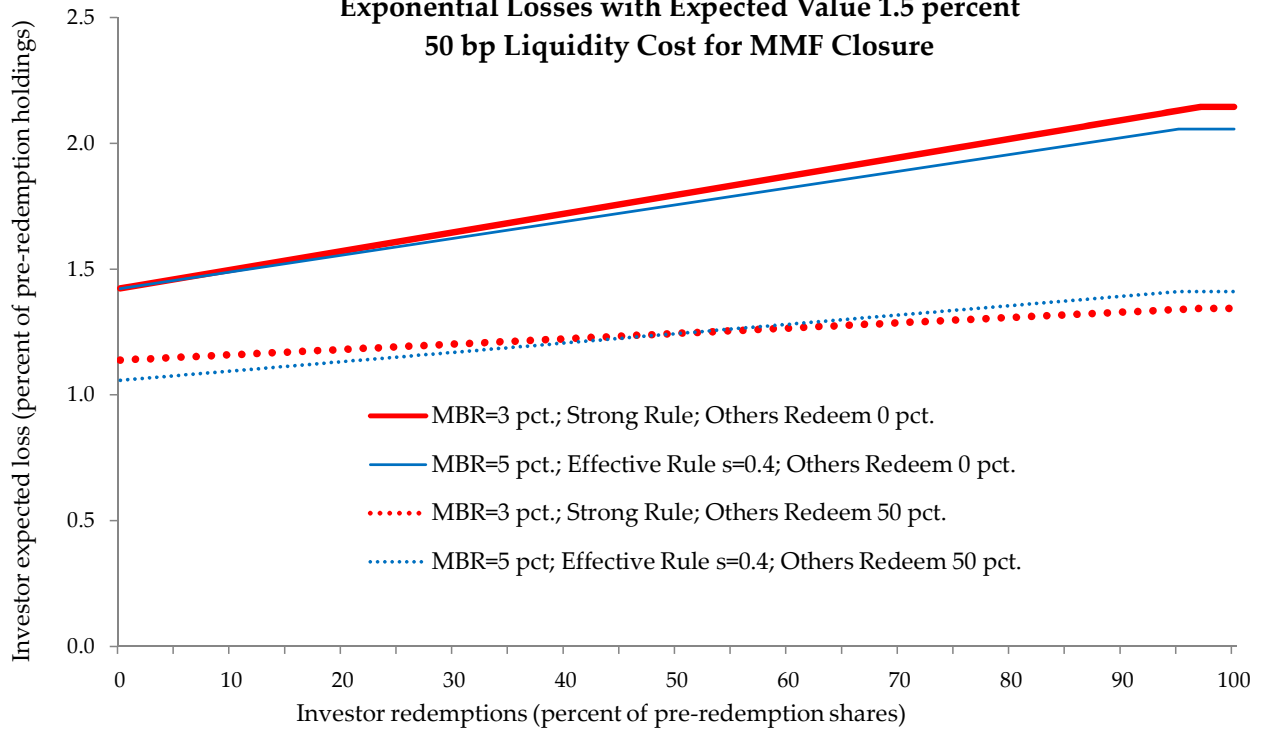
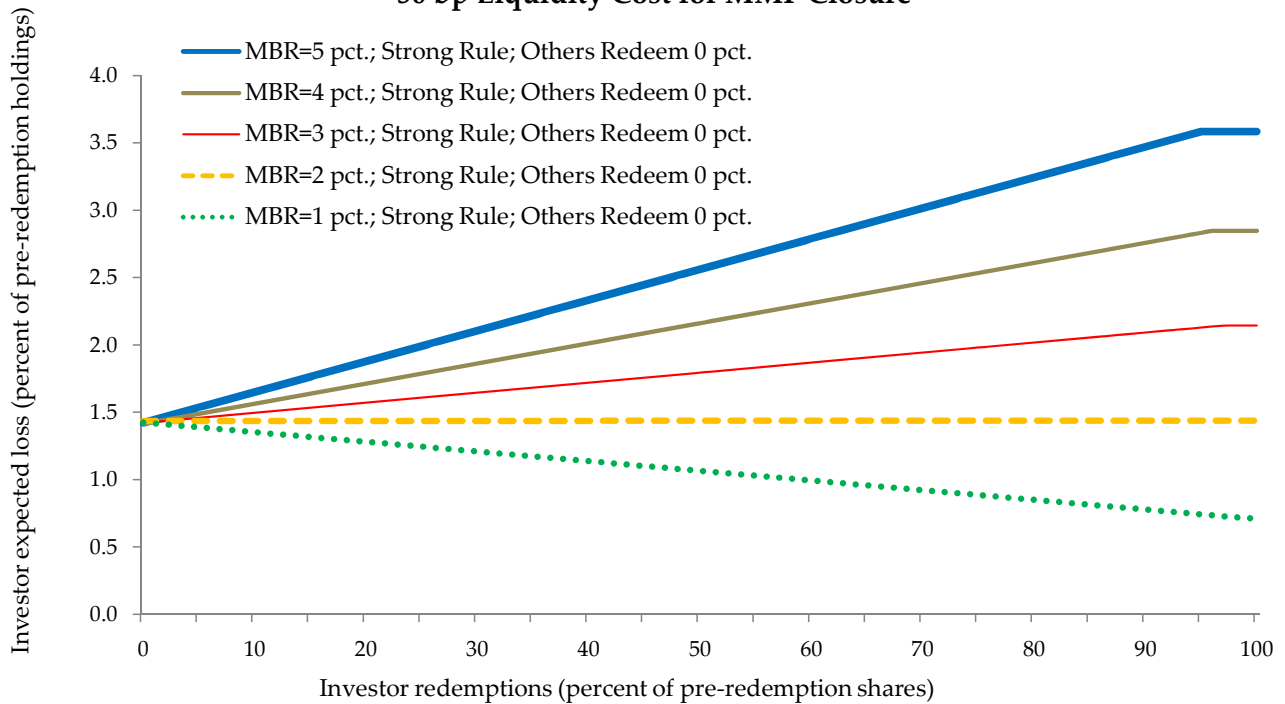


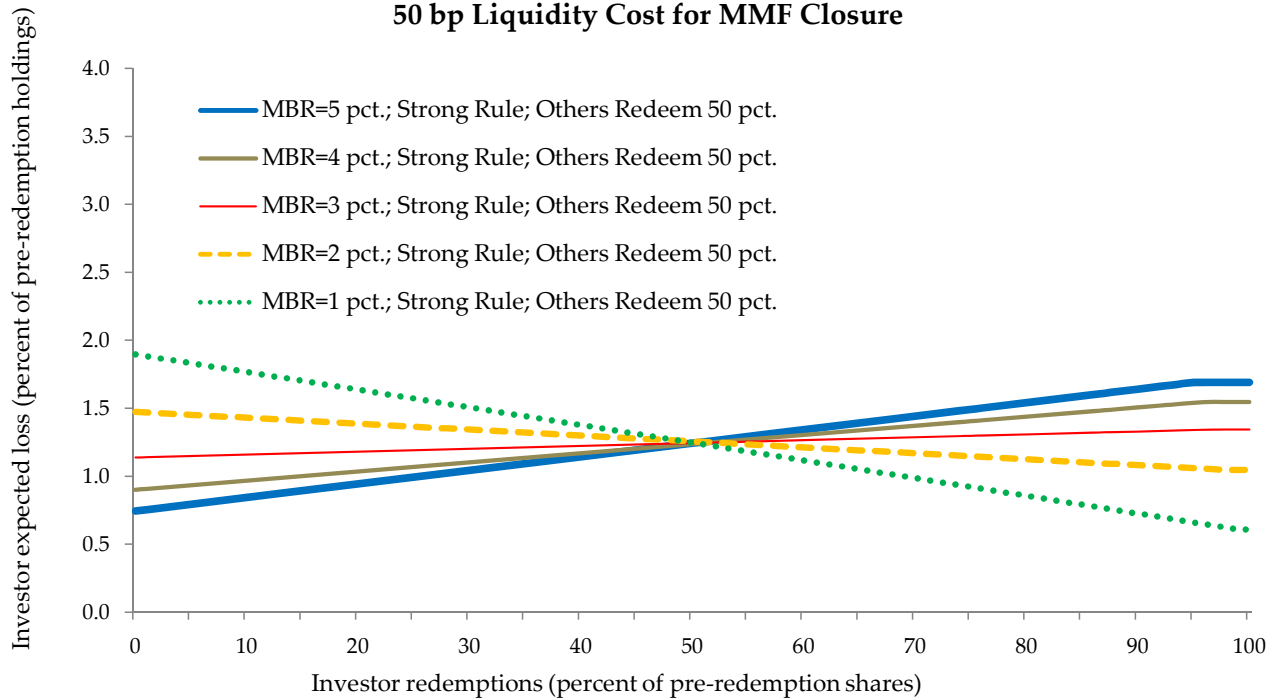
Chart 17. Investor Expected Loss with a 50 bp Buffer
Exponential Losses with Expected Value 1.5 percent
50 bp Liquidity Cost for MMF Closure



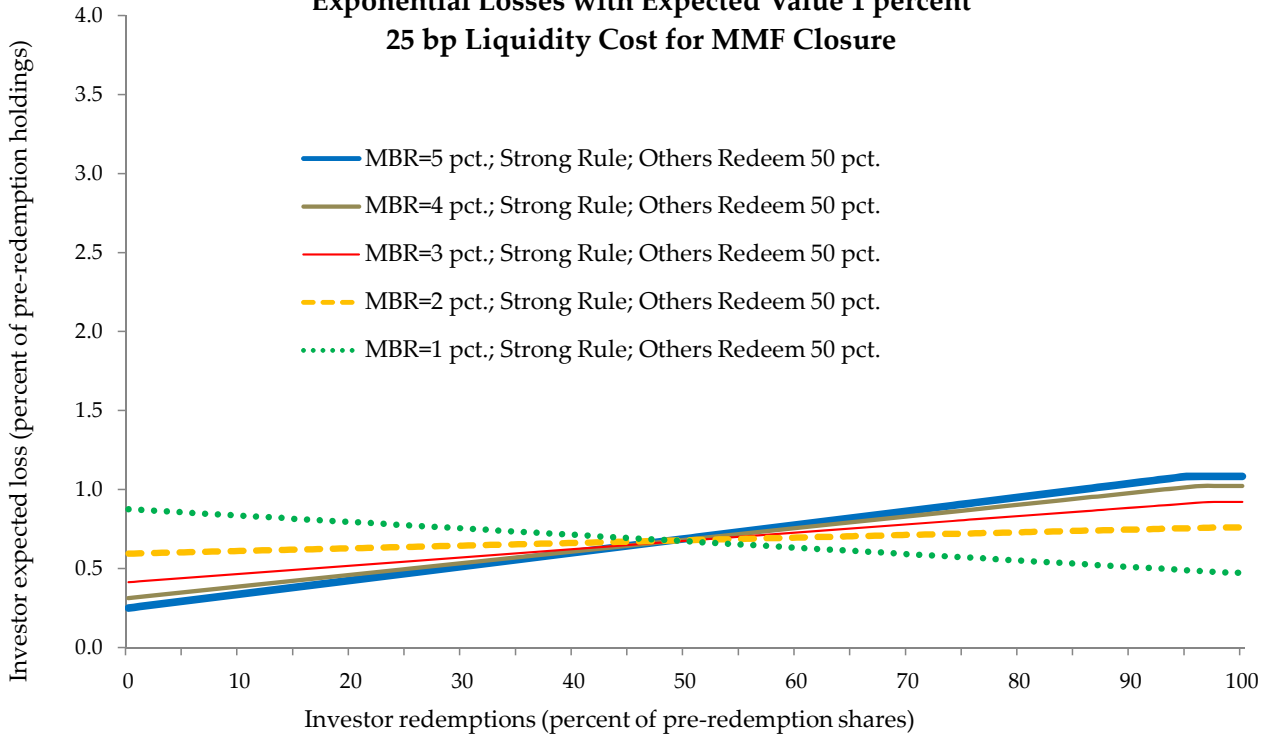
**Chart 18. Investor Expected Loss with a 50 bp Buffer
Exponential Losses with Expected Value 1.5 percent
50 bp Liquidity Cost for MMF Closure**



**Chart 19. Investor Expected Loss with a 50 bp Buffer
Exponential Losses with Expected Value 1.5 percent
50 bp Liquidity Cost for MMF Closure**



**Chart 20. Investor Expected Loss with a 50 bp Buffer
Exponential Losses with Expected Value 1 percent
25 bp Liquidity Cost for MMF Closure**



**Chart 21. Investor Expected Loss with a 50 bp Buffer
Exponential Losses with Expected Value 2 percent
100 bp Liquidity Cost for MMF Closure**

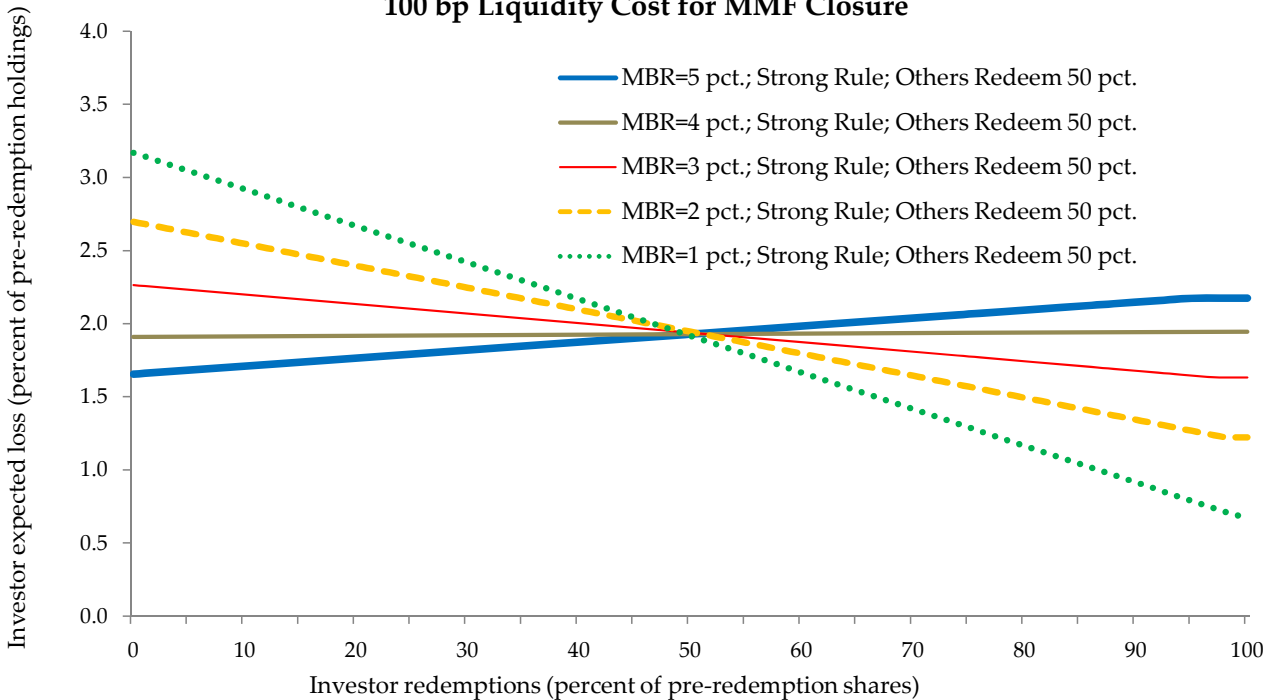


Chart 22. Investor Expected Loss with a 50 bp Buffer
Exponential Losses with Expected Value 1.5 percent
50 bp Liquidity Cost for MMF Closure

