Economic Quarterly—Volume 98, Number 4—Fourth Quarter 2012—Pages 231–254

Some Theoretical Considerations Regarding Net Asset Values for Money Market Funds

Huberto M. Ennis

On Tuesday, September 16, 2008, the day after Lehman Brothers filed for bankruptcy, the Reserve Primary Fund, a large prime money market fund, announced that it would not be able to redeem investors' funds one for one. The fund had "broken the buck" mainly due to losses on its holdings of Lehman's debt instruments. In the days that followed, outflows from prime money funds spiked, with investors withdrawing, in the space of a week, approximately \$300 billion—roughly 15 percent of total assets invested in these funds at the time (Financial Stability Oversight Council 2012). By Friday of that week, the U.S. Treasury and the Federal Reserve would decide to implement several major interventions aimed at stabilizing the money market funds industry. While outflows did, in fact, slow down in the following weeks, money funds continued divesting large amounts of commercial paper and other assets for some time.

The interventions announced by the U.S. Treasury and the Federal Reserve on September 19, 2008, were broad and unprecedented. The Temporary Guarantee Program adopted by the Treasury Department guaranteed that shareholders of those funds opting to participate would receive the fund's stable net asset value (NAV) per share were the fund to suspend redemptions and fully liquidate. At the same time, the

I would like to thank Todd Keister, Jeff Lacker, Jon Lecznar, Ned Prescott, Zhu Wang, and Alex Wolman for comments on an earlier draft. All errors and imprecisions are of course my exclusive responsibility. The views expressed in this article are those of the author and do not necessarily represent the views of the Federal Reserve Bank of Richmond or the Federal Reserve System. E-mail: huberto.ennis@rich.frb.org.

Federal Reserve created the Asset-Backed Commercial Paper Money Market Mutual Funds Liquidity Facility that was used to extend central bank credit to banks buying high-quality asset-backed commercial paper from money market funds (see Duygan-Bump et al. [2013]).

Money market funds (or, money funds, for short) are open-end mutual funds that invest in short-term high-credit-quality debt instruments such as commercial paper, large certificates of deposit, Treasury bonds, and repurchase agreements. Most money funds maintain a stable redemption value of shares, usually set at a value equal to one, and pay dividends that reflect the prevailing short-term interest rates. As of September 2012, there were 632 money market funds in the United States with total assets under management of approximately $2.9 trillion. In comparison, deposits at banking institutions amount to about $11 trillion. So, the size of the U.S. money market fund industry is significant.

SEC rule 2a-7 pursuant to the Investment Company Act of 1940 provides the regulatory framework for these funds. The rule permits funds to use the amortized cost method of valuation to compute net asset values and allows the funds to round such value to the nearest 1 percent.[1] The possibility of stable net asset values is a consequence of these provisions. At the same time, the rule puts limitations on the type of assets that the funds can hold: Funds must hold low-risk investment instruments with remaining maturity no longer than a given maximum date.

Within the broader category of money market funds, there are different sub-categories based on the main investments taken by the funds. Prime money funds hold predominantly private debt instruments. Government funds, instead, are restricted to invest only in government-issued securities. Prime money funds tend to be more exposed to credit risk (Rosengren 2012) and they are the ones that experienced serious financial distress during the second half of 2008.

In February 2010, partly as a response to the problems with prime money funds during the crisis, the Securities and Exchange Commission (SEC) adopted amendments to rule 2a-7 intended to make money funds more resilient and less likely to break the buck. The changes tightened restrictions on the amount of risk that money funds can assume and, for the first time, required that money funds maintain liquidity buffers to help them withstand sudden demands for redemptions. The new

[1] The amortized cost method allows the funds to value assets at their acquisition cost rather than market value, and interest earned on the asset is accrued uniformly over the maturity of the asset (adjusting for amortization of any premium or accretion of any discount involved upon purchase).

rules also enhanced information disclosure by funds and provided a framework for the liquidation of funds that break the buck and suspend redemptions.

Even after the wide-ranging revisions of rule 2a-7 in 2010, many policymakers and interested parties believe that a more comprehensive reform of the money funds industry is still necessary. In November 2012, the Financial Stability Oversight Council (FSOC) made public a set of proposed recommendations to the SEC for further reform (Financial Stability Oversight Council 2012). The Council proposed three different avenues for reform. The first alternative is to remove the valuation and pricing provisions in rule 2a-7 and to require money market funds to have a floating NAV that reflects the market value of their assets.

The second alternative is to require funds to maintain a buffer of assets in excess of the value implied by a fixed (and stable) NAV on outstanding shares. This buffer would be combined with a minimum balance at risk—in certain circumstances a small percentage of each investor's shares would be made available for redemption only on a delayed basis (see McCabe et al. [2012] for a detailed analysis of the minimum balance at risk idea). Finally, the third proposal is to require funds to hold a risk-based buffer and combine it with requirements on portfolio diversification, liquidity, and disclosure.[2]

To assess the Council's proposals, or any other reform proposal, it seems crucial first to be able to discern what is the ultimate function that money funds perform in the economy and how appropriate regulations depend on that. There are (at least) two possible ways to think about this issue. On one hand, some observers have argued that money funds provide useful maturity transformation by issuing claims (shares) that can be redeemed on demand while, at the same time, investing in longer-term financial instruments. Even though the funds' portfolios are concentrated in relatively short-term instruments, funds stand ready to redeem shares on demand and, hence, are exposed to a maturity mismatch and the threat of illiquidity.

On the other hand, it may be that the main role of money funds is to manage the portion of investors' portfolios intended to be allocated to relatively short-term money market instruments. In other words, according to this view, money funds are expert "cash" managers and, for this reason, it is efficient for investors to delegate to them the administration of part of their short-term and liquid investment strategy.

[2] See the FSOC document for a thorough description and evaluation of the reform proposals (Financial Stability Oversight Council 2012). The document also provides a good summary of the institutional details of the U.S. money market funds industry.

Assessing which of the two alternative views best describes the economic value associated with money funds is important for choosing the appropriate design of a regulatory framework. In particular, how redemption values should be computed often depends on this assessment. The aim of this article is to illustrate this point by presenting and comparing the implications of using different methods for computing NAVs in two very simple models that capture, in a stark way, the two aforementioned views about the function of money funds.

The first model is a version of the canonical maturity transformation framework introduced by Diamond and Dybvig in 1983. We find that, to the extent that NAV regulations are designed in a way that still allow funds to fulfill their basic function, then illiquidity and potential instability are likely to remain an integral feature of the money fund business. Furthermore, from this standpoint, computing appropriate market-sensitive NAVs requires an estimation of the amount of withdrawals that the fund can be expected to face. This process of anticipation is especially difficult because it involves predicting economic behavior that depends on agents' expectations about the decisions of others.[3]

The second model maintains many of the structural features of the first model, but is modified so that the motives investors have to deposit money with the fund are different. In particular, investors no longer derive value from maturity transformation but, instead, they rely on the funds exclusively to manage their investments.[4] In this case, we find different implications relative to the first model. Computing NAVs that accurately reflect market valuations is perfectly compatible with the role played by the funds and can actually make the funds more stable. The model also illustrates how a wave of withdrawals from a poorly performing fund may just be the way that the system has to implement the best possible allocation of resources. Trying to stop that process would, in fact, be detrimental to economic efficiency.

Obviously, it is hard to determine which is the main function that money funds are performing in the economy, or even if they are essential organizations to pursue the highest attainable welfare of society. This article considers two candidate functions, one at a time. However, it is certainly possible that money funds perform, at least to a certain extent, these and potentially other functions simultaneously. Sorting these issues out is essentially an empirical undertaking,

[3] Chen, Goldstein, and Jiang (2010, Appendix A) study a different, yet related model of a mutual fund where the redemption strategies of agents are also interdependent in equilibrium and can generate the conditions for fund instability.

[4] The recent article by Parlatore Siritto (2013) also studies a model where the main function of money funds is to manage the assets of investors.

beyond the scope of our study. The objective in this article is, instead, rather theoretical. The point we want to illustrate is that once one has taken a stand on the answer to the empirical question, some theoretical implications follow that can help guide the design of an appropriate regulatory policy for money funds.

In principle, the models we present could be extended and modified to evaluate the other reform proposals currently being considered. For example, to understand the implications of requiring a buffer of assets one would need to take a stand on the way the buffer is being funded and model the objectives of the agents providing such funding. While this is potentially a productive activity, it would complicate the models in a way that would reduce the clarity of the results related to NAV policies. For this reason, we choose to limit our discussions to the NAV proposals.[5]

Before turning to the models, we should mention here that there is, in fact, a third commonly held perspective on the role of money funds in the economy, which we will not discuss in this article. The money funds industry developed and grew briskly in the 1970s, a period when banks were subject to strict interest rate ceilings imposed by regulation. These restrictions on the ability of banks to pay competitive rates did not apply to money funds and allowed money funds to become a natural alternative to banks (see Rosen and Katz [1983] for example). Even though the restrictions have been mostly removed now, funds may still be a vehicle for regulatory arbitrage to the extent that they are not subject to strict capital requirements and other regulations faced by banks.

The rest of the article is organized as follows. In the next two sections, we study two alternative frameworks that can be used to think about the problem of setting the appropriate redemption value of shares in a mutual fund. The first model, presented in Section 1, considers the case in which the role of the fund is to perform a maturity transformation function. The second model, in which the fund is just an investment vehicle that performs no essential maturity transformation, is the subject of Section 2. We close the article in Section 3 with a brief conclusion.

[5] Another aspect left unexplored in this article is the possibility of contingent support from an institutional sponsor when the fund experiences financial distress. Sponsor support has played a significant role in the recent history of U.S. money market funds (Rosengren 2012). For a theoretical analysis of the issue, see Parlatore Siritto (2013).

1. MATURITY TRANSFORMATION

The canonical framework for studying maturity transformation in financial economics is the Diamond and Dybvig (1983) model of banking. A way to obtain desirable allocations in such an environment is to allow for an institutional arrangement that resembles a mutual fund. In this section, we analyze the implications of this theory for the determination of the fund's net asset value.[6]

The Model

There is a continuum of agents of mass 1. Agents are risk averse and each owns one unit of resources at the beginning of time. Time is denoted by $t = 0, 1$. Agents are homogeneous ex ante, but in period 0 a proportion q of the agents gets a preference shock and needs to consume at that time to be able to get any utility. We call these agents impatient and the $1 - q$ remaining agents, patient. Patient agents are indifferent about consuming at time 0 or 1. There is a productive technology that returns $R > 1$ units of resources in period 1 per unit of resources (not consumed and) invested in period 0. Resources can be taken out of the production technology during period 0 at a one-for-one basis (one unit per unit invested); in other words, there are no liquidation "costs" from interrupting the production process at an early stage.

A Benchmark Optimal Allocation

Since $R > 1$, there is a clear benefit from delaying consumption in this economy. For this reason, it is generally optimal to have patient agents consume only in period 1. Impatient agents, however, must consume in period 0.

Consider the solution (c_0^*, c_1^*) to the following planning problem:

$$\max_{c_0, c_1} qu\left(c_0\right) + \left(1 - q\right) u\left(c_1\right) \tag{PP1}$$

subject to

$$\left(1 - q\right) c_1 = R\left(1 - qc_0\right).$$

We take such a solution as a benchmark optimal allocation in this environment. It is the allocation that maximizes the sum of the total

[6] There is an extensive literature dedicated to the study of possible extensions of the Diamond-Dybvig model (see, for example, Freixas and Rochet [2008]). We use the simplest version of the model that allows us to illustrate the general points we are trying to make. Studying the implications for money funds of extensions of the model in various directions is a potentially fruitful activity. We consider this section a first step in that direction.

utility of both groups of agents, patient and impatient, subject to the resource constraint. In this constraint, $1-qc_0$ is the amount of resources left after making a payment of value c_0 to each of the q impatient agents. This amount remains invested in the productive technology and is multiplied by the return R after waiting until period 1. In period 1, the resulting resources are divided between the remaining $1-q$ patient agents and each of them gets an amount equal to c_1.

When investors' coefficient of relative risk aversion is greater than one it can be shown that

$$1 < c_0^* < c_1^* < R.$$

The thing to notice here is that patient and impatient agents share the return from the productive investment in the optimal allocation. This is a form of insurance. Impatient agents get more than their initial resources even though the productive investment has not yielded any returns at the time that these agents wish to consume. This insurance is possible because only a proportion of the agents is expected to be impatient.

Institutions: An Open-End Mutual Fund

There are two main categories of mutual funds: those that are open-end and those that are closed-end. Open-end mutual funds stand ready to redeem shares held by investors at an announced net asset value. Closed-end mutual funds, instead, issue a fixed number of shares that in principle trade in a securities market but do not redeem shares on demand. Money market funds in the United States are predominantly open-end funds. Given the focus of our study, we restrict attention to this arrangement in the main body of the article. The reasons for the prevalence of open-end funds is the subject of active academic research (see, for example, Stein [2005]). We do not address the issue here but we present a brief analysis in the Appendix of how a closed-end fund would work in this environment.[7]

Suppose that at the beginning of time agents form an open-end mutual fund and deposit their endowment with the fund. The fund then invests the resources and sets dividend payments and a NAV that determines how much an agent is entitled to withdraw from the fund at each time.

[7] There are many complex issues associated with the economics of closed-end mutual funds. For a survey of the subject see Lee, Shleifer, and Thaler (1990). Cherkes, Sagi, and Stanton (2008) is an interesting recent contribution.

One way for the fund to implement the optimal allocation (c_0^*, c_1^*) is to set a NAV equal to 1 and assign $c_0^* - 1$ new shares to each investors in period 0 in the form of a dividend payment. At that point, then, each agent has in their account c_0^* shares of the fund. If only the proportion q of agents that need to consume early decide to withdraw from the fund, then total withdrawals from the fund equal qc_0^* and there will be enough resources to pay the rest (a proportion $1 - q$) of the agents an amount equal to c_1^* in period 1. Since $c_1^* > c_0^*$, an investor that expects these payments and does not need to consume early will be willing to wait to withdraw. For this reason, the optimal allocation is a possible outcome associated with this mutual fund scheme.

As is well-known from the bank-run literature, when withdrawals from the fund happen sequentially, there is another possible outcome associated with this scheme (see Diamond [2007] for a simple exposition). Given that c_0^* is greater than unity, if all agents attempt to withdraw at time 0 then the fund would not have enough resources to cover all the required payments. As a result, if agents expect that all other agents will attempt to withdraw from the fund, then they also have incentives to try to withdraw, creating a situation that would resemble a run on the fund.[8]

It is also well-known from the bank-run literature that a scheme that allows the suspension of redemptions after q withdrawals will be able to costlessly rule out the run equilibrium. In reality, money funds can *and have* asked the SEC to authorize them to suspend redemptions after experiencing a wave of withdrawals. However, the authorization is usually granted under the assumption that the fund will fully liquidate and terminate operations after that. To the extent that the requirement of full liquidation still imposes costs on the fund, the suspension becomes less effective in limiting the incidence of runs.

In the model, the possibility of runs arises because, after the fund has distributed the new shares as dividends, if all agents are expected to want to withdraw from the fund at time 0, then the current value of fund assets is not sufficient to justify a NAV equal to 1. In particular, at time 0 total assets in the fund have a current (liquidation) value of 1. Agents, however, own $c_0^* > 1$ shares which, with a NAV of 1, entitle them to total time-0 payments that are greater than the current (liquidation) value of assets (one unit). An obvious solution to this problem

[8] The fact that withdrawals take place sequentially during time 0 implies that the fund initially makes payments without knowing the total number of time-0 withdrawals that will ultimately happen. If the fund would be able to observe the total number of withdrawal requests before making any actual payments, then it is easy to show that the fund would adjust the value of those payments in such a way that runs could not happen in equilibrium.

is not allowing the fund to allocate new shares in the form of dividends before the actual returns are realized. However, the "early" dividends are essential for implementing the benchmark optimal allocation when the NAV is set to equal 1.[9]

In general, however, the fund may not want to value assets at their liquidation value (i.e., using a NAV equal to 1). Suppose, instead, that the fund sets a NAV equal to the *future discounted value* of the cash flow from the assets (FDV for short). If the manager of the fund (or some regulator) looks at the assets currently in the fund and disregards the withdrawal issue, following FDV would require setting a NAV equal to $\frac{R}{1+r}$, where r is an appropriate discount rate.

Since we are considering a situation without discounting, one possibility would be to take $r = 0$. In this case, the fund's NAV will be set to equal R. We know, however, that if agents withdrawing at time 0 get a payment equal to R, then the optimal allocation will not be implemented (since $c_0^* < R$). Furthermore, if q agents get R in period 0, then there will not be enough resources to pay R or c_1^* to those agents withdrawing (and consuming) at time 1. If withdrawals from the fund happen sequentially, the only optimal withdrawal strategy for *all* investors under these payments is to try to withdraw early in a situation resembling a run.

Given that the rate of return on investment between $t = 0$ and $t = 1$ is equal to R, another possibility would be to use $1 + r = R$ as the appropriate discounting to compute the FDV. In this case, then, the fund's NAV will be set to equal unity and again, without an early distribution of shares in the form of dividends, the optimal allocation would not be obtained. An attractive aspect of setting this value for the NAV is that the unique equilibrium in this case is for only impatient agents to withdraw at $t = 0$. While this conveys a sense of stability to the fund, it is also the case that impatient agents consume only one unit (not c_0^*) in this situation and, hence, the fund no longer performs the maturity transformation function that was the purpose of its creation.

It is unclear the extent to which money funds in reality are able to make higher payments to investors in anticipation of future expected returns. In the model, implementing a value of c_0 greater than 1 requires such anticipation. Money funds may not be performing the type of maturity transformation suggested by this model. We will consider an alternative model in the next section.

[9] Initially each agent owns one share with a NAV equal to 1. As impatient agents need to consume $c_0^* > 1$ to conform with the benchmark optimal allocation, an entitlement of extra shares needs to be assigned to agents in period 0 so that impatient agents can actually consume an amount greater than 1 (c_0^*) at the appropriate time.

Even if the model in this section is the relevant one, it could be that due to legal (or "best practice") restrictions, money funds do not perform the function described here. For example, suppose that the law requires that the fund pays dividends only after returns have been realized and always sets the NAV at the current liquidation value of the assets. In that case, the fund would set a NAV equal to 1 in period 0 and the payments would be given by $c_0 = 1$ and $c_1 = R$. This payment scheme, again, makes the fund immune to runs even when withdrawals are restricted to happen in a sequential manner.

The main insight thus far is that the maturity transformation function may involve a tradeoff between efficiency and stability. Some schemes result in a system that is immune to runs but does not provide beneficial insurance to impatient agents. Other schemes transfer resources appropriately among agents but make funds open to instability. The setting of the NAV plays a crucial role in the design of these schemes.

Variable Liquidation Terms

Suppose the fund is not able to liquidate and recover the invested resources one for one at time 0. Instead, the fund can only get ξ per each unit initially invested and later liquidated during period 0. In principle, the value of ξ may depend on the amount x being liquidated early. That is, ξ is a function of x.

An optimal arrangement is one that delivers the consumption allocation (c_0^*, c_1^*) obtained by solving the following problem:

$$\max_{c_0,c_1,x} qu\left(c_0\right) + \left(1-q\right)u\left(c_1\right) \tag{PP2}$$

subject to

$$\begin{aligned} qc_0 &= x\xi\left(x\right), \\ \left(1-q\right)c_1 &= R\left(1-x\right). \end{aligned}$$

Here, the first constraint indicates that to make a payment of c_0 to each of the q impatient agents, the fund needs to liquidate x units of investment, which allows it to obtain $x\xi\left(x\right)$ units of resources at time 0 when the payments to impatient agents need to occur. After liquidating x units of resources, $1-x$ units are left in the productive technology and, hence, result in $R\left(1-x\right)$ available resources at time 1. The second constraint, then, says that these resources will be used to pay an amount c_1 to each of the $1-q$ patient agents.

It is easy to see that if $\xi(x) = R$ for all x, then $c_0^* = c_1^* = R$. In this case, a NAV equal to R per share implements the optimal allocation.[10] However, if $\xi(x) < R$ for some x then it becomes less obvious how to compute an appropriate NAV. For example, if $\xi(x) = \widetilde{\xi} < R$ for all x then $c_0^* < c_1^* < R$ and a fund trying to implement the best arrangement for its investors could need to set a NAV that would expose it to instability. The benchmark situation we studied before is the particular case when $\widetilde{\xi} = 1$.

When funds liquidate, they usually sell assets in the market. It is often argued that the price of the assets may depend on how much is being liquidated. In our simple framework, liquidation at time 0 does not involve market prices but rather the direct technological costs of liquidating productive investment. Still, using the flexibility of the function ξ we can consider some cases that produce valuable insights about the more complex situation in which market prices play a role during liquidation. In particular, consider the case in which $\xi(x) = R$ as long as $x \leq q$ and $\xi(x) = \widetilde{\xi} < R$ if x is greater than q. Here, again, the appropriate NAV would depend on the expected number of withdrawals. Suppose that the fund expects to have q withdrawals. Then, using a NAV equal to R allows the fund to implement the allocation $c_0^* = c_1^* = R$ with only impatient agents withdrawing from the fund at time 0.

However, if unexpected extra withdrawals were to happen (that is, if more than q agents decide to withdraw at time 0), the NAV would have to be drastically adjusted. Evidently, a crucial issue is how soon in the withdrawal process would the fund realize that withdrawals will be higher than q. If this realization comes after the first q withdrawals have already happened, then the fund will have to adjust the NAV at that point. The appropriate value of the NAV would depend on how many more withdrawals are expected after the first q. Suppose that after seeing that withdrawals continue beyond the first q the fund expects $q' > q$ withdrawals. Then, setting a NAV equal to $\widetilde{\xi}$ would make the fund solvent but would destroy any insurance possibilities that the fund could still try to exploit given that q' is expected to be lower than 1.

This extension of the model captures in a stylized manner the technological (or market-based) costs that are often associated with the

[10] Notice here that when $\xi(x) = R$ for all x, the fund has the ability to come up with resources immediately at no cost. For each unit of resources that the fund invests in the productive technology, it can get R units immediately, without waiting or bearing any risk. For this reason, the case of $\xi(x) = R$ seems of limited applicability for understanding actual real life investment situations.

early liquidation of an investment position. The analysis clearly illustrates that liquidation costs, in interaction with expectations about the number of early withdrawals, significantly complicate the setting of an appropriate NAV.

Portfolio Choice: Adding a Liquid Asset

Suppose now that in the setup just studied the liquidation value of the productive technology is $\xi(x) = \widetilde{\xi} < 1$ for all x. This situation may seem peculiar since some *costly* liquidation is taking place even though it is completely predictable. In other words, given that the fund is expecting at least q redemptions, it would be better to invest some resources in an asset that, while less productive, avoids any significant liquidation costs (i.e., a more liquid asset).

To address this issue, we extend the previous setup to include an alternative technology that returns, per unit invested at the beginning of time 0, one unit of resources at any time. Then, an optimal arrangement would produce the allocation that solves the following problem:

$$\max_{c_0, c_1, \gamma, x} qu(c_0) + (1-q)\, u(c_1) \tag{PP3}$$

subject to

$$\begin{aligned} qc_0 &= \gamma + x\widetilde{\xi}, \\ (1-q)\, c_1 &= R(1-\gamma-x), \end{aligned}$$

where γ is the portion invested in the liquid asset and x, again, is the amount liquidated at time 0 of the fund's investment in the productive technology, $1-\gamma$. As before, the two constraints are resource constraints on payments at time 0 and 1, respectively. The first constraint shows that the investment γ in the liquid asset is fully used to make payments to impatient agents. In the second constraint, total unliquidated productive investment is now equal to $1-\gamma-x$. Multiplying this amount by $R > 1$, we obtain the total available resources at time 1 that can be used to make payments of value c_1 to each of the $1-q$ patient agents.

When $\widetilde{\xi} < 1$ and the fund expects that exactly q agents will withdraw at $t = 0$, it is optimal to choose $x^* = 0$ and $\gamma^* = qc_0^*$. Furthermore, the optimal values of c_0 and c_1 are given by the same c_0^* and c_1^* obtained in the benchmark optimal allocation (problem PP1). The perfect predictability of the number of withdrawals, combined with the fund's access to a liquid asset, implies that costly liquidation never happens.

How should the fund compute its NAV at time 0? Here, again, combining the payment of early dividends with a NAV equal to 1 would be consistent with obtaining the optimal allocation as an equilibrium outcome. The alternative approach based on calculating a FDV with a discount rate $r = 0$ would result in a value of the NAV equal to $\gamma^* 1 + (1 - \gamma^*) R$. While the FDV method is often considered natural, it is easy to show that in this case the implied NAV is greater than c_0^* and, hence, it would provide too much consumption to those agents withdrawing in period 0 (relative to the optimal allocation).[11]

The fact that the fund can perfectly predict the amount of withdrawals is important and may be considered unrealistic. Uncertainty over q significantly complicates the calculations. To gain some perspective on this issue, consider a situation where the fund was expecting q withdrawals but instead $\widetilde{q} > q$ withdrawals happen. After making the first q payments the fund would have to reassess the rest of its planned payments. Suppose that after making the first q payments the fund immediately discovers that the number of withdrawals will be $\widetilde{q} > q$. Then, the optimal continuation payments would solve the following problem:

$$\max \left(\widetilde{q} - q\right) u\left(c_0'\right) + \left(1 - \widetilde{q}\right) u\left(c_1'\right) \tag{PP4}$$

subject to

$$\left(\widetilde{q} - q\right) c_0' = x\widetilde{\xi},$$

$$\left(1 - \widetilde{q}\right) c_1' = R\left(1 - \gamma^* - x\right).$$

The first constraint indicates that for the fund to be able to make a payment of value c_0' to $\widetilde{q} - q$ agents in period 0 it will have to liquidate an amount x of productive investment that, given liquidation costs, results in $x\widetilde{\xi}$ available resources. It is important to realize here that the fund has already made q payments of size c_0^*, and since $\gamma^* = qc_0^*$, there are no more liquid assets available to make extra payments in period 0. The second constraint (over payments in period 1) is similar to that in the previous problem. Let us denote by $c_0'^*$ and $c_1'^*$ the solution to problem PP4.[12]

Setting the appropriate continuation NAV in this case is again a difficult issue. Note that there are only $(1 - \gamma^*)$ units of the asset left at the fund after the initial q withdrawals. These assets can be liquidated at a rate of $\widetilde{\xi} < 1$ and the fund has to still make $1 - q$ payments. In

[11] We know that $c_0^* < c_1^*$, $c_0^* = \gamma^*/q$, and $c_1^* = R(1 - \gamma^*) / (1 - q)$. Then, we have that $\gamma^*/q < R(1 - \gamma^*) / (1 - q)$, which can be rearranged to $\gamma^* + (1 - \gamma^*) R > \gamma^*/q = c_0^*$.

[12] We do not discuss here whether the fund managers would have the incentives at this point to redesign payments so as to maximize the remaining investors' utility. Perhaps reputational issues could be brought to bear in explaining a behavior of the fund in line with that suggested by the optimal continuation payments studied here.

principle, using current values of the assets, the fund would set a NAV equal to $(1-\gamma^*)\,\widetilde{\xi}/(1-q)$ and it can be shown that $c_0'^*$ is actually greater than this number. The reason for the discrepancy between the optimal continuation payment $c_0'^*$ and the NAV computed using current valuations is essentially the same as we discussed before: The fund does not expect to have to liquidate all assets (as long as $\widetilde{q} < 1$) and, as a consequence, it can still provide some insurance (maturity transformation) to the agents requesting early redemptions. In the optimal continuation, the fund's payments to these agents are such that they receive a portion of the returns coming from the productive investment that will be held to maturity.

This last extension of the model shows that when the fund holds a portfolio of investments, some more liquid than others (as it would want to do, given that it expects some withdrawals to happen early and some to happen late), the standard methods for computing NAVs again may fail to deliver the most desirable allocations. In summary, then, setting appropriate values for NAVs within the maturity transformation paradigm often involves a tradeoff between efficiency and stability. This is the case in the simplest version of the model and it remains true even when we consider liquidation costs and a non-trivial portfolio choice available to the fund.

2. INVESTMENT MANAGEMENT

In this section, we study a model in which the mutual fund performs the function of investment management. The underlying justification is an assumption that the fund can administer the allocation of funds to productive activities more efficiently than individual investors. For this reason, then, investors delegate management functions to the fund by investing directly in it. The model is again very simple. We attempt to stay as close as possible to the formal analysis of the previous section but introduce some modifications that produce a different perspective on the recent experiences with money funds.

The Model

There is a mass 1 of risk averse agents and each of them own one unit of resources at the beginning of time. Time is again given by $t = 0, 1$. Different from the model in the previous section, here all agents are patient (that is, they are indifferent between consuming at either time 0 or 1). There is a risky productive technology that returns a random amount R of resources in period 1 per unit of resources invested in period 0. The value of R gets realized after investment in this risky technology has taken place. However, resources can be removed from the risky productive technology at any time during period 0 on a one-for-one basis. Agents can also invest in an alternative riskless technology at any time during period 0 that returns a fix gross return $R_z > 1$ in period 1 per unit of resources invested in period 0. Call z the amount invested in this alternative riskless technology.

A Benchmark Optimal Allocation

Since z can be decided after observing the realization of R, it is optimal to make z a function of R. The optimal allocation of resources solves the following planning problem:

$$\max_{c(R),z(R)} E\left[u\left(c\left(R\right)\right)\right], \tag{PP5}$$

subject to

$$c\left(R\right) = R\left[1 - z\left(R\right)\right] + R_z z\left(R\right)$$

and

$$0 \leq z\left(R\right) \leq 1 \qquad \text{for all } R.$$

The expectation in the objective function is taken with respect to the random variable R. The first constraint is a resource constraint that must hold pointwise, for each possible value of R. It says that consumption is equal to the return on the portfolio of investment implied by $z\left(R\right)$. The second constraint reflects natural non-negativity requirements on the amount invested in each of the two technologies.

Let us denote by $z^*\left(R\right)$ the optimal investment strategy implied by the solution to this problem. We have that $z^*\left(R\right) = 1$ whenever $R < R_z$ and $z^*\left(R\right) = 0$ when $R > R_z$. If $R_z = R$, then the value of z^* is not pinned down by this problem and it is irrelevant for payoffs. Just for concreteness assume that $z^*\left(R_z\right) = 0$.

Institutions: An Investment Fund

Since all agents are equally exposed to the underlying uncertainty in the environment, risk-sharing is no longer a reason for them to pool resources in a fund. Assume, however, that only the fund has the necessary infrastructure (expertise) to be able to invest in the technology with random return R. Agents have to decide whether to invest in the fund before the value of R is realized. Let e be the amount of the initial resources that each agent decides to keep outside the fund. Hence, the amount $1 - e$ of resources is invested in the fund.

Once the value of R is realized and observed, agents may want to withdraw some of the resources initially invested in the fund. At that time, the fund calculates a NAV and allows withdrawals according to that value. Suppose R can take a finite number of possible values. We use the subindex $j \in J$ to indicate the different values of R, where J is a finite set. Let p_j be the probability that $R = R_j$ for each $j \in J$ and, of course, $\Sigma_{j \in J} p_j = 1$. Denote by h_j and z_j the NAV set by the fund and the amount that an agent withdraws from the fund, respectively, when $R = R_j$. Then, the optimization problem faced by an investor is the following:

$$\max_{e, \{c_j, z_j\}_{j \in J}} \sum_{j \in J} p_j u(c_j) \tag{IP}$$

subject to

$$c_j = R_j (1 - e - z_j) + R_z (h_j z_j + e)$$

and $0 \leq z_j \leq 1 - e$ for all $j \in J$, and $0 \leq e \leq 1$. Agents initially invest $1 - e$ at the fund and then withdraw z_j after they discover that returns will be equal to R_j. The shares z_j withdrawn from the fund are valued at a NAV equal to h_j and, hence, the total amount withdrawn equals $h_j z_j$. Agents re-invest this amount in the alternative riskless technology, together with the previously invested amount e. Hence, total consumption equals the sum of resources obtained from the fund, $R_j (1 - e - z_j)$, and from the riskless technology, $R_z (h_j z_j + e)$.

The Case of a Fixed NAV Equal to One

Since the fund can physically liquidate investment one for one, setting $h_j = 1$ for all j is feasible. When $R_j > R_z$ for some $j \in J$ and the fund sets $h_j = 1$ for all j, agents will be willing to invest all their endowment in the fund at the beginning of time. To see this, define $z'_j = z_j + e$ for all $j \in J$ and note that now we can write $c_j = R_j \left(1 - z'_j\right) + R_z\, z'_j$ since $h_j = 1$ for all j. Given that we still have the constraint $z_j \leq 1 - e$ as a requirement, choosing $e = 0$ relaxes the domain constraints on z_j and, consequently, can only improve the solution to the agent's problem. In particular, note that when $h_j = 1$ and $e = 0$ the problem of the agent is the same as the planning problem for the benchmark optimal allocation (PP5), but where now $z(R_j) = z_j$ stands for withdrawals from the fund in state j. Parallel to the solution of problem (PP5), then, whenever R_j is less than R_z the optimal value of z_j equals 1 and agents withdraw all their investments from the fund. Even though this event could look like a run on the fund, it is actually part of the process involved in obtaining an optimal allocation of resources.

This result provides an interesting perspective on some proposals to reform the regulatory framework for money market funds. Specifically, some reform proposals are designed to provide investors with a disincentive to withdraw from a troubled fund. The objective is to reduce the incidence of runs. However, we see here that limiting the ability of investors to reallocate resources at certain points in time could stand in the way of economic efficiency.

Note that we have considered only the case when investment in the fund actually constitutes a risky alternative for the agents. It is often the case, however, that money funds are considered a relatively safe investment alternative. It would not be hard to modify the model so that R_z is random and R is a fixed (safe) return. While the results have a similar flavor, some of the interpretations may not be as natural. For example, investors would want to withdraw from the fund at those times when R_z is relatively high. In other words, run-like episodes in relatively safe funds would tend to be associated with "good times" (high returns) for investors.

Variable Liquidation Terms

So far, we have studied a situation where the fund can liquidate investment one for one. More generally, suppose that the fund can obtain resources equal to ξ_j per unit liquidated of the risky productive technology, with $j \in J$. To simplify the calculations in what follows, assume that $J = \{L, H\}$ with $R_H > R_L$ and $p_L = p$ (so that $1 - p$ is the probability that $R = R_H$).

An optimal arrangement in this case produces an allocation that solves the following problem:

$$\max_{e,\{c_j,z_j\}_{j=L,H}} pu\left(c_L\right) + (1-p)\, u\left(c_H\right) \tag{PP6}$$

subject to

$$c_j = R_j \left(1 - e - z_j\right) + R_z \left(\xi_j z_j + e\right)$$

and $0 \leq z_j \leq 1 - e$ for $j = L, H$, and $0 \leq e \leq 1$.

In principle, the liquidation values could be independent of the observed value of R. When $\xi_L = \xi_H = 1$, problem (PP6) is equivalent to problem (IP) with $h_j = 1$ for all j. Then, when $R_H > R_z$, it is optimal to set e equal to zero (recall that e must be chosen before the realization of R can be observed). More generally, however, when $\xi_L = \xi_H = \xi$ for some value of $\xi \in (0, 1)$ and $R_L < R_z$, it is possible to have an optimal value of e that is different from zero. There are two cases to consider, depending on whether ξR_z is greater or less than R_L.

When $\xi R_z < R_L$, it is never optimal to liquidate investments in the funds, and the expressions for consumption are given by:

$$\begin{aligned} c_L &= R_L + (R_z - R_L)\, e, \\ c_H &= R_H - (R_H - R_z)\, e. \end{aligned} \tag{NL}$$

It is clear here that there is a tradeoff involved in choosing the optimal value of e. Investing more in the fund (lower e) increases consumption when returns are high (when $R = R_H$) but decreases consumption when returns are low (when $R = R_L < R_z$). For some parameter values the optimal value of e is positive.

When $\xi R_z > R_L$, it is optimal to liquidate investments when the realization of R is known to be equal to R_L. Given this, the expressions for consumption are now given by:

$$\begin{aligned} c_L &= \xi R_z + (R_z - \xi R_z)\, e, \\ c_H &= R_H - (R_H - R_z)\, e. \end{aligned} \tag{FL}$$

Notice the similarities with respect to the previous expressions, (NL). As a result, it is not hard to see that a similar logic applies and that for certain parameter values the way to balance the tradeoff of returns is to choose an interior (positive) value of e.

It is important to realize here that, given the information constraints implied by the environment, this situation reflects ex ante efficient choices. However, when $R = R_L$, costly liquidation takes place. This liquidation may be regarded as a regrettable outcome ex post but it should be understood that trying to avoid it through regulation could be detrimental to ex ante welfare.

Even though we do not model explicitly a market for assets we can use the model, as in the previous section, to help us think about a situation in which the fund is liquidating assets by selling them (potentially at a discount) in the market. To this end, let us consider the case in which ξ_j is positively correlated with R_j. One particular, simple version of this correlation is when $\xi_j = \xi R_j$ for $j = L, H$. This assumption implies that the liquidation value of assets reflects immediately the deterioration in prospective future returns, as one would expect would happen in a market. We turn to the study of this case next.

First, it is easy to see that if $\xi R_z > 1$ then it is always optimal to set $z_L = z_H = 1 - e$ and liquidate all investments from the fund immediately after making them. This seems an implausible situation, mainly due to the stark timing in the model. Hence, we will proceed here under the assumption that $\xi R_z \leq 1$.

When $\xi R_z < 1$ it is optimal to set $z_L = z_H = 0$ and the expressions for consumption are the same as those labeled (NL) above. As before, then, the choice of e reflects a tradeoff between lower returns in good times and higher returns in bad times.[13]

Comparing the problem for the optimal arrangement, (PP6), with the problem of the private investor, (IP), we can see that by setting $h_j = \xi R_j$ for $j = L, H$ the fund would be able to provide the agents with the optimal contract. Under this arrangement, agents do not liquidate any of their investments in the fund, regardless of the state of asset returns. That is, agents choose $z_L = z_H = 0$ and the fund never experiences a wave of withdrawals.

The key to understanding this result is to note that when the return R is expected to be low, the NAV set by the fund immediately adjusts to reflect the lower valuation of the fund's assets. By the time the investors get a chance to withdraw, the losses are already reflected in the withdrawal values. There is no way in which withdrawing from the fund can be used by investors as a way to "escape" the expected losses associated with the low returns from the fund's assets.

Delays in Adjusting the NAV

Suppose, as before, that $\xi_j = \xi R_j$ for $j = L, H$. Now, however, assume that the fund is not able to immediately adjust the NAV when the news about the returns of the assets are first revealed. As an example, suppose that the fund initially sets an (unconditional) redemption value of shares h equal to one (before any information about returns have been revealed) and that the fund is only able to adjust h after q investors have had an opportunity to withdraw from the fund.[14]

The payments to the first q investors are now given by:

$$c_L = R_L (1 - e - z_L) + R_z z_L + R_z e,$$

$$c_H = R_H (1 - e - z_H) + R_z z_H + R_z e,$$

and it is optimal for these investors to set $z_L = 1 - e$ and $z_H = 0$. In other words, those investors that are able to withdraw from the fund at a NAV equal to 1 will withdraw all their investments when the return on the assets is expected to be low and will leave all their investments in the fund if the return on the assets is expected to be high.

When $R = R_L$, after the first q agents have redeemed their shares, the fund will be able to reset its NAV. At that point, the fund would

[13] Under constant relative risk aversion, it is easy to show that the amount invested in the fund $1 - e$ is increasing in the average return R and decreasing on the (mean-preserving) variance of R.

[14] This timing can perhaps be motivated by thinking of a gradual process of diffusion of information, whereby only some agents find out that returns will be low before the fund is able to (or willing to) adjust redemption values.

have already liquidated $s = q(1-e)/\xi R_L$ units of the initial $(1-e)$ investments and the payoff to the remaining investors would have to be recalculated. In particular, if the fund sets a NAV equal to ξR_L, the payoff to these agents from withdrawing from the fund equals

$$\xi R_L \frac{1-e-s}{1-q} R_z.$$

The payoff from not withdrawing equals

$$R_L \frac{1-e-s}{1-q}.$$

Given that $\xi R_z < 1$, these agents will prefer not to withdraw.

This example illustrates how delays in updating the NAV of an investment fund may create the conditions for an initial rush of withdrawals resembling a run, which only stops after the NAV has been appropriately adjusted. Within the context of this interpretation about the nature of money funds, *floating* NAVs that adjust every time an investor has an opportunity to withdraw could be helpful in reducing fund instability.

At this point, it is natural to ask why delays in the adjustment of NAVs would happen. Current regulation allows money funds not to reflect in their redemption value deviations from the market value of their assets as long as they are small (fewer than 50 basis points). Furthermore, it seems possible that announcing changes in redemption values that were otherwise expected to be relatively constant would raise awareness and doubts among investors. If fund managers perceive a threshold-like effect from making these announcements they would have incentives to delay them on the hope that new information arrives and reverts the negative news previously received.

3. CONCLUSION

Money market funds experienced considerable distress in 2008 during the U.S. financial crisis. Their resiliency was questioned again in 2011 during the European sovereign crisis (see Chernenko and Sunderam [2012] and Rosengren [2012]). Currently, a generalized concern exists that the instability of money funds may have systemic consequences (Financial Stability Oversight Council 2012). For these reasons, there is a heated ongoing debate about the appropriate reform of the regulatory framework that applies to these funds.

In this article, we have presented two models that represent, in a stylized manner, two possible alternative interpretations of the economic function fulfilled by money funds. In both models, money funds may experience waves of withdrawals that resemble runs. The frameworks, however, are not flexible enough to address systemic concerns such as contagion and economy-wide disruptions triggered by the troubles in the money funds industry. Still, some important insights about fund stability and regulation arise from the analysis. One of the main lessons of the article is that the appropriate regulation of money market funds depends on the stand taken with respect to the fundamental economic function performed by the funds.

In particular, if money funds are mainly providers of maturity transformation services, then the setting of the redemption value of shares needs to take into account the optimal insurance component involved in this kind of arrangement. Extreme versions of floating net asset values may undermine this function, just as narrow banking tends to undermine the maturity transformation function of banks. Perhaps some instability is inextricably associated with maturity transformation, and trying to completely rule out instability translates into ruling out any degree of maturity transformation. Under this view, stable money funds can, in effect, be redundant institutions.

However, in the second model we presented in this article, we took on the interpretation that money funds are instead investment managers that are able to access, select, and implement beneficial asset-allocation strategies for their investors. Under this view, money funds do not perform any maturity transformation. We learned that in this case a timely adjustment of the fund's redemption value of shares (such as a floating NAV) may be conducive to stability and is compatible with the fund's intended function. To a certain extent, then, alternative reform-proposals involving NAVs indirectly reflect different perspectives about the main function that money funds perform in the economy.

APPENDIX

In this appendix we study an arrangement resembling a closed-end fund in the environment presented in Section 1. We can interpret this arrangement as a version of the financial intermediation system proposed by Jacklin (1987).

Suppose that at the beginning of time, investors form a fund that issues shares in exchange for investors' endowment. The fund, then, invests in a productive technology with return R. The value of each share is set to equal 1 and each share pays a dividend d_t at $t = 0, 1$. In other words, each share represents the right to a dividend stream.

At time $t = 0$ investors holding a share receive the dividend d_0 and a market for ex-dividend shares opens. *Redemptions of shares are not allowed* at time $t = 0$ (i.e., it is a closed-end fund).

Clearly, the q impatient agents will want to sell their shares. If the fund sets $d_0^* = qc_0^*$ and $d_1^* = R(1 - qc_0^*)$, we have that market clearing in the shares market is given by

$$(1 - q)\, d_0^* = vq,$$

where v is the price of a share and $(1 - q)\, d_0^*$ is the total amount of resources in the hands of patient agents that can be used to buy the q shares of the impatient agents. The equilibrium price is given by $v^* = (1 - q)\, c_0^*$. Note that, for each share, patient agents pay $(1 - q)\, c_0^*$ and receive in the following period $R(1 - qc_0^*)$. Since $R(1 - qc_0^*) = (1 - q)\, c_1^* > (1 - q)\, c_0^*$, patient agents want to buy the shares at the price v^*. Patient agents, as a group, then consume d_1^* since they own all the shares in period $t = 1$ and each of them consume

$$\frac{d_1^*}{1 - q} = \frac{R}{1 - q}(1 - qc_0^*) = c_1^*.$$

Impatient agents consume $d_0^* + v^*$ (the dividend plus the proceeds from selling the shares) and we have that

$$d_0^* + v^* = qc_0^* + \frac{(1 - q)\, qc_0^*}{q} = c_0^*.$$

We see here, then, that a closed-end fund could also implement the optimal allocation in this environment. In fact, this arrangement would make the fund immune to runs. The reasons for why funds choose to be open-end were left unmodeled in this article. See Stein (2005) for a general discussion of this issue and for a possible explanation.

REFERENCES

Chen, Qi, Itay Goldstein, and Wei Jiang. 2010. "Payoff Complementarities and Financial Fragility: Evidence from Mutual Fund Outflows." *Journal of Financial Economics* 97 (August): 239–62.

Cherkes, Martin, Jacob Sagi, and Richard Stanton. 2008. "A Liquidity-Based Theory of Closed-End Funds." *Review of Financial Studies* 22 (April): 257–97.

Chernenko, Sergey, and Adi Sunderam. 2012. "Frictions in Shadow Banking: Evidence from the Lending Behavior of Money Market Funds." Fisher College of Business Working Paper 2012-4 (September).

Diamond, Douglas W. 2007. "Banks and Liquidity Creation: A Simple Exposition of the Diamond-Dybvig Model." Federal Reserve Bank of Richmond *Economic Quarterly* 93 (Spring): 189–200.

Diamond, Douglas W., and Philip H. Dybvig. 1983. "Bank Runs, Deposit Insurance, and Liquidity." *Journal of Political Economy* 91 (June): 401–19.

Duygan-Bump, Burcu, Patrick M. Parkinson, Eric S. Rosengren, Gustavo A. Suarez, and Paul S. Willen. 2013. "How Effective Were the Federal Reserve Emergency Liquidity Facilities? Evidence from the Asset-Backed Commercial Paper Money Market Mutual Fund Liquidity Facility." *Journal of Finance* 68 (April): 715–37.

Financial Stability Oversight Council. 2012. "Proposed Recommendations Regarding Money Market Mutual Fund Reform." Washington, D.C.: U.S. Department of the Treasury (November).

Freixas, Xavier, and Jean-Charles Rochet. 2008. *Microeconomics of Banking*. Cambridge, Mass.: The MIT Press.

Jacklin, Charles. 1987. "Demand Deposits, Trading Restrictions, and Risk Sharing." In *Contractual Arrangements for Intertemporal Trade*, edited by E. Prescott and N. Wallace. Minneapolis: University of Minnesota Press, 26–47.

Lee, Charles M. C., Andrei Shleifer, and Richard H. Thaler. 1990. "Anomalies. Closed-End Mutual Funds." *Journal of Economic Perspectives* 4 (Fall): 153–64.

McCabe, Patrick E., Marco Cipriani, Michael Holscher, and Antoine Martin. 2012. "The Minimum Balance at Risk: A Proposal to Mitigate the Systemic Risks Posed by Money Market Funds." Federal Reserve Bank of New York Staff Report No. 564 (July).

Parlatore Siritto, Cecilia. 2013. "The Regulation of Money Market Funds: Adding Discipline to the Policy Debate." Manuscipt, New York University.

Rosen, Kenneth T., and Larry Katz. 1983. "Money Market Mutual Funds: An Experiment in Ad Hoc Deregulation: A Note." *Journal of Finance* 38 (June): 1,011–7.

Rosengren, Eric S. 2012. "Money Market Mutual Funds and Financial Stability." Speech given at Federal Reserve Bank of Atlanta 2012 Financial Markets Conference, Stone Mountain, Ga., April 11.

U.S. Securities and Exchange Commission. 2010. "Money Market Fund Reform: Final Rule." Available at www.sec.gov/rules/final/2010/ic-29132.pdf.

Stein, Jeremy C. 2005. "Why Are Most Funds Open-End? Competition and the Limits of Arbitrage." *Quarterly Journal of Economics* 120 (February): 247–72.