Runs on Money Market Mutual Funds *

By LAWRENCE SCHMIDT, ALLAN TIMMERMANN AND RUSS WERMERS†

Draft: February 22, 2016

We study daily money market mutual fund flows at the individual share class level during September 2008. This fine granularity of data facilitates new insights into investor and portfolio holding characteristics conducive to run risk in cash-like asset pools. Empirically, we find that cross-sectional flow data observed during the week of the Lehman failure are consistent with key implications of a simple model of coordination with incomplete information and strategic complementarities. Similar conclusions follow from daily models fitted to capture dynamic interactions between investors with differing levels of sophistication within the same money fund, holding constant the underlying portfolio.
JEL: G01, G21, G23
Keywords: Money market mutual funds; bank runs; strategic complementarities.

* We thank the Editor (Martin Eichenbaum) and three anonymous referees whose insightful comments led to significant improvements of the manuscript. We are also grateful to the participants and discussants at numerous conferences and seminar presentations over the last few years for many helpful comments and suggestions. In particular, we also thank Franklin Allen, Sean Collins, Peter Crane, Marco Cipriani, Doug Diamond, Engelbert Dockner, Emily Gallagher, Stefano Giglio, Philip Gharghori, Itay Goldstein, Richard Green, Zhiguo He, Woodrow Johnson, Asaf Manela, Konstantin Milbradt, Chris Parsons, Giacomo Rondina, Amir Sufi, Philip Schnabl, Chester Spatt, Laura Starks, Sascha Steffen, Luke Taylor, Paula Tkac, and Harald Uhlig for their feedback.

† Schmidt: Department of Economics, University of Chicago, Chicago, IL, ldwschmidt@uchicago.edu. Timmermann: Rady School of Management, University of California, San Diego, La Jolla, CA, atimmermann@ucsd.edu. Wermers: Smith School of Business, University of Maryland at College Park, College Park, MD, wermers@umd.edu.

Runs on financial institutions have long been a subject of academic and regulatory interest, and are widely thought to have important welfare consequences. Although a large part of the theoretical literature on runs has focused on commercial banks, recent studies such as Allen, Babus, and Carletti (2009) and Gennaioli, Shleifer, and Vishny (2013) consider the broader issue of instability in the so-called shadow banking system. During the recent financial crisis, many types of intermediated asset pools suffered run-like behavior, e.g., ETFs, asset-backed SIVs, and ultrashort-duration bond funds. Especially vulnerable were vehicles with cash-like liabilities, for which the liquidity mismatch became magnified during the crisis: creditors demanded unusually high-frequency access to their cash, while the liquidity of assets plunged. The crisis highlighted that run-like behavior can occur in a far broader set of pooled vehicles than bank deposits.

This paper brings unique evidence to the study of run-like behavior in pooled investment vehicles by studying the crisis in money market mutual funds (MMMFs, henceforth) following the bankruptcy of Lehman Brothers on September 15, 2008. During the year prior to September 2008, Kacperczyk and Schnabl (2013) find that the riskiness of asset holdings of MMMFs increased dramatically, and that the dispersion in yields across MMMFs increased from less than 30 to more than 150 bps/year. To be sure, these fundamental shifts increased the cross-sectional heterogeneity in MMMF riskiness, making the crisis that unfolded in September 2008 a unique testbed within which to study runs in pooled investments that are similar in structure, but have very different risks and yields, as well as having a different investor clientele.

On September 16, a single money market mutual fund that held just over 1% of its portfolio in Lehman commercial paper, the Reserve Primary Fund, "broke the buck," that is, marked the net asset value of the fund below the \$1 book value per share that investors normally expect as their redemption value; billions of dollars in investor redemptions occurred almost immediately.[1] The following day, run-like behavior spread to many other MMMFs that cater to institutional investors.

[1]This fund held, as of August 31, 2008, \$64.5 billion in total net assets, of which 0.82% was invested in Lehman commercial paper (all due during October 2008) and 0.39% in Lehman medium-term notes (due March 20, 2009). A press release from the Reserve states that the Primary Fund honored redemptions prior to 3 p.m. Eastern Time on September 16 at \$1 per share, but closed the day at \$0.97 per share.

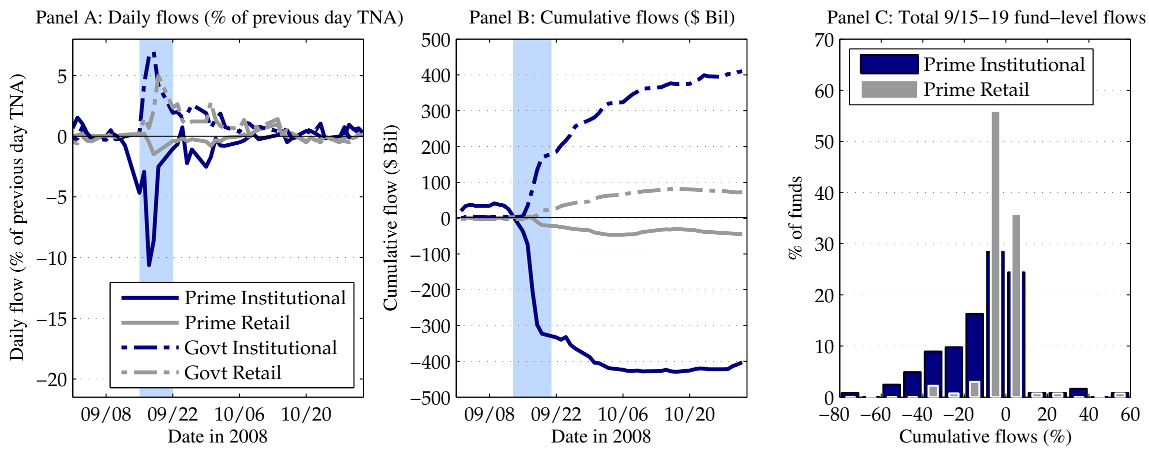

FIGURE 1. MONEY MARKET MUTUAL FUND FLOWS IN SEPTEMBER-OCTOBER 2008

Note: This figure plots summary statistics for fraction flows (one-day difference in log assets) to/from money market mutual fund share classes in different aggregate categories during September and October 2008. For each category, Panel A displays the daily percentage flow, while Panel B plots the cumulative flow (in billions of dollars) over this period. Panel C presents a histogram of the cross-sectional distribution of fund-level percentage change in total assets under management for share classes within each category during September 15-19 (i.e., flows during the MMMF crisis week following the failure of Lehman Brothers, which is indicated by blue shading).

Panel A of Figure 1 shows daily percentage flows to (from) MMMF categories during September and October 2008. Outflows from prime institutional share classes amount to more than 10% of assets on a single fateful day, September 17, 2008. Movements out of prime retail share classes were far more subdued. In contrast, MMMFs holding U.S. Government-backed securities (mainly Treasuries and agencies) experienced strong inflows, as investors sought the liquidity of the U.S. Government market as part of a "flight-to-safety". On a cumulated basis, as shown in Panel B, the flows out of prime institutional share classes amounted to $400bn during the first two weeks of the crisis. Moreover, these massive category-level outflows were far from equally distributed across funds. To demonstrate this, we calculate the Lehman week percentage change in total assets under management ("flow") for each fund separately within the prime institutional and prime retail categories. Panel C of Figure 1 shows evidence of extremely high dispersion in flows from prime institutional funds over the crisis week, with some funds experiencing a withdrawal greater than 50% of assets during the week, with others experiencing either modest withdrawals or even large inflows. In contrast, the vast majority of prime retail funds experienced only minor withdrawals or deposits.[2]

Even a cursory examination of the MMMF crisis strongly suggests that investors feared that fund managers were exposed to a common risk–a sudden and sharp reduction in the liquidity of commercial paper backed by financial institutions–even though they had largely diversified away most issuer-specific and sector-specific idiosyncratic risks. We view the MMMF crisis of September 2008 as a unique natural experiment–after a major, unexpected shock–that enables a study of the mechanism driving investor runs from a broad segment of similar intermediated portfolios.

Our daily data on flows to different institutional share classes within the same prime (non-government) MMMF portfolio allow us to study the interaction between redemptions of institutional investors in share classes with lower expense ratios (larger investment minimums) and institutional investors in share classes with higher expense ratios (smaller minimums), both residing within the same fund and holding pro-rata shares in the same portfolio of securities (thus, keeping constant the fundamental quality of these holdings as well as any subsidization, explicit or implicit, by the fund advisor). Since the largest institutional investors have more "skin in the game" and, presumably, have access to more resources (analysts and data) than smaller institutional investors, we would expect them to be both more attentive and better informed than their smaller counterparts about the quality of the portfolio that they hold together.

We use these features of our data to shed light on theoretical models of financial intermediation, a key feature of which is the interaction between fundamentals and investors' strategic behavior, and how this interaction depends on investor information. In Morris and Shin's (1998) benchmark model of regime change, agents exogenously receive private and public signals about the strength of fundamentals. The common public signal acts as an additional coordination device for agents' actions, affecting the probability of regime changes (Goldstein and Pauzner, 2005). Weak fundamentals trigger a run, but the run would not have occurred were it not for agents' self-fulfilling beliefs about the

[2]We note that a broader run might have occurred, had the crisis played out longer, or if the Treasury, the Federal Reserve, and some fund advisors had not stepped in with implicit or explicit assurances of backup support. For instance, Wachovia announced, on September 17, that it would support three of its money funds that were vulnerable to excessive outflows.

actions of other agents, which amplify the effect of a bad draw of fundamentals.[3] Such interactions between strategic behavior and fundamentals makes empirical identification of strategic complementarities very difficult (Goldstein, 2013).

Our paper develops a simple, static coordination game with strategic complementarities and information asymmetries to derive a set of predictions on the relationship between investor sophistication and fund flows. These predictions are that, first, within a particular MMMF, sophisticated investor outflows in reaction to a negative shock to portfolio fundamentals are greater than unsophisticated investor outflows. Second, holding constant the level of sophistication of each of a fund's investor types, outflows following such a shock are weakly increasing in the fraction of more sophisticated investors (within the same fund). Third, after such a shock, the *difference* between the outflows of sophisticated and unsophisticated investor types should be enhanced, the larger the fraction of sophisticated investors (and, thus, the stronger the strategic complementarities).

Using expense ratios at the share class level as a proxy for investor sophistication, since expense ratios are strongly (and negatively) correlated with the required minimum dollar investment level in MMMFs, we empirically test these theoretical predictions with a set of cross-sectional regressions. We find evidence to support all three predictions, consistent with strategic externalities playing an important role in the September 2008 run episode. Moreover, we find that (better-informed) sophisticated investors respond to the actions of their (less-informed) unsophisticated counterparts within the same MMMF. Large-scale institutional investors redeem more strongly in response to the redemptions of smaller-scale institutions (in the same MMMF) when these small investor redemptions represent a larger share of fund assets, consistent with large investors reacting to the magnitude of strategic complementarities posed by small investor actions.

Several key themes differentiate our analysis from that of prior studies. First, while previous studies document that institutional and retail investors behaved differently dur-

[3] Angeletos and Werning (2006) allow the public signal to be generated endogenously as an equilibrium outcome from a financial market. When public signals are endogenous, there can be regions with multiple equilibria. Angeletos, Hellwig, and Pavan (2007) present an extension of their baseline model, where agents receive noisy signals about the size of attacks during the previous period. This can generate "snowballing" effects similar to those in the herding literature (e.g., Chari and Kehoe, 2003; Gu, 2011), where large attacks can be immediately followed by additional attacks, as investors update their beliefs based upon the actions of other investors.

ing the financial crisis, we exploit investor heterogeneity to identify that strategic complementarities within and between these investor types were in play, using theoretically motivated tests. We find that the context of the mixture of institutional and retail investors, as well as the quality of portfolio fundamentals, drove a large dispersion in the way that institutional investors reacted. In some cases, their reaction was quite muted; in other cases, their reaction posed an existential threat to the MMMF from which they redeemed. We find that a higher concentration of less sophisticated investors–which includes both small-scale institutions and retail investors–within a fund weakens complementarities, and mitigates the strategic incentives of the most sophisticated investors to run.

Second, we study the dynamics of fund flows at the daily interval, and use the multiple share class structure of the MMMF data to study the daily interaction between institutional investors of different levels of sophistication.[4] Third, we use a "diff-in-diff" approach that compares flows in ultra-short bond funds with those in MMMFs during the week following the Lehman default. In these empirical tests, we find evidence that flows out of same-complex prime MMMFs are larger than flows out of ultra-short funds that invest in comparable securities, consistent with strategic complementarities being stronger for fixed net asset value (NAV) (share price) MMMFs than for comparable floating NAV ultra-short funds. Fourth, while prior studies estimate average effects, we use quantile panel regressions to study the distribution of flows conditional on observables, allowing us to quantify the degree of heterogeneity in investor outflows for funds with similar asset holdings and investor characteristics.

More broadly, our analysis contributes to several empirical literatures on financial crises. First, we present new evidence of strategic complementarities, consistent with results in recent work by Chen et al. (2010), Qian and Tanyeri (2013), and Hertzberg et al. (2011), albeit in a different institutional setting and identified via different methods. Second, we provide evidence on the link between risk-taking and run-like behavior, supporting work on banking panics by Gorton (1988), Schumacher (2000), Martinez-Peria

[4]Previous studies such as McCabe (2010) and Kacperczyk and Schnabl (2013) do not exploit the multiple share class structure of MMMFs, nor propose formal tests for strategic complementarities.

and Schmukler (2001), and Calomiris and Mason (2003). Third, similar to Kelly and Ó Gráda (2000), Ó Gráda and White (2003), and Iyer and Puri (2012), we identify investor characteristics which are linked to runs, and consider dynamic interactions between different types of investors in a strategic setting. Finally, we contribute to the literature on the recent financial crisis, particularly Gorton and Metrick (2012), Covitz et al. (2013), Acharya et al. (2013), and Schroth et al. (2014), who provide evidence of run-like behavior by financial intermediaries in the repo and asset-backed commercial paper markets.[5]

Our analysis proceeds as follows. Section I provides details of the institutional background to our analysis. Section II presents our simple theoretical model, while Section III introduces our data and tests the predictions from the theoretical model. Section IV compares the run on MMMFs to that experienced by ultra-short bond funds during the financial crisis. Section V analyzes the dynamic interactions between investor types through vector autoregressions fitted to daily flows, Section VI quantifies the effect of different factors on the cross-sectional distribution of flows, while Section VII concludes.

I. Institutional background

Traditional commercial bank deposit accounts and MMMFs are similar in some respects (e.g., the presumption of dollar-in-dollar-out), but quite different in others (e.g., no explicit deposit guarantees and vastly different regulatory structures, including disclosure requirements). Like other mutual funds, MMMFs are regulated under the Investment Company Act of 1940 and its various amendments (henceforth, ICA). However, they operate under a special provision of the ICA, Rule 2a-7, which allows them to value investor shares at the "amortized cost" or "book value" of assets—an accounting-based rather than a market-based principle—that is, shares are valued at the purchase price of securities minus computed premium or discount, amortized over the securities' remaining life. This provision of the ICA allows MMMFs to maintain a constant $1.00 per

[5]A key source of fragility in these markets, which differs from the money market, is rollover risk; see, e.g., He and Xiong (2012a,b). However, MMMFs are one of the primary sources of demand for these assets, and, thus, inherit rollover risk on the asset side of their balance sheets; a corporation unable to roll its paper will default on its maturing paper. Parlatore (2014) shows how strategic complementarities in MMMF sponsors' support decisions can create incentives for MMMFs within a sponsor's umbrella to run on the asset (e.g., repo and ABCP) markets.

share net asset value. For investors, this fixed-value has many advantages.[6]

Like banks, MMMFs seek to offer highly liquid liabilities, while holding less liquid assets. To be sure, this liquidity mismatch is much less extreme for MMMFs, but still raises the possibility that a MMMF might become liquidity-constrained and unable to meet redemption requests.[7] These risks have been controlled differently in banks and MMMFs. Banks are required to maintain capital, and depositors are insured up to a certain level, but banks may generally hold highly illiquid assets (e.g., 30-year mortgages), hold assets that may be lower-rated or difficult to rate or price, and employ leverage. MMMFs, in contrast, under Rule 2a-7, must hold only assets that are (normally) highly liquid with high credit quality, and generally may not use leverage.[8]

Due to the presence of significant fixed costs of operation, it is common for multiple share classes with different levels of expense ratios to coexist in a single MMMF–a feature which we exploit heavily in our empirical tests to follow. Apart from potentially different expense ratios, these shareclasses always enjoy identical (pro-rata) claims on the fund's assets. Thus, redemptions in one share class–should they result in a reduction in portfolio liquidity for a fund–negatively impact remaining investors in all share classes of that fund equally. Those share classes with low expense ratios require high investment minimums, which allow only large-scale (and, presumably, more sophisticated and attentive) investors to buy them; those with high expense ratios are populated with smaller-scale investors. Thus, funds with multiple share classes enable us to compare the behavior of players who differ in the precision of their information about fundamentals

[6]This provision allows retail investors to use their MMMFs for transactions purposes, such as paying bills and settling securities trades, without worrying about daily fluctuations in MMMF balances. A constant $1.00 NAV also allows many kinds of institutions (e.g., state and local governments) to hold their liquid balances in MMMFs, since their charters or laws generally disallow liquidity fund investments in variable NAV products; many industrial corporations have similar restrictions on investments of their excess cash balances. And, for both retail and institutional investors, a constant $1.00 NAV vastly simplifies tax accounting by eliminating the need to track the capital gains and losses that arise with a variable-NAV mutual fund.

[7]This issue can also arise with long-term mutual funds. Mutual funds are required to offer investors the ability to redeem their shares on a daily basis at the fund's end-of-day NAV. It is at least theoretically possible that requests for redemptions could outstrip a fund's ability to liquidate its underlying portfolio in order to satisfy those redemptions. This possibility is more meaningful for bond mutual funds, such as during a financial crisis if liquidity were to dry up in certain fixed income instruments (e.g., emerging market bonds). See Chen, Goldstein, and Jiang (2010) for empirical evidence of strategic behavior by investors in long-term mutual funds consistent with investors perceiving this liquidity risk and, more recently, Goldstein, Jiang and Ng (2015) for evidence of liquidation externalities generated by strategic complementarities among corporate bond funds.

[8]Over the years prior to 2008, the provisions of Rule 2a-7 have been tightened to further reduce systemic risks (see Collins and Mack, 1994).

and/or the behavior of other investors, holding constant all fund characteristics –notably portfolio quality and fund management company characteristics (e.g., the ability to subsidize fund losses to prevent a run).

II. A simple theoretical model

Next, we develop a set of robust predictions using a theoretical setting which captures several of the features of the MMMF market described above. We do so in the context of a simple coordination game with regime change, then discuss several testable implications for investor redemption behavior in funds with multiple shareclasses. Our model introduces a simple form of heterogeneity–uninformed versus informed investors-into a representative example from the class of global games models which have been used to study run-like behavior.

A. Basic structure

Our setup and exposition follows studies such as Morris and Shin (2001) and Angeletos and Werning (2006), except that we will allow for a fraction of uninformed agents that monitor the market less closely, and, as such, are unaware of a potential run as it is developing.[9] The model's status quo–maintaining a NAV of \$1 per share–will be abandoned if a sufficiently large fraction of agents withdraw from the fund.

As is standard in the global games literature, we assume that there is an exogenous fundamental θ, which nature draws from a prior distribution $\theta \sim N(\theta_0, \sigma_0)$. A continuum of ex-ante identical agents (indexed by i) decide whether to run on the fund based on noisy signals about θ. As is standard, we normalize the sign of θ so that higher values indicate stronger fundamentals; as θ increases, the fund becomes less vulnerable to runs.

After observing noisy signals about θ, each agent i chooses between two possible actions, a_i: withdraw from the fund ($a_i = 1$) or maintain the existing investment ($a_i = 0$). We normalize the payoff from not attacking to 0, and assume that the payoff from attacking is $1-c$ if the status quo is abandoned and $-c$ otherwise. Defining the aggregate

[9]Morris and Shin (2001) develop a version of the Goldstein and Pauzner (2005) model that follows a very similar structure as the model developed here. An analogous extension to the one here yields the same cross-sectional predictions.

action $A = \int_0^1 a_i di$, the status quo is abandoned when $A > \theta$. The payoff of agent i is

$$U(a_i, A, \theta) = a_i(\mathbf{1}\{A > \theta\} - c), \tag{1}$$

where $\mathbf{1}$ is an indicator which equals 1 when the status quo is abandoned. The interesting feature of the game is the presence of a coordination motive. Each agent's expected payoff from playing $a_i = 1$ increases as the mass of other attacking agents, A, increases.

In contrast to Angeletos and Werning (2006), we assume that agents do not always monitor conditions in money markets. For simplicity, we assume that measure $\mu > 0$ of the agents ("informed agents") update their beliefs about market conditions. These agents receive informative signals about the fundamental, θ, and choose actions to best respond to beliefs about the behavior of other informed agents. Each informed agent receives two signals: a private signal, $x_i = \theta + \sigma_x \xi_i$, where $\xi_i \sim N(0,1)$ and i.i.d. across agents, and an exogenous public signal, $z = \theta + \sigma_z \varepsilon$, where $\varepsilon \sim N(0,1)$ captures common noise. The remaining measure $1 - \mu$ of agents receive a completely uninformative signal about fundamentals and play the strategy $a_i = 0$.[10]

We focus on monotone equilibria, defined as perfect Bayesian equilibria in which each informed agent attacks if and only if her private signal is below a threshold $x^*(z)$, which may depend on the public signal, z. Given the assumed behavior of uninformed agents, the aggregate attack size is:

$$A(\theta, z) = P[\text{agent } i \text{ is informed}] \cdot P[x < x^*(z)|\theta] = \mu \cdot \Phi\left[\sqrt{\alpha_x}(x^*(z) - \theta)\right], \tag{2}$$

where $\Phi(\cdot)$ is the standard normal cdf, and $\alpha_x \equiv 1/\sigma_x^2$ is the precision of the private signal. Define the precisions $\alpha_0 \equiv 1/\sigma_0^2$ and $\alpha_z \equiv 1/\sigma_z^2$ analogously. The equilibrium is defined by two objects. The first, $\theta^*(z)$, implicitly defined by $A(\theta, z) = \theta$, satisfies

$$x^*(z) = \theta^*(z) + \frac{1}{\sqrt{\alpha_x}} \Phi^{-1}\left[\frac{1}{\mu}\theta^*(z)\right]. \tag{3}$$

[10]Implicitly, this amounts to an assumption that, in equilibrium, the prior distribution of θ–the unconditional distribution of fundamentals–assigns sufficiently high probability that the fund will not break the buck such that an agent's expected utility is positive when she receives an uninformative signal. We provide a sufficient condition for this behavior to be optimal in our proof of Proposition 1 in the online appendix.

$\theta^*(z)$ is the maximum value of the fundamental such that the status quo is abandoned when the realization of the public signal is z. The second object is an indifference condition. The marginal agent, who receives a signal $x^*(z)$, should be indifferent between withdrawing and remaining invested in the fund. Hence, $x^*(z)$ solves $P[\theta \leq \theta^*(z)|x,z] = c$. An informed agent's posterior belief about $\theta \sim N(\frac{1}{\alpha^*}(\alpha_x \cdot x + \alpha_z \cdot z + \alpha_0 \cdot \theta_0), \alpha^*)$, where $\alpha^* \equiv \alpha_x + \alpha_z + \alpha_0$, so the indifference condition is

$$\Phi\left[\sqrt{\alpha^*}\left(\theta^*(z) - \frac{\alpha_x}{\alpha^*}x^*(z) - \frac{\alpha_z}{\alpha^*}z - \frac{\alpha_0}{\alpha^*}\theta_0\right)\right] = c. \tag{4}$$

Combining (3-4) yields a fixed point condition in θ^* defining the threshold $\theta^*(z)$

$$\Phi^{-1}\left[\frac{1}{\mu}\theta^*\right] - \frac{\alpha_z + \alpha_0}{\sqrt{\alpha_x}}\theta^* = \sqrt{1 + \frac{\alpha_z}{\alpha_x} + \frac{\alpha_0}{\alpha_x}} \cdot \Phi^{-1}[1-c] - \frac{\alpha_z}{\sqrt{\alpha_x}}z - \frac{\alpha_0}{\sqrt{\alpha_x}}\theta_0. \tag{5}$$

The key to our identification argument is that, all else constant, increases in μ make the left hand side of (5) strictly smaller, while the expression on the right hand side stays the same. Proposition 1, which characterizes the effect of μ on $x^*(z)$, follows from (5).

PROPOSITION 1: *For all θ and z, $x^*(z)$ is bounded below and above by functions $\underline{x}^*(z,\mu)$ and $\overline{x}^*(z,\mu)$, respectively, both of which are increasing in μ. The equilibrium is unique for all z and $\underline{x}^*(z,\mu) = \overline{x}^*(z,\mu)$ iff $\sqrt{2\pi\alpha_x} > (\alpha_z + \alpha_0)\mu$. In the unique case, the probability that an informed agent attacks is weakly increasing in μ.*

Proposition 1, proved in an Online Appendix, states that, all else constant, increases in the measure of informed agents μ raise the probability that an informed agent attacks whenever the equilibrium is unique. To see this, note that when the expression on the left hand side of (5) is increasing near a fixed point $\theta^*(z)$, increases in μ will increase the threshold $\theta^*(z)$ which triggers a run. From (3), an increase in $\theta^*(z)$ increases the threshold for the private signal of the marginal agent, $x^*(z)$. Therefore, coordination on the status quo becomes more difficult to sustain, and more agents will choose to run as μ increases. When $\sqrt{2\pi\alpha_x} > (\alpha_z + \alpha_0)\mu$, the left hand side of (5) is globally increasing, the equilibrium is always unique, and our comparative static is sharp.

When $\sqrt{2\pi\alpha_x} < (\alpha_z + \alpha_0)\mu$, there exist values of z such that multiple x^* and θ^* solve

(3-5). For these signal realizations, there are three possible equilibria, and the model can only place upper and lower bounds on the fraction of agents who choose to attack. Our proposition states that both of these bounds shift upwards in response to increases in μ. In this limited sense, our comparative static is robust to potential multiplicity of equilibria.

Our setup is intentionally kept simple. In an online appendix we discuss how similar comparative statics emerge in versions of the bank run games of Goldstein and Pauzner (2005) and He and Manela (forthcoming) when they are augmented with a measure of uninformed agents. Alternative models such as Hellwig and Veldkamp (2009) endogenize the acquisition of information but involve similar trade-offs and will imply similar comparative statics results, i.e., adding a fraction of investors who maintain the status quo weakens the complementarities and thus makes it easier for the strategic investors to coordinate on the status quo.

B. Testable predictions

For purposes of solving the model, it is sufficient to know the fraction of investors who receive informative signals each period. Now, we introduce a simple way of interpreting the flows from funds with multiple shareclasses, which generates several testable predictions. As a simple illustration, assume there are two types of investors, and each fund has two shareclasses. Large, sophisticated investors invest in shareclass S, and smaller, unsophisticated investors invest in shareclass U. If monitoring money market conditions is primarily associated with fixed costs, then type S investors will face lower monitoring costs per dollar invested. As such, we would expect them to monitor conditions more frequently. Accordingly, we will assume that type S and U investors receive informative signals with probabilities p and q, respectively, where $0 \leq q < p \leq 1$.

Now, imagine that we observe data from a cross-section of funds, which are indexed by j. These funds receive i.i.d. draws of the fundamental, θ_j, and are ex-ante heterogeneous in the fraction of assets under management, ω_j, owned by shareclass S. Then, after the failure of Lehman, measure $\omega_j p + (1 - \omega_j)q$ receive informative signals and make strategic choices. Therefore, ω_j is proportional to the (asset-weighted) average level of investor sophistication for a given fund. Our tests only require that we observe cumulative actions of investors in each shareclass, as well as empirical proxies for ω_j.

Our model makes three key predictions which are testable at the shareclass level. Suppose that the conditions are satisfied such that a unique equilibrium exists. Then, we can uniquely define $A^*(\theta, z, \omega)$ as the expected action of an agent who has received an informative signal in a fund with fundamental θ, public signal z, and fraction ω of investors in shareclass S. Conditional on θ and z, the following comparative statics hold (note that "outflow," as used below, is defined as the net of investor dollar sales minus purchases, all divided by lagged total net assets of a share class):

1) Within funds, outflows from shareclass S, $p \cdot A^*(\theta, z, \omega)$, are larger than outflows from shareclass U, $q \cdot A^*(\theta, z, \omega)$.
2) Since $A^*(\theta, z, \omega_2) - A^*(\theta, z, \omega_1) \geq 0$ for any $\omega_2 \geq \omega_1$, expected outflows for each type of shareclass are weakly increasing in the fraction of sophisticated investors, ω. Moreover, the marginal effect of changing ω on expected outflows is higher for type S than for type U investors.
3) Within funds, the difference in outflows between shareclass S and shareclass U, $(p - q) \cdot A^*(\theta, z, \omega)$, is increasing in ω.

Since these comparative statics hold conditional on θ and z, they also hold unconditionally. Note that predictions 1 and 3 hold *within* funds and are, therefore, identifiable in regressions with fund fixed effects.

III. Testing the theory: Evidence of strategic behavior from investor share class flows

Compared with commercial banks, our data on money market mutual funds (MMMFs) are unique in several dimensions that allow unprecedented insights into the mechanism of runs through a high level of granularity in both the cross-sectional and time domains.[11] This section first introduces our data, before proceeding with a set of cross-sectional empirical tests of the predictions from the simple model of Section II.

[11] In addition, during the week of the crisis, MMMFs carried no explicit insurance, further making them a unique subject for the study of run-like behavior in cash-like pooled vehicles where investors derive liquidity insurance from each other.

A. Data

Our data are purchased from iMoneyNet, a company that collects daily information from over 2,000 U.S. registered MMMFs that invest primarily in U.S. short-term, dollar-denominated debt obligations, and cover the period from February to December 2008. A comparison with statistics from the Investment Company Institute indicates that iMoneyNet covers about 93.5% of the dollar value of the entire U.S. prime MMMF universe–funds that can invest in assets such as repurchase agreements and commercial paper–as of the end of 2008, which is the main focus of this study. These MMMFs are offered to retail as well as institutional investors, through different share classes, which are pro-rata claims on cashflows from the fund portfolio's security holdings. The iMoneyNet data consist of fund investment objective, fund family/adviser (i.e., "complex") identity, share class type (i.e., retail vs. institutional), daily total net assets by share class, portfolio average maturity of (fixed-income) holdings, and weekly sector breakdown (e.g., commercial paper vs. repurchase agreement) of portfolio holdings. Importantly, the data include share class-level expense ratios. These data are especially crucial for our study, as they allow us to identify investor "skin in the game," which is likely to be highly related to investor sophistication and attentiveness.

Further, note that an institution subscribing to iMoneyNet can easily determine outflows from a share class that occurred during the prior day, before that institution makes its investment decision for the current day. Therefore, each institution can condition its decision to invest/disinvest on information about the actions of all other institutions (and retail investors) during any prior days.[12]

Also, since the Reserve Primary Fund is widely viewed by market participants as initiating the crisis through its "breaking-the-buck" announcement, we consider flows to this fund as occurring exogenous to those of other funds. Two other small MMMFs within the Reserve complex, Liquid Performance and Primary II, did not hold Lehman debt (at

[12]Fund total net assets (TNA) and other variables are available to subscribers of iMoneynet, and we expect that larger scale ("skin in the game") investors subscribe to the database and actively analyze it. Closing-day TNA data is available by 8 a.m. the following morning, which allows a quick calculation of prior-day flows (thanks to Pete Crane, publisher of Money Fund Intelligence, for this information).

least as of August 31, 2008), but did experience some outflows, as well, that were likely tied to a very negative investor perception about the Reserve brand.[13] In order to exclude any potential "brand effect," we exclude observations from all funds within the Reserve complex from our analysis.

Further information on our dataset, including univariate summary statistics, are available in a supplemental online appendix.

B. Expense Ratios and Account Sizes

Our empirical analysis uses the expense ratio charged to a given share class as a proxy for the level of sophistication of investors in that share class. Table 1 empirically documents that lower expense ratio share classes are generally only available to investors with larger amounts of money to invest. Specifically, for each institutional share class, i, for which data on the minimum required investment amount is available as of September 12, 2008 (the Friday prior to the Lehman default), we regress the log of minimum dollar investment on expense ratio (ER). These coefficients are reported in Panel A of Table 1. The specifications in Panel B replace the continuous measure ER with a dummy variable which equals 1 when the expense ratio is less than 35 basis points, and zero, otherwise. We report results separately for prime shareclasses, government shareclasses, and both prime and government shareclasses pooled in the same regression.[14]

Table 1 shows that there is a strongly negative and highly significant relation between minimum investment and expense ratio across shareclasses, and that this finding is robust to whether or not a fund fixed effect is included. Moreover, this finding holds separately for prime funds and for government funds, as well as when we pool both types. Using the continuous ER measure (Panel A), a one bp/year increase in the ER is associated with a 5% decrease in the minimum investment for prime shareclasses. Panel B indicates that prime institutional share classes with an ER less than 35 bp have dollar investment

[13] Some uncertainty surrounded the holdings of all three Reserve funds, as their August 31, 2008 portfolio holdings were not made public until October 29, 2008.

[14] It is important to note that so-called "institutional" share classes are sometimes offered to retirement plans (e.g., 401(k) plans), which are ultimately controlled by individual investors. Based on industry sources, we estimate that less than 1% (by value) of prime institutional share classes charging no more than 35 bps/yr in expense ratio consist of majority retirement accounts, while they account for almost 10% of institutional share classes with an expense ratio greater than 35 bps/yr. Thus, almost all retirement accounts are included in our definition (using 35 bps/yr) of "less sophisticated institutions" throughout this paper, and they are a minority even in that category.

Panel A: Continuous expense ratio measure

	Prime only		Government only		Prime and Government	
Shareclass expense ratio (pp/yr)	-6.05***	-5.23***	-7.64***	-4.91***	-6.80***	-5.04***
	(-2.668)	(-3.052)	(-4.683)	(-3.934)	(-4.975)	(-4.977)
R^2	0.055	0.907	0.102	0.876	0.075	0.892
Fund Dummies	No	Yes	No	Yes	No	Yes
Number of Observations	226	226	265	265	491	491
Degrees of Freedom	224	132	263	175	489	308

Panel B: Discrete expense ratio measure

	Prime only		Government only		Prime and Government	
Dummy: Expense ratio $<$ 35 bp/yr	2.68***	2.06***	3.01***	1.64***	2.83***	1.81***
	(3.290)	(3.293)	(5.736)	(3.955)	(6.071)	(5.142)
R^2	0.055	0.905	0.086	0.868	0.069	0.887
Fund Dummies	No	Yes	No	Yes	No	Yes
Number of Observations	226	226	265	265	491	491
Degrees of Freedom	224	132	263	175	489	308

TABLE 1—REGRESSIONS OF MINIMUM INVESTMENT LEVELS ON SHARECLASS EXPENSE RATIOS

Note: This table presents evidence that lower expense ratios are associated with higher investment minimums in the cross-section of institutional share classes of MMMFs with available data on minimums, measured as of September 12, 2008. In Panel A, each cell in the table provides the coefficient and t-statistic from a univariate regression of the log of the minimum dollar investment level of a given shareclass on its expense ratio, measured in bp/year). Each column corresponds to different subsets of the institutional shareclass population: prime institutional shareclasses, government institutional shareclasses, and both populations pooled, with and without fund fixed effects, respectively. In Panel B, we replace the continuous expense ratio with a dummy variable which equals one when the expense ratio is less than 35 bp/yr, and zero, otherwise.

minimums which are, on average, almost 7 times ($\exp(2.06) - 1 \approx 6.8$) larger than other shareclasses, suggesting, again, that the lowest ERs are found in share classes requiring the largest (by far) minimum investment. This evidence is consistent with more sophisticated investors residing in low ER institutional shareclasses, in the sense that they hold larger accounts, have more "skin in the game," and, thus, greater incentives to continuously monitor the fundamentals of the MMMFs in which they invest.

C. Cross-Sectional Tests

Expanding our analysis to use the share class expense ratio (ER)–which, unlike investment minimums, is available for all share classes in iMoneyNet–as a (negatively signed) proxy for investor sophistication, we next conduct a set of simple tests of the theoretical predictions of our model. To this end, we compute value-weighted percentage outflows (investor redemptions minus purchases divided by total net assets) of all share classes within one of three bins (retail share classes, and institutional share classes with an ER above or less than or equal to 35 bps/year) in individual prime funds during September 12-19, 2008. Then, Figure 2 plots the average of these bin values across all prime funds,

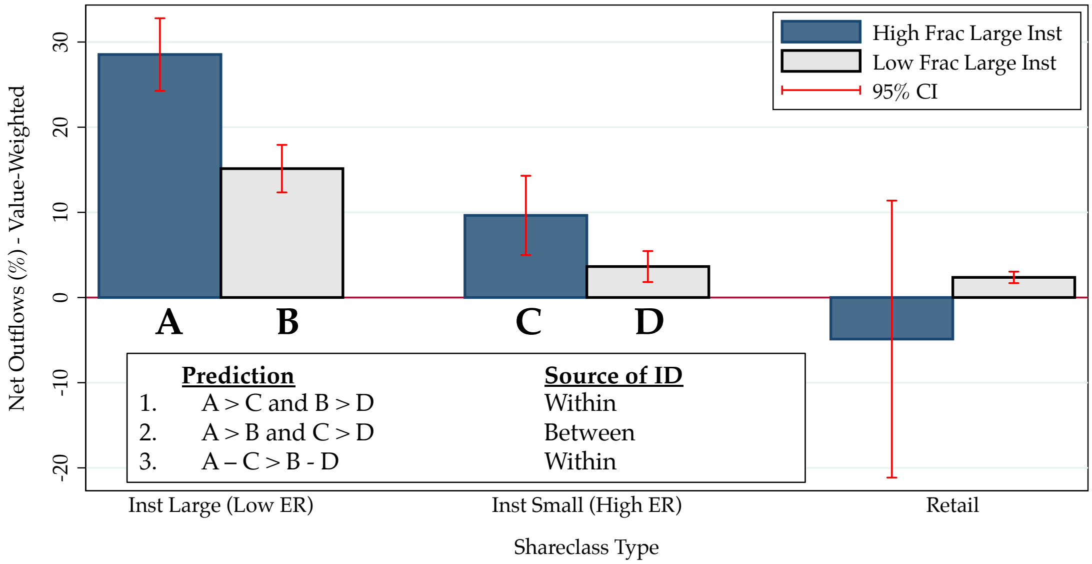

FIGURE 2. SUMMARY OF TESTS OF THEORETICAL PREDICTIONS

Note: This graph shows net outflows (the percentage change in total fund assets under management multiplied by -1) by shareclass type, across all prime MMMFs from September 15 to September 19th. First, each share class within each prime fund is placed in one of three bins: (1) retail share classes, and institutional share classes having an expense ratio (2) above or (3) no more than 35 bps/yr. Next, each of these three bins is subdivided, according to whether or not the share class is a claim on a prime fund with a high fraction–at least 75%–of its investment dollars (across all share classes, retail and institutional) represented by institutional share classes with an expense ratio no more than 35 bps/yr. Value-weighted average net outflows with 95% confidence intervals (CI), are shown for each of the six bins.

separately for funds having a high vs. low fraction of large institutional share classes relative to total prime fund assets (with the cut-off at 75%, by value, of large institutional share classes), ω.[15] For purposes of generating the figure, we classify all institutional shareclasses with expense ratios less than or equal to 35 bps/year (essentially the cross-sectional mean of 34 bps/year) as sophisticated, while the remaining ones are classified as unsophisticated.[16]

The first prediction from our model of Section II is that, within funds, outflows from sophisticated investors (S) should exceed outflows from unsophisticated investors (U). This implies that outflows from low ER share class investors should exceed those from high ER share class investors, both among funds dominated by large institutional investors (A>C in Figure 2) and within funds dominated by small institutional investors

[15]The 75% figure is chosen so as to try to balance the number of funds represented in columns A-D, but the results are very robust to using a different cutoff.
[16]The findings are robust to using other values such as 25 or 45 bp/yr.

(B>D). This prediction is strongly supported by the data.

The second prediction is that, among share classes with similar levels of investor sophistication, outflows following a negative shock to fundamentals should be larger when the share class is a claim on a fund with a higher fraction of sophisticated investors (ω). Our empirical test for low ER share classes (A>B) and for high ER share classes (C>D) also strongly supports this prediction.

The third prediction is that the difference in outflows between within-fund high and low sophistication investors (low ER and high ER share classes, respectively) is increasing in the fraction of sophisticated investors, ω, that reside within the fund. In our graph, this amounts to A-C>B-D. Again this prediction is supported by the data.

Tables 2 and 3 provide more formal tests of these predictions. In Table 2, we consider Prediction 1 by regressing share class-level percentage flows (here, an outflow is negatively signed) over September 15-19 on share class expense ratio.[17] We employ, in the various specifications, either no fixed effect, a fund fixed effect, or a fund complex fixed effect.[18] We find (models 1 to 3) that the share class expense ratio is positively and significantly associated with one-week investor flow, regardless of whether fund or complex fixed effects are included. To illustrate, using regression 3, a prime institutional share class with an ER that is one standard deviation (20 bps/year) above that of another share class (in the same fund) carrying an average level (34 bps/year) of ER is predicted to experience an outflow during the crisis week that is 13.2% of assets lower.

This finding is further confirmed in models 4 to 6, where we adopt a non-parametric approach. Here, we separate prime institutional share classes into three bins based on their ER. The base category (model intercept) represents funds with an ER above 35 bps/year. These specifications show that the lowest ER segment (charging no more than 15 bps/year) exhibits a much higher level of outflows than the next segment (between

[17]Throughout the paper, with the exception of Table 1, share class-level and fund-level regressions only include observations with at least \$100 million in assets, respectively. This approach generally excludes small-scale funds that have a limited number of investors. Results are unchanged, albeit a bit noisier, if we lower this threshold to \$10 million.

[18]A fund fixed effect allows us to hold constant the quality of the portfolio, as well as the implied insurance provided by any potential subsidization by the management company as we examine the reaction of different share classes; alternatively, the coarser complex fixed effect provides a (less precise) control for fundamentals and implied insurance, while retaining the effect of differences in clientele across different funds (which increases the statistical power to detect the effect of clientele differences on flows).

	(1)	(2)	(3)	(4)	(5)	(6)
Share class expense ratio	10.89***	11.14***	12.42***			
	(5.181)	(4.153)	(3.153)			
Indicators:						
$15 < ER \leq 35$				-15.59***	-21.29***	-20.85***
				(-4.865)	(-5.122)	(-2.682)
$ER \leq 15$				-34.88***	-36.09***	-38.97**
				(-3.913)	(-2.894)	(-2.395)
Dummies	None	Complex	Fund	None	Complex	Fund
Clustering	Fund	Complex	Fund	Fund	Complex	Fund
N	258	258	258	258	258	258
Degrees of Freedom	256	191	135	255	190	134
R2	0.057	0.369	0.565	0.065	0.381	0.575

TABLE 2—REGRESSIONS OF PRIME INSTITUTIONAL SHARE CLASS-LEVEL FLOWS ON EXPENSE RATIO

Note: To test model Prediction 1, this table presents estimated coefficients from OLS regressions of the change in logged prime institutional share class assets under management (i.e., flows as a fraction of lagged assets under management, $\times 100$) from September 15 to September 19th on that share class's expense ratio (ER), which has been normalized by its cross-sectional standard deviation. Depending on the specification, we include no fixed effects, complex fixed effects, or fund fixed effects. The final three columns replace the continuous share class ER with two dummy variables for different ranges of ER; the omitted category includes all prime institutional shareclasses with an ER >35 bp/yr (t-statistics in parentheses).

15 and 35 bps/year). Specifically, model 4 indicates that share classes having an ER between 15 and 35 bps/year experience an outflow of 14% of assets, while those with an ER below 15 bps/year experience an outflow of 29.5% of assets**,** both relative to share classes with an ER above 35 bps/year (which experience an outflow of 8% during that week). Relative outflows when ER<15 bps/year only mildly increase to 30% and 32% of assets when complex and fund fixed-effects are included, respectively (models 5 and 6). These findings, again, confirm Prediction 1 from Section II: the most sophisticated investors redeem much more strongly during the crisis week.

Table 3 presents cross-sectional regressions of the cumulative "Lehman-week" change in logged assets (i.e., fraction flows, multiplied by 100), at the individual share class level (for "sophisticated" share classes only), on the dollar percentage of sophisticated investors (ω) at the fund-level and five control variables: shareclass expense ratio; fund size (log assets); fund risk, proxied by gross 7-day SEC yield; fund liquid asset share;[19]

[19]Our definition of liquid assets starts with that of Duygan-Bump, et al. (2013), which equals the sum of the portfolio weights of investments in repo, Treasury, and U.S. Agency securities. Asset holding data is available weekly from iMoneyNet, so we interpolate between weekly values using a method which is similar to the one in Strahan and Tanyeri (forthcoming). We estimate the fraction of portfolio holdings that mature each day by dividing the percentage of the portfolio that matures within the next week by five, the number of trading days in a week. We then compare the cumulative change in assets under management since the value was last updated with that day's investor flow (change in TNA). If the share of maturing assets exceeds net redemptions, we assume that the liquid asset share remains unchanged. If redemptions exceed maturing assets–which is extremely rare prior to the start of the crisis period–we assume that the

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
%Soph (25 bp/yr)	-42.5**							
	(-2.49)							
%Soph (30 bp/yr)		-36.6***						
		(-2.81)						
%Soph (35 bp/yr)			-34.0***			-12.9	-50.6***	
			(-3.04)			(-1.04)	(-3.79)	
%Soph (50 bp/yr)				-31.4***				
				(-3.13)				
%Soph (150 bp/yr)					-28.9***			
					(-2.94)			
%Soph (35 bp/yr)$\times$ ER							17.9***	31.5***
							(4.00)	(3.34)
Shareclass expense ratio	8.0	16.6	15.9**	14.9***	10.4***	5.7	-2.8	-10.0*
	(0.51)	(1.56)	(1.99)	(4.02)	(4.92)	(1.28)	(-1.24)	(-1.84)
Log total fund assets	-12.8***	-10.9***	-10.4***	-8.3***	-7.5***	-1.5	-7.5***	
	(-2.96)	(-2.83)	(-3.29)	(-3.26)	(-3.39)	(-0.66)	(-3.52)	
Average gross yield	-0.5	-1.6	-2.0	-4.6	-3.4	-6.8	-3.0	
	(-0.07)	(-0.25)	(-0.33)	(-0.72)	(-0.65)	(-1.25)	(-0.57)	
Liquid asset share	5.2	3.1	3.1	0.3	1.6	-0.8	1.3	
	(0.80)	(0.54)	(0.58)	(0.07)	(0.36)	(-0.17)	(0.28)	
Fund business	-0.5	-1.1	-1.2	-1.6	-1.7	-2.0	-1.3	
	(-0.12)	(-0.26)	(-0.36)	(-0.57)	(-0.66)	(-0.90)	(-0.54)	
Sample selection	Sophisticated Institutional: ER $\leq$ cutoff					$ER > 35$ bp/yr	All	All
Fixed effects	None	None	None	None	None	None	None	Fund
N	118	138	161	215	258	97	258	258
Degrees of freedom	111	131	154	208	251	90	250	134
R^2	0.080	0.088	0.094	0.110	0.118	0.076	0.128	0.582

TABLE 3—REGRESSIONS OF PRIME INSTITUTIONAL SHARE CLASS-LEVEL FLOWS ON FUND AND INVESTOR CHARACTERISTICS

Note: To test model Prediction 2, this table presents estimated coefficients from OLS regressions of the change in logged share class-level assets under management (i.e., flows as a fraction of lagged assets under management ($\times 100$)) for each prime institutional money market share class from September 15 through September 19th on expense ratio (ER); % large-scale institutions (%Soph; defined as the fraction of total fund investment dollars owned by investors of institutional shareclasses with $ER \leq$ a particular cutoff–25 bp for the first column); logged fund-level size; average annualized 7-day SEC gross yield over the six month period prior to the crisis (March-August 2008); liquid asset share, estimated daily by computing dollar proportion Treasury and U.S. agency securities plus repo investments, plus (estimated) maturing securities, minus net redemptions; and fund business, defined as one minus proportion (by value) of fund complex aggregate mutual fund assets that are represented by prime institutional share classes. The bottom five variables in the table have each been divided by their most recently available cross-sectional standard deviations to normalize. Different columns correspond with different cutoffs and different sample selection criteria. To test Prediction 3 of our theoretical model, the final two columns add an interaction term between shareclass expense ratio and $\%Soph$ at the fund-level. Note that the specification in the last column includes fund fixed effects. t-statistics are indicated in parentheses.

and proportion (by assets) of the complex's mutual fund business other than its prime institutional money funds. This last variable is also used by Kacperczyk and Schnabl (2013), as a proxy for the tendency of a sponsor to avoid taking risk in its prime MMMFs, in order to protect its remaining fund business.[20]

difference between redemptions and maturing assets is met by selling liquid securities. Results are similar if we use the weekly liquid asset share from iMoneyNet in place of our daily "real-time" variable.

[20]Specifically, Kacperczyk and Schnabl (2013) find a (negative) relation between this variable and the willingness of the management company to take higher risks in their prime money funds. While Average Gross Yield should pick up risk-taking, it might do so with measurement error since it is measured over a several-month lagged period; thus, Fund Business might partially be capturing this unmeasured risk.

To measure the sensitivity of our classification scheme to the threshold used to separate sophisticated from unsophisticated institutional investors, columns 1-5 consider a range of expense ratio cutoff values (25, 30, 35, 50, and 150 bps/year) in deciding which share classes to consider sophisticated, and, thus, to include in the regression. We point out that all institutional share classes have an expense ratio below 150 bps/year, so Regression 5 identifies all institutional share classes as "sophisticated".

Recall that Prediction 1 of our model is that outflows during the crisis week, all else equal, are expected to be largest in share classes with the lowest expense ratios. This is supported by the positive coefficients on "Share Class Expense Ratio" in columns 1-5 which are statistically significant in those specifications that do not truncate share class observations (i.e., retain a larger sample size) when the ER variable exceeds low thresholds (columns 3-5).[21] The negative coefficient on (logged) fund size shows that, for large funds, the same outflow in percentage terms is associated with a larger liquidation, in dollar terms, with the potential for a larger (negative) price impact on securities sold. The coefficient on the yield variable is negative, as we would expect–funds holding riskier assets experience larger redemptions–but is imprecisely estimated. Finally, coefficients on liquidity and fund business are insignificant, indicating that, relative to substantial differences in clientele, these considerations are of secondary importance to investors, perhaps because they are viewed as being inadequate to stem a run should a large fraction of sophisticated investors decide to redeem.[22] Interestingly, in univariate regressions of fund-level flows on yield, liquidity, or fund business, each of these variables is significant at either the 5% or 10% level, and the signs are as expected. However, the explanatory power of these variables in multivariate regressions is quite low, indicating that cross-sectional heterogeneity in the strength of complementarities in funding liquidity across funds may have been somewhat larger than the heterogeneity in portfolio risk (market liquidity) across funds.

Columns 1-5 test the second prediction from our model, namely that outflows increase

[21] To verify this, we estimated a version of specification (3) in Table 3 that includes all shareclasses, and found that the coefficient on the share class expense ratio is positive and highly statistically significant.

[22] For instance, the median prime fund, on September 12, 2008, holds only 17% highly liquid assets.

in the fraction of sophisticated investors that reside in a particular money fund. For example, the first model measures the dollar percentage of the prime fund that is financed by institutional share classes carrying an ER, at most, of 25 bps/year. Across these models, consistent with theory we find (independent of the cutoff used to classify sophisticated institutional shareclasses) a strongly negative and highly significant relation between fund flows and the fraction of sophisticated investors residing in a fund. For example, model 1 predicts that a fund having 40% sophisticated investors (defined as $ER \leq 25$ bps/year) will have outflows from those sophisticated investors that are 8.1% of assets greater (all else equal) than a fund having 20% sophisticated investors. Moreover, the magnitude of this coefficient increases monotonically (in absolute value) as the threshold decreases, across models, from 150 bps/year to 25 bps/year, indicating that our proxy (ER) captures increasingly sophisticated investors as it decreases.

Our model implies that the relation between flows from less sophisticated investors and the fraction of sophisticated investors should be weaker than the relation for flows from more sophisticated investors. To test this feature, Column 6 estimates the same specification for institutional shareclasses with $ER \geq 35$ bps. Indeed, the magnitude of the coefficient on the fraction of sophisticated investors is smaller (and no longer statistically significant) than in columns 1-5, although its sign remains negative.

The third prediction of our model is that the difference in outflows between within-fund high and low sophistication investors is increasing in the fraction of sophisticated investors that reside within the fund. Columns 7-8 of Table 3 test this prediction by including an interaction term between the total fraction of a money fund's assets held by sophisticated investors and the sophistication of a given share class of that fund, proxied (negatively) by the expense ratio. The positive and highly significant coefficient on this interaction term strongly indicates that higher levels of sophisticated (i.e., low expense ratio) investors in a given money fund lead to larger differences between the predicted behavior of sophisticated and unsophisticated investors.[23]

[23] We perform a number of additional exercises in order to verify the robustness of the results of Table 3. In particular, we show that our results are not particularly sensitive to the inclusion of additional control variables, as well as nonlinear transformations of the existing control variables. We also verify that our findings with the interaction terms (specifications 7-8 in Table 3) hold for other choices of the cutoff separating sophisticated from unsophisticated types. Given space

	Panel A: Prediction 1				Panel C: Predictions 1-3			
	9/15	9/15-16	9/15-17	9/15-18	9/15	9/15-16	9/15-17	9/15-18
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Shareclass expense ratio	1.6**	3.0**	7.8***	10.1***	-1.2	-2.7	-6.1*	-8.4*
	(2.35)	(2.26)	(3.01)	(3.26)	(-1.12)	(-1.40)	(-1.69)	(-1.97)
% Soph (35 bp/yr) × ER					4.1**	8.0**	19.5***	26.1***
					(2.21)	(2.48)	(3.48)	(3.94)
Specification	All institutional, Fund F.E. (N=258)				All institutional, Fund F.E. (N=258)			
R^2	0.542	0.538	0.436	0.521	0.552	0.555	0.461	0.546

	Panel B: Predictions 1-2				Panel D: Predictions 1-3			
	9/15	9/15-16	9/15-17	9/15-18	9/15	9/15-16	9/15-17	9/15-18
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
% Soph (35 bp/yr)	-6.4***	-10.0***	-19.2***	-27.8***	-8.9***	-13.9***	-29.6***	-44.4***
	(-2.67)	(-3.12)	(-3.93)	(-4.46)	(-2.94)	(-3.48)	(-4.80)	(-5.39)
% Soph (35 bp/yr) × ER					3.4***	4.9***	11.1***	15.6***
					(2.92)	(3.23)	(4.73)	(4.91)
Shareclass expense ratio	2.4	3.3	13.1**	15.7**	-1.0	-1.1	-2.3*	-3.2*
	(1.16)	(1.03)	(2.51)	(2.39)	(-1.64)	(-1.53)	(-1.88)	(-1.96)
Log total fund assets	-1.5***	-2.6***	-5.7***	-8.4***	-1.0***	-1.9***	-4.0***	-6.0***
	(-2.82)	(-2.74)	(-3.51)	(-4.13)	(-2.78)	(-2.78)	(-3.54)	(-4.11)
Specification	ER $\leq$ 35 bp/yr, Controls, No F.E. (N=161)				All institutional, Controls, No F.E. (N=258)			
R^2	0.062	0.072	0.107	0.152	0.088	0.115	0.129	0.165

TABLE 4—REGRESSIONS OF SHARECLASS-LEVEL FLOWS ON FUND- AND INVESTOR CHARACTERISTICS: CUMULATIVE FLOWS OVER THE COURSE OF LEHMAN WEEK

Note: This table presents estimated coefficients (t-statistics) from OLS regressions of the change in the log of shareclass-level assets under management (i.e., flows as a fraction of lagged assets under management) for prime institutional money market funds ($\times 100$), cumulated over parts of the period from September 15-18, 2008, on expense ratio, % large-scale institutions (%Soph; calculated with 35 bp/yr cutoff), as well as the log of fund size. Regressions in panels B and D also control for average annualized gross yield, liquid asset share, and fund business; these (insignificant) coefficients are suppressed. See notes to Table 3 for more detailed variable descriptions. Expense ratio (ER) and size are divided by their cross-sectional standard deviations. Different columns denote different sample periods. Panels C and D add an interaction term between shareclass expense ratio and $\%Soph$ at the fund-level.

We note, here, that the liquidity mismatch of a MMMF's balance sheet is captured by the variable %Soph, on the liability side, and by Liquid Asset Share, on the asset side. The regression results in Table 3 indicate that funding (liability side) liquidity is a much more significant driver of investor redemption behavior than market (asset side) liquidity. However, this is very likely due to the implications of funding liquidity for asset liquidity–changes in funding liquidity (i.e., increases in %Soph) must be met with changes in asset liquidity (i.e., increases in Liquid Asset Share). In the sense that abrupt changes in funding liquidity lead changes in asset liquidity, we should expect the former to be a more significant predictor of investor behavior, and it is.

Our tests in Tables 2-3 examined the determinants of shareclass-level flows cumulated over the entire week following Lehman's failure. Table 4 provides additional insights

considerations, these results are reported in a supplemental Online Appendix.

about dynamics by demonstrating how these cross-sectional differences evolved over the course of the week. Each panel of the table has four columns. Column one reports coefficients for a regression whose dependent variable is the shareclass-level flow for Monday, September 15th; the dependent variable from column two includes flows from September 15th and 16th, and so forth. Different panels correspond to different specifications from Tables 2 and 3: Panel A corresponds to specification 3 of Table 2, while Panels B-D correspond to specifications 3, 8, and 7 of Table 3, respectively.

The results are qualitatively similar across all three panels: The cross-sectional differences emphasized in Tables 2-3, which are consistent with all three predictions of the model, are already present on September 15th, and these differences become more pronounced as the week progresses. This pattern is consistent with the most sophisticated investors being among the early movers, particularly when complementarities at the fund level are strong.[24]

Our finding that large investors are more prone to run contrasts with the evidence in Chen et al. (2010) who, for a sample of equity mutual funds, find that large investors are less likely to run. In our setting, institutional investors are assumed to be atomistic but better-informed about fundamentals, making them more likely to run. In contrast, Chen et al. (2010) model a large investor that internalizes the negative effects of its future actions–weakening complementarities–when the primary source of complementarities is the price impact of future redemptions. While these findings contrast, note that the two studies' settings are markedly different. First, MMMFs' fixed NAV structure and potential sponsor support create additional first-mover advantages which are not present for equity mutual funds. MMMF payoffs are specifically designed so as to be informationally-insensitive, making price discovery difficult, whereas equity prices are readily available. Prior to the Lehman episode, larger MMMF investors had a much stronger incentive to invest in information acquisition technologies, and, as such, were much more likely to be

[24]These findings are consistent with He and Manela's (forthcoming) dynamic model, in which investors, upon hearing a rumor that a bank may become illiquid, optimally wait for a time before running on the bank. Crucially, the equilibrium waiting time decreases when other agents are more likely to be informed, implying that the cross-sectional differences in run-like behavior should grow (as we see in the data) during the course of the crisis. As we discuss in the Online Appendix, a version of their model also implies similar comparative statics to those from our static model.

aware of the potential gap between book and market values relative to smaller investors, strengthening complementarities. Second, even though institutional investors' money market accounts could be very large, typically no single institutional investor dominates a single MMMF. The differences in model assumptions and empirical results highlight the importance of the market setting for investor incentives and suggests that there is room for further empirical and theoretical work on the subject.

A caveat is in order before proceeding. Our analysis uses investor clientele for identification of strategic complementarities. Other (omitted) variables could be correlated with variation in clientele across funds, and so, ultimately, we cannot rule out alternative explanations of our findings related to the influence of unobserved asset quality (beyond the proxies for asset quality that we already include, such as yield and liquidity) or unobserved investor characteristics (beyond our proxy of expense ratio). For instance, it is plausible that investors are more likely to receive a bad signal about the quality of asset holdings in funds with high fractions of sophisticated investors, and, so, redeem more aggressively from such funds–vs. our hypothesis that sophisticated investors are reacting strategically to each other. We note, however, that such increased information, under this mechanism, would have to be available to sophisticated investors, and not to the unsophisticated–otherwise, our within-fund tests would still control for the information environment. Moreover, this alternative mechanism does not explain why (as we will show shortly) sophisticated investor flows should depend on unsophisticated investor actions. Thus, we believe that our hypotheses regarding strategic complementarities are more reasonable, given our empirical results, than this alternative hypothesis.

IV. The role of fixed NAV: Comparison with ultra-short (variable NAV) bond funds

The cross-sectional tests of Table 3 explore strategic complementarities by changing the type of investor (large-scale "sophisticated" vs. smaller-scale "unsophisticated" institutional investors), while controlling for portfolio fundamentals (quality of assets), as proxied by the lagged 7-day portfolio gross yield. A current issue of contention among academics, regulators, and the investment management industry is whether the fixed share price of MMMFs is a significant contributor to the potential for investor runs.[25] Indeed, the latest round of SEC regulatory changes, effective October 2016, requires all prime institutional shareclasses to be offered for purchase and sale (by the management company) at a "floating NAV," defined as a per-share price that reflects an assessment by the management company of the daily market valuation of portfolio assets.[26]

The regulatory point of view is that a fixed $1 per share creates an enhanced first-mover advantage in response to an adverse shock to credit quality. In this setting, the ideal test for the effect of the fixed NAV structure on investors' redemption behavior would be to compare the actions of similar investors which hold mutual funds with similar credit risk exposures, where one group of investors has a floating NAV while the other has a fixed NAV. Since no MMMFs carried a variable share price in 2008, we examine the sector of mutual funds with holdings closest to prime MMMFs, and carrying a variable price: ultra-short bond funds. We analyze those ultra-short funds closest to prime MMMFs: ultra-short funds that hold both government and non-government debt (i.e., we exclude ultra-short government-only funds). Ultra-short funds do not fix their share price so, according to this point of view, we would expect complementarities between investors to be more keen in MMMFs than in ultra-short bond funds during the September 2008 market run since a floating (market) share price is designed to internalize a greater portion of the (negative) externalities of running from a fund as its fundamentals deteriorate.

We propose a simple test which provides suggestive evidence for the effect of the fixed NAV on investor redemption behavior. Our key identifying assumption is that the change in credit quality of the ultra-short portfolios was weakly larger relative to the change experienced by MMMFs operated by the same fund management company. Then, under the hypothesis that the fixed NAV structure has no effect on investors' incentives to redeem quickly, we would expect redemptions in ultra-short funds to be weakly larger than redemptions from comparable MMMF investors.

[25]See, for example, Gordon and Gandia (2014), SEC (2014) www.sec.gov/rules/final/2014/33-9616.pdf, and Fidelity Management and Research (2013) https://www.fidelity.com/about-fidelity/corporate/nancy-prior-speech-striking-the-right-regulatory-balance-for-money-market-mutual-funds.

[26]See www.sec.gov/rules/final/2014/33-9616.pdf. As another acknowledgement of the enhanced first-mover advantage among institutions, prime institutional shareclasses will no longer be allowed to co-exist as claims on a fund (a portfolio) with retail shareclasses.

	(1)	(2)	(3)	(4)
Complex Prime Retail MMF flow	0.53***		0.60***	
	(4.109)		(3.493)	
Complex Prime MMF flow		0.34**		0.18*
		(2.145)		(1.869)
Institutional shareclass dummy	0.34	-1.22	0.38	0.15
	(0.093)	(-0.280)	(0.074)	(0.028)
Expense ratio	1.65	2.64	2.26	2.01
	(1.127)	(1.585)	(0.870)	(0.685)
Annualized yield	0.30	0.87	-0.02	0.75
	(0.213)	(0.632)	(-0.013)	(0.460)
Log total fund assets	-6.29***	-5.82***	-7.31***	-4.82*
	(-3.608)	(-2.979)	(-2.786)	(-1.936)
Sample selection	All Ultrashort funds		Ultrashort funds > $10 MM	
Number of observations	45	50	36	39
R^2	0.302	0.234	0.273	0.132

TABLE 5—REGRESSIONS OF SHARE CLASS-LEVEL FLOWS FOR SAME-COMPLEX ULTRA-SHORT BOND FUNDS ON FLOWS FOR PRIME MMMFS

Note: This table presents estimated coefficients from OLS regressions of share class-level flows as a fraction of lagged assets under management (i.e., the change in log assets under management, adjusted for changes in market value $\times 100$) for prime ultra-short bond fund institutional share classes during the month of September 2008 on flows over the same period from same-complex prime retail MMMF (models 1 and 3) or same-complex prime retail plus prime institutional MMMF share classes (models 2 and 4); in each case, these same-complex MMMF flows are aggregated to the complex level. In addition, each regression includes a dummy variable which equals one for an institutional share class, the shareclass expense ratio, the log of fund size (total net assets under management as of August 31, 2008) and its annualized gross yield (fund interest distributions over the previous 12 months, divided by current share price); the latter three variables have been divided by their cross-sectional standard deviations. All variables (including money market flows) are constructed using the CRSP mutual fund database.

Ultra-short bond funds serve a clientele consisting almost exclusively of retail investors.[27] Thus, a perfect comparison to our prime institutional MMMF share classes, which consist of much larger proportions of truly non-retail investors, is not possible. However, we compare flows during and after the crisis week of 2008 between ultra-short funds and both retail and institutional prime MMMFs to gain insights about the role of a floating share price in generating fragility. Here, we conduct such an analysis using monthly flow data and other fund characteristics from the Center for Research in Security Prices (CRSP) for both ultra-short funds and prime MMMFs.[28]

[27]Closer inspection of the ultra-short funds data suggests that approximately 85% of assets under management in institutional shareclasses are held either directly or indirectly by retail investors. Retail investors are able to purchase lower-cost institutional share classes in ultra-short funds through, for example, their 401(k) plans. We thank the Investment Company Institute for this information.

[28]We use CRSP MMMF data, instead of iMoneynet data because CRSP provides a unique identifier which allows us to link MMMFs with ultra-short funds managed within the same fund complex. Since ultra-short funds that hold only government securities are not separately identified by the Lipper classification in CRSP, we manually identify those funds that also invest in non-government securities in order to provide a better comparator group to prime MMMFs, which also invest outside government securities. CRSP asset data are only available on a monthly basis; accordingly, we focus on explaining investor flows over the entire month of September 2008. Monthly flows (as a fraction of beginning-of-month assets) for a given share class during September are estimated as $flow_{Sep} = \ln(tna_{Sep}) - \ln(tna_{Aug}) - r_{Sep}$, where tna_{Sep} and r_{Sep} denote total net assets under management at the end of September 2008 and (continuously compounded) monthly return during that month for a given shareclass.

Table 5 reports the results from regressing crisis-month share class level flows to ultra-short prime bond funds on same-complex flows to prime MMMF retail share classes (models 1 and 3) or to same-complex prime retail plus prime institutional share class MMMF flows (models 2 and 4).[29] The regressions include a dummy equaling one for institutional (vs. retail) ultra-short shareclasses, as well as (normalized) controls for the expense ratio, yield, and ultra-short fund size (aggregated across all share classes). The rightmost two columns exclude very small ultra-short bond fund share classes ($<$ \$10 million in assets) from the analysis as a robustness test. In the regressions of ultra-short flows on retail MMMF flows (1 and 3), the estimated slope coefficient is 0.53 to 0.6, and is (statistically) significantly different from both zero and one; this reflects that flows to or from ultra-short share classes have the same direction (generally outflows), but roughly half the amplitude of same-complex prime retail MMMF share classes during the crisis week. Models 2 and 4 show that the amplitude of ultra-short share class flows is even more muted, compared to all prime MMMF share classes within the same complex. (As expected, the institutional ultra-short dummy has little relevance, since most investors in this class are, in reality, retail investors).

To give a broader sense of the fundamental shock and the resulting price pressure on the shares of ultra-short funds, Panel A of Figure 3 shows the average cumulative holding period return for ultra-short funds during each day of September 2008. Ultra-short funds experienced three consecutive daily returns more than three standard deviations below zero, with cumulative losses of -75 bps during the Lehman week. In contrast, MMMFs (excluding funds managed by the Reserve complex) experienced modest gains from earned interest. Note, from Panel B, that the floating value of ultra-short funds reduced outflows from these funds during the month of September.

By October 2008, MMMF shares were backstopped through several programs initiated by the Federal Reserve and the U.S. Treasury, which significantly reduced the scope for complementarities.[30] In contrast, no backup programs were implemented explicitly for

[29]We compare flows within the same complex as a rough control for the implicit potential subsidization by the complex during a crisis and for the quality of portfolio holdings.

[30]Specifically, programs were announced on September 19 (Treasury), October 7 and 21 (Federal Reserve). Further, the SEC allowed MMMFs to price their portfolio holdings at amortized cost for a period (when quotes on commercial

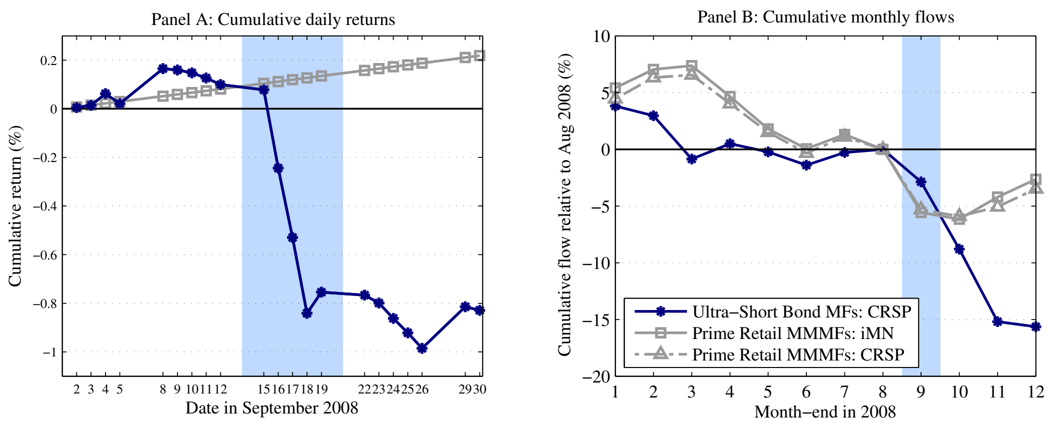

FIGURE 3. CUMULATIVE RETURNS AND FLOWS FOR MMMF AND ULTRA-SHORT BOND FUNDS

Note: Panel A shows the daily return of prime MMMFs and ultra-short bond funds, averaged (value-weighted) across all funds each day, then cumulated from the beginning of the month to each day during September 2008. Panel B shows the monthly flow (i.e., the change in log of assets under management, adjusted for investment returns) each month during 2008, averaged (value-weighted) across all prime MMMF retail shareclasses (using both the iMoneynet dataset and the CRSP dataset, which match closely) and all ultra-short bond funds, then cumulated over consecutive months in 2008. We normalize the cumulative flow to zero as of the end of August 2008.

ultra-short bond funds (although they may have indirectly benefitted from the commercial paper backup programs). This is consistent with the longer-term cumulative flow patterns observed in Panel B, which show that ultra-short bond funds experienced less-severe outflows during September 2008, but continued to experience moderate outflows (and losses) during October and November. By contrast, outflows from prime MMMFs largely ended by early October.

Differences between MMMFs and ultra-short bond funds prevent us from drawing too strong conclusions from these results about the role of a variable share price in reducing the importance of payoff complementarities. While ultra-short funds normally hold riskier assets than MMMFs (thus, the variable share price), they might not have participated as heavily in the degradation of asset quality in the financial commercial paper market during September 2008. With this caveat, our results point to fixed share prices as likely amplifying the run on prime MMMFs triggered by the deterioration in the fundamentals during the Lehman week. Further, as is studied in greater detail for corporate bond funds by Goldstein et al. (2015), there is still scope for complementarities even

paper were generally regarded as unreliable), which potentially helped to forestall more MMMFs from breaking the buck.

with a floating NAV. Investors still have an incentive to redeem early given the low degree of secondary market liquidity and the potential lack of reliable market prices for use in NAV calculations. Moreover, the new regulations will effectively segment the MMMF market between institutional and retail investors, which, according to our model above, will potentially strengthen complementarities among institutional investors.

V. Dynamic Interactions between Investor Types

As noted earlier, a unique feature of our data is that we observe the daily flows of shareclasses associated with investor types of different levels of sophistication. In this section, we use our data to study whether the actions of unsophisticated investors affect the actions of sophisticated investors in a dynamic setting.

To measure flow dynamics, we estimate a model that captures daily interactions between sophisticated and unsophisticated investors within the same MMMF. As above, we partition institutional shareclasses as sophisticated or unsophisticated based on the level of their expense ratios, while retail shareclasses are always classified as unsophisticated.[31] To maximize the number of paired observations in our sample while retaining consistency across funds, institutional shareclasses are classified as sophisticated if their expense ratio (1) falls below the median expense ratio within the institutional shareclasses of that fund; and (2) is no greater than 35 bps/year.[32]

Table 6 reports results from regressions that use daily (fraction) flow from low expense ratio (sophisticated institutional) shareclasses, Low_{it}, aggregated within a given MMMF, as the dependent variable. We regress this variable on its own one-day lag, Low_{it-1}, as well as one-day lagged flow from aggregated (within a given MMMF) high expense ratio (unsophisticated institutional and all retail) shareclasses, $High_{it-1}$, and interactions of these variables with $\%High_{it-1}$ and $\%High_{it-2}$, the one- and two-day lagged investment value of unsophisticated investors (institutional plus all retail) in a given prime

[31]Funds with a single share class are excluded from the analysis, which leads to a reduction in the sample size from 123 to 64 funds. The excluded share classes have similar yield, liquidity and flow standard deviations but tend to be smaller with a lower concentration of sophisticated investors (higher ERs) than funds with multiple share classes.

[32]In a few cases, two shareclasses within the same fund have the same expense ratio. Prior to generating the ordering, we pool the assets under management for such shareclasses. In cases with an odd number of institutional shareclasses, we code the median shareclass as sophisticated only if its expense ratio is no greater than 35 bps.

Variable	Early Crisis (1)	Peak Crisis (2)	Lehman Week (3)	Early Crisis (4)	Peak Crisis (5)	Lehman Week (6)
$Low_{i,t-1}$	0.20	0.18	0.25**	0.17	0.19	0.28**
	(1.32)	(1.16)	(2.35)	(1.18)	(1.13)	(2.40)
$Low_{i,t-1} \times \%High_{i,t-2}$	-0.25	0.35	0.03	-0.42	0.15	-0.21
	(-0.62)	(0.54)	(0.07)	(-1.24)	(0.26)	(-0.73)
$High_{i,t-1}$	-0.04	-0.32**	-0.16*	-0.03	-0.27*	-0.12
	(-0.38)	(-2.26)	(-1.77)	(-0.34)	(-1.83)	(-1.17)
$High_{i,t-1} \times \%High_{i,t-2}$	0.85**	1.22**	1.02**	0.46	1.05**	0.67*
	(2.19)	(2.23)	(2.42)	(1.46)	(2.53)	(1.99)
$High_{i,t}$				-0.23**	-0.10	-0.19
				(-2.40)	(-0.78)	(-1.57)
$High_{i,t} \times \%High_{i,t-1}$				1.58***	1.54**	1.64***
				(4.75)	(2.47)	(3.13)
$\%Low_{i,t-1}$	0.83	-9.38***	-4.66***	0.21	-12.98***	-7.21***
	(0.46)	(-2.99)	(-2.93)	(0.13)	(-3.60)	(-4.56)
Average yield$_{i,t-1}$	0.22	-0.76	-0.33	0.08	-0.81	-0.32
	(0.47)	(-0.64)	(-0.41)	(0.19)	(-0.73)	(-0.49)
Log total fund assets$_{i,t-1}$	-0.88**	-0.15	-0.61	-0.90**	-0.10	-0.58
	(-2.22)	(-0.17)	(-1.01)	(-2.22)	(-0.11)	(-1.04)
Avg institutional expense ratio$_{i,t-1}$	1.02	2.59	2.69**	0.67	1.90	2.15**
	(1.65)	(1.59)	(2.35)	(1.08)	(1.28)	(2.35)
Liquid asset share$_{i,t-1}$	0.53	1.89*	1.38*	0.58	1.79	1.45**
	(1.16)	(1.67)	(1.95)	(1.29)	(1.61)	(2.29)
Fund business$_{i,t-1}$	-0.24	0.01	-0.12	-0.26	0.04	-0.06
	(-0.79)	(0.02)	(-0.20)	(-0.91)	(0.05)	(-0.12)
N	320	190	318	320	190	318
R^2	0.166	0.279	0.263	0.233	0.370	0.348

TABLE 6—DETERMINANTS OF DAILY FLOWS FROM SOPHISTICATED (LOW ER) SHARECLASSES

Note: For each fund, we separate prime institutional share classes into two categories, based on their expense ratios. The first category, "Low," consists of share classes which have expense ratios that are lower than the median expense ratio (across all institutional share classes within a given fund). All remaining share classes are included in the "High" category, including all retail share classes. The value of shares outstanding is then aggregated across all Low share classes and, separately, across all High share classes within a given fund. (Funds with a single share class are excluded from this analysis.) For each fund and date, we calculate the first difference in the log of aggregate value within each category (i.e., fraction flow), which we denote by $Low_{i,t}$ and $High_{i,t}$. The table presents the coefficients from panel regressions with $Low_{i,t}$ as the dependent variable on $Low_{i,t-1}$ and $High_{i,t-1}$, estimated for three different subperiods in 2008: 9/10-9/16 "Early Crisis", 9/17-9/19 "Peak Crisis", and 9/15-9/19 "Lehman Week", respectively. We multiply $Low_{i,t}$ and $High_{i,t}$ by 100 to express them in log percentage points. We also include interaction variables between, for example, $High_{i,t-1}$ and $\%High_{i,t-2}$, which is defined as two-day lagged fraction of total MMMF value within a MMMF represented by "High" (both institutional and all retail) ER shareclasses. Control variables, described in the notes to Table 3, have been divided by their (cross-sectional) standard deviations for ease of interpretation. All specifications also include unreported time dummies. Standard errors are clustered at the fund level. R^2 reports the overall R^2. t-statistics are reported in parentheses.

MMMF as a fraction of fund TNA.[33] These interaction terms are key to our analysis: they test whether sophisticated investors react differently to the behavior of unsophisticated investors when the latter comprise a high fraction of the fund's assets.

Finally, we include five control variables (yield, log fund size, and average institutional expense ratio within a given MMMF, liquid asset share, and fund business, defined as described previously for Table 3) in our regression. Average institutional expense ratio

[33] $High_{it-1}x\%High_{it-2}$ is the contribution to total lagged fund flows coming from unsophisticated investors.

is used to control for differences in clientele across different MMMFs.

To explore differences in the dynamics of flows over potentially very different market regimes, we estimate the model over two separate sub-periods, namely, the "early-crisis" period (Wednesday, September 10 to Tuesday, September 16, 2008) and the "peak crisis" period (Wednesday, September 17 to Friday, September 19).[34] We also provide estimates over the full "Lehman week" (September 15 to September 19).

First, consider the coefficient on the interaction term $High_{it-1} \times \%High_{it-2}$, which captures the concentration of unsophisticated investors ($\%High_{it-2}$) in amplifying the reaction of sophisticated investors to the outflows of the unsophisticated, $High_{it-1}$ (Prediction 2 from our model of Section II). In both subsamples, and in the entire Lehman week, this coefficient is positive, close to one, and highly significant. These estimates show very different responses of sophisticated investors to prior-day flows of unsophisticated investors across funds with different concentrations of unsophisticated investors. To illustrate, during the Lehman week as a whole, in response to a 5% outflow by unsophisticated investors, a prime MMMF that consists of a 50% concentration of unsophisticated money experiences a one-day (expected) outflow that is 2% of assets greater, among sophisticated investors, relative to a prime MMMF with a 10% concentration (exp(1.02x-.05x.4)-1).

These results suggest that the actions of sophisticated and unsophisticated investors become more coordinated in funds where the latter group plays a greater role. To our knowledge, these dynamic interactions between heterogenously-informed investors have not been modeled formally in the literature. However, we find results which are intuitively consistent with complementarities. When agents' payoffs depend on the actions of others and there are complementarities, agents want their actions to be positively correlated with the aggregate action, so that they run when other investors run. This implies that, in funds dominated by unsophisticated investors, sophisticated investors optimally wish to observe the actions of unsophisticated investors before timing their withdrawals,

[34]We partition in this manner because the news that the Reserve Primary Fund "broke the buck" due to its Lehman holdings first became known by the public late in the day on September 16. By the following morning, the media had widely circulated this information.

which is what we find.

We note here that one might interpret the positive coefficient on the interaction term in an alternative way that has little to do with strategic complementarities. That is, if funds sell their highest quality assets first to meet redemptions of any kind, lagged outflows from low and high expense ratio shareclasses should have the same effect on Low_{it}–the idea being that sophisticated investors are merely reacting to worsening fundamentals, which might be inferred from observing one-day lagged outflows. We provide some evidence counter to this alternative mechanism: in Table 6, the coefficients on $Low_{it-1} \times \%High_{it-2}$ and $High_{it-1} \times \%High_{it-2}$ having the same magnitude, but opposite signs. Thus, it appears that outflows are viewed, by sophisticated investors, not in isolation, but in the context of the concentration of unsophisticated investors in the MMMF.[35]

We further note that, consistent with Prediction 2 of our model, the coefficient on $\%Low_{it-1}$ is statistically significant and economically large during the peak crisis and Lehman week. During the peak crisis, the coefficient on this variable suggests cumulative three-day redemptions about 25% higher for funds consisting purely of sophisticated investors, compared with funds consisting almost entirely of unsophisticated investors.

To put this into perspective, in the dynamic model of Angeletos, et al. (2007), informed investors update their beliefs based on a noisy signal about the size of previous attacks, and the observation of whether the regime (in our case, the probability of maintaining the $1 NAV) survived the previous-period attack. Large prior attacks, ceteris paribus, reveal negative information about fundamentals, making future attacks more likely.[36] However, an institution's survival from previous attacks may also result in upwardly revised beliefs among investors about the strength of fundamentals. Our empirical results on informed investors' response to prior-day redemptions of other informed investors reflects the relative importance of these two channels (the existence of large attacks vs.

[35] Another interesting dynamic is modeled by He and Manela (forthcoming), who show that less-informed agents may optimally wait and withdraw after early movers (who are, in equilibrium, better-informed) after learning that the bank's portfolio liquidity may have been impaired. That is, less-informed investors wait until the marginal cost of waiting (imposed by the risk of additional better-informed investors running) equals the marginal benefit of waiting (through earning interest). In unreported tests, we find some evidence of unsophisticated investors reacting to the lagged outflows of sophisticated investors, however, this effect is weaker than the reaction of sophisticated shown in Table 6.

[36] If prior attacks also serve to weaken fundamentals (e.g., by forcing funds to sell their most liquid asset holdings), future attacks become even more likely in the presence of large prior attacks.

the survival of funds after large attacks). Overall, our evidence suggests that attacks by sophisticated investors were viewed negatively during the crisis week, and did not result in upwardly revised beliefs about fundamentals (even though no other funds, besides the Reserve Primary fund, officially broke the buck).[37]

Turning to the control variables, liquid asset share has a statistically significant effect on flows during the peak crisis period, and during the entire Lehman week, but not during the early crisis, suggesting that concerns about asset liquidity became much more keen as the crisis unfolded. As in prior tables of this paper, we also find some evidence that sophisticated investors are more likely to run in larger funds, although this evidence is strong only during the early crisis period. Funds with larger average institutional expense ratios, and therefore less sophisticated investors (relative to other funds), on average experienced weaker redemptions. This finding reinforces that the level of sophistication of an investor plays a key role in the tendency to redeem during the crisis.

Finally, we recognize that low ER investors could also respond to the same-day behavior of high ER investors.[38] Redemption requests at MMMFs are placed throughout the day. While most funds redeem all shares at the end of the day, some institutional share classes allow redemptions at various points during the day. Regardless, investors submit redemption requests throughout the day, even if their redemptions are not honored until 4 p.m. Eastern Time. Thus, it is plausible that large investors could be "tipped off" about the behavior of small investors in plenty of time to redeem their own shares on the same day.[39] Columns 4-6 in Table 6 account for this effect by adding *contemporaneous* flows from small-scale investors, $High_{i,t}$, along with an interaction term, $High_{i,t} \times \%High_{i,t-1}$. We find only a modest effect of the contemporaneous value of $High_{it}$ on Low_{it}.[40] In contrast, the contemporaneous interaction term

[37] We note here, however, that several funds ventured close to breaking the buck, and may have avoided doing so only because their fund boards were reluctant to make this move even though fundamentals indicated that the buck had already, in reality, been broken. Others avoided breaking the buck through subsidizations by their advisors.

[38] Angeletos and Werning (2006) present a model in which a subset of (better informed) investors receive a noisy signal about redemptions of other (less informed) investors. They consider this public signal in conjunction with their own private signals prior to making their decisions.

[39] On September 19, 2008, Ameriprise sued the Reserve Management Company, alleging that larger institutions were "tipped off" about the Primary Fund's holdings of Lehman on Monday, September 15, but not smaller investors. Note that the information from the August 31, 2008 portfolio report was not filed with the SEC until October 29, 2008.

[40] Admittedly, the simultaneity of $Low_{i,t}$ and $High_{i,t}$ makes this test somewhat more crude, particularly with respect

$High_{i,t} \times \%High_{i,t-1}$ is significant and economically large. Sophisticated and unsophisticated investor actions are apparently more likely to be coordinated precisely when our model predicts that they have the strongest incentives to coordinate.

VI. Nonlinearities and Magnitudes: Evidence from Dynamic Quantile Regressions

We next adopt a panel quantile regression approach that further allows us to identify the fundamental characteristics of a fund that may make it more susceptible to run-like behavior by investors and helps identify what role the size and/or sophistication level of a fund's investor base play in increasing its exposure to run-like risk. Our quantile approach may also be more robust to potential multiplicity of equilibria.[41] We see the primary purpose of this analysis as helping to uncover the magnitude of cross-sectional heterogeneity, both before and after controlling for observable characteristics. Second, in the absence of a liquid secondary market, optimal policy and welfare calculations are likely to be sensitive to the *ex-post* distribution of outflows across funds.[42]

Figure 4 plots the 10^{th}, 50^{th}, and 90^{th} quantiles of the daily percentage change in total assets under management ("flow") for each fund within prime institutional and prime retail share classes (each dollar-aggregated within the fund). During the period ending on Friday, September 5, 2008, one week prior to the Lehman bankruptcy, the distribution of flows across prime institutional share classes (Panel A) is fairly tight. However, the distribution widens during the following two weeks; massive redemptions are highly concentrated among a small subset of funds. For instance, the 10^{th} percentile, on September 17, experiences outflows greater than 15% of prior-day total net assets. In contrast, the median fund experiences an outflow of about 2% on the same day.

to interpreting the coefficient on $High_{i,t}$. However, the main coefficient of interest is the interaction term.

[41]Global games models with sufficiently precise private information result in a unique equilibrium and so allow for clear predictions and comparative statics results; our Proposition 1 is an example of this. However, these conditions may not hold, leaving a role for multiplicity and for nonfundamental sources of volatility to affect outcomes. See Echenique and Komunjer (2009) and Angeletos and Pavan (2012) for further details about this robustness property.

[42]This is true even if all funds face the same *ex-ante* run risk, particularly if there is non-trivial heterogeneity in portfolio holdings. This would be the case if, for example, fund-level redemption costs are a convex function of outflows. If runs are concentrated at a small number of funds, those assets that are overweighted by these funds are likely to face substantial selling pressure, potentially leading to fire sales, followed by further outflows. These fire sales could also create lucrative opportunities for funds which do not experience excessive ouflows, since yields on the distressed assets would rise due to these frictions, creating an additional incentive for investors to move their money in the same direction as other investors. Not only might redemptions be absorbed more easily in the former, relative to the latter case, but there is also less likelihood of amplification through future withdrawals.

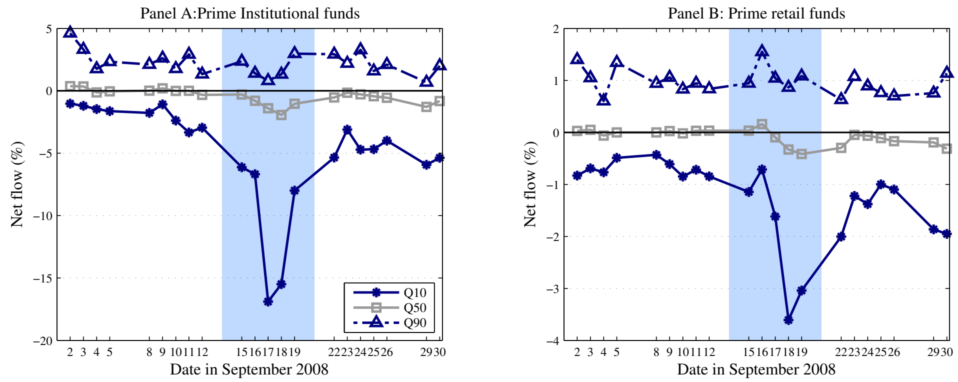

FIGURE 4. QUANTILES OF DAILY FLOW DISTRIBUTIONS BY CATEGORY DURING SEPTEMBER 2008

Note: This figure plots the 10th, 50th, and 90th daily quantiles of log (fund-level) aggregated prime institutional share class total net assets (i.e., daily fraction flows) in Panel A, and prime retail in Panel B, during September 2008. The week following the failure of Lehman Brothers, September 15-19, is shaded.

While the pattern of flow quantiles for prime retail share classes (Panel B) is qualitatively similar to that of prime institutional share classes, retail investors exhibited a much more muted response to the crisis. Furthermore, peak outflows among retail funds occurred on September 18, a day later than peak outflows among institutional funds. Those few retail investors who redeem only do so after information from the popular media becomes widely available (on September 17-18) about the evolving crisis.[43]

Next, we estimate a dynamic model for the distribution of flows, as summarized by these 10th, 50th, and 90th percentiles of daily changes in fraction flows, given observable characteristics, using dynamic panel quantile regressions.[44] In this way, we can determine whether fund and/or investor characteristics differentially influence flows in different parts (e.g., the center vs. tails) of the conditional cross-sectional distribution. Other than lagged fund flows, our model includes the following covariates: %sophisticated investors (prime institutional share classes with ER$\leq$35 bps/year), logged flow standard deviation, average yield, and logged fund total assets.[45] We estimate the relevant pa-

[43]See, for example, "Money Market Funds Enter a World of Risk," New York Times, September 17, 2008.

[44]For simplicity, we now refer to the aggregate of prime institutional share classes as a prime institutional "fund," but the readers should be reminded that a fund can consist of both prime institutional and prime retail share classes.

[45]To account for asymmetries in fund flows, our model allows lagged fund flows above the cross-sectional median flows to have a differential effect, relative to lagged fund flows below the median. Moreover, to keep the specification simple (as quantile regressions already use a large number of degrees-of-freedom), we employ only four covariates, but our results are robust to including other covariates in the model. Furthermore, for robustness, our results winsorize the lagged flow variables and the flow standard deviation at their 2% values within each tail.

Variable	Value	Cumulative Flow Quantile 1%	5%	10%	50%	90%
	$f(\bar{x})$	-51.25	-41.21	-35.61	-17.09	1.09
% Sophisticated	$f(\bar{x}+\sigma_x)$	-62.70	-51.39	-45.28	-22.38	-0.14
	$f(\bar{x}-\sigma_x)$	-40.88	-32.01	-27.37	-12.50	3.09
	Difference	-21.82 ***	-19.37 ***	-17.91 ***	-9.88 ***	-3.23
	p-value	[0.004]	[0.003]	[0.002]	[0.001]	[0.184]
	p-value vs. median	[0.025]	[0.022]	[0.021]	-	[0.041]
Average gross yield	$f(\bar{x}+\sigma_x)$	-55.05	-44.44	-38.75	-19.16	0.30
	$f(\bar{x}-\sigma_x)$	-48.15	-38.10	-32.69	-14.91	1.98
	Difference	-6.90 *	-6.34 **	-6.06 **	-4.25 **	-1.68
	p-value	[0.050]	[0.040]	[0.028]	[0.011]	[0.223]
	p-value vs. median	[0.178]	[0.187]	[0.176]	-	[0.139]
Log flow std. dev.	$f(\bar{x}+\sigma_x)$	-62.68	-50.13	-43.50	-19.07	11.61
	$f(\bar{x}-\sigma_x)$	-41.75	-33.24	-28.80	-14.47	-2.05
	Difference	-20.94 ***	-16.89 ***	-14.70 ***	-4.60 **	13.66 **
	p-value	[0.000]	[0.000]	[0.000]	[0.014]	[0.013]
	p-value vs. median	[0.000]	[0.000]	[0.000]	-	[0.000]
Log fund total assets	$f(\bar{x}+\sigma_x)$	-56.90	-46.34	-40.68	-20.60	-0.01
	$f(\bar{x}-\sigma_x)$	-45.64	-35.91	-30.69	-13.25	3.04
	Difference	-11.26 **	-10.43 **	-9.99 ***	-7.35 ***	-3.05
	p-value	[0.038]	[0.019]	[0.009]	[0.000]	[0.158]
	p-value vs. median	[0.246]	[0.239]	[0.225]	-	[0.089]

TABLE 7—MARGINAL EFFECTS OF FUND CHARACTERISTICS ON CUMULATIVE FLOW QUANTILES

Note: This table shows the impact of explanatory variables on cumulative flow distributions (as a percentage of initial assets) for prime institutional share classes (aggregated to the fund level) for the September 15-19 period. These estimates are obtained by simulating from an estimated dynamic quantile panel regression model for daily flows that is further described in an appendix. Columns report the 1st, 5th, 10th, 50th, and 90th quantiles of the cumulative flow distributions, respectively. We begin by fixing each of the explanatory variables at its average, assuming that the initial value of lagged flows equals the prime institutional category average. Then, we report the impact on the simulated flow distribution of adding and subtracting one standard deviation to each explanatory variable, as well as p-values for a test of whether the difference in the simulated quantiles is statistically significant, obtained by using the bootstrapped distribution of parameter estimates from our model, as well as the p-value of whether the marginal effect is significantly different at a given quantile, relative to the marginal effect at the median (using the bootstrapped distribution).

rameters using the sequential semi-parametric method of Schmidt and Zhu (2015), and calculate standard errors using simple bootstrap procedures.[46] As in Table 6, we allow the coefficients of the dynamic model to change between early and peak crisis periods. Further details and estimates for the model are provided in the Online Appendix.

We use our fitted model to simulate the impact of perturbations in initial observable characteristics on one-week (cumulative) flow distributions.[47] Table 7 presents a summary of this simulation, which provides two primary insights. First, as shown in the first row of Table 7, regardless of fund characteristics, there is substantial heterogeneity in

[46]We use a clustered bootstrap, where we construct our bootstrap-sampled data by drawing a complete fund time series with replacement, which allows for arbitrary serial correlation in the residuals within funds.

[47]We fix initial values for the fund characteristics, yielding initial conditional quantile forecasts, and simulate from a parametric distribution which satisfies the conditional quantile restrictions. Substituting this simulated draw into the law of motion from the estimated model generates conditional quantile estimates for the next day. Iterating forward, we are able to trace out the (simulated) distribution of cumulative flows over the course of the crisis. Further details, along with some additional graphical results, are in the Online Appendix.

fund flows even for funds with similar observable characteristics. Second, we find much stronger evidence (statistically and economically) of a nonlinear dependency of flows on investor characteristics in left-tail quantiles (1% and 5%), relative to the median or right tail quantiles (50% and 90%). For example, at the 1 percentile of flows (the largest outflows), a one standard deviation perturbation upward in the percentage of a prime institutional fund owned by sophisticated investors is expected to produce an outflow that is 22% of assets greater than that of a one standard deviation perturbation downward; at the median, this sensitivity is a much lower 10% of assets; a statistical test that the difference between the impact of %Sophisticated between these two quantiles is zero exhibits a p-value less than 1%. Compared to these effects, the nonlinear impact of the yield and fund total asset variables are somewhat muted. Given the backward-looking nature of the reported yield, it is likely that such information quickly became stale during the fast-moving Lehman week. Finally, the (lagged) flow standard deviation has a modest effect on the median (-2.8%), but increases the spread of the distribution.

VII. Conclusion

This paper studies prime money market mutual fund (MMMF) run-like behavior during the crisis period following the Lehman default of September 2008. We find that run behavior was especially pronounced among prime MMMF share classes that cater predominantly to very large-scale institutional investors. We use data on different share classes within the same fund to identify differences in the flow behavior of sophisticated (low expense ratio) vs. less sophisticated (high expense ratio) investors, keeping portfolio holdings and other fund characteristics, such as implied sponsor subsidization, constant. Our empirical tests show strong evidence consistent with strategic complementarities playing an important role in investor actions during the crisis.

In addition to these tests, our paper provides a set of new empirical findings which shed light on theoretical models of runs on financial intermediaries. First, we show that runs can develop and evolve in a matter of days, and that they involve important feedback effects from past flows on following-day flows. Second, we show that it is difficult to identify, ex-ante, which funds are subject to runs, although there is clear evidence that large scale investors are keenly aware of the quality of the fund holdings, the financial strength of the fund advisor, and the characteristics of other investors in the same funds.

We draw several policy conclusions from our findings. First, the presence of retail share classes and small-scale share classes weaken the strategic incentives for large-scale institutional investors to redeem their shares. This may have implications for the latest round of regulatory reforms, which require both a floating share price (NAV) and no commingling of portfolio assets between an institutional and a retail share class. Our results suggest that, even if floating net asset values weaken complementarities, the segregation of institutional money might strengthen the strategic complementarities among institutional investors; the size of this effect will depend on the mix of sophisticated vs. unsophisticated institutions within a given fund that results after the segregation is implemented. Indeed, our empirical results on ultra-short bond funds, which already carry floating share prices, indicate that such complementarities may persist in prime institutional MMMFs after October 2016.

Finally, our paper brings some suggestions for future research, some of which may carry implications for commercial banks. That is, our findings suggest that further theoretical work is needed on the stability of pooled funds with a large and time-varying maturity and/or liquidity mismatch between assets and liabilities. Most theoretical models tend to consider a static environment, where a single financial institution interacts with investors in isolation, whereas the run-like episodes of the recent financial crisis simultaneously affected multiple financial institutions. Our results shed additional light on the role of market-wide conditions in affecting run-like behavior at individual institutions, which may inform future theoretical work that more fully models the feedback between markets and pooled vehicles that are both under stress.

REFERENCES

Abreu, D., and M. Brunnermeier, 2003, Bubbles and Crashes. *Econometrica* 71, 173-204.

Acharya, V.V., P. Schnabl, and G. Suarez, 2013, Securitization without risk transfer. *Journal of Financial Economics* 107, 515-536.

Allen, F., A. Babus, and E. Carletti, 2009, Financial Crises: Theory and Evidence. Working paper.

Angeletos, G-M, C. Hellwig, and A. Pavan, 2007, Dynamic Global Games of Regime Change: Learning, Multiplicity, and the Timing of Attacks. *Econometrica* 71, 711-756.

Angeletos, G-M, and A. Pavan, 2012, Selection-free Predictions in Global Games with Endogenous Information and Multiple Equilibria. Working Paper.

Angeletos, G-M, and I. Werning, 2006, Crises and Prices: Information Aggregation, Multiplicity, and Volatility. *American Economic Review* 96, 1720-1736.

Calomiris, C., Mason, J.R., 2003, Fundamentals, panics, and bank distress during the depression. *American Economic Review* 93, 1615-1647.

Chari, V.V., and P.J. Kehoe, 2003, Hot Money. *Journal of Political Economy* 111(6), 1262-1292.

Chen, Y., 1999, Banking Panics: The Role of the first-come, first-served Rule and Information Externalities. *Journal of Political Economy* 107, 946-968.

Chen, Q., I. Goldstein, and W. Jiang, 2010, Payoff Complementarities and Financial Fragility: Evidence from Mutual Fund Outflows. *Journal of Financial Economics* 97, 239-262.

Chernenko, S., and A Sunderam, 2014, Frictions in Shadow Banking: Evidence from the Lending Behavior of Money Market Funds. *Review of Financial Studies*, 27 (6), 1717-1750.

Collins, S.S., and P.R. Mack, 1994, Avoiding Runs in Money Market Mutual Funds: Have Regulatory Reforms Reduced the Potential for a Crash. Working Paper 94-14, Federal Reserve Board Finance and Economics Discussion Series.

Covitz, D., N. Liang, and G. Suarez, 2013, The Evolution of a Financial Crisis: Panic in the Asset-Backed Commercial Paper Market. *Journal of Finance* 68 (3), 815-848.

Duygan-Bump, B., Parkinson, P., Rosengren, E., Suarez, G.A., and P. 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* 62 (2), 715-737.

Echenique, F., and I. Komunjer, 2009, Testing Models with Multiple Equilibria by Quantile Methods. *Econometrica* 77, 1281-1297.

Gennaioli, N., A. Shleifer, and R.W. Vishny, 2013, A Model of Shadow Banking. *Journal of Finance* 68, 1331-1363.

Goldstein, I., 2013, Empirical Literature on Financial Crises: Fundamentals vs. Panic, In *The Evidence and Impact of Financial Globalization*, edited by Gerard Caprio (Elsevier, Amsterdam), 523-534.

Goldstein, I., H. Jiang, and D.T. Ng, 2015, Investor Flows and Fragility in Corporate Bond Funds. Working paper, Wharton School of Business.

Goldstein, I., and A. Pauzner, 2005, Demand-deposit contracts and the probability of bank runs. *Journal of Finance* 60, 1293–1328.

Gordon, J. and C. Gandia, 2014, Money Market Funds Run Risk: Will Floating Net Asset Value Fix the Problem? *Columbia Business Law Review* 314, 751-781.

Gorton, G., and A. Metrick, 2012, Securitized banking and the run on repo. *Journal of Financial Economics* 104, 425-451.

Gu, C., 2011, Herding and bank runs. *Journal of Economic Theory* 146, 163-188.

He, Z., and A. Manela, forthcoming, Information Acquisition in Rumour-Based Bank Runs. *Journal of Finance*

He, Z., and W. Xiong, 2012a, Dynamic debt runs. *Review of Financial Studies* 25, 1799-1843.

He, Z., and W. Xiong, 2012b, Rollover Risk and Credit Risk. *Journal of Finance* 67, 391-429.

Hellwig, C., and L. Veldkamp, 2009, Knowing What Others Know: Coordination Motives in Information Acquisition. *Review of Economic Studies* 76, 223-251.

Hertzberg, A., Liberti, J.M., Paravisini, D., 2011, Public information and coordination: evidence from a credit registry expansion. *Journal of Finance* 66, 379-412.

Iyer, R., and M. Puri, 2012, Understanding bank runs: The importance of depositor-bank relationships and networks. *American Economic Review* 102, 1414-1445.

Kacperczyk, M., and P. Schnabl, 2013, How Safe are MMMFs? *Quarterly Journal of Economics*, 1073-1122.

Kelly, M., and C. Ó Gráda, 2000, Market Contagion: Evidence from the Panics of 1854 and 1857. *American Economic Review* 90(5), 1110-1124.

Martinez-Peria, M.S., Schmukler, S.L., 2001, Do depositors punish banks for bad behavior? Market discipline, deposit Insurance, and banking crises. *Journal of Finance* 56, 1029-1051.

McCabe, P., 2010, The Cross Section of Money Market Fund Risks and Financial Crises. Working Paper.

Morris, S. and H.S. Shin, 1998, Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks. *American Economic Review* 88, 587-597.

Morris, S. and H.S. Shin, 2001, Rethinking Multiple Equilibria in Macroeconomic Modeling. *NBER Macroeconomics Annual* volume 15, pages 139-182.

Ó Gráda, C., and E. White, 2003, The Panics of 1854 and 1857: A View from the Emigrant Industrial Savings Bank. *Journal of Economic History* 63(1), 213-240.

Parlatore, C., 2014, Fragility in MMMFs: Sponsor Support and regulation. Mimeo.

Qian, M. and B. Tanyeri, 2013, Litigations and Mutual Fund Runs. Unpublished working paper, National University of Singapore.

Schmidt, L. and Y. Zhu, 2015, Quantile Spacings: A Simple Method for the Joint Estimation of Multiple Quantiles. Mimeo, UChicago.

Schroth, E., G. Suarez, and L.A. Taylor, forthcoming, Dynamic debt runs and financial fragility: Evidence from the 2007 ABCP crisis. *Journal of Financial Economics*.

Schumacher, L., 2000, Bank runs and currency run in a system without a safety net: Argentina and the 'tequila' shock. *Journal of Monetary Economics* 46, 257-277.

Strahan, P. and B. Tanyeri, Once Burned, Twice Shy? MMMF Responses to a Systemic Liquidity Shock. *Journal of Financial and Quantitative Analysis*, forthcoming.

Appendix (For Online Publication Only)

Appendix A: Background Information

This appendix provides a timeline of events of the crisis of September 2008 and describes details of the iMoneynet database.

A.1. Key Money Market Events of September-October 2008

Numerous traumatic economic events had occurred since August 2007, putting considerable pressure on MMMFs. From August 2007 to August 2008, several unregulated liquidity pools used by institutional investors failed, both in the U.S. and elsewhere. This led to vast inflows to MMMFs, as these institutional investors turned to the tighter regulatory provisions required of MMMFs under Rule 2a-7, and, perhaps, to the implicit backup of sponsors for their MMMFs in a time of peril. It is very likely that this vast inflow of money believed there was little chance that a systematic risk event would significantly impact the mutual fund industry, setting the stage for the widespread impact of a common extreme risk event, as modeled by Gennaioli et al. (2013).

Then, the Federal Government declined to assist a reeling Lehman Brothers, which failed on September 15, 2008. On September 16, 2008, the Reserve Primary Fund (which held about $750 million in commercial paper issued by Lehman Brothers) "broke the buck." Immediately, prime MMMFs began to see vast outflows, and they struggled to sell securities to meet these redemptions. On Friday, September 19, 2008, the U.S. Treasury offered a guarantee to MMMFs in exchange for an "insurance premium" payment. On that same day, the Federal Reserve announced The Asset-Backed Commercial Paper Money Market Mutual Fund Liquidity Facility to provide funding to U.S. depository institutions and bank holding companies to finance purchases of high-quality asset-backed commercial paper (ABCP) from money market mutual funds under certain conditions. This program was set up to assist MMMFs holding such paper to meet redemption demands and to promote liquidity in the ABCP market and money markets, more generally. This program began operations on September 22, 2008, and was closed on February 1, 2010.

In addition, in response to the growing difficulty of corporations in rolling over their short-term commercial paper, the Fed announced The Commercial Paper Funding Facility on October 7, 2008, followed by additional details on October 14, 2008. This program took effect on October 27, 2008, and was designed to provide credit to a special purpose vehicle that would purchase three-month commercial paper from U.S. issuers.

On October 21, 2008, the Federal Reserve announced yet another program, The Money Market Investor Funding Facility (MMIFF). The MMIFF was a credit facility provided by the Federal Reserve to a series of special purpose vehicles established by the private sector. Each SPV was able to purchase eligible money market instruments from eligible investors using financing from the MMIFF and from the issuance of ABCP. Eligible Assets included certificates of deposit, bank notes and commercial paper with a remaining maturity of at least seven days and no more than 90 days.

Further, the SEC took a number of actions, perhaps the most important being to allow MMMFs to price their underlying securities at amortized cost at a point during the crisis when quotes on commercial paper were generally regarded as unreliable. Following these developments, investors continued to redeem shares from prime MMMFs, but at a diminishing rate. By the end of October 2008, the MMMF crisis was essentially over.

A.2. Further Details on iMoneynet Database

Our daily MMMF data from iMoneyNet cover the period February 2008 to June 30, 2009, and include data on funds that no longer exist. We estimate that these data cover about 93.5% of prime MMMFs in existence at the end of 2008. We approximate daily fund share class flows as the daily fraction change in share class total net assets.[48] From the perspective of a subscriber to iMoneyNet–mainly those who invest in MMMFs–a day's flow data for a share class are available well before 4 p.m. Eastern Time on the following day. Thus, an institution subscribing to iMoneyNet can easily view outflows from a fund that occurred during the prior day, before making its decision for the current day.

Since the Reserve Primary Fund is widely viewed by market participants as initiating the crisis through its "breaking-the-buck" announcement, we consider flows to this fund as developing exogenous to that of other funds. Other MMMFs within the Reserve complex likely held Lehman and experienced simultaneous outflows, as well. Thus, we exclude observations from all funds within the Reserve complex from our analysis.

One other detail of our data construction is worth noting. Our fund business variable, which is intended to proxy for sponsor reputational concerns emphasized in Kacperczyk and Schnabl (2013), is calculated using a slightly different data set. Kacperczyk and Schnabl calculate total complex mutual fund assets by combining data from a variety of sources, but the primary source is the CRSP mutual fund database. We use a measure of the same quantity which was kindly provided by the Investment Company Institute.

Panel A of Table A1 presents univariate summary statistics for prime institutional share classes as of September 12, 2008. Specifically, we show the mean, standard deviation, and a range of quantiles for the main covariates used in our models for prime institutional share classes (aggregated to the fund level) during the week of September 15-19, 2008.

Panel B reports cross-sectional correlations between covariates, measured at the share class level, as of that week. Most notable is the strong negative correlation between yield and liquid asset share (-0.66), which reflects that less liquid assets tend to earn higher yields. Also, we find that the fraction of large-scale investors is negatively correlated with shareclass expense ratio. Most of the remaining covariates are only moderately correlated, with all other pairwise correlations being lower than 36% in absolute value.

[48]Almost all money fund dividends are reinvested in the same money fund share class, so distributions (and their passive reinvestments) have a negligible effect on our estimates of flows.

Panel A: Univariate Summary Statistics

				Quantiles			
Variable	N	Mean	Std. Dev.	5%	25%	75%	95%
% Sophisticated (35 bp/yr cutoff)	258	69.6	35.6	0.00	44.0	97.3	100.0
Shareclass total net assets ($ bil)	258	4.82	8.83	0.12	0.46	4.69	24.50
Fund total net assets ($ bil)	258	18.97	23.05	0.54	3.14	26.57	62.26
Shareclass expense ratio (%/yr)	258	0.34	0.20	0.12	0.18	0.45	0.72
Average gross yield (%/yr)	258	2.94	0.15	2.67	2.88	3.04	3.14
Liquid asset share (%)	258	18.57	14.22	2.00	8.00	23.00	49.00
Fund business (%)	258	72.4	18.3	33.1	63.4	84.9	98.9
Std. dev. of daily log flows (%)	258	2.50	1.41	0.55	1.70	2.99	5.23
Cumulative flow 9/15-19 (%)	258	-12.8	23.7	-56.4	-21.5	0.1	8.9

Panel B: Pairwise Correlation Matrix

	% Soph	Log TNA	ER	Yield	Liquid	Fund Bus.	Log SD
Log total fund assets	0.27*						
Shareclass expense ratio (%/yr)	-0.47*	-0.10					
Average gross yield (%/yr)	-0.02	0.24*	0.08				
Liquid asset share (%)	0.05	-0.22*	-0.05	-0.66*			
Fund business (%)	-0.14*	-0.28*	0.03	-0.01	0.09		
Log flow standard deviation	0.36*	-0.09	-0.26*	-0.08	0.11	-0.06	
Cumulative flow 9/15-19 (%)	-0.27*	-0.31*	0.32*	-0.18*	0.20*	0.11	-0.10

TABLE A1—SUMMARY STATISTICS FOR PRIME INSTITUTIONAL FUNDS AS OF 9/15/2008

Note: This table presents summary statistics for the cross section of money market funds in the Prime Institutional category within the iMoneyNet database as of September 12, 2008, the Friday before the failure of Lehman Brothers. Variables are as follows: % Sophisticated, our measure of the % large-scale institutions, defined as the fraction of total fund investment dollars owned by investors of institutional shareclasses with $ER \leq 35$ bp/yr; shareclass and fund total net assets, in billions of dollars; shareclass expense ratio, in percentage points; Average gross yield is the average daily value of the (annualized) 7-day gross yield, in percentage points, from March through August, 2008; liquid asset share, a "real-time" estimate of liquid assets available as a fraction of total net assets, which is calculated by comparing an estimate of maturing assets with net redemptions; fund business, defined as one minus the proportion (by value) of fund complex aggregate mutual fund assets that are represented by prime institutional share classes; the standard deviation of daily changes in log institutional total assets over the period March-August 2008, which is winsorized by 2% in either tail; and Lehman week cumulative flows, defined as the percentage change in share class-level assets under management (i.e., flows as a fraction of lagged assets under management), in percentage points, for each prime institutional money market share class from September 15 through September 19th on expense ratio (ER). Panel A presents distributional statistics, while Panel B presents pairwise correlations between variables from Panel A or their log transformations, where stars indicate significance at the 5% level.

Of particular note is the fact that our measure of the percentage of sophisticated investors has only a -2% correlation with average gross yield and a 5% correlation with liquid asset share, our two controls for portfolio risk, neither of which is statistically significant.

Consistent with other regression results in the paper, the last row of Panel B shows pairwise correlations between cumulative crisis-week flows and each of these characteristics. These correlations, when significant, generally have the signs we would expect: funds with a higher concentration of sophisticated investors, more total assets under management, lower expense ratios, higher yields, lower liquidity, and lower fund business were all associated with significantly higher outflows during the crisis week.[49]

[49]Note, however, the standard errors in Table A1 are not corrected for heteroskedasticity or within-fund covariances.

APPENDIX B: SUPPLEMENTARY ANALYSES

This appendix contains a number of supplementary analyses to complement the main results from the text. Section B.1 presents the proof of Propositon 1. Section B.2 discusses how the key comparative statics from Proposition 1 can emerge in alternative modeling enviornments. Section B.3 presents several regressions of fund-level outflows on a number of fund characteristics, placing an emphasis on the role of average yield, portfolio liquidity, and fund business in predicting investor behavoir. Finally, Sections B.4 and B.5 discuss the robustness of the results from Tables 3 and 6 in the main text to several alternative modeling choices.

B.1. Proof of Proposition 1

The proof of our Proposition 1 is quite similar to Proposition 1 in Angeletos and Werning (2006), itself closely related to Morris and Shin (2001). As such, we emphasize the key aspects of the proof and the minor differences in equlibrium conditions resulting from uninformed agents and refer the reader to these papers for further technical details.

As discussed in the main text, an equilibrium is defined by two objects–a threshold on the fundamental θ^* and a threshold private signal $x^*(z)$ for the marginal agent. Recall that we combined the two optimality conditions into a single fixed point condition (equation 3) in the main text, which depends on θ^* (the corresponding threshold signal $x^*(z)$ is pinned down by equation (2) in the main text):

$$\text{(B1)}\quad \Phi^{-1}\left[\frac{1}{\mu}\theta^*\right] - \frac{\alpha_z+\alpha_0}{\sqrt{\alpha_x}}\theta^* \equiv \Gamma(\theta,\mu) = \sqrt{1+\frac{\alpha_z}{\alpha_x}+\frac{\alpha_0}{\alpha_x}}\cdot\Phi^{-1}[1-c] - \frac{\alpha_z}{\sqrt{\alpha_x}}z - \frac{\alpha_0}{\sqrt{\alpha_x}}\theta_0.$$

Note that the right hand side of (B1) does not depend on θ^*. Since Γ varies from $-\infty$ to ∞, at least one solution will always exist, and the equilibrium will be unique for all public signals z whenever $\Gamma(\theta,\mu)$ is strictly increasing in θ. $\frac{\partial\Gamma}{\partial\theta} = \frac{1}{\mu\phi(\mu^{-1}\theta^*)} - \frac{\alpha_z+\alpha_0}{\sqrt{\alpha_x}}$, which, from the curvature of the normal density, is bounded below by $\frac{\sqrt{2\pi}}{\mu} - \frac{\alpha_z+\alpha_0}{\sqrt{\alpha_x}}$ and thus is strictly positive when $\sqrt{2\pi\alpha_x} > (\alpha_z+\alpha_0)\mu$ (the condition from the proposition).

Next, we derive our key comparative static for the equilibrium objects θ^* and x^*. Since (B1) holds with equality, the implicit function theorem implies that the partial derivative of a particular element of the equilibrium correspondence with respect to μ is

$$\left.\frac{\partial\theta^*}{\partial\mu}\right|_{\theta^*,\mu} = \frac{\mu^{-2}[\phi(\mu^{-1}\theta^*)]^{-1}}{[\mu\phi(\mu^{-1}\theta^*)]^{-1} - \frac{\alpha_z+\alpha_0}{\sqrt{\alpha_x}}},$$

which is positive whenever the denominator ($\partial\Gamma/\partial\theta$) is positive. As discussed above, $\partial\Gamma/\partial\theta > 0$ when the equilibrium is unique. Also, given that $\lim_{\theta\downarrow-\infty}\Gamma(\theta,\mu) = -\infty$ and $\lim_{\theta\uparrow\infty}\Gamma(\theta,\mu) = \infty$, the denominator will be positive in the neighborhood of the lowest and highest possible equilibrium θ^*, which we will denote by $\underline{\theta}^*(z,\mu)$ and $\overline{\theta}^*(z,\mu)$, respectively, so $\frac{\partial\underline{\theta}^*}{\partial\mu} > 0$ and $\frac{\partial\overline{\theta}^*}{\partial\mu} > 0$. The fact that $\partial\underline{x}^*/\partial\mu > 0$ and $\partial\overline{x}^*/\partial\mu > 0$ follows immediately from applying the implicit function theorem to the indifference condition (4).

Finally, we discuss the optimality of the behavior of uninformed agents. For simplicity, we have simply assumed that these agents choose to play $a_i = 0$. However, it will be optimal for them to do so provided that the prior probability that the regime fails is less than c. This probablilty is well-defined without additional assumptions when the equilibrium is unique; otherwise, we can bound it above by

$$\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \mathbf{1}[\theta \leq \overline{\theta}^*(z,\mu)]\pi(\theta,z)d\theta dz,$$

where $\pi(\theta, z)$ is the prior (bivariate normal) density of the fundamental θ and public signal z. For sufficiently high θ_0, this condition should be satisfied. However, there could exist regions of the parameter space for which their assumed behavior is not consistent with utility maximization. When thinking about our particular context, our assumption that these agents play $a_i = 0$ appears reasonable given that essentially no major MMMF had broken the buck prior to the Lehman episode.

B.2. Robustness of main theoretical predictions to alternative modeling assumptions

We next briefly discuss how similar comparative statics to our Proposition 1 emerge in versions of the bank run games of Goldstein and Pauzner (2005) and He and Manela (forthcoming) when they are augmented with a measure of uninformed agents. Our discussion of these papers is intended only to highlight the basic structure of each model and, in particular, how to derive our key comparative statics. We refer the reader to the papers discussed below for additional details and discussion.

We begin with the Goldstein and Pauzner (2005) model, which adds noisy signals to an environment similar to the Diamond and Dybvig (1988) bank run game, yielding sharp comparative statics. Morris and Shin (2001) present a stylized version of their model which has a very similar structure to ours. Agents must decide whether to withdraw early or maintain their investments based on a noisy signal about the return on a risky asset in a setting with payoff complementarities.[50] The complementarities emerge from the fact that the expected log return on the risky investment is assumed to be a decreasing function of the fraction of agents who withdraw early.

As in our model from Section II, the equilibrium in the bank run game involves a threshold on the private signal received by a marginal agent who is indifferent between attacking and maintaining the status quo. This threshold, which is defined in terms of the posterior mean of an informed agent's beliefs about fundamentals in Morris and Shin's analysis, satisfies a fixed point condition. When private information is sufficiently informative relative to other sources of common (prior/public) information, this threshold is unique. When we extend the model to assume that a fraction of agents receive uninformative signals, it is easy to show that increasing the fraction of uninformed agents moves

[50]Morris and Shin (2001) do not consider a public signal although their model extends naturally to a case with noisy public signals. Note, however, that the sufficient condition on the relative signal precision guaranteeing uniqueness will be different.

the threshold upwards, which makes runs less likely to occur. An analogous argument to the one from Section II.B implies that the same three testable predictions derived from the regime change game of Section II also emerge in this context.

More generally, many static global games models of bank runs/regime change involve fixed point equations similar to equation (5) in the main text. As is discussed in Goldstein (2013), these models often partition the support of the fundamental θ into three regions: (i) a region in which the regime fails regardless of any agents' behavior (in our model: $\theta \in (-\infty, 0]$), (ii) a region in which the regime fails but could have survived were agents able to coordinate on the status quo (in our model: $\theta \in (0, \theta^*]$), (iii) a region in which the regime survives (in our model: $\theta \in [\theta^*, \infty)$). By definition, adding a measure of uninformed agents will not affect the boundaries of the first region. However, the basic mechanism of our model is preserved: adding a measure of uninformed agents who play the status quo weakens complementarities, making it easier for agents to coordinate and shifting the boundary (θ^*) separating region (ii) from region (iii) downwards. This is likely to have a similar effect on the fixed point (and, in turn, the optimal strategies of informed agents) in other contexts when analogously extended.

Next, we turn to the dynamic model of He and Manela (forthcoming). He and Manela incorporate uncertainty about the liquidity of a bank's investments and endogenous information acquisition into the asynchronous awareness framework of Abreu and Brunnermeier (2003). At a (stochastic) point in time, a rumor that a bank may be illiquid begins spreading among depositors with a Poisson arrival rate (β, the "rumor-spreading rate"). Agents, upon hearing the rumor, are unable to observe whether or not the bank is truly illiquid nor do they observe the point at which the rumor began to circulate. If the bank is illiquid, it fails if a sufficiently high fraction of investors withdraw their deposits within an exogenously specified period of time (the "awareness window"), generating complementarities. Agents, who choose whether to deposit with a bank at each point in time, additionally have the opportunity to acquire costly private signals about the bank's liquidity after hearing the rumor.

He and Manela's framework includes informed (heard the rumor) and uninformed (did not hear the rumor) agents, and the uninformed agents maintain their existing deposits with the bank.[51] The key parameter of interest is the rumor spreading rate β. As β increases, a higher fraction of agents are aware that the bank may be illiquid at a given point in time, increasing the potential number of agents who may attack the bank. He and Manela show that increases in β decrease the time it takes for an illiquid bank to fail. In particular, informed agents wait less time on average before withdrawing their deposits, implying that expected outflows are higher at each point in time.

The key mechanism which generates our main testable predictions is actually quite similar to this comparative static for β. Informed agents need to forecast the behavior

[51] Analogous to our discussion at the end of our proof of Proposition 1 above, a simple technical condition ensures that it is optimal for agents who do not hear the rumor to behave in this manner. Moreover, among those who have heard the rumor, there is a subset of individuals who receive private signals which perfectly reveal the bank's liquidity state. This additional dimension in which agents are heterogenously informed is of secondary importance for our discussion here.

or other agents when deciding how to respond to potentially bad news about liquidity. Complementarities are stronger when more agents also received the bad news, increasing fragility when the bank is illiquid.[52] Therefore, the rumor spreading rate in He and Manela plays a similar role to the fraction of informed agents μ from our static model in Section II. Moreover, a similar argument to that of Section II.B yields the same predictions for cross-sectional flows from funds with multiple shareclasses. These predictions follow if we simply reinterpret β as the average rate at which the rumor spreads across all depositors, then assume that the rumor spreads more quickly among sophisticated relative to unsophisticated investors. The key element of the argument is that, since all agents behave in the same way once they have heard the rumor, an informed agent's optimal strategy only depends on the overall rate at which the rumor spreads in the population.[53]

B.3. Regressions of Fund-level flows on fundamental and investor characteristics

Table B1 regresses cumulative Lehman week flows on the variables featured in the cross-sectional regressions of Table 3. Of particular interest is the explanatory power of "fundamental" measures such as average gross yield, portfolio liquid asset share, and the fund business variable (a proxy for a given MMMF sponsor's reputational considerations, which is likely to indicate a higher willingness to support a troubled fund) in the absence of the other controls. Because all of these variables are constant for all shareclasses within the same fund, we pool all institutional shareclasses to the fund-level before calculating flows.

Panel A presents the estimated coefficients from regressions on various combinations of these three variables. In univariate specifications, each of the variables is significant at either the 5% or 10% level, and the signs are as expected. Higher yield, lower liquidity, and lower sponsor reputational concerns all increase predicted fund-level outflows. In multivariate regressions, the sponsor reputation variable remains significant at the 10% level. Yield and liquidity remain highly significant when separately included in a regression along with fund business; however, these two variables are collinear and are both insignificant in the last specification in column 6. It is also worth noting that the explanatory power of these regressions is somewhat low–the maximum R^2 is 6.8% in the cross-section.

Panel B adds four additional variables–average institutional expense ratio, % large scale investors, log fund size, and the log flow standard deviation–to the specifications from Panel A. The latter three of these additional variables are statistically significant, and their associated coefficients are quantitatively large. One can also observe a dramatic

[52]He and Manela also show that increasing β increases the private incentives for agents to acquire information, a mechanism which further strengthens complementarities.

[53]To apply the comparative statics from He and Manela (forthcoming) in support of our argument, we abstract away from differences in information acquisition costs across different types of investors after they hear the rumor. Allowing these costs to vary across different types of informed agents, while plausible, would complicate the analysis considerably, and we conjecture that it will have little impact on the qualitative predictions emphasized here. We leave the study of such an extension for further research.

Panel A: Specifications without Investor Characteristics

Variable	(1)	(2)	(3)	(4)	(5)	(6)
Average gross yield	-5.04**			-4.97**		-3.08
	(-2.180)			(-2.172)		(-0.875)
Liquid asset share		4.86**			4.83**	2.78
		(2.590)			(2.477)	(0.811)
Fund business			3.63*	3.52*	3.59	3.54
			(1.722)	(1.677)	(1.652)	(1.649)
N	123	123	123	123	123	123
R^2	0.041	0.039	0.022	0.062	0.060	0.069

Panel B: Specifications with Investor Characteristics

Variable	(1)	(2)	(3)	(4)	(5)	(6)
Average gross yield	-3.80*			-3.78*		-2.23
	(-1.852)			(-1.828)		(-0.716)
Liquid asset share		3.79*			3.78*	2.34
		(1.880)			(1.872)	(0.717)
Fund business			-0.66	-0.53	-0.53	-0.50
			(-0.388)	(-0.302)	(-0.301)	(-0.282)
% Soph (35 bp/yr)	-14.25***	-15.44***	-15.55***	-14.53***	-15.72***	-15.05***
	(-2.920)	(-3.205)	(-2.981)	(-2.737)	(-3.000)	(-2.766)
Log total fund assets	-10.53***	-10.57***	-11.60***	-10.68***	-10.72***	-10.51***
	(-5.757)	(-6.269)	(-6.460)	(-5.679)	(-6.294)	(-5.948)
Avg institutional expense ratio	-1.19	-1.83	-2.29	-1.31	-1.95	-1.50
	(-0.623)	(-1.053)	(-1.189)	(-0.642)	(-1.051)	(-0.738)
std_logflowstdev_wins	-5.10***	-5.24***	-5.01**	-5.17***	-5.30***	-5.29***
	(-2.689)	(-2.839)	(-2.560)	(-2.662)	(-2.799)	(-2.740)
N	123	123	123	123	123	123
R^2	0.366	0.366	0.345	0.366	0.367	0.371

TABLE B1—DETERMINANTS OF FUND-LEVEL INSTITUTIONAL FLOWS DURING LEHMAN WEEK

Note: This table presents estimated coefficients from OLS regressions of the change in the log of fund-level assets under management (i.e., flows as a fraction of lagged assets under management) for prime institutional money market funds ($\times 100$) from September 15 through September 19th on average annualized 7-day SEC gross yield over the six month period prior to the crisis (March-August 2008); liquid asset share, estimated daily by computing dollar proportion Treasury and U.S. agency securities plus repo investments, plus (estimated) maturing securities, minus net redemptions; and fund business, defined as one minus proportion (by value) of fund complex aggregate mutual fund assets that are represented by prime institutional share classes. Panel B adds expense ratio, % large-scale institutions (% Soph; defined as the fraction of total fund investment dollars owned by investors of institutional shareclasses with expense ratios (weakly) below a cutoff–25 bp for the first column.), log of fund size (total assets under management as of September 12, 2008), and the standard deviation of daily log flows, computed from March-August, 2008. All coefficients except for % Soph have been divided by their cross-sectional standard deviations.

increase in R^2 after these variables are included. Yield and liquid asset share remain marginally significant in columns 1-2 and 4-5, but the investor characteristics subsume the explanatory power of the fund business variable.

It is worth noting that these results are not at all incompatible with the basic story emphasized in Kacperczyk and Schnabl (2013). If MMMFs may have been treated almost as perfect substitutes by large-scale institutional investors, small differences in portfolio risk may have been sufficient to generate fairly substantial differences in clientele. Kacperczyk and Schnabl present convincing evidence that fund management companies with substantial reputational concerns took less portfolio risk in the period prior to the Lehman episode, which is also fully consistent with our finding from Table A1 of a significantly negative correlation between our fund business variable and the fraction of large-scale investors. We find that, during the peak of the crisis, these clientele measures appear to possess greater explanatory power than portfolio risk measures, suggesting that cross-sectional heterogeneity in the strength of complementarities across funds may have been somewhat larger than the heterogeneity in portfolio risk across funds.

B.4. Robustness exercises for specifications in Table 3

For the sake of brevity, our regressions in Table 3 only tested prediction 3 of our theoretical model (involving the interaction terms) using a 35 bp/yr cutoff. Table B2 verifies that similar results hold for other choices of the cutoff, both with and without fund fixed-effects.

We have also added other potential control variables such as the standard deviation of daily log fund flows (columns 2 and 7 in Table B3), calculated using data from March-August 2008, and squared versions of the controls and found the results to be robust to these added controls (columns 4 and 8 in Table B3).

Our final covariate is the proportion of total MMMF assets for a complex (e.g., Fidelity), represented by prime institutional share classes (PIPERC, reported in columns 3 and 7 in Table B3) which can be interpreted as a proxy that (negatively) represents the ability of the fund management company to subsidize its funds during the crisis week.[54] PIPERC is not a perfect proxy for implied complex subsidization, as it may also proxy for the percentage of sophisticated investors at the complex level and, thus, for the strength of strategic complementarities. This alternative interpretation might influence our results if sophisticated investors monitor not only their own funds, but also other same-complex MMMFs. Again we find that the results in Table 3 are robust to the inclusion of this control.

In unreported results, we also verify that our results are not sensitive to our choice to calculate logarithmic, rather than arithmetic, cumulative flows. Extremely similar results, which are available upon request, obtain if we instead use arithmetic flows.

[54]We also note that PIPERC is (negatively) related to our variable, Fund Business, which is patterned after a similar variable introduced by Kacperczyk and Schnabl (2013). In their paper, funds whose sponsors have a lower level of other mutual fund business take on more risk before the crisis, and, therefore, suffer larger outflows because of this risk-taking. Although our Average Gross Yield variable is designed to proxy for risk-taking, it is entirely possible that it does not do so without measurement error, and, thus, that Fund Business may be capturing some of the effect outlined by Kacperczyk and Schnabl (2013).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
%Soph (25 bp/yr)	-42.6***							
	(-3.20)							
%Soph (25 bp/yr)× ER	13.6***				27.3**			
	(3.20)				(2.54)			
%Soph (30 bp/yr)		-45.5***						
		(-3.46)						
%Soph (30 bp/yr)× ER		15.0***				27.4**		
		(3.51)				(2.49)		
%Soph (50 bp/yr)			-55.3***					
			(-3.97)					
%Soph (50 bp/yr)× ER			18.5***				32.3***	
			(4.24)				(3.16)	
%Soph (150 bp/yr)				-58.6***				
				(-3.67)				
%Soph (150 bp/yr)× ER				17.7***				33.4**
				(3.47)				(2.18)
Shareclass expense ratio	0.2	-0.9	-4.6*	-5.9	-5.8	-6.3	-15.1*	-19.1
	(0.11)	(-0.40)	(-1.79)	(-1.65)	(-0.87)	(-0.89)	(-1.96)	(-1.51)
Log total fund assets	-6.0***	-6.6***	-7.8***	-7.9***				
	(-2.79)	(-3.15)	(-3.59)	(-3.44)				
Average gross yield	-2.9	-3.2	-3.0	-3.5				
	(-0.54)	(-0.60)	(-0.57)	(-0.68)				
Liquid asset share	1.4	1.0	1.2	1.3				
	(0.32)	(0.23)	(0.27)	(0.28)				
Fund business	-1.3	-1.2	-1.8	-1.9				
	(-0.52)	(-0.50)	(-0.69)	(-0.71)				
Sample	All Inst	All Inst	All Inst	All Inst	All Inst	All Inst	All Inst	All Inst
Fixed effects	None	None	None	None	Fund	Fund	Fund	Fund
N	258	258	258	258	258	258	258	258
Degrees of freedom	250	250	250	250	134	134	134	134
R2	0.127	0.127	0.124	0.124	0.579	0.579	0.576	0.570

TABLE B2—CROSS-SECTIONAL REGRESSIONS OF SHARECLASS-LEVEL FLOWS ON FUND- AND INVESTOR CHARACTERISTICS - ROBUSTNESS OF SPECIFICATIONS (7) AND (8) OF TABLE 3 USING DIFFERENT CUTOFFS

Note: This table presents estimated coefficients from OLS regressions of the change in the log of shareclass-level assets under management (i.e., flows as a fraction of lagged assets under management) for prime institutional money market funds ($\times 100$) from September 15th through September 19th on expense ratio, % large-scale institutions (%Soph; defined as the fraction of total fund investment dollars owned by investors of institutional shareclasses with expense ratios (weakly) below a cutoff–25 bp for the first column.), and an interaction term between shareclass expense ratio and $\%Soph$ at the fund-level, as well as the log of fund size (total assets under management as of September 12, 2008) and its average annualized gross yield in the six month period prior to the crisis (March-August 2008), liquid asset share, estimated daily by computing dollar proportion Treasury and U.S. agency securities plus repo investments, plus (estimated) maturing securities, minus net redemptions; and fund business, defined as one minus proportion (by value) of fund complex aggregate mutual fund assets that are represented by prime institutional share classes. All control variables except % Soph are divided by their cross-sectional standard deviations. Different columns use different cutoffs for % Soph. Columns 5-8 include fund fixed effects.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
% Soph (35 bp/yr)	-34.0***	-29.2***	-26.5**	-30.4***	-50.6***	-51.8***	-42.3***	-45.7***
	(-3.04)	(-2.75)	(-2.50)	(-2.64)	(-3.79)	(-3.83)	(-3.80)	(-3.64)
% Soph (35 bp/yr) $\times$ ER					17.9***	21.7***	16.0***	16.8***
					(4.00)	(3.96)	(4.00)	(3.88)
Shareclass expense ratio	15.9**	15.4**	13.6*	12.3	-2.8	-5.8*	-1.3	-2.2
	(1.99)	(2.04)	(1.79)	(1.61)	(-1.24)	(-1.95)	(-0.63)	(-0.94)
Log total fund assets	-10.4***	-12.1***	-11.3***	-61.4**	-7.5***	-9.1***	-8.4***	-54.3**
	(-3.29)	(-3.73)	(-3.42)	(-2.24)	(-3.52)	(-3.95)	(-3.73)	(-2.38)
Average gross yield	-2.0	-1.0	-1.5	-75.1	-3.0	-2.6	-3.1	-29.3
	(-0.33)	(-0.16)	(-0.24)	(-1.16)	(-0.57)	(-0.50)	(-0.60)	(-0.46)
Liquid asset share	3.1	3.5	2.6	-8.7	1.3	1.8	0.6	-9.7
	(0.58)	(0.66)	(0.54)	(-1.23)	(0.28)	(0.40)	(0.15)	(-1.29)
Fund business	-1.2	-1.9	-3.8	-13.0	-1.3	-1.9	-3.5	-14.2*
	(-0.36)	(-0.58)	(-0.80)	(-1.19)	(-0.54)	(-0.76)	(-1.01)	(-1.70)
Log flow standard deviation		-9.0*				-6.3*		
		(-1.96)				(-1.97)		
Complex % prime institutional			-5.2				-4.5	
			(-1.11)				(-1.28)	
Average gross yield^2				1.8				0.7
				(1.18)				(0.41)
Log total fund assets^2				4.8**				4.3**
				(1.99)				(2.10)
Liquid asset share^2				1.9				2.1
				(1.56)				(1.45)
Fund business^2				1.7				1.9*
				(0.96)				(1.71)
Sample	Sophisticated institutional: ER $\leq$ 35 bp				All institutional shareclasses			
N	161	161	161	161	258	258	258	258
Degrees of freedom	154	153	153	150	250	249	249	246
R^2	0.094	0.118	0.102	0.120	0.128	0.143	0.134	0.150

TABLE B3—CROSS-SECTIONAL REGRESSIONS OF SHARECLASS-LEVEL FLOWS ON FUND- AND INVESTOR CHARACTERISTICS - ROBUSTNESS TO ADDITIONAL CONTROLS

Note: This table presents estimated coefficients from OLS regressions of the change in the log of shareclass-level assets under management (i.e., flows as a fraction of lagged assets under management) for prime institutional money market funds $(\times 100)$ from September 15th through September 19th on expense ratio, % large-scale institutions (%Soph; defined as the fraction of total fund investment dollars owned by investors of institutional shareclasses with expense ratios (weakly) below a cutoff–25 bp for the first column.), and an interaction term between shareclass expense ratio and $\%Soph$ at the fund-level, as well as the log of fund size (total assets under management as of September 12, 2008) and its average annualized gross yield in the six month period prior to the crisis (March-August 2008), liquid asset share, estimated daily by computing dollar proportion Treasury and U.S. agency securities plus repo investments, plus (estimated) maturing securities, minus net redemptions; and fund business, defined as one minus proportion (by value) of fund complex aggregate mutual fund assets that are represented by prime institutional share classes. Specification 2 adds a control for the standard deviation of daily fund-level flows, computed from March-August 2008, while Specification 3 adds a control for the percent of complex total assets under management in prime institutional shareclasses ("PIPERC", measured as of September 12, 2008). All control variables except % Soph are divided by their cross-sectional standard deviations. Specification 4 includes squared values of the original control variables. The right four columns columns add an interaction term between shareclass expense ratio and $\%Soph$ at the fund-level.

Variable	Early Crisis (1)	Peak Crisis (2)	Lehman Week (3)	Early Crisis (4)	Peak Crisis (5)	Lehman Week (6)
$Low_{i,t-1}$	0.20	0.13	0.21*	0.17	0.14	0.24**
	(1.31)	(0.82)	(1.88)	(1.19)	(0.86)	(2.07)
$Low_{i,t-1} \times \%High_{i,t-2}$	-0.23	0.31	0.04	-0.41	0.14	-0.18
	(-0.58)	(0.55)	(0.13)	(-1.19)	(0.26)	(-0.67)
$High_{i,t-1}$	-0.04	-0.31**	-0.17*	-0.03	-0.26*	-0.13
	(-0.38)	(-2.14)	(-1.79)	(-0.35)	(-1.76)	(-1.21)
$High_{i,t-1} \times \%High_{i,t-2}$	0.77**	1.13*	0.90*	0.43	0.98**	0.61
	(2.18)	(1.72)	(1.90)	(1.41)	(2.04)	(1.64)
$High_{i,t}$				-0.23**	-0.08	-0.18
				(-2.42)	(-0.68)	(-1.50)
$High_{i,t} \times \%High_{i,t-1}$				1.55***	1.42**	1.53***
				(4.58)	(2.49)	(3.18)
$\%Low_{i,t-1}$	1.35	-7.85**	-3.16	0.51	-11.57***	-6.00***
	(0.69)	(-2.29)	(-1.58)	(0.28)	(-3.17)	(-3.59)
Average yield$_{i,t-1}$	0.45	0.28	0.58	0.20	0.00	0.32
	(0.94)	(0.20)	(0.53)	(0.52)	(0.00)	(0.39)
Log total fund assets$_{i,t-1}$	-1.22***	-1.81*	-1.94***	-1.08***	-1.37	-1.52**
	(-2.92)	(-1.83)	(-2.73)	(-2.74)	(-1.31)	(-2.52)
Avg institutional expense ratio$_{i,t-1}$	0.54	0.43	0.91	0.44	0.23	0.91
	(0.74)	(0.26)	(0.76)	(0.63)	(0.15)	(0.94)
Liquid asset share$_{i,t-1}$	0.51	2.08*	1.53**	0.56	1.97*	1.56**
	(1.07)	(1.85)	(2.12)	(1.21)	(1.72)	(2.32)
Fund business$_{i,t-1}$	-0.52	-1.22	-1.14	-0.42	-0.89	-0.78
	(-1.28)	(-1.02)	(-1.18)	(-1.27)	(-0.83)	(-1.07)
Log flow std. dev.$_{i,t-1}$	-0.35	-2.22**	-1.81**	-0.15	-1.88**	-1.34**
	(-0.88)	(-2.03)	(-2.01)	(-0.40)	(-2.08)	(-2.12)
Complex PIPERC$_{i,t-1}$	-0.57	-2.49	-2.04	-0.31	-1.85	-1.42*
	(-1.11)	(-1.64)	(-1.64)	(-0.74)	(-1.59)	(-1.74)
N	320	190	318	320	190	318
R^2	0.171	0.316	0.292	0.234	0.394	0.363

TABLE B4—DETERMINANTS OF DAILY FLOWS FROM SOPHISTICATED (LOW ER) SHARECLASSES - ROBUSTNESS TO ADDITIONAL CONTROLS

Note: For each fund, we separate prime institutional share classes into two categories, based on their expense ratios. The first category, "Low," consists of share classes which have expense ratios that are lower than the median expense ratio (across all institutional share classes within a given fund). All remaining share classes are included in the "High" category, including all retail share classes. The value of shares outstanding is then aggregated across all Low share classes and, separately, across all High share classes within a given fund. (Funds with a single share class are excluded from this analysis.) For each fund and date, we calculate the first difference in the log of aggregate value within each category (i.e., fraction flow), which we denote by $Low_{i,t}$ and $High_{i,t}$. The table presents the coefficients from panel regressions with $Low_{i,t}$ as the dependent variable on $Low_{i,t-1}$ and $High_{i,t-1}$, estimated for three different subperiods in 2008: 9/10-9/16 "Early Crisis", 9/17-9/19 "Peak Crisis", and 9/15-9/19 "Lehman Week", respectively. We multiply $Low_{i,t}$ and $High_{i,t}$ by 100 to express them in log percentage points. We also include interaction variables between, for example, $High_{i,t-1}$ and $\%High_{i,t-2}$, which is defined as two-day lagged fraction of total MMMF value within a MMMF represented by "High" (both institutional and all retail) ER shareclasses. Relative to the baseline specification from Table 6, we add a control for the standard deviation of daily fund-level flows, computed from March-August 2008 and winsorized by 2% in either tail, and a control for the percent of complex total assets under management in prime institutional shareclasses ("PIPERC", measured as of September 12, 2008). Control variables, described in the notes to Table 3, have been divided by their (cross-sectional) standard deviations for ease of interpretation. All specifications also include unreported time dummies. Standard errors are clustered at the fund level. t-statistics are reported in parentheses.

B.5. Robustness exercises for specifications in Table 6

Table B4 undertakes a similar set of robustness checks for the time-series VAR specification in Table 6. Specifically, in Table B4 we include log flow standard deviation and complex PIPERC as additional control variables to the model. At 0.77, 1.13, and 0.90 for the early crisis, peak crisis and Lehman week, respectively, the estimated coefficients of the lagged interaction term ($High_{i,t-1} \times \%High_{i,t-2}$, reported in columns 1-3 in Table B4) are only marginally lower than the estimates in Table 6 (0.85, 1.22, and 1.02). Moreover, At 1.55, 1.42 and 1.53, the coefficients on the contemporaneous interaction term ($High_{it} \times \%High_{i,t-1}$, reported in columns 4-6 of Table B4) during the early crisis, peak crisis and Lehman week are very similar to those reported in Table 6 (1.58, 1.54, and 1.64, respectively).

Similarly, the results are robust to not applying a 2% winsorization to the lagged flows. Here, the estimates on the lagged interaction term ($High_{i,t-1} \times \%High_{i,t-2}$) during the three subperiods become (0.73, 1.13, and 0.83) (previously 0.85, 1.22, and 1.02), while the estimated coefficients on the contemporaneous interaction term ($High_{it} \times \%High_{i,t-1}$) change to 1.58, 1.34, and 1.47 (previously 1.58, 1.54, and 1.64) and remain highly statistically significant.

Appendix C: Quantile Regression Methodology and Coefficient Estimates

This Appendix introduces the panel quantile regression methodology used in Section VI, explains how we estimate the models, and presents some estimation results and robustness tests.

C.1. Methodology

We focus on modeling three quantiles, namely the 10th, 50th (median) and 90th, of the flow distribution, conditional on a vector of observable variables. Three quantiles is the minimum number sufficient to allow for *heterogeneity* and *asymmetry* in the flow distributions. In this way, we can determine whether fund and/or investor characteristics differentially affect funds in different parts (e.g., the center vs. tails) of the conditional cross-sectional distribution.

We adopt the following specification for conditional quantiles of a variable $Y_{i,t}$:

$$\begin{aligned} &Y_{i,t} = f_0(X_{i,t}, \beta) + \epsilon^0_{i,t} = X'_{i,t}\beta_0 + \epsilon^0_{i,t} && P[\epsilon^0_{i,t} < 0|X_{i,t}] = 0.5 \\ \text{(C1)}\quad &Y_{i,t} = f_1(X_{i,t}, \beta) + \epsilon^1_{i,t} = X'_{i,t}\beta_0 - \exp[X'_{i,t}\beta_1] + \epsilon^1_{i,t} && P[\epsilon^1_{i,t} < 0|X_{i,t}] = 0.1 \\ &Y_{i,t} = f_2(X_{i,t}, \beta) + \epsilon^2_{i,t} = X'_{i,t}\beta_0 + \exp[X'_{i,t}\beta_2] + \epsilon^2_{i,t} && P[\epsilon^2_{i,t} < 0|X_{i,t}] = 0.9. \end{aligned}$$

The functions $f_0(\cdot)$, $f_1(\cdot)$, and $f_2(\cdot)$ represent the median, 10^{th}, and 90^{th} percentiles of the distribution of $Y_{i,t}$ given $X_{i,t}$, respectively.To facilitate interpretation of the results, we anchor the model around the conditional median, governed by β_0, of the flow distribution. We then add (or subtract) spreads, governed by β_1 and β_2, that quantify the difference between the effect of covariates on funds in the left or right tails of the cross-sectional distribution of fund flows. We first look at the effect of covariates on the median, then separately consider any additional effects on these spreads of an exponential affine functional form. This guarantees that the conditional quantiles never cross and yields an internally consistent dynamic model.

As our specification is relatively new, some discussion of how to interpret parameters is in order. β_0 governs the effect of $X_{i,t}$ on the median level of flows and affects the other conditional quantiles as well. Since β_0 shifts the entire distribution of flows, its interpretation is quite similar to an OLS regression coefficient. We refer to these terms as "median exposures" or "common exposures", though we emphasize that these coefficients affect all quantiles symmetrically. β_1 captures the additional effect of covariates on the left tail of the flow distribution–the spread between the median and the 10^{th} percentile. For ease of exposition, we refer to this distance as a fund's "left tail exposure."[55]

Our model for the conditional quantiles has an additional interpretation which is useful in a panel context. Partitioning the vector $X_{i,t} = [W'_{i,t}, Z'_t]'$, where $W_{i,t}$ is a vector of

[55] If $\beta_1^{(j)}$ and $X_{i,t}^{(j)}$ are the j^{th} elements of β_1 and $X_{i,t}$, respectively, then a one unit increase in $X_{i,t}^{(j)}$ generates a $\beta_1^{(j)}$ percent increase in the left tail exposure for a given fund. β_2 governs a fund's right tail exposure, defined analogously. From a fund's perspective, increases in left tail exposure are "bad" (indicating higher downside risk) while increases in right tail exposure are "good".

fund-specific characteristics and Z_t is a vector of time-specific factors, our model is

$$\begin{aligned}\text{(C2)}\quad f_0(X_{i,t},\beta) &= W'_{i,t}\lambda_0 + Z'_t\gamma_0 \equiv W'_{i,t}\lambda_0 + \alpha_{0,t}\\ f_1(X_{i,t},\beta) &= W'_{i,t}\lambda_0 - \exp[W'_{i,t}\lambda_1]\exp[Z'_t\gamma_1] \equiv W'_{i,t}\lambda_0 - \exp[W'_{i,t}\lambda_1]\alpha_{1,t}\\ f_2(X_{i,t},\beta) &= W'_{i,t}\lambda_0 + \exp[W'_{i,t}\lambda_2]\exp[Z'_t\gamma_2] \equiv W'_{i,t}\lambda_0 + \exp[W'_{i,t}\lambda_2]\alpha_{2,t},\end{aligned}$$

where $\lambda = [\lambda'_0, \gamma'_0, \lambda'_1, \gamma'_1, \lambda'_2, \gamma'_2]'$. Here $(\alpha_{0,t}, \alpha_{1,t}, \alpha_{2,t})'$ is a vector of time-specific shocks. $\alpha_{0,t}$ is a shock that shifts the distribution for all funds symmetrically. $\alpha_{1,t}$ and $\alpha_{2,t}$ scale up or down the left and right tail exposures, respectively.[56] This specification makes sense in our application, given the important interactions between market-wide events (e.g., declines in liquidity) and investor behavior. We also allow the coefficients to change over different subperiods.[57]

Before going further, we introduce some terminology to ease the exposition in the discussion that follows. Our specification in (C2) enables us to compare different quantiles of the conditional flow distribution, holding conditioning variables, $W_{i,t}$, fixed. A "median fund" is not particularly lucky or unlucky when compared with funds with similar observable characteristics, experiencing flows that are relatively close to $f_0(X_{i,t}, \beta)$. In contrast, a "left tail fund" is relatively unlucky, experiencing flows relatively close to $f_1(X_{i,t}, \beta)$, while a "right tail fund" is relatively lucky. Relative to peers with similar observables, left tail funds are most likely to have experienced run-like behavior, so we are particularly interested in comparing left tail funds with different values of $W_{i,t}$.[58]

RECURSIVE ESTIMATION PROCEDURE

To see how we estimate the parameters of the model in (C1), it is helpful to rewrite the data generating process for $Y_{i,t}$ as

$$\text{(C3)}\qquad Y_{i,t} = X'_{i,t}\beta_0 - D_{i,t}\exp[X'_{i,t}\beta_1]\eta_{i,t} + (1 - D_{i,t})\exp[X'_{i,t}\beta_2]\eta_{i,t},$$

where $\eta_{i,t}$ is a nonnegative random variable with $P[\eta_{i,t} < 1|X_{i,t}] = 0.8$ and $D_{i,t}$ is a Bernoulli random variable which equals 1 with probability 0.5.[59] The left tail and right tail exposures, $\exp(X'_{i,t}\beta_1)$ and $\exp(X'_{i,t}\beta_2)$, are analogous to "semi-variances", where β_1 and β_2 separately govern the variance of bad and good shocks, respectively. If $\beta_1 = \beta_2$, this is consistent with a simple mean-variance model where the variance is a loglinear function of $X_{i,t}$. This alternative way of writing the DGP also mirrors the manner in which we estimate the relevant parameters.

Our analysis uses the method of Schmidt and Zhu (2015) which sequentially esti-

[56]This multiplicative structure gives λ_1 and λ_2 useful factor-loading interpretations. If $W'_{i,t}\lambda_1 = 0$, a fund's left tail exposure is equal to the aggregate shock; as $W'_{i,t}\lambda_1$ increases, the sensitivity to the aggregate shock increases.

[57]For example, perhaps investors put a heavy weight on the riskiness of a fund's holdings during the early stages of a crisis, but place less weight on this during later stages.

[58]In a number of cases, a variable has a strong effect on a fund's left tail exposure while having a minor effect on its median and right tail exposures; thus, changes in $W_{i,t}$ have little effect on flows from median or right tail funds, but they make a large difference for left tail funds.

[59]The conditional quantile restrictions hold since if $P(\eta_{i,t} < 1|X_{i,t}) = 0.8$, $P(Y_{i,t} < X'_{i,t}\beta_0 - \exp[X'_{i,t}\beta_1]|X_{i,t}) = P(D_{i,t} = 1|X_{i,t}) \times P(\eta_{i,t} > 1|X_{i,t}) = 0.5 \times (1 - 0.8) = 0.1$.

mates the parameters of interest using a series of standard linear quantile regressions. Specifically, we first estimate β_0 using standard linear quantile regression. Next, we estimate β_1 and β_2 by splitting the sample into two halves based on the signs of the residuals and performing an additional linear quantile regression on the log of these residuals. Using the positive residuals, we can estimate β_2. To see why this works, note that if $Y_{i,t} - X'_{i,t}\beta_0 > 0$, $Y_{i,t} - X_{i,t}\beta_0 = \exp[X'_{i,t}\beta_2]\eta_{i,t}$. Taking logs, we get that $\log[Y_{i,t} - X'_{i,t}\beta_0] = X'_{i,t}\beta_2 + \log\eta_{i,t}$. Given our assumption that $P[\eta_{i,t} < 1|X_{i,t}] = P[\log\eta_{i,t} < 0|X_{i,t}] = 0.8$, the transformed model satisfies the standard assumptions for linear quantile regression. To get feasible estimators, β_0 is replaced with $\hat{\beta}_0$, the initial estimate from the quantile regression for the median. An analogous procedure works for the absolute value of the negative residuals, enabling us to estimate β_1.

MULTI-PERIOD FLOW SIMULATIONS

We next explain how we simulate from the dynamic model to study the relationship between explanatory variables and cumulative flows during the crisis period. These simulations are used to generate Table 7 in the main text. We begin by fixing each of the explanatory variables at its average, while the initial value of lagged flows is assumed to be equal to the category average, i.e., $Y_{i,\tau-1} = \bar{Y}_{\tau-1}$, where τ is the first date in the simulation. Next, we take one of the elements of $X_{i,\tau}$ and add or subtract one standard deviation.

Our method for simulating a single daily flow mirrors the data generating process as described in Equation (C3). Given the model parameters, it is straightforward to draw $Y_{i,t}$ given $X_{i,t}$ by drawing a Bernoulli random variable, $D_{i,t}$, along with $\eta_{i,t}$, whose distribution remains to be specified. We assume that $\eta_{i,t}$ is distributed as an exponential random variable with rate parameter $-\log 0.2$, which ensures that $P(\eta_{i,t} < 1) = 0.8$. Figure B1 demonstrates that the distribution fits the data quite well; kernel density estimates of the fitted residuals, $\hat{\eta}_{i,t}$, are essentially indistinguishable from the parametric density, for both positive and negative residuals. We calculate cumulative flows by summing up the simulated $\{Y_{i,t}\}_{t=\tau}^{\tau+h}$.

We update several elements of $X_{i,t}$, given a simulated value of $Y_{i,t-1}$. The first is $Y_{i,t-1} - \bar{Y}_{t-1}$, which we calculate by subtracting off the actual cross-sectional mean from the data. Second, we update $LOGTNA$ by adding $Y_{i,t-1}$. Third, we update $LIQUIDRT$ by assuming that any redemptions in excess of maturing assets (estimated using the average cross-sectional weighted average maturity) are met by selling liquid assets. Then, given $X_{i,t}$, we simulate $Y_{i,t}$. Iterating back and forth, we trace out the path of cumulative flows.

For each set of initial $X_{i,\tau}$, we simulate 50,000 total sample paths for cumulative flows. We calculate the 1st, 5th, 10th, 50th, and 90th quantiles of the set of simulated paths, respectively. In addition, using the bootstrapped distribution (10,000 replications) of parameter estimates, we compute two statistical tests. The first tests whether the marginal effect of the variable of interest on cumulative flows is significantly different from zero. The second tests whether the difference between the marginal effect at the

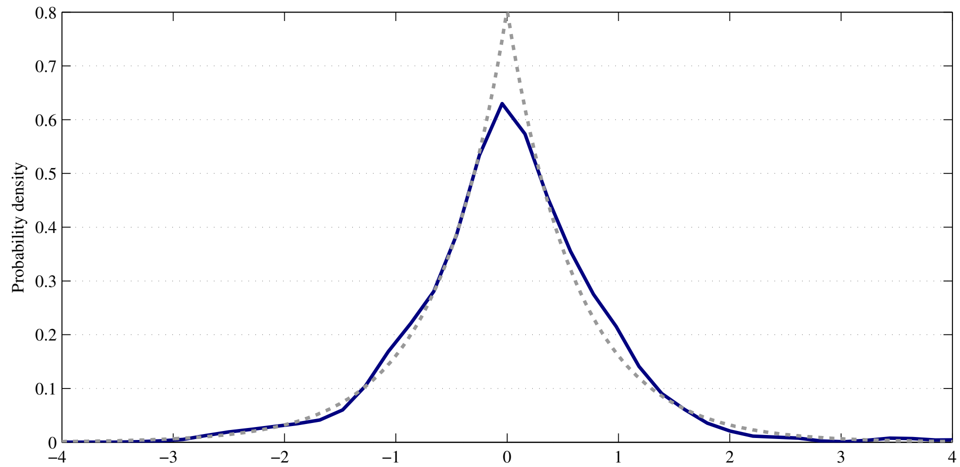

FIGURE C1. STANDARDIZED RESIDUAL DENSITY FROM BASELINE MODEL - INSTITUTIONAL FUNDS

Note: This figure plots the empirical density of the standardized residuals from the estimated model in Table C1, $\hat{\eta}_{i,t}$, which is generated using a kernel smoother. The dashed line plots the density of a Laplace-distributed random variable which has been normalized to satisfy the conditional quantile restriction which is assumed when estimating the model.

median and a different quantile differs from zero.

C.2. Coefficient estimates

Table C1 presents our quantile regression estimates for the panel of Prime Institutional share classes. The dependent variable is the daily change in log (fund-level) aggregated share class total net assets. In our discussion to follow, for simplicity, we often refer to the aggregate of prime institutional share classes as a prime institutional "fund," but the reader should be reminded that, strictly speaking, a fund can consist of both prime institutional and prime retail share classes.

As was the case in Table 6, we split the sample into, and allow the model coefficients, β_0, β_1, and β_2 to change, over the early crisis and peak crisis subperiods, and each column of the table presents estimates for a specific subperiod. Panel A presents our estimates of β_0, the coefficients governing the conditional median. Panels B and C present our estimates of β_1 and β_2, the coefficients governing left and right tail exposures, respectively. We express the dependent variable in log percentage points, and divide all characteristics other than lagged flows by their cross-sectional standard deviations.

We briefly summarize the main results from Table C1. At the median, the coefficient on the fraction of sophisticated investors increases in absolute magnitude from -15 bps to -39 bps as we move from the early to the peak crisis. Both coefficients are statistically significant but the associated magnitudes are relatively small. Yields and log flow standard deviations are negatively correlated with median flows during the peak crisis but not in the early crisis, suggesting that riskier portfolios were associated with higher outflows during the peak crisis. Median outflows were also bigger for the largest funds both during

Variable	Panel A Common (Median) Exposure		Panel B Left Tail Exposure		Panel C Right Tail Exposure	
	Early Crisis	Peak Crisis	Early Crisis	Peak Crisis	Early Crisis	Peak Crisis
% Sophisticated$_{i,t-1}$	-0.0015 **	-0.0039 **	0.2889 ***	0.3703 **	0.1158	0.0170
	[0.022]	[0.024]	[0.008]	[0.037]	[0.332]	[0.319]
Average gross yield$_{i,t-1}$	-0.0007	-0.0054 ***	0.0962 *	0.1207	0.0142	0.1284
	[0.185]	[0.000]	[0.080]	[0.160]	[0.451]	[0.153]
Log flow std. dev.$_{i,t-1}$	-0.0007	-0.0046 **	0.5358 ***	0.4072 ***	0.4927 ***	0.5022 ***
	[0.208]	[0.034]	[0.000]	[0.006]	[0.000]	[0.002]
Log total fund assets$_{i,t-1}$	-0.0024 ***	-0.0092 ***	0.0441	0.1694	0.0052	0.2571 **
	[0.000]	[0.000]	[0.226]	[0.232]	[0.346]	[0.039]
$y_{i,t-1} - \bar{y}_{t-1} > 0$	0.1093 *	0.3573 **				
	[0.098]	[0.016]				
$y_{i,t-1} - \bar{y}_{t-1} < 0$	0.2643 ***	0.4513 ***				
	[0.003]	[0.001]				
$\lvert y_{i,t-1} - \bar{y}_{t-1} \rvert$			0.0600 *	-0.0136	0.0949 **	0.0729 *
			[0.087]	[0.501]	[0.042]	[0.057]
N	615	367	615	367	615	367
Pseudo-R^2 (50,10,90)	0.053	0.185	0.284	0.328	0.155	0.049

TABLE C1—FUND-LEVEL PANEL QUANTILE REGRESSION COEFFICIENTS - PRIME INSTITUTIONAL

Note: This table presents the coefficients from estimating equation (C1) via quantile regression using the recursive method in Schmidt and Zhu (2015). The dependent variable ($y_{i,t}$) is the daily log difference in fund-level assets under management for prime institutional funds, in percentage points (i.e., $\times$ 100). Panel A, on the left, reports β_0, which controls the conditional median and shifts all quantiles symmetrically. Panel B, in the middle, reports β_1, which governs the width of the left tail (the distance between the median and the 10th percentile). Panel C, on the right, reports β_2, which controls the width of the right tail (the distance between the 90th percentile and the median). All three sets of coefficients are allowed to vary over two different periods in 2008: 9/10-9/16 Early Crisis and 9/17-9/19 Peak Crisis, respectively. More detailed variable descriptions may be found in Table A1. In addition to the coefficients in the table, models include time dummies to capture the common shocks, $\alpha_{0,t}$, $\alpha_{1,t}$, and $\alpha_{2,t}$. Numbers in brackets are one-sided bootstrapped p-values clustered at the fund level. With the exception of lagged flows, all variables are divided by their (cross-sectional) standard deviations.

the early and late crisis. Finally, median outflows become more strongly serially correlated during the peak crisis compared to the early crisis period and the responsiveness to lagged outflows is somewhat higher (relative to inflows) in both periods.

The estimates for the left tail exposures show that the shift of the flow distribution to the left is accompanied by a widening in the left tail. Specifically, the coefficients on the fraction of sophisticated investors, at 37% and 29% in the early and peak crisis, respectively, are economically large and highly significant in both crisis periods. Thus, a one standard deviation increase in the fraction of sophisticated investors increases the distance between the median and 10th percentile of the flow standard deviation (i.e., increased outflows) by 37% during the peak crisis. Clearly nonlinearities in how flows are related to the fraction of sophisticated investors become very important during the crisis. We observe similarly large effects for the log flow standard deviation during the early and peak crisis periods. In contrast, the yield and fund size variables only have modest effects on the left tail of the flow distribution.

Finally, the estimates for the right tail exposures suggest that funds with higher flow standard deviations also had a greater likelihood of experiencing inflows on a given day during the crisis, suggesting that funds that had more volatile flows prior to the crisis also experienced more volatility (i.e., higher second moments) during the crisis. Also, large

Variable	Value	Cumulative Flow Quantile 1%	5%	10%	50%	90%
	$f(\bar{x})$	-50.96	-38.23	-32.30	-14.75	3.28
% Sophisticated	$f(\bar{x}+\sigma_x)$	-64.42	-48.60	-41.19	-18.93	2.18
	$f(\bar{x}-\sigma_x)$	-39.36	-29.87	-25.22	-11.13	5.06
	Difference	-25.06 ***	-18.73 ***	-15.97 ***	-7.81 ***	-2.88
	p-value	[0.002]	[0.001]	[0.001]	[0.001]	[0.397]
	p-value vs. median	[0.006]	[0.005]	[0.004]	-	[0.026]
Average gross yield	$f(\bar{x}+\sigma_x)$	-58.27	-44.21	-37.63	-17.84	0.96
	$f(\bar{x}-\sigma_x)$	-43.64	-32.94	-27.71	-11.86	5.56
	Difference	-14.62 **	-11.26 **	-9.92 **	-5.97 **	-4.60
	p-value	[0.046]	[0.033]	[0.026]	[0.011]	[0.266]
	p-value vs. median	[0.121]	[0.123]	[0.125]	-	[0.191]
Log flow std. dev.	$f(\bar{x}+\sigma_x)$	-63.20	-46.63	-39.02	-16.37	10.25
	$f(\bar{x}-\sigma_x)$	-42.22	-32.01	-27.14	-13.13	0.24
	Difference	-20.98 ***	-14.62 ***	-11.88 ***	-3.24 *	10.00 ***
	p-value	[0.002]	[0.002]	[0.002]	[0.085]	[0.002]
	p-value vs. median	[0.001]	[0.001]	[0.001]	-	[0.000]
Log fund total assets	$f(\bar{x}+\sigma_x)$	-55.17	-42.56	-36.18	-17.40	3.61
	$f(\bar{x}-\sigma_x)$	-45.93	-34.08	-28.31	-11.74	4.35
	Difference	-9.23 *	-8.47 **	-7.87 **	-5.65 ***	-0.75
	p-value	[0.052]	[0.032]	[0.024]	[0.007]	[0.547]
	p-value vs. median	[0.129]	[0.112]	[0.108]	-	[0.032]

TABLE C2—MARGINAL EFFECTS OF FUND CHARACTERISTICS ON CUMULATIVE FLOW QUANTILES - COEFFICIENTS ESTIMATED ON 20TH, 50TH, 80TH QUANTILES

Note: This table shows the impact of explanatory variables on cumulative flow distributions (as a percentage of initial assets) for prime institutional share classes (aggregated to the fund level) for the September 15-19 period. These estimates are obtained by simulating from an estimated dynamic quantile panel regression model for daily flows. Whereas the dynamic panel regression model presented in Table 7 is estimated for the 10th, 50th, and 90th quantiles, this table repeats the exercise where the model is instead estimated for the 20th, 50th, and 80th quantiles. Columns report the 1st, 5th, 10th, 50th, and 90th quantiles of the cumulative flow distributions, respectively. We begin by fixing each of the explanatory variables at its average, assuming that the initial value of lagged flows equals the prime institutional category average. Then, we report the impact on the simulated flow distribution of adding and subtracting one standard deviation to each explanatory variable, p-values for a test of whether the difference in the simulated quantiles is statistically significant, obtained by using the bootstrapped distribution of parameter estimates from our model, as well as the p-value of whether the marginal effect is significantly different at a given quantile, relative to the marginal effect at the median (using the bootstrapped distribution).

funds were particularly likely to experience inflows during the peak crisis.

C.3. Alternative quantile estimation results

The quantile regression results presented in the main text are generated by estimating a daily model for the conditional 10th, 50th, and 90th quantiles of the cross-sectional flow distribution. Given that our choice of the 10th and 90th percentiles for purposes of estimating comparative statics in the left and right tails is somewhat arbitrary, we show that similar results obtain when we instead estimate a model for the conditional 20th and 80th percentiles. These results, which are tabulated in Table C2, are qualitatively and quantitatively similar to those reported in Table 7.