

Who Monitors the Mutual Fund Manager, New or Old Shareholders?

WOODROW T. JOHNSON*

ABSTRACT

This study tests whether mutual fund shareholders continue to trade in response to fund returns after they make their initial investment in fund shares. It decomposes the relationship between fund returns and shareholder flow in a large, proprietary panel of all shareholder transactions in one mid-size no-load mutual fund family. Results show that both new and old shareholders buy shares during periods of good returns; however, shareholder outflow is essentially unrelated to fund returns. This lack of a return-sell relationship is not driven by locked-in pension assets, shareholders' ignorance of ongoing fund returns, or embedded capital gains. However, there is evidence that exchanges between equity funds in the family are more correlated with returns of the destination fund than with returns of the origination fund. This may indicate that flow between equity mutual funds is driven by shareholders buying new funds rather than selling old funds.

I. Introduction

Agency problems are pervasive in economics. The literature analyzes a wide range of tools that stakeholders in various organizations deploy to protect their interests from self-serving managers. In the context of open-end mutual funds, Fama and Jensen (1983) suggest that the traditional tools are relatively unimportant because most of the shareholder-manager agency conflicts are resolved through shareholders' transactions. Their argument is based on two key characteristics

*University of Oregon. I can be reached via wtj@uoregon.edu, T (541) 346-3558, and F (541) 346-3341. Correspondence should be sent to Finance Department, 1208 University of Oregon, Eugene, OR 97403. I thank Diane Del Guercio, Wei Jiang (the WFA discussant), Wayne Mikkelson, Megan Partch, Jon Reuter, Erik Sirri (the AFA discussant), Paula Tkac, and seminar participants at the American Finance Association meetings, Arizona State University, Brigham Young University, the Pacific Northwest Finance Conference, the University of Oregon, and the Western Finance Association meetings for helpful comments and the Q Group for financial support. [The latest version of this document may be downloaded from http://lcb1.uoregon.edu/wtj/.](http://lcb1.uoregon.edu/wtj/)

of mutual funds that differentiate funds from industrial companies and other organizations. First, shareholders directly affect the amount of assets the manager controls through their buys and sells. Second, the manager's compensation is proportional to the fund's total net assets and is, therefore, largely determined by shareholders' individual buys and sells.¹ In other words, Fama and Jensen (1983) argue that the agency problem in mutual funds can be solved if shareholders "reward" the manager with inflow after he posts good returns and "punish" the manager with outflow after he posts poor returns.

The literature provides a strong link between shareholders' transactions and the manager's incentives: Ippolito (1992), Patel, Zeckhauser, and Hendricks (1994), Gruber (1996), and others show that fund-level net shareholder flow is positively correlated with lagged fund returns. More recent research has just started to explore the fund-level relationship between gross shareholder flow and fund returns. Edelen (1999), Bergstresser and Poterba (2002), Goetzmann and Massa (2003), O'Neal (2004), and others suggest that although gross inflow is related to past returns, gross outflow is not. This is a puzzling empirical regularity because monitoring in the Fama-Jensen sense should be strong for sells.

This paper builds upon the existing gross-flow studies by decomposing, for the first time, the return-flow relationship *within* the fund using shareholder-level data that links together individual shareholders' transactions through time, trade by trade.² The central question of interest is whether shareholders are equally responsive to returns after they make their initial investment in fund

¹More precisely, the fund's expense ratio includes a proportional management fee that compensates the investment advisor for its services. The investment advisor presumably chooses a compensation contract for the fund manager that aligns the advisor's and manager's incentives.

²There are other papers that use shareholder-level mutual fund data in other contexts. Goetzmann and Massa (2002) identify momentum and contrarian shareholders in a few Fidelity index mutual funds. Johnson (2004) measures shareholders' investment horizons in one mutual fund family. Niehaus and Shrider (2006) study how shareholders choose which fund in their portfolio to sell using data from a full-service broker/dealer. Ivković and Weisbenner (2006) explore behavioral issues that influence shareholder redemptions in one brokerage house.

shares as they are at account opening. The first contribution of this paper is a comparison of shareholders' account-opening buys with their post-opening buys. The second contribution is a series of tests that drill down to see exactly why shareholders' sells are unrelated to returns.

Studying trading differences between “new” shareholders' account-opening buys and “old” shareholders' subsequent transactions will shed light on the incentives of fund managers. For example, if old shareholders neither buy nor sell in response to ongoing returns, the manager could choose investment policies designed to attract new shareholders—in an attempt to increase fund size and his compensation—even if those policies are costly for old shareholders. Along these lines, Barclay, Pearson, and Weisbach (1998) argue that fund managers make excessive distributions at the expense of old shareholders in an attempt to be more attractive to new shareholders. Christofferson and Musto (2002) suggest that the manager can profitably raise the fees that old shareholders pay in an existing fund while simultaneously opening a clone fund with lower fees (and a correspondingly higher expected return) for new shareholders.

Understanding trading differences between new and old shareholders will also shed light on some of the frictions shareholders face when they trade. For example, if it is costly for shareholders to pay attention to ongoing performance after making their initial investment in fund shares, shareholders' account-opening buys would be more sensitive than their subsequent buys and sells to returns. The analysis may also indicate whether old shareholders have access to alternative investments when making additional contributions or whether they are constrained to invest through their previously-chosen fund (see Del Guercio and Tkac (2002)).

This paper exploits a proprietary database that includes a panel of all shareholder transactions within and across all funds in one mid-size actively managed no-load mutual fund family over a six-year period. The data include comprehensive information about each transaction. Additionally,

the data connect each transaction to the shareholder who placed it. This allows transactions to be linked through time and across observable characteristics such as the account's tax status.

The evidence shows that shareholders respond to returns when they place buy orders. In a manner consistent with rewarding the fund manager for good realized performance, both new and old shareholders buy more shares after periods of high fund returns than they do after periods of low fund returns. Old shareholders generate the majority of the number of buy transactions—even after excluding automatic transactions and reinvested fund distributions—but new shareholders are responsible for the majority of the fund's dollar-weighted inflow.

Consistent with prior research, outflow is remarkably constant across all levels of fund returns in the aggregated data: shareholders neither increase their sells during periods of poor returns nor decrease their sells during periods of good returns. These results suggest that the risk of losing assets does not incentivize the fund manager to work hard. For example, if the manager is content with the fund's size—that is, if he is willing to forgo the possibility of large inflow in response to superior future returns—he may choose to reduce his current workload without affecting his current compensation by indexing (part of) the fund's portfolio.

Why are shareholders who are anxious to buy in response to good returns unwilling to sell in response to poor returns? Gruber (1996) conjectures that shareholder flow will not respond to poor returns if it consists of pension accounts that are locked into an inferior menu of funds. The data reject this possibility because the fund family does not have a significant amount of pension assets.

A second hypothesis is that shareholders are unaware of returns after opening their accounts—perhaps they sell only in response to idiosyncratic liquidity shocks—because it is too costly for them to pay attention to ongoing returns. The data reject this hypothesis because shareholders,

after opening their accounts, buy shares in response to ongoing returns. In fact, new and old shareholders' buys are highly correlated (the contemporaneous monthly and quarterly correlation coefficients are 0.89 and 0.94, respectively), suggesting that they use the same signals to buy shares. Nevertheless, the return-buy relationship is stronger for new shareholders than it is for old shareholders.

A third potential explanation of the buy-sell asymmetry is that accrued tax liabilities make taxable shareholders unwilling to remove assets in response to poor returns. This hypothesis is also rejected by the data: the return-sell relationship for tax-deferred households is, at best, only marginally stronger than that for taxable households.

Fourth, shareholders might move assets from one fund to another in search of better returns. For example, they might rank the universe of funds each quarter and trade into the fund at the top of the list.³ Under this model, shareholders do not sell because their current fund performed poorly. Instead, they sell because another fund performed better. To explore this possibility, exchanges between equity funds in the fund family are linked together. Although the evidence is not strong due to the small size of the exchange subsample, it is consistent with the hypothesis that shareholders care more about destination fund returns than about origination fund returns. More broadly, this may indicate that flow between equity mutual funds is driven by shareholders buying new funds rather than selling old funds.

The final test shows that sample selection is important when assessing the return-sell relationship. Ivković and Weisbenner (2006) use a new methodology to suggest that mutual fund shareholders, as a group, sell shares in response to poor returns. Their result is puzzling because it is

³A branch of the mutual fund literature suggests the shareholders can earn abnormal returns by actively trading funds. For example, Hendricks, Patel, and Zeckhauser (1993) report that substantial gains can be earned by trading funds every quarter. Bollen and Busse (2005) use daily returns over quarterly horizons to show that the top decile of funds in one quarter deliver excess returns the next.

inconsistent with not only this study but also other studies such as Bergstresser and Poterba (2002) and O'Neal (2004). This paper, therefore, proposes and tests an alternative explanation of their result: their database consists of a particular subset of mutual fund shareholders with a particular trading pattern that is more return-sell sensitive than is the average shareholder. When this study restricts itself to the same type of shareholders and transactions they use, it replicates their finding: shareholders sell in response to poor returns.

Taken together, the overall results of this paper suggest that shareholder flow is an incomplete monitoring mechanism. Although aggregate shareholder inflow rewards the manager after periods of good performance, aggregate shareholder outflow does not punish him after periods of poor performance. Alternative agency control mechanisms, especially those that are responsive to poor performance, must be important in the mutual fund marketplace (see Agrawal and Knoeber (1996) and Almazan, Brown, Carlson, and Chapman (2004)). These may include career and reputational concerns of the manager and the fund family (see Chevalier and Ellison (1999), Farnsworth (2003), and Gervais, Lynch, and Musto (2005)). Also, the fund's board of directors might be more active during periods of poor performance than it is during periods of good performance (see Khorana, Tufano, and Wedge (2006)).

Although the incentive feature of shareholder flow in Fama and Jensen (1983) is not present in every model of shareholder flow, most papers suggest that shareholders should both buy and sell shares in response to fund returns. For example, Berk and Green (2004) present a rational model of shareholder flow in a world with neither asymmetric information nor moral hazard. Fund returns signal the ability of the manager, and shareholders react accordingly: "at each point in time... [shareholder assets] flow to and from each fund so that the expected excess return... is zero" (page 1275). Although shareholder flow does not incentivize the portfolio manager in this

model, shareholders still need to be vigilant in watching for poor returns so they can sell accordingly. Lynch and Musto (2003) present a model with very different implications. In their world, the fund family fires individual managers who post poor returns, effectively breaking the link between poor past performance and future returns. Thus, shareholders do not sell in response to poor returns even though they buy in response to good returns.

From a policy perspective, it is noteworthy that retail households sell poor returns while other shareholders do not because many commentators suggest that households are unsophisticated shareholders that need regulatory protection. This paper suggests that households, as a group, are not as passive as previously suggested. In fact, retail households are the most performance-sensitive group considered in this paper.

The remainder of the discussion is organized as follows. Section II describes the database used in the analysis. Section III reports the return-flow relationship for both net and gross shareholder flow. Section IV tests whether old shareholders chase ongoing returns. Section V explores alternative motivations for selling shares. Section VI concludes.

II. Data

The data for this study were supplied, generously, by an anonymous mutual fund family. The family is an open-end, no-load complex with fees and policies that are standard in the industry. It is well above the median fund family in terms of total assets under management (i.e., this family is *not* small), and it sponsors approximately ten actively managed funds, including both equity and fixed-income funds. The equity funds have similar investment objectives; moreover, none is a sector or specialty fund. All fixed-income accounts are excluded from the analysis because they may be traded differently from the equity accounts. Fund shares have been distributed geographically—

in terms of both number and value of accounts—in a way that closely mirrors the distribution of wealth in the United States, with the exception of a disproportionately large presence in the investment advisor’s home state.

The fund family provided an electronic copy of its database for the period between fall 1994 and summer 2000.⁴ The database contains three main files: shareholders (well over fifty thousand), transactions (just under one million), and funds (around ten). The shareholder file includes registration information for each account. The transaction file includes all shareholder transactions in each fund. The fund file includes a complete history of net asset values (NAVs) and distributions for each fund.

Fund distributions are removed from the database because they are unrelated to the research question. In particular, shareholders choose whether or not to reinvest distributions when they open their accounts.⁵ Even if returns affect this initial choice, they probably do not affect whether shareholders change the reinvestment option going forward. As a practical matter, essentially no shareholders change their reinvestment option after account opening.

In making the distinction between new and old shareholders, it might be helpful to identify all accounts owned by each shareholder throughout the entire sample period. Unfortunately, data limitations make it impossible to do this consistently. This may not be a big handicap because the fund family suggests that only 10–15% of shareholders own more than one account.

⁴This database is also used by Johnson (2004). Although he drops from his analysis all shareholders who opened their accounts before fall 1994, this study does not. Although pre-fall-1994 transactions are unavailable for these left-censored accounts, their characteristics (including account-opening date) are available.

⁵Most shareholders choose to reinvest dividends. The account-, transaction-, and dollar-weighted dividend reinvestment rates are 97.1%, 97.7%, and 95.9%, respectively. These rates are similar to those found in the mutual fund industry as a whole. For example, Bergstresser and Poterba (2002) report that the dollar-weighted reinvestment rate is over 93% in equity mutual funds.

A. Bookkeeping Arrangements

Shareholders have traditionally purchased no-load mutual fund shares directly from the fund. However, many shareholders now choose to interact with an intermediary (such as a mutual fund supermarket) that collects the transactions of its customers and passes them through to the fund. The trading technology provided by these intermediaries is generally superior to that which is available to the non-intermediated shareholder. It includes the ability to open multiple mutual fund accounts without completing additional paperwork and the ability to get same-day pricing on asset flow across mutual fund family boundaries. Thus, these shareholders are predicted to trade differently: either shareholders with high preferences for trade self-select the intermediary (a selection effect) or the intermediary's superior trading technologies encourage otherwise identical shareholders to trade differently (a treatment effect).

Some (but not all) of the intermediaries establish "omnibus house accounts" with the fund. Under this bookkeeping arrangement, the fund does not see the actual transactions placed by the underlying shareholders. Instead, the fund receives a daily report of aggregated gross flow from the intermediary. This arrangement obscures the trading behavior of the underlying shareholders. For this reason, these accounts are dropped from the database.

B. Shareholder Flow

The unit of observation throughout the analysis is shareholder flow which is computed from individual shareholder transactions. To be consistent with the existing literature, transactions are aggregated to the monthly level and are scaled by lagged total net assets (TNA).

The quarterly and yearly aggregation periods used in prior research are infeasible in this study because the database contains only one fund family over fewer than six years. Unreported re-

sults from daily and weekly robustness checks are qualitatively similar to the reported monthly results. However, the highly partitioned subsets used in some of the following regressions are less meaningful under higher frequency aggregation periods because the proportion of periods with no shareholder flow increases as the partition gets finer.

This study focuses on comparing the return-flow relationship across different groups of shareholders (say, groups A and B) by regressing their flow on lagged excess returns. One possible way to scale the data is to aggregate the dollars traded by each shareholder i and divide this sum by the shareholders' lagged TNA as follows:

$$\begin{aligned}\text{flow}_{n,t}^A &= \frac{\sum_{i \in A} \text{dollar}_{i,n,t}}{\text{TNA}_{n,t-1}^{A \cup B}} \\ \text{flow}_{n,t}^B &= \frac{\sum_{i \in B} \text{dollar}_{i,n,t}}{\text{TNA}_{n,t-1}^{A \cup B}},\end{aligned}\tag{1}$$

for each fund n and month t . However, it is hard to compare return-flow sensitivities across shareholder groups if one group is larger than the other group.⁶ For this reason, this study uses an alternative methodology that scales each group's aggregated transactions by its own lagged TNA as follows:

$$\begin{aligned}\text{flow}_{n,t}^A &= \frac{\sum_{i \in A} \text{dollar}_{i,n,t}}{\text{TNA}_{n,t-1}^A} \\ \text{flow}_{n,t}^B &= \frac{\sum_{i \in B} \text{dollar}_{i,n,t}}{\text{TNA}_{n,t-1}^B}.\end{aligned}\tag{2}$$

This methodology obviously requires that $\text{TNA}^{A \cup B}$ can be decomposed into TNA^A and TNA^B .

⁶Suppose all shareholders are identical, trading at the same time and in the same quantity. If group A contains twice as many shareholders as group B does, it would *mechanically* exhibit a stronger return-flow relationship in a regression of flow on returns.

Whenever this is not possible, the second-best scaling must be used:

$$\begin{aligned} \text{flow}_{n,t}^A &= \frac{\sum_{i \in A} \text{dollar}_{i,n,t}}{\text{TNA}_{n,t-1}^{A \cup B}} \\ \text{flow}_{n,t}^B &= \frac{\sum_t \sum_{i \in A} \text{dollar}_{i,n,t}}{\sum_t \sum_{i \in B} \text{dollar}_{i,n,t}} \cdot \frac{\sum_{i \in B} \text{dollar}_{i,n,t}}{\text{TNA}_{n,t-1}^{A \cup B}}. \end{aligned} \quad (3)$$

This approach rescales the aggregated transactions of shareholder group B to be the same size as that for group A . For example, if lifetime flow from group A is double that from group B , this scaling will simply double flow from group B . Differences in regression coefficients across groups A and B will reflect, therefore, differences in the timing of the dollar transactions and not simply aggregate differences in their magnitudes.

C. *Representativeness*

The database contains extraordinarily detailed data that can be used to decompose the within-fund return-flow relationship. However, this benefit comes at a price: the data set covers only one mutual fund family. This raises the question of whether this family's return-flow relationship is different from that of other families. To address this issue, three sets of results are presented in this subsection that can be directly compared with the existing literature. Taken together, this evidence (and evidence presented in Section III) suggests that this family's aggregated shareholder flow is not atypical. There is no compelling reason to believe that the shareholders in this mutual fund family are systematically different in their within-fund monitoring from shareholders in other mutual fund families. Nevertheless, only future research can definitively address this issue.

First, the monthly relationship between shareholder flow and lagged fund returns is shown in Figure 1. The monthly fund returns are sorted into ten deciles. The average flow in the next month is calculated for each of the ten deciles. The graph reveals a mildly convex association between returns and net flow that is consistent with prior research (see Sirri and Tufano (1998, Figure 1)).

The analysis is repeated separately for inflow and outflow. It shows that the net-flow relationship is entirely driven by inflow—the outflow graph is comparatively flat (see O’Neal (2004, Figures 4–5)). This suggests that even though shareholders buy in response to high fund returns, they sell for liquidity or other reasons that are unrelated to the returns of their fund.

Second, the abnormal monthly time t shareholder flow is regressed on aggregates time t flow to funds in its same style group, its own time $t - 1$ flow, its own time $t - 1$ return, its change in Jensen’s alpha from $t - 2$ to $t - 1$, and its change in Jensen’s alpha from $t - 2$ to $t - 1$ squared. The unreported estimated coefficients from this market model between November 1996 and October 1999 are all within the interquartile range of a similar regression run on each of the 3,388 unique funds in Morningstar’s domestic equity category with the exception of the coefficient on lagged returns (79th percentile) and lagged shareholder flow (77th percentile).⁷ This relatively high sensitivity to returns should bias this study *against* the lack of response to returns that is documented in the following sections.

Third, daily net shareholder flow in this fund family is similar to that in the broader fund industry. For example, Greene and Hodges (2002) report that the mean daily net flow in their sample of TrimTabs mutual funds between February 2, 1998 and March 31, 2000 is -0.01% . They report that the 25th, 50th, and 75th percentile of net flow is -0.06% , -0.01% , and 0.04% , respectively. The median (mean) daily net flow for this mutual fund family over the same time period is within their interquartile range: -0.041% (-0.044%).

⁷Paula Tkac kindly provided these calculations. See Del Guercio and Tkac (2003, Table 1).

III. Do Shareholders Symmetrically Buy and Sell Returns?

If shareholders actively follow the performance of the mutual fund manager, net shareholder flow will be positively correlated with returns. Moreover, gross shareholder flow will also be linked to managerial performance: buys should be higher in good times than in bad times, and sells should be higher in bad times than in good times. In the framework of Fama and Jensen (1983), it is especially important that sells respond to periods of poor returns.

A. *Net Flow*

The mutual fund literature measures fund performance many different ways, ranging from the simple (raw returns) to the complex (four-factor alphas). This study uses the fund's excess returns which is defined to be the difference between the returns of the fund and its benchmark index (as listed in the fund's prospectus). This measure helps mitigate concerns about the time-series dynamics of raw returns in up or down years and it will be highly correlated with the the fund's performance relative to its peers (because they track the same index). Unreported robustness checks show that the main results of the study are not sensitive to this choice. For example, qualitatively similar results are found using just fund returns or fund returns with benchmark returns.

Table 1 presents monthly OLS regression results from two specifications of the return-flow relationship for net flow in all accounts.⁸ The first specification includes six lagged excess returns, fund dummies, and a constant. The second specification adds concurrent excess returns. Estimates are multiplied by 100; therefore, they can be interpreted as the basis point change in fund size for a 1% change of the independent variables. The estimated standard errors are robust to heteroskedastic disturbances.⁹

⁸Early studies that look at the return-flow relationship using net shareholder flow include Ippolito (1992), Patel, Zeckhauser, and Hendricks (1994), and Gruber (1996).

⁹The standard errors are not adjusted for autocorrelated disturbances by clustering on the fund (see Molton (1986))

Specification 2 tabulates the main results. It shows that there is a positive, statistically significant relationship between net shareholder flow and returns for every lag. These results are consistent with the hypothesis that shareholders, in aggregate, monitor the fund manager. They buy relatively more when results are relatively good, and they buy relatively less when results are relatively poor. These results are consistent with prior research.

B. Gross Flow

The net-flow model implicitly assumes that shareholders' buying and selling decisions are symmetric. However, a significant amount of the fund's gross shareholder flow crosses each day—even more crosses each month—which suggests that there is an asymmetry in how shareholders buy and sell shares. To explore this issue, net flow is decomposed into gross inflow and outflow. Inflow and outflow are separately regressed on excess returns of the fund, fund dummies, and a constant. Zellner's Seemingly Unrelated Regression (SUR) models are estimated instead of OLS models. Although SUR produces the same estimates as OLS, SUR allows for the across-equation hypothesis tests that are essential to this study.

Table 2 presents two sets of monthly regression results that parallel those from the previous table.¹⁰ The first specification includes six lagged excess returns, fund dummies, and a constant. The second specification adds concurrent excess returns. Estimates are multiplied by 100; therefore, they can be interpreted as the basis point change in fund size for a 1% change of the independent variables.

because the number of regressors exceeds the number of mutual funds in the database (after all, each regression includes fund dummies). Although clustered standard errors can be computed, doing so violates the asymptotic theory used to justify their calculation.

¹⁰Although the bulk of the return-flow mutual fund literature is grounded in net shareholder flow, a handful of recent studies use gross shareholder flow data. Two well-known examples are Edelen (1999) and Bergstresser and Poterba (2002).

Specification 2 presents the main results. It shows that outflow is, essentially, orthogonal to returns. This contrasts sharply with the strong relationship that exists for inflow. These shareholders are eager to reward the manager for good performance, but they are unwilling to punish him for poor performance.¹¹ (These results are consistent with prior research.) Is this buy-sell asymmetry driven by shareholders' failure to monitor the fund manager after their initial investment in fund shares, or are other forces—such as taxes—affecting shareholders' propensity to sell?

IV. Do Old Shareholders Chase Returns?

A common view in the literature is that the fund's old shareholders are passive and unresponsive to returns while new shareholders are highly performance sensitive (see, for example, Gruber (1996) and Zheng (1999)). If this characterization were correct, the fund manager could exploit the fund's old shareholders in the absence of other monitoring mechanisms. Barclay, Pearson, and Weisbach (1998) argue that fund managers make excessive distributions at the expense of old shareholders in an attempt to be more attractive to new shareholders. Christofferson and Musto (2002) suggest that the manager can profitably raise the fees that old shareholders pay in an existing fund while simultaneously opening a clone fund with lower fees (and a correspondingly higher expected return) for new shareholders.

This section tests whether old fund shareholders respond to returns when buying or selling shares. The first set of tests examines whether shareholders monitor the manager after account opening in the same way they monitor at account opening. The second set attempts to identify subsets of shareholders that have above-average sensitivities to ongoing returns.

¹¹Because the mutual fund industry is growing rapidly, funds that are not keeping up are being punished: they lose their relative ranking according to assets under management, and they are unable to capture larger economies of scale in, for example, fund distribution and other shareholder services. I thank Sean Collins for making this point at the AFA meetings.

A. Automatic Transactions

The mutual fund offers automatic investment and withdrawal plans (AIPs and AWP, respectively) to its shareholders whereby the fund automatically debits or credits the shareholder's bank account at a prespecified frequency, typically monthly. Unreported results show that more than ten percent of shareholders in non-omnibus house accounts establish an automatic plan at account opening.

Table 3, Panel A reports the distribution of automatic and non-automatic transactions in non-omnibus house accounts. Results shows that AIP buys comprise 46.76% of the total number of all buys. However, AIP buys tend to be small compared with non-AIP buys. This is evident in the fact that AIP buys aggregate to only 1.99% of the total dollar buys in the fund. The table shows that AWP sells are an even smaller part of the fund's total outflow.

AIP and AWP transactions are removed from the database because they are a *de minimis* part of the fund's total flow and they are probably unrelated to ongoing fund returns: only a handful of accounts establish an automatic plan after account opening or terminate one before account closure. Automatic transactions do not appear in any of the following tables in this study.

B. New and Old Shareholders' Buys

Inflow is decomposed into new shareholders' account-opening buys and old shareholders' post-opening buys. These variables are separately regressed on excess returns of the fund, fund dummies, and a constant. It is hypothesized that new shareholders are more sensitive than old shareholders to returns due to differences in the investment opportunities available to new and old shareholders. New shareholders are, presumably, unconstrained and choose the fund when it outperforms a group of peer funds (see Gruber (1996)). Old shareholders, however, might be locked

into their fund and buy shares irrespective of ongoing fund returns (see Del Guercio and Tkac (2002)). For example, old shareholders' ongoing monitoring costs might outweigh any potential benefit of active trade.

Table 3, Panel B compares the distribution of new and old shareholders' buys. It shows that 61.72% of the total number of buys come from old shareholders. On a dollar-weighted basis, however, old shareholders' buys are smaller than new shareholders' buys, comprising only 37.11% of inflow.

Table 4 presents the two-equation SUR results. Specification 1 scales buys by total net assets (see equation 1), and specification 2 rescales old shareholders' buys to be as large as new shareholders' buys (see equation 3). For each specification, the first column reports estimated coefficients and standard errors for the new buys, the second column reports results for the old buys, and the third column reports p -values from the hypothesis tests that the corresponding new and old coefficients are equal.

As expected, new shareholders actively monitor the manager. They open more accounts when the fund returns are high than they open when fund returns are low. The results for old shareholders are similar: they buy shares when the fund performs well. This suggests that old shareholders are not passive, buying only to satisfy idiosyncratic liquidity needs. Instead, they monitor ongoing fund returns—rewarding the fund manager for posting good returns—in a manner similar to new shareholders. Unreported results show that the contemporaneous correlation between new and old shareholders' buys is 0.94 over calendar quarters and 0.89 over calendar months. These results may indicate that old shareholders, rather than being locked into their current fund, face an investment opportunity set that is similar to that for new shareholders. If so, this would support the assertion by Fama and Jensen (1983) that shareholder flow is a sufficient monitoring mechanism

for the mutual fund industry.

Although both new and old shareholders buy in response to returns, the new shareholders' response is stronger: their (scaled) estimated coefficients are more than double those from old shareholders in specification 2. Hypothesis tests reported in Table 4 show that all of these differences are statistically significant. These differences may indicate that individual shareholders are less sensitive to returns after making their initial investment in fund shares than they are at account opening, perhaps because the ongoing returns do not provide much incremental information for old shareholders. Alternatively, shareholder heterogeneity might drive the results: after account opening, one group buys returns with the same intensity as at account opening while the other group buys randomly. These hypotheses are explored next.

C. *Shareholders Who Open after High or Low Returns*

One group of shareholders that may be particularly sensitive to ongoing returns is those who opened their accounts during periods of good returns. "High" shareholders joined the fund when excess returns were above the median; "low" shareholders joined the fund when excess returns were below the median. *Future* transactions from high shareholders are predicted to be positively correlated with ongoing returns (they have revealed the fact that they trade in response to good returns) while *future* transaction from low shareholders are predicted to be either uncorrelated with ongoing returns (if they do not pay attention) or negatively correlated with ongoing returns (if they are contrarian). Time-series concerns about entry in "up" or "down" market years are mitigated by the fact that all returns are calculated relative to the fund's benchmark.

This subsection considers two high-low partitions of lagged return. The first one is based on one-month returns at account opening, and the second one is based on cumulative six-month returns at

account opening. Table 3, Panels C–D show that high shareholders execute both more *post-opening* buys and more *post-opening* sells than low shareholders do. High shareholders' transactions tend to be slightly smaller than those from low shareholders. For example, one-month high shareholders place 57.36% of the number of post-opening buys, but these buys aggregate to only 54.14% of the fund's post-opening inflow.

Table 5 presents the four-equation SUR results for high and low shareholders (account-opening buys are necessarily excluded from these regressions). For inflow, columns one and two report estimated coefficients and standard errors while column three reports p -values from the hypothesis tests that the corresponding coefficients are equal. Outflow is reported in a parallel fashion in columns 4–6. Flow is scaled by shareholder type (see equation 2). Panel A reports results from one-month returns while Panel B reports results from cumulative six-month returns.

The inflow results show that high shareholders continue to chase returns after joining the fund while low shareholders appear to ignore ongoing returns. However, the high-low differences are statistically significant for only two return lags in Panel B (note that the low shareholders have very small R^2 s and very large standard errors).

The outflow evidence is contrary to expectations for both shareholder types. High shareholders ignore returns when selling shares while low shareholders sell when returns are poor. In Panel B, four of the seven high-low differences are statistically significant, and the magnitude of the low shareholders' coefficients are especially large (they are the largest outflow coefficients in this study).

This novel decomposition of shareholders supports the hypothesis of shareholder heterogeneity in ex ante return preferences, but the documented buy-sell asymmetry for both high and low shareholders is puzzling. Although it is unclear what motivates the low shareholders to sell poor returns so strongly (after all, they joined the fund when it performed poorly), a possible explanation of the

high shareholders' asymmetry is tested below in Subsection V.B.

D. Trading Frequency

Shareholders also differ in the frequency of their transactions. They choose both the number of transactions to place after opening (for example, one or ten) and the time between transactions (for example, one month or one year). This subsection partitions shareholders in real time along both dimensions to test whether the shareholder's *prior* trading history is correlated with *future* return sensitivities, ignoring, as always, both automatic transactions and fund distribution. On the one hand, frequent transactions may reflect high idiosyncratic liquidity needs (no-load mutual funds are a low-cost investment vehicle for meeting such needs). If so, future transactions from these shareholders would be uncorrelated with returns. On the other hand, frequent transactions may identify shareholders who actively follow the fund's performance. These shareholders' future transactions would be correlated with returns.

Table 3 shows how *post-opening* transactions and flow are split between frequent and infrequent shareholders.¹² Panel E presents statistics for the number-of-prior-transactions partition. It shows that 76% of the number of buys come from shareholders who have placed at least two prior transactions. The number of sells are approximately evenly split across this shareholder grouping with 54% of them belonging to shareholders who have not placed at least two prior transactions. The dollar-weighted results indicate that shareholders who place many trades make small buys and

¹²Neither of the partitions is forward looking. For the number-of-prior-transactions partition, shareholders all start in the few-transactions group. Each shareholder stays in that group until he has made two post-opening transactions, at which point he immediately and permanently switches to the many-trades group. For the time-since-last-transaction partition, all shareholders start in the recent group upon account opening. Shareholders stay in the recent group until six months have passed without an additional transaction. At that point, they switch to the distant group and stay there until they trade again. Thus, shareholders can repeatedly switch between the recent and distant groups. They can also remain in one group for an extended period of time. For both partitions, it is arbitrarily assumed that accounts opened before fall 1994 have not traded since account opening because the database does not contain transactions before fall 1994.

large sells relative to shareholders who place few trades. Panel F presents statistics for the time-between-transactions partition. It shows that just over 80% of the number of buys occur within six months of the account's previous transaction while just over 35% of the number of sells occur within six months of the account's prior transactions. The dollar-weighted results indicate that shareholders who have traded recently make small buys and large sells relative to shareholders who have not traded recently.

Table 6 presents the four-equation SUR results for transaction frequency (account-opening buys are excluded from these regressions). For inflow, columns one and two report estimated coefficients and standard errors while column three reports p -values from the hypothesis tests that the corresponding coefficients are equal. Outflow is reported in a parallel fashion in columns 4–6. Flow is scaled by partition type (see equation 2). Panel A presents results for the number-of-prior-transactions partition (at least two; at most one), and Panel B presents results for the time-since-last-transaction partition (less than six months; more than six months).

Panel A shows that transactions from shareholders who have two or more prior transactions (after account opening) do not respond strongly to ongoing returns when either buying or selling. Even though their inflow and outflow estimated coefficients are economically large, they are statistically weak (both of the R^2 s are around 3%, 12 of the 14 return coefficients are insignificant, and both hypothesis tests that either the inflow coefficients or outflow coefficients are jointly equal to zero cannot be rejected). However, transactions from accounts with fewer than two prior transactions (after account opening) are correlated with returns for both inflow (lags 1–6) and outflow (lags 0–2). Because the many-transaction group has large standard errors—the largest in this study for both inflow and outflow—the reported many-few hypothesis tests are not rejected except for lag 3 in the inflow equation (the p -value is 9%). Nevertheless, the results are consistent with the

notion that multiple transactions are driven by shareholder's idiosyncratic liquidity needs and not fund returns.¹³

Panel B shows that shareholders with recent transactions (i.e., within the previous six months) essentially ignore ongoing returns when buying and selling (unreported tests again fail to reject the hypothesis that either the inflow coefficients or outflow coefficients are jointly equal to zero). This occurs despite the concern that recent transactions proxy for recent good returns: to the extent shareholders chase good returns, recent transactions could proxy for recent good returns because most post-opening transactions are buys. The evidence also shows that buys (but not sells) from shareholders who have not traded recently are strongly correlated with ongoing returns. Thus, this panel suggests that trades placed within six months of prior trades are motivated by idiosyncratic liquidity needs rather than ongoing returns. Another important implication is that even shareholders who have not traded for a while continue to monitor the fund, buying additional shares when the fund does well. But neither shareholder group is willing to sell shares when the fund performs poorly.

E. Complete and Partial Liquidations

A final place to look for evidence of post-opening trading heterogeneity is between the two types of sells shareholders can place: those that completely liquidate the account and those that partially liquidate the account. The previous evidence that frequent traders ignore returns suggests that complete liquidations are more likely than partial liquidations to be motivated by poor returns.

Table 7 separately regresses complete and partial sells on lagged excess returns. The overall

¹³In unreported robustness checks, the regressions are repeated for groups based on 3-or-more to 12-or-more prior transactions. The outflow and few-transaction inflow results are very similar across each of these additional specifications. However, the many-transaction group results differ in four of the ten cases: in groups 3–5, the positive return-buy relationship is statistically stronger than it is in the reported two-or-more specification.

evidence suggests that the return-flow relationship does not differ much across these two types of sells. However, the weak evidence that does exist suggests that complete liquidations are more sensitive than partial liquidations to returns (two lags are statistically significant in the former case while none is statistically significant in the later case).

V. What Motivates Sells?

Why are shareholders who are anxious to buy in response to good returns unwilling to sell in response to poor returns? This section considers three factors that may affect the observed return-sell relationship. The first one is accrued taxes. Shareholders with embedded capital gains may be reluctant to remove assets in response to poor returns. The second factor is reinvesting behavior. Shareholders may sell in order to reinvest in better-performing funds rather than to get out of a poorly-performing fund. The final one is a quick look at shareholder sampling issues.

A. Taxes

If taxes were an important deterrent, the return-sell relationship in tax-deferred accounts should be stronger than that in taxable accounts. To test this hypothesis, this subsection focuses on the relatively homogeneous subset of non-intermediated retail households (as defined in Subsection II.A) in order to avoid confounding tax effects with other clientele effects than might arise in a broader sample of account types. This subset excludes, for example, trusts, college endowments, institutions, and, of course, all shareholders who invest through omnibus accounts.

Table 3, Panel G decomposes retail household flow. The results indicate that 79.22% of buys are taxable and 77.95% of sells are taxable. Although the tax-deferred transactions are fewer in number, they tend to be larger in value: they comprise 29.55% of dollar-weighted buys and 27.37%

of dollar-weighted sells.

Table 8 presents the four-equation SUR results. For inflow, columns one and two report estimated coefficients and standard errors while column three reports p -values from the hypothesis tests that the corresponding coefficients are equal. Outflow is reported in a parallel fashion in columns 4–6.

The outflow results show that for every return lag, taxable and tax-deferred households have a statistically identical response (the inflow results are not of direct interest, but they do indicate that taxable shareholders respond more strongly than tax-deferred shareholders to good returns). However, the tax-deferred shareholders have more statistically significant coefficients than the taxable shareholders have, and the hypothesis test that the return coefficients are jointly equal to zero is rejected for tax-deferred outflow but not for taxable outflow. Thus, the evidence on whether tax lock-in affects the buy-sell asymmetry is mixed. The differences that exist are not strong, suggesting that other factors must be important. In particular, taxes do not explain most of the buy-sell asymmetry.

B. Exchanges

Some shareholders might actively trade mutual funds, moving from one fund to another in search of better returns. For example, they might rank the universe of funds each quarter and trade into the fund at the top of the list (see Hendricks, Patel, and Zeckhauser (1993) and Bollen and Busse (2005)). Under this model, shareholders do not sell because their current fund performed poorly. Instead, they sell because another fund performed better. The present data are not ideally suited to test this model because most sells leave the mutual fund family: whether and where the proceeds are reinvested is unknown. Nevertheless, some sells are exchanged into other funds in the

family. This subsection compares fund returns on both sides of these transactions to test whether shareholders' exchanges move away from poorly performing funds or towards better-performing funds. The results should be interpreted with care because exchanges make up a small part of the fund's total transaction database

All daily shareholder exchanges between distinct equity funds in the family are extracted from the database of shareholder transactions in non-omnibus house accounts. Returns of the destination and origination funds are compared for each transaction for horizons ranging from one day to one year. Table 9 presents the results. The first column tabulates the proportion of exchange transactions in which the destination fund returns are positive. The second column tabulates the proportion of exchange transactions in which the destination returns exceed the origination returns. The results in Panels A and B are transaction- and dollar-weighted, respectively.

The results show that shareholders exchange into funds that have better track records than their prior funds. For example, about two-thirds of exchanges are made into funds that outperformed the origination fund over the last one- to three-month horizons. The results are most pronounced over longer-term horizons, but they are also evident over each of the last five trading days. The dollar-weighted results are essentially the same as the transaction-weighted results, suggesting that buy-motivated exchanges are about the same size as the other exchanges.

As an additional test, exchanges are aggregated into monthly gross inflow and outflow. These flows are regressed on lagged excess returns in a two-equation SUR framework. The unreported results are weak—due to the small sample—but they are consistent with the notion that exchanges are motivated by the good performance of the destination fund rather than the poor performance of the origination fund.

C. *Shareholder Sampling Issues*

In a recent paper, Ivković and Weisbenner (2006) suggest that mutual fund shareholders, as a group, sell in response to poor returns. This finding is puzzling because it contradicts not only this study but also many other studies of the return-flow relationship that use gross shareholder flow data such as Bergstresser and Poterba (2002) and O’Neal (2004). It is important to resolve this inconsistency because their paper is based on a commonly-studied database: investors in one discount brokerage house between 1991 and 1996 (see, for example, Barber and Odean (2000)).

Ivković and Weisbenner (2006) argue that their finding is driven by methodological differences: they use raw fund returns to explain outflow while other studies use risk-adjusted returns. However, they are unable to test whether another difference is important: they analyze only retail household investors that self-selected into a particular discount-brokerage house while other studies include each fund’s entire shareholder population. They report that the quarterly correlation between the brokerage flow and estimated flow from CRSP data is only 0.50, suggesting that the shareholders in their sample do not trade in tandem with the broader population of mutual fund shareholders.

This paper is uniquely suited to show whether part of their finding might be caused by the type of shareholders analyzed: this database includes all shareholders (both households and non-households) that held or traded shares in the fund, irrespective of the channel through which they came to the fund. Table 10 shows that *within-fund* shareholder heterogeneity is important and that it probably explains part of their result. In other words, their documented return-sell relationship may be unique to the particular type of shareholders they study.

The first specification repeats, for reference purposes, the original gross-flow model from Table 2 (the coefficients are not identical because automatic transactions are excluded from this version). The second specification, a proper subset of the first, consists of just retail households (the coeffi-

coefficients in this table cannot be derived from Table 8 because that table used type-specific scaling). The results show that the household-only subset of shareholders is more sensitive than the entire population of fund shareholders to poor returns when selling: the R^2 increases; the point estimates of the individual coefficients are now all negative, three of them are statistically significant, and the hypothesis test that the coefficients are jointly equal to zero is rejected with a p -value of 0.0928.

Another important—but subtle—data difference between this paper and Ivković and Weisbenner (2006, page 6) is that they chose to exclude particular households and transactions from their database: shareholders who make post-opening buys in the same fund and all transactions that follow the first sell in the same fund. In effect, Ivković and Weisbenner (2006) skew their database toward shareholders who do not place many within-fund trades. The evidence presented above in Table 6 suggests that this might introduce an important bias.

Specification 3 tabulates the return-sell relationship for their particular subset of households (“IW Households”) by purging similar shareholders from the household database. The evidence indicates that the IW Households are more return-sell sensitive than the overall shareholder population. More importantly, it also shows that IW Households are more return-sell sensitive than the Household Only shareholders: the R^2 increases; the point estimates are larger, more of them are significant, and the hypothesis test that the coefficients are jointly equal to zero is rejected with a p -value less than 0.0000.

In summary, these results cast some doubt on the return-sell methodological contribution of Ivković and Weisbenner (2006). At least part of their result appears to be an artifact of the particular type of shareholders they chose to study (which is motivated by their tax focus—they need to calculate the shareholder’s tax basis).

VI. Conclusion

Shareholders of open-end mutual funds hold a noteworthy option that is not available in most organizations studied by financial economists: they can add or remove assets at any time and at a fair price. This ability, if properly exercised, could substitute for regulatory oversight and alternative forms of governance (see Fama and Jensen (1983)). In order to understand whether fund shareholders are adequate monitors of fund managers, this study tests whether shareholders continue to respond to returns after they make their initial investment in fund shares.

Results show that “new” and “old” shareholders have a similar, positive response to lagged returns when buying fund shares. Additionally, the contemporaneous correlation of inflow from new and old shareholders is 0.94 over calendar quarters and 0.89 over calendar months. These results indicate that old shareholders are not locked into their current fund, buying additional shares without regard to ongoing returns. Instead, it appears that new and old shareholders face similar investment opportunity sets: when new shareholders buy shares, so, too, do old shareholders.

In stark contrast to the buy results, the sell evidence shows that returns do not affect the shareholders’ decision to remove assets from the fund. This asymmetric buy-sell relationship is puzzling. First, why are shareholders unwilling to use the information in poor returns as a signal to sell fund shares when they are anxious to use the information in good returns as a signal to buy fund shares? Second, if poor returns do not cause sells, what does? The data reject the following three potential explanations of the buy-sell asymmetry: pension accounts are locked into an inferior menu of funds; shareholders do not pay attention to the fund’s ongoing returns and, hence, do not know when the fund performs poorly; and accrued taxes make shareholders with embedded gains unwilling to remove assets. The data do not reject a fourth explanation: shareholders sell their current fund in order to reinvest the proceeds in a better-performing fund. This explanation is consistent

with the hypothesis that shareholders, rather than placing transactions in isolation, strategically sell assets in order to buy other, better-performing assets. Exploring shareholders' transitions from one asset to another as they dynamically update their portfolios may be a profitable avenue for future research.

There is also evidence of shareholder heterogeneity. For example, despite the fact that shareholders, in the aggregate, do not sell in response to poor returns, some shareholder groups do. In particular, the results show that retail households sell when the fund performs poorly. This (and another data sampling issue) may explain why Ivković and Weisbenner (2006) report that mutual fund shareholders, as a group, sell poor returns while other studies report the opposite result.

This paper also demonstrates that transactions from shareholders who trade frequently are not related to returns. Instead, these shareholders' trades might be motivated by liquidity needs that are unrelated to fund performance. This may not be surprising because no-load mutual funds are a low-cost investment vehicle through which shareholders can ameliorate small or recurring liquidity shocks. This contrasts with households who frequently trade stocks—their trades appear to be motivated by returns (see, for example, Barber and Odean (2000)).

Taken together, the results suggest that shareholder flow is an incomplete monitoring mechanism. Although aggregate shareholder inflow rewards the manager after periods of good returns, aggregate outflow does not punish him after periods of poor returns. Future research could focus on how other monitoring mechanisms (such as boards of directors) respond to periods of poor returns.

Of course, it is possible that portfolio managers are adequately incentivized through inflow. Because the mutual fund industry is growing rapidly, funds that are not keeping up are being punished: they lose their relative ranking according to assets under management, and they are

unable to capture larger economies of scale in, for example, fund distribution and other shareholder services.¹⁴

This paper studies the return-flow relationship in mutual funds, shedding new light on how fund shareholders monitor fund managers. However, the findings are grounded in the returns and flow of one mutual fund family. This raises the question of whether the documented return-flow relationships are different from those found in other families. Only future research can definitively address this concern. Nevertheless, the aggregate behavior of shareholders in this family is remarkably similar to that reported in the existing literature, including Bergstresser and Poterba (2002, page 410), Del Guercio and Tkac (2003, Table 1), Greene and Hodges (2002, Table 3), and Sirri and Tufano (1998, Figure 1). This suggests that the main results might generalize to shareholders in other mutual fund families. In fact, the database used in this study is tilted slightly toward more return-sensitive shareholders than are found in the typical fund (see Del Guercio and Tkac (2003)). Thus, the main results of this paper might actually be understated.

¹⁴I thank Sean Collins for making this point.

References

- Agrawal, Anup, and Charles R. Knoeber, 1996, Firm performance and mechanisms to control agency problems between managers and shareholders, *Journal of Financial and Quantitative Analysis* 31, 377–397.
- Almazan, Andres, Keith C. Brown, Murray Carlson, and David C. Chapman, 2004, Why constrain your mutual fund manager?, *Journal of Financial Economics* 73, 289–321.
- Barber, Brad, and Terrance Odean, 2000, Trading is hazardous to your wealth: The common stock investment performance of individual investors, *Journal of Finance* 55, 773–806.
- Barclay, Michael, Neil Pearson, and Michael Weisbach, 1998, Open-end mutual funds and capital-gains taxes, *Journal of Financial Economics* 49, 3–43.
- Bergstresser, Daniel, and James Poterba, 2002, Do after-tax returns affect mutual fund inflows?, *Journal of Financial Economics* 63, 381–414.
- Berk, Jonathan B., and Richard C. Green, 2004, Mutual fund flows and performance in rational markets, Working paper, Haas School of Business, University of California at Berkeley.
- Bollen, Nicolas P. B., and Jeffery A. Busse, 2005, Short-term persistence in mutual fund performance, *Review of Financial Studies* 18, 569–597.
- Chevalier, Judith, and Glenn Ellison, 1999, Career concerns of mutual fund managers, *Quarterly Journal of Economics* 114, 389–432.
- Christofferson, Susan E. K., and David K. Musto, 2002, Demand curves and the pricing of money management, *Review of Financial Studies* 15, 1499–1524.

- Del Guercio, Diane, and Paula A. Tkac, 2002, The determinants of the flow of funds of managed portfolios: Mutual funds vs. pension funds, *Journal of Financial and Quantitative Analysis* 37, 523–557.
- , 2003, Star power: The effect of Morningstar ratings on mutual fund flows, Working paper.
- Edelen, Roger M., 1999, Investor flows and the assessed performance of open-end mutual funds, *Journal of Financial Economics* 53, 439–466.
- Fama, Eugene F., and Michael C. Jensen, 1983, Separation of ownership and control, *Journal of Law and Economics* 26, 301–325.
- Farnsworth, Heber, 2003, Reputational effects in portfolio management, Working paper, Washington University.
- Gervais, Simon, Anthony W. Lynch, and David K. Musto, 2005, Fund families as delegated monitors of money managers, *Review of Financial Studies* 18, 1139–1169.
- Goetzmann, William N., and Massimo Massa, 2002, Daily momentum and contrarian behavior of index fund investors, *Journal of Financial and Quantitative Analysis* 37, 375–390.
- , 2003, Index funds and stock market growth, *Journal of Business* 76, 1–28.
- Greene, Jason T., and Charles W. Hodges, 2002, The dilution impact of daily fund flows on open-end mutual funds, *Journal of Financial Economics* 65, 131–158.
- Gruber, Martin J., 1996, Another puzzle: The growth in actively managed mutual funds, *Journal of Finance* 51, 783–810.

- Hendricks, Darryll, Jayendu Patel, and Richard Zeckhauser, 1993, Hot hands in mutual funds: Short run persistence of relative performance, 1974–1988, *Journal of Finance* 48, 93–130.
- Ippolito, Roger, 1992, Consumer reaction to measures of poor quality: Evidence from the mutual fund industry, *The Journal of Law and Economics* 35, 45–70.
- Ivković, Zoran, and Scott Weisbenner, 2006, “Old” money matters: The sensitivity of mutual fund redemption decisions to past performance, Working paper, College of Business, University of Illinois at Urbana-Champaign.
- Johnson, Woodrow T., 2004, Predictable investment horizons and wealth transfers among mutual fund shareholders, *Journal of Finance* 59, 1979–2012.
- Khorana, Ajay, Peter Tufano, and Lei Wedge, 2006, Board structure, mergers, and shareholder wealth: A study of the mutual fund industry, Working Paper.
- Lynch, Anthony W., and David K. Musto, 2003, How investors interpret past returns, *Journal of Finance* 58, 2033–2058.
- Molton, Brent R, 1986, Random group effects and the precision of regression estimates, *Journal of Econometrics* 32, 385–397.
- Niehaus, Greg, and David G. Shrider, 2006, Redemption behavior by load mutual fund investors: Evidence on simplifying heuristics and framing, Working paper.
- O’Neal, Edward S., 2004, Purchase and redemption patterns of US equity mutual funds, *Financial Management* 33, 63–90.

Patel, Jayendu, Richard J. Zeckhauser, and Darryll Hendricks, 1994, Investment flows and performance: Evidence from mutual funds, cross-border investments, and new issues, in Ryuzo Sato, Richard M. Levich, and Rama V. Ramachandran, ed.: *Japan, Europe, and the International Financial Markets: Analytical and Empirical Perspectives* . chap. 4, pp. 51–72 (Cambridge University Press: New York, NY).

Sirri, Erik R., and Peter Tufano, 1998, Costly search and mutual fund flows, *Journal of Finance* 53, 1589–1622.

Zheng, Lu, 1999, Is money smart? A study of mutual fund investors' fund selection ability, *Journal of Finance* 54, 901–933.

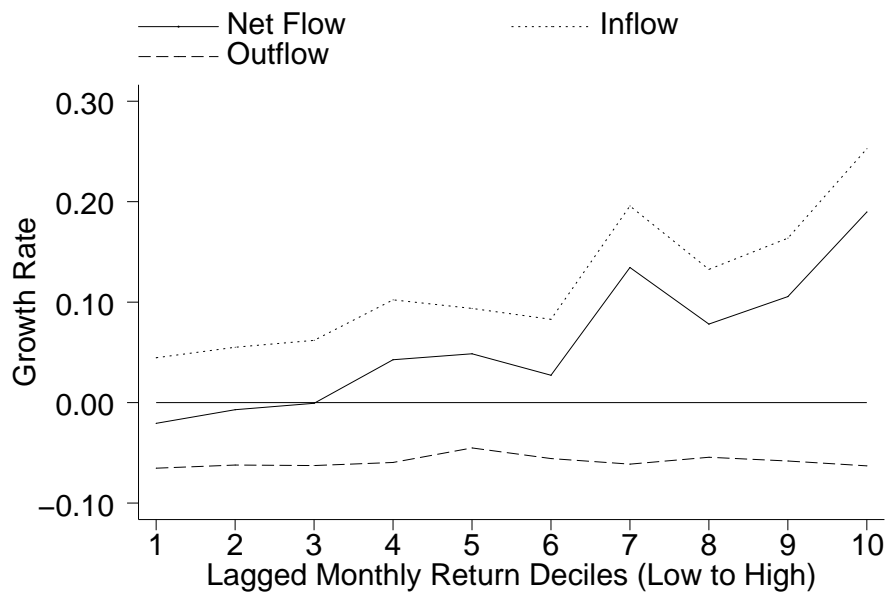


Figure 1. Return-Flow Relationship. Presents the monthly relationship between shareholder flow and lagged fund returns from all equity funds in one anonymous mutual fund family between fall 1994 and summer 2000. Fund distributions, whether reinvested or not, are excluded.

Table 1
Net Shareholder Flow

This table presents two OLS regression models of monthly net shareholder flow on excess returns of the fund (fund minus benchmark), fund dummies, and a constant from all equity funds in one anonymous mutual fund family between fall 1994 and summer 2000. Flow is aggregated from daily non-omnibus shareholder transactions (excluding fund distributions) and scaled by the shareholders' lagged total net assets. All estimates are multiplied by 100; they can be interpreted as the basis point change in fund flow for a 1% change in the independent variable. Robust standard errors are presented below their coefficients. Fund dummies and constants are included in each specification; however, their estimates are not reported. Statistical significance at ten, five, and one percent is denoted with one, two, and three asterisks, respectively.

Return Lags	Specification 1 Net Flow	Specification 2 Net Flow
Month 0		43.989**
		19.056
Month -1	48.729***	53.868***
	17.861	17.773
Month -2	52.971***	55.335***
	18.497	18.458
Month -3	52.721***	61.033***
	17.447	19.513
Month -4	48.376***	56.136***
	16.742	18.433
Month -5	53.780***	60.447***
	16.574	17.740
Month -6	69.517***	69.183***
	18.871	18.299
R^2	0.153	0.177

Table 2
Gross Shareholder Flow

This table presents two two-equation SUR models of monthly gross shareholder flow on excess returns of the fund (fund minus benchmark), fund dummies, and a constant from all equity funds in one anonymous mutual fund family between fall 1994 and summer 2000. Flow is aggregated from daily non-omnibus shareholder transactions (excluding fund distributions) and scaled by the shareholders' lagged total net assets. All estimates are multiplied by 100; they can be interpreted as the basis point change in fund flow for a 1% change in the independent variable. Standard errors are presented below their coefficients. Fund dummies and constants are included in each specification; however, their estimates are not reported. Statistical significance at ten, five, and one percent is denoted with one, two, and three asterisks, respectively.

Return Lags	Specification 1		Specification 2	
	Inflow	Outflow	Inflow	Outflow
Month 0			39.757***	-4.232
			15.293	3.747
Month -1	47.649***	-1.079	52.294***	-1.574
	15.789	3.831	15.703	3.848
Month -2	51.811***	-1.160	53.947***	-1.387
	15.783	3.830	15.617	3.827
Month -3	46.327***	-6.394*	53.840***	-7.194*
	15.967	3.875	16.039	3.930
Month -4	50.873***	2.497	57.886***	1.750
	16.932	4.109	16.947	4.152
Month -5	53.271***	-0.509	59.297***	-1.150
	18.368	4.457	18.297	4.483
Month -6	66.211***	-3.306	65.909***	-3.274
	18.510	4.492	18.290	4.481
pseudo R^2	0.171	0.132	0.191	0.136

Table 3
Transactions vs. Flow

This table compares shareholder transactions with shareholder dollar flow from all equity funds in one anonymous mutual fund family between fall 1994 and summer 2000. Panel A compares automatic with non-automatic transactions; Panel B compares new with old shareholders' buys and complete with partial liquidations; Panels C–D compare shareholders who joined the fund during periods of high excess fund returns (high) with those who joined during periods of low excess fund returns (low) over the last one or six months; Panel E compare shareholders who have traded 2 or more times (many) with those who have not (few); Panel F compares shareholders whose last transaction (including account opening) happened within the previous six months (recent) with those whose last transaction did not (distant); and Panel G compares taxable with tax-deferred shareholders. This table excludes fund distributions, whether reinvested or not, and transactions from omnibus house accounts. Panels B–G exclude automatic transactions, and Panels C–F exclude the account-opening buy. The results are pooled across all funds in the family.

A. Automatic and Non-automatic Transactions		
Trade Type	Transaction Weighted	Dollar Weighted
Buys		
Automatic (AIPs)	46.76%	1.99%
Non-Automatic	53.24%	98.01%
Sells		
Automatic (AWPs)	0.72%	0.03%
Non-Automatic	99.28%	99.97%
B. Transaction Type		
Trade Type	Transaction Weighted	Dollar Weighted
Buys		
New Shareholders	38.28%	62.89%
Old Shareholders	61.72%	37.11%
Sells		
Complete Liquidations	75.01%	71.71%
Partial Liquidations	24.99%	28.29%
C. Shareholders Who Open after High or Low Returns (Lag 1)		
Trade Type	Transaction Weighted	Dollar Weighted
Buys		
High	57.36%	54.14%
Low	42.64%	45.86%
Sells		
High	53.25%	51.95%
Low	46.75%	48.05%
D. Shareholders Who Open after High or Low Returns (Lags 1–6)		
Trade Type	Transaction Weighted	Dollar Weighted
Buys		
High	68.19%	65.00%
Low	31.81%	35.00%
Sells		
High	67.08%	67.30%
Low	32.92%	32.70%
E. Number of Prior Post-Opening Transactions		
Trade Type	Transaction Weighted	Dollar Weighted
Buys		
Many	75.80%	65.92%
Few	24.20%	34.08%
Sells		
Many	46.37%	55.76%
Few	53.63%	44.24%
F. Time Since Last Transaction		
Trade Type	Transaction Weighted	Dollar Weighted
Buys		
Recent	81.34%	71.26%
Distant	18.66%	28.74%
Sells		
Recent	36.69%	43.07%
Distant	63.31%	56.93%
G. Taxable and Tax-Deferred Households		
Trade Type	Transaction Weighted	Dollar Weighted
Buys		
Taxable	79.22%	70.45%
Tax-Deferred	20.78%	29.55%
Sells		
Taxable	77.95%	72.63%
Tax-Deferred	22.05%	27.37%

Table 4
New and Old Shareholders' Buys

This table presents two two-equation SUR model of monthly gross partitioned shareholder flow on excess returns of the fund (fund minus benchmark), fund dummies, and a constant from all equity funds in one anonymous mutual fund family between fall 1994 and summer 2000. Partitioned flow is aggregated from daily non-omnibus shareholder transactions (excluding fund distributions and automatic transactions) and scaled by the shareholders' lagged total net assets. "New" consists of account-opening buys, and "old" consists of post-opening buys. The second specification rescales old flow to be the same size as new flow (see equation 3). For the first model, estimated coefficients and standard errors are multiplied by 100 and are presented in columns 1 and 2. The return coefficients can be interpreted as the basis point change in fund flow for a 1% change of returns. Column 3 presents p -values from hypothesis tests that the coefficients are equal across the partition. Parallel results for the second model are presented in the columns 4–6. Fund dummies and constants are included in each specification; however, their estimates are not reported. Statistical significance at ten, five, and one percent is denoted with one, two, and three asterisks, respectively.

Return Lags	Unadjusted Flow			Rescaled Flow		
	New	Old	Equal?	New	Old	Equal?
Month 0	37.778***	2.164	0.002***	37.778***	4.591	0.001***
	12.980	3.274		12.980	6.864	
Month -1	43.875***	8.682***	0.002***	43.875***	18.853***	0.015**
	13.328	3.362		13.328	7.048	
Month -2	45.782***	8.292**	0.001***	45.782***	17.065**	0.005***
	13.255	3.344		13.255	7.009	
Month -3	48.013***	5.871*	0.000***	48.013***	13.148*	0.001***
	13.613	3.434		13.613	7.198	
Month -4	48.170***	9.792***	0.002***	48.170***	20.802***	0.014**
	14.383	3.628		14.383	7.606	
Month -5	52.630***	6.665*	0.001***	52.630***	15.916*	0.002***
	15.529	3.917		15.529	8.212	
Month -6	57.233***	8.472**	0.000***	57.233***	16.928**	0.001***
	15.523	3.916		15.523	8.209	
pseudo R^2	0.196	0.107		0.196	0.138	

Table 5

Shareholders Who Open after High or Low Returns

This table presents two four-equation SUR models of monthly gross partitioned shareholder flow on excess returns of the fund (fund minus benchmark), fund dummies, and a constant from all equity funds in one anonymous mutual fund family between fall 1994 and summer 2000. Partitioned flow is aggregated from daily non-omnibus shareholder transactions (excluding fund distributions and automatic transactions) and scaled by type-specific lagged total net assets (see equation 2). "High" shareholders joined the fund when lagged excess fund returns were above the median; "low" shareholders joined the fund when lagged excess fund returns were below the median. Panel A calculates returns over the prior one month at account opening while Panel B calculates returns over the prior six-month period at account opening. This static taxonomy necessarily excludes each account's opening transaction. Estimated coefficients and standard errors are multiplied by 100 and are presented in columns 1, 2, 4, and 5. The return coefficients can be interpreted as the basis point change in fund flow for a 1% change of returns. Columns 3 and 6 present p -values from hypothesis tests that the coefficients are equal across the partition. Fund dummies and constants are included in each specification; however, their estimates are not reported. Statistical significance at ten, five, and one percent is denoted with one, two, and three asterisks, respectively.

Panel A. Shareholder Sort Based on One-Month Lagged Returns						
Return Lags	Gross Inflow			Gross Outflow		
	High	Low	Equal?	High	Low	Equal?
Month 0	9.606**	16.532	0.694	2.025	-13.630**	0.055*
	4.331	17.704		5.599	6.105	
Month -1	14.584***	7.666	0.701	4.922	-8.178	0.117
	4.447	18.178		5.749	6.269	
Month -2	12.624***	22.797	0.571	-0.591	-14.048**	0.106
	4.423	18.079		5.718	6.235	
Month -3	17.518***	22.582	0.784	1.890	-17.902***	0.021**
	4.543	18.567		5.872	6.403	
Month -4	13.920***	19.687	0.767	-0.062	-4.727	0.605
	4.800	19.618		6.205	6.766	
Month -5	19.104***	17.711	0.947	10.590	-4.984	0.110
	5.182	21.180		6.699	7.304	
Month -6	12.032**	36.476*	0.245	-1.303	-9.950	0.375
	5.180	21.172		6.696	7.302	
pseudo R^2	0.152	0.047		0.099	0.124	
Panel B. Shareholder Sort Based on Six-Month Lagged Returns						
Return Lags	Gross Inflow			Gross Outflow		
	High	Low	Equal?	High	Low	Equal?
Month 0	13.990*	5.475	0.693	0.321	-11.815**	0.134
	7.651	20.911		6.707	5.507	
Month -1	15.247*	-18.533	0.128	-2.541	-10.682*	0.328
	7.871	21.512		6.900	5.666	
Month -2	15.796**	-17.762	0.128	-0.657	-15.139***	0.080*
	7.825	21.385		6.859	5.632	
Month -3	21.444***	-19.026	0.074*	12.935*	-30.888***	0.000***
	8.041	21.975		7.048	5.787	
Month -4	23.238***	-15.182	0.109	-4.707	-13.208**	0.344
	8.497	23.223		7.449	6.116	
Month -5	22.675**	42.630*	0.440	13.054	-5.045	0.062*
	9.167	25.053		8.035	6.598	
Month -6	17.718*	-26.603	0.091*	-3.493	-20.560***	0.082*
	9.289	25.387		8.143	6.686	
pseudo R^2	0.129	0.031		0.103	0.161	

Table 6
Trading Frequency

This table presents two four-equation SUR models of monthly gross partitioned shareholder flow on excess returns of the fund (fund minus benchmark), fund dummies, and a constant from all equity funds in one anonymous mutual fund family between fall 1994 and summer 2000. Partitioned flow is aggregated from daily non-omnibus shareholder transactions (excluding fund distributions and automatic transactions) and scaled by type-specific lagged total net assets (see equation 2). In Panel A, shareholder transactions are partitioned based on the shareholder's number of prior post-opening transactions. "Many" shareholders have placed two or more prior transactions after account opening while "few" shareholders have placed no more than one prior transaction after account opening. In Panel B, shareholder transactions are partitioned based on the time of the shareholder's prior transaction (including time since the account opening). "Recent" transactions were placed within six-months while "distant" transactions were not. These *dynamic* taxonomies necessarily exclude each account's opening transaction. Individual shareholders can repeatedly move between the recent and distant groups, but they can only move once from the few group to the many group. Estimated coefficients and standard errors are multiplied by 100 and are presented in columns 1, 2, 4, and 5. The return coefficients can be interpreted as the basis point change in fund flow for a 1% change of returns. Columns 3 and 6 present *p*-values from hypothesis tests that the coefficients are equal across the partition. Fund dummies and constants are included in each specification; however, their estimates are not reported. Statistical significance at ten, five, and one percent is denoted with one, two, and three asterisks, respectively.

Panel A. Number of Prior Post-Opening Transactions						
Return Lags	Gross Inflow			Gross Outflow		
	Many	Few	Equal?	Many	Few	Equal?
Month 0	32.196	-0.527	0.168	-3.411	-5.101**	0.900
	24.331	1.873		13.830	2.247	
Month -1	24.877	5.144***	0.417	3.860	-6.354***	0.460
	24.939	1.919		14.176	2.303	
Month -2	33.625	3.730*	0.216	3.640	-3.891*	0.584
	24.791	1.908		14.092	2.289	
Month -3	47.198*	4.995**	0.090*	-11.937	-0.766	0.430
	25.547	1.966		14.522	2.359	
Month -4	47.824*	5.667***	0.111	14.709	-3.375	0.229
	27.099	2.086		15.404	2.502	
Month -5	27.269	6.014***	0.451	2.344	0.288	0.898
	28.946	2.228		16.453	2.673	
Month -6	41.668	7.057***	0.219	-3.944	-2.200	0.913
	28.883	2.223		16.418	2.667	
pseudo R^2	0.035	0.119		0.026	0.171	
Panel B. Time Since Last Transaction						
Return Lags	Gross Inflow			Gross Outflow		
	Recent	Distant	Equal?	Recent	Distant	Equal?
Month 0	-5.002	11.601*	0.014**	-8.045	1.273	0.328
	4.835	6.212		7.901	5.498	
Month -1	5.106	15.573**	0.133	11.030	-2.853	0.156
	4.964	6.378		8.113	5.645	
Month -2	-2.296	17.563***	0.004***	-5.628	0.339	0.540
	4.937	6.343		8.068	5.614	
Month -3	-10.269**	20.509***	0.000***	-21.270**	6.575	0.005***
	5.070	6.514		8.286	5.766	
Month -4	1.206	19.735***	0.014**	10.547	-0.120	0.313
	5.357	6.883		8.755	6.092	
Month -5	1.480	15.162**	0.092*	1.550	-0.840	0.834
	5.784	7.431		9.453	6.577	
Month -6	3.611	18.161**	0.073*	-3.586	2.824	0.574
	5.782	7.429		9.449	6.575	
pseudo R^2	0.091	0.135		0.113	0.075	

Table 7
Complete and Partial Liquidations

This table presents two two-equation SUR model of monthly gross partitioned shareholder flow on excess returns of the fund (fund minus benchmark), fund dummies, and a constant from all equity funds in one anonymous mutual fund family between fall 1994 and summer 2000. Partitioned flow is aggregated from daily non-omnibus shareholder transactions (excluding fund distributions and automatic transactions) and scaled by the shareholders' lagged total net assets. "Complete" consists of sells that liquidate the account, and "partial" consists of sells that do not liquidate the account. The second specification rescales partial flow to be the same size as complete flow (see equation 3). For the first model, estimated coefficients and standard errors are multiplied by 100 and are presented in columns 1 and 2. The return coefficients can be interpreted as the basis point change in fund flow for a 1% change of returns. Column 3 presents *p*-values from hypothesis tests that the coefficients are equal across the partition. Parallel results for the second model are presented in the columns 4–6. Fund dummies and constants are included in each specification; however, their estimates are not reported. Statistical significance at ten, five, and one percent is denoted with one, two, and three asterisks, respectively.

Return Lags	Unadjusted Outflow			Rescaled Outflow		
	Complete	Partial	Equal?	Complete	Partial	Equal?
Month 0	-5.140*	0.940	0.091*	-5.140*	3.770	0.203
	2.840	2.320		2.840	6.520	
Month -1	-3.420	1.860	0.152	-3.420	5.110	0.236
	2.920	2.380		2.920	6.690	
Month -2	-3.250	1.860	0.164	-3.250	7.720	0.125
	2.900	2.370		2.900	6.650	
Month -3	-6.920**	-0.280	0.078*	-6.920**	1.450	0.255
	2.980	2.430		2.980	6.830	
Month -4	-1.920	3.700	0.157	-1.920	10.870	0.099*
	3.150	2.570		3.150	7.220	
Month -5	-0.710	-0.420	0.946	-0.710	-1.420	0.933
	3.400	2.770		3.400	7.800	
Month -6	-2.810	-0.440	0.581	-2.810	0.070	0.731
	3.400	2.770		3.400	7.790	
pseudo R^2	0.110	0.075		0.110	0.090	

Table 8

Taxable Households and Tax-Deferred Households

This table presents one four-equation SUR model of monthly gross partitioned shareholder flow on excess returns of the fund (fund minus benchmark), fund dummies, and a constant from all equity funds in one anonymous mutual fund family between fall 1994 and summer 2000. Partitioned flow is aggregated from daily non-omnibus shareholder transactions (excluding fund distributions and automatic transactions) and scaled by type-specific lagged total net assets (see equation 2). “Taxable” consists of flow from households that establish taxable accounts. “Tax-Deferred” consists of flow from households that establish tax-deferred accounts. This table excludes all households that invest through an intermediary such as a mutual fund supermarket. It also excludes all non-households such as trusts, college endowments, and other institutions. Estimated coefficients and standard errors are multiplied by 100 and are presented in columns 1, 2, 4, and 5. The return coefficients can be interpreted as the basis point change in fund flow for a 1% change of returns. Columns 3 and 6 present p -values from hypothesis tests that the coefficients are equal across the partition. Fund dummies and constants are included in each specification; however, their estimates are not reported. Statistical significance at ten, five, and one percent is denoted with one, two, and three asterisks, respectively.

Return Lags	Gross Inflow			Gross Outflow		
	Taxable	Tax-Deferred	Equal?	Taxable	Tax-Deferred	Equal?
Month 0	59.801*** 20.782	39.596*** 15.038	0.044**	-4.293 3.168	-5.178* 3.138	0.802
Month -1	75.267*** 21.340	51.107*** 15.441	0.019**	-5.727* 3.253	-11.353*** 3.222	0.120
Month -2	73.502*** 21.223	53.515*** 15.356	0.051*	-4.000 3.235	-6.493** 3.205	0.488
Month -3	81.153*** 21.796	56.556*** 15.771	0.020**	-0.068 3.322	-6.029* 3.291	0.107
Month -4	78.497*** 23.029	58.504*** 16.663	0.073*	-3.506 3.510	-1.633 3.477	0.631
Month -5	84.408*** 24.863	64.858*** 17.991	0.104	0.826 3.790	-2.910 3.754	0.375
Month -6	90.271*** 24.854	63.922*** 17.984	0.028**	-2.569 3.789	-5.453 3.753	0.494
pseudo R^2	0.200	0.204		0.124	0.131	

Table 9
Equity Exchange Transactions

This table presents the proportion of exchange transactions that occur during periods of high returns from all equity funds in one anonymous mutual fund family between fall 1994 and summer 2000. This analysis excludes exchanges between accounts in the same fund, automatic investment or withdrawal plan transactions, all fund distributions, and all omnibus exchanges. The first column compares the destination fund returns with zero. The second column compares the destination fund returns with the origination fund returns. Panel A weights each transaction equally. Panel B weights each transaction by its dollar value.

A. Transaction Weighted		
Return Lags	Destination is Positive	Destination Exceeds Origination
Day 0	51.0%	51.3%
Day -1	55.0%	55.1%
Day -2	53.9%	56.2%
Day -3	53.0%	54.5%
Day -4	53.2%	52.5%
Day -5	52.1%	53.5%
Month -1	74.1%	65.8%
Quarter -1	73.1%	66.9%
Half -1	79.3%	77.4%
Year -1	93.3%	75.9%
B. Dollar Weighted		
Return Lags	Destination is Positive	Destination Exceeds Origination
Day 0	50.9%	51.1%
Day -1	54.5%	54.9%
Day -2	54.5%	56.3%
Day -3	53.1%	54.1%
Day -4	54.2%	51.0%
Day -5	53.4%	52.1%
Month -1	74.1%	65.3%
Quarter -1	71.0%	66.9%
Half -1	79.5%	77.0%
Year -1	92.1%	74.1%

Table 10
Shareholder Sampling Issues

This table presents three two-equation SUR models of monthly gross shareholder flow on excess returns of the fund (fund minus benchmark), fund dummies, and a constant from all equity funds in one anonymous mutual fund family between fall 1994 and summer 2000. Flow is aggregated from daily non-omnibus transactions (excluding fund distributions and automatic transactions) and scaled by lagged TNA of the respective shareholder group. The first specification includes all shareholders; the second specification includes only retail households; and the third specification includes only the type of retail households studied by Ivković and Weisbenner (2006). All estimates are multiplied by 100; they can be interpreted as the basis point change in fund flow for a 1% change in the independent variable. Standard errors are presented below their coefficients. Fund dummies and constants are included in each specification; however, their estimates are not reported. Statistical significance at ten, five, and one percent is denoted with one, two, and three asterisks, respectively.

Return Lags	All Shareholders		Households Only		IW Households Only	
	Inflow	Outflow	Inflow	Outflow	Inflow	Outflow
Month 0	39.942***	-4.232	52.600***	-4.777*	46.009**	-8.564***
	15.257	3.747	18.608	2.676	17.876	2.638
Month -1	52.555***	-1.572	66.371***	-7.363***	64.510***	-13.243***
	15.667	3.848	19.107	2.748	18.355	2.709
Month -2	54.073***	-1.389	66.616***	-4.911*	60.597***	-8.469***
	15.581	3.827	19.002	2.733	18.255	2.694
Month -3	53.882***	-7.201*	72.526***	-2.370	66.215***	-6.558**
	16.002	3.930	19.515	2.807	18.748	2.767
Month -4	57.961***	1.741	71.283***	-2.717	63.968***	-8.011***
	16.907	4.152	20.620	2.965	19.809	2.923
Month -5	59.292***	-1.167	77.356***	-0.674	66.503***	-4.924
	18.254	4.483	22.262	3.202	21.386	3.156
Month -6	65.709***	-3.281	81.324***	-3.649	85.583***	-4.497
	18.247	4.481	22.254	3.200	21.379	3.155
pseudo R^2	0.192	0.136	0.203	0.156	0.198	0.186