J. Finan. Intermediation 19 (2010) 143–168

Contents lists available at ScienceDirect

J. Finan. Intermediation

journal homepage: www.elsevier.com/locate/jfi

Who incentivizes the mutual fund manager, new or old shareholders? ☆

Woodrow T. Johnson [1]

US Securities and Exchange Commission, 100 F Street, NE, Washington, DC 20549-9360, United States

ARTICLE INFO

Article history:
Received 27 July 2008
Available online 16 September 2009

Keywords:
Mutual funds
Shareholder flow
Return-flow relationship
Smart money

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 midsize 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 related more strongly to returns of the destination fund than to 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. Supermarket shareholders are smart insofar as they exchange into funds that subsequently outperform their prior funds during their individual holding periods.

☆ I thank Diane Del Guercio, Wei Jiang, Wayne Mikkelson, Megan Partch, Jon Reuter, Lenore Robbins, Erik Sirri, 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. Special thanks go to the anonymous mutual fund family that provided proprietary data for this study (which I received in my capacity as a graduate student prior to joining the Commission staff) and to its employees who helped in its interpretation.

[1] The Securities and Exchange Commission, as a matter of policy, disclaims responsibility for any private publication or statement by any of its employees. The views expressed herein are those of the author and do not necessarily reflect the views of the Commission or of its staff.

E-mail address: JohnsonW@sec.gov

1042-9573/$ - see front matter © 2009 Elsevier Inc. All rights reserved.
doi:10.1016/j.jfi.2009.09.003

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 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.[2] In other words, Fama and Jensen (1983) argue that the agency problem in mutual funds can be solved if shareholders incentivize the manager to work hard by "rewarding" him with inflow after he posts good returns and "punishing" him with outflow after he posts poor returns.[3]

The mutual fund literature documents a strong link between shareholders' transactions and the manager's incentives: Ippolito (1992), Patel et al. (1994), Gruber (1996), and others show that fund-level net shareholder flow is positively correlated with lagged fund returns. However, 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), Del Guercio and Tkac (2008), O'Neal (2004), and others suggest that although gross inflow is related to past returns, gross outflow is not. The asymmetric response to lagged returns is puzzling.

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.[4] The central question of interest is whether shareholders are equally responsive to returns after they make their initial investment in fund 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 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, Brown et al. (1996) show that managers might alter portfolio risk in a way that harms old shareholders in an effort to increase their own compensation. Barclay et al. (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

[2] More 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.

[3] In some principal-agent models, principals acquire private information when they assess the agent. In the case of common stocks, for example, investors would not trade on the basis of public information because the stock price already incorporates that information. In the case of mutual funds, however, it may be possible to gain from trading on public information because funds always trade at NAV. In particular, managerial quality is never incorporated in the price of open-end mutual funds (see Gruber (1996, p. 807) and Berk and Green (2004)).

[4] There 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.

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 midsize 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, however, 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?[5] 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.[6] 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 the returns of the new fund that is entered than about the returns of the old fund that is exited. More

[5] Most of the literature's evidence regarding the sensitivity of outflow to past returns is cross sectional: shareholders do not get out of the worst funds each year. The time series evidence presented in this paper and elsewhere shows a different result: shareholders do not get out of a particular fund when it performs poorly. The time series evidence is economically interesting for at least two reasons. First, it shows an empirical asymmetry in shareholder behavior that contrasts with the symmetry found in many theoretical models. Second, it suggests that shareholder outflow does not incentivize the fund manager to work hard from one period to the next. In fact, it may even encourage the manager to change his behavior through time in a way that hurts old shareholders (see, for example, Brown et al. (1996)).

[6] A branch of the mutual fund literature suggests the shareholders can earn abnormal returns by actively trading funds. For example, Hendricks et al. (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 in the next quarter.

broadly, this may indicate that flow between equity mutual funds is driven by shareholders buying new funds rather than selling old funds. Comparing the performance of the two funds *after* the exchange over *individual* holding periods, only supermarket shareholders demonstrates the "smart money" effect of Gruber (1996) and Zheng (1999).

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 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 incentive-alignment mechanism, contrary to Fama and Jensen (1983). 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 et al. (2004)). These may include career and reputational concerns of the manager and the fund family (see Chevalier and Ellison (1999), Farnsworth (2003), and Gervais et al. (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 et al. (2007)).

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, p. 1275) 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". Although shareholder flow does not incentivize the portfolio manager in their 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 portfolio 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.

1. 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.[7] 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 status

[7] This 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.

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.[8] Even if returns affect this initial choice, they probably do not affect whether shareholders change the reinvestment option going forward. As a practical matter, shareholders do not 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.

1.1. 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 trade quickly between different mutual fund families. 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, omnibus house accounts are dropped from the database.

1.2. 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 robustness checks from daily and weekly aggregation periods are qualitatively similar to the reported monthly results. However, the highly partitioned subsets used in some of the following regressions are less meaningful under weekly or daily 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:

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

[8] Most 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.

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.[9] For this reason, this study uses an alternative methodology that scales each group's aggregated transactions by its own lagged TNA as follows:

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

This methodology obviously requires that $\text{TNA}^{A\cup B}$ can be decomposed into TNA^{A} and TNA^{B}. Whenever this is not possible, the second-best scaling must be used:

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

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 the 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.

1.3. Representativeness

The database contains extraordinarily detailed data that can be used to decompose the within-fund return-flow relationship for the first time. 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 2) 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 trading 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 Fig. 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, Fig. 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, Figs. 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 aggregated 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 3388 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).[10] This relatively high sensitivity to returns should bias this study *against* the lack of response to returns that is documented in the following sections.

[9] Suppose 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.

[10] Paula Tkac kindly provided these calculations. See Del Guercio and Tkac (2008, Table 1).

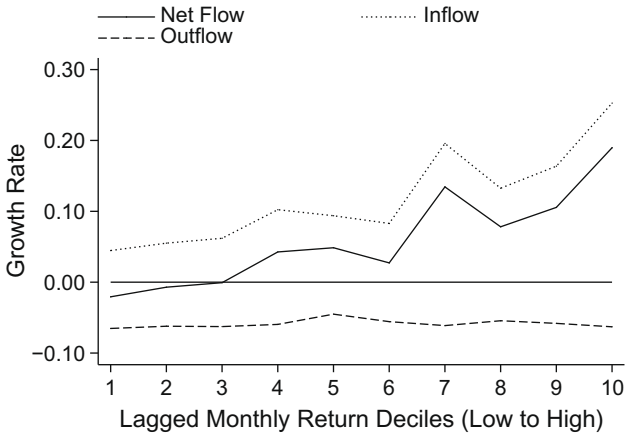

Fig. 1. Return-flow relationship. Presents the monthly relationship between shareholder flow and lagged fund returns from all equity funds in one anonymous mutual fund family between fall 1994 and summer 2000. Fund distributions, whether reinvested or not, are excluded.

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 833 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%).

2. 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.

2.1. 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 are 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 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 raw fund returns by themselves or raw fund returns along with benchmark returns.

Table 1 presents monthly OLS regression results from two specifications of the return-flow relationship for net flow in all accounts.[11] 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%

[11] Early studies that look at the return-flow relationship using net shareholder flow include Ippolito (1992), Patel et al. (1994), and Gruber (1996).

Table 1
Net shareholder flow.

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

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 so 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. Statistical significance at 10, 5, and 1% is denoted with 1, 2, and 3 asterisks, respectively.

change of the independent variables. The estimated standard errors are robust to heteroskedastic disturbances.[12]

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' transactions, in aggregate, incentivize the fund manager. Shareholders 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.

2.2. Gross flow

The net-flow model implicitly assumes that shareholders' buying and selling decisions are symmetric. To explore whether this is realistic, 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, the SUR framework facilitates 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.[13] The first specification includes six lags of 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.

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. These results are consistent with prior research. Is this buy-sell asymmetry driven by

[12] The standard errors are not adjusted for autocorrelated disturbances by clustering on the fund (see Molton (1986)) 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.

[13] Although 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).

Table 2
Gross shareholder flow.

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

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 so 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. Statistical significance at 10, 5, and 1% is denoted with 1, 2, and 3 asterisks, respectively.

shareholders' failure to pay attention to the fund after their initial investment in fund shares, or are other forces—such as taxes—affecting shareholders' propensity to sell?

3. 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 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 mechanisms that align incentives. Brown et al. (1996) show that managers might alter portfolio risk in a way that harms old shareholders in an effort to increase their own compensation. Barclay et al. (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 respond to fund returns after account opening in the same way they respond at account opening. The second set attempts to identify subsets of shareholders that have above-average sensitivities to ongoing returns.

3.1. Automatic transactions

The mutual fund offers automatic investment and withdrawal plans (AIPs and AWPs, respectively) to its shareholders whereby the fund automatically debits or credits the shareholder's bank account at a prespecified frequency, typically monthly. More than 10% of the shareholders in non-omnibus house accounts establish an automatic plan.

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 panel shows that AWP sells are an even smaller part of the fund's total outflow.

Table 3
Transactions vs. flow.

Trade type	Equally weighted (%)	Dollar weighted (%)
A. Automatic and non-automatic transactions		
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		
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)		
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)		
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		
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		
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*		
Buys		
Taxable	79.22	70.45
Tax-deferred	20.78	29.55

Table 3 (*continued*)

Trade type	Equally weighted (%)	Dollar weighted (%)
Sells		
Taxable	77.95	72.63
Tax-deferred	22.05	27.37

This table compares equally- and dollar-weighted shareholder 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 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.

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.

3.2. 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 to test whether new shareholders are more sensitive than old shareholders to returns. New shareholders are generally 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, the cost of paying attention to ongoing fund returns 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 Eq. (1)), and specification 2 rescales old shareholders' buys to be as large as new shareholders' buys (see Eq. (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' buys incentivize the manager. They open more accounts when the fund returns are good than they open when fund returns are poor. 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 pay attention to 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 an important incentive-alignment mechanism for the mutual fund industry.

Although both new and old shareholders buy in response to returns, the new shareholders' response is stronger: the rescaled flow results reported in specification 2 show that the new shareholders' estimated coefficients are more than double those from old shareholders, and the reported hypothesis tests show that each of these differences is statistically significant at 5%. These differences may indicate that individual shareholders are less sensitive to returns after making their initial

Table 4
New and old shareholders' buys.

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	

This table presents two two-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 the shareholders' lagged total net assets (see Eq. (1)). "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 Eq. (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 columns 4–6. Statistical significance at 10%, 5%, and 1% is denoted with 1, 2, and 3 asterisks, respectively.

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.

3.3. 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 while "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* transactions 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 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 esti-

Table 5
Shareholders who open after high or low returns.

Return lags	Gross inflow			Gross outflow		
	High	Low	Equal?	High	Low	Equal?
Panel A. Shareholder sort based on one-month lagged returns						
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						
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	

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 Eq. (2)). "High" shareholders joined the fund when lagged excess fund returns were above the median, and "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 respective partition. Statistical significance at 10%, 5%, and 1% is denoted with 1, 2, and 3 asterisks, respectively.

mated 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 Eq. (2)). Panel A reports results from sorts based on one-month returns while Panel B reports results from sorts based on 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. This is the expected result. However, the high-low differences are statistically significant for only two return lags in Panel B. Note that the low shareholders have very large standard errors and small R^2s.

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 Section 4.2.

3.4. 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 distributions. 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 ongoing performance. If so, these shareholders' future transactions would be correlated with returns.

Table 3, Panels E–F show how *post-opening* transactions and flow are split between frequent and infrequent shareholders.[14] 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 large sells relative to shareholders who place few trades. Panel F presents statistics for the time-since-last-transaction 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, excluding, of course, account-opening buys. 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 Eq. (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^2s are around 3%, 12 of the 14 return coefficients are insignificant, and both hypothesis tests that the inflow coefficients or the 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

[14] Neither 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.

Table 6
Trading frequency.

Return lags	Gross inflow			Gross outflow		
	Many	Few	Equal?	Many	Few	Equal?
Panel A. Number of prior post-opening transactions						
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	
	Recent	Distant	Equal?	Recent	Distant	Equal?
Panel B. Time since last transaction						
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	

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 Eq. (2)). In Panel A, shareholder transactions are partitioned based on the shareholder's number of prior post-opening transactions. "Many" shareholder 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 respective partition. Statistical significance at 10%, 5%, and 1% is denoted with 1, 2, and 3 asterisks, respectively.

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.[15]

[15] In 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.

Panel B shows that shareholders with recent transactions (i.e., within the previous six months) essentially ignore ongoing returns when buying and selling; moreover, two unreported tests again fail to reject the hypothesis that either the inflow coefficients or the 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 transactions placed within six months of prior transactions 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 pay attention to the fund's ongoing returns, buying additional shares when the fund does well. But neither shareholder group is willing to sell shares when the fund performs poorly.

3.5. Complete vs. partial liquidations

An additional 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 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: complete liquidations have more significant coefficients and a slightly larger R^2.

Table 7
Complete vs. partial liquidations.

Return lags	Unadjusted outflow			Rescaled outflow					
	Complete	Partial	Equal?	Complete	Partial	Equal?			
Month 0	−5.140	*	0.935	0.091	*	−5.140	*	3.772	0.203
	2.844		2.317			2.844		6.516	
Month −1	−3.418		1.860	0.152		−3.418		5.113	0.236
	2.920		2.379			2.920		6.690	
Month −2	−3.245		1.863	0.164		−3.245		7.725	0.125
	2.904		2.366			2.904		6.654	
Month −3	−6.919	**	−0.285	0.078	*	−6.919	**	1.449	0.255
	2.983		2.430			2.983		6.833	
Month −4	−1.924		3.704	0.157		−1.924		10.872	0.099
	3.151		2.567			3.151		7.220	
Month −5	−0.710		−0.419	0.946		−0.710		−1.419	0.933
	3.402		2.772			3.402		7.795	
Month −6	−2.811		−0.437	0.581		−2.811		0.072	0.731
	3.401		2.771			3.401		7.792	
Pseudo R^2	0.110		0.075			0.110		0.090	

This table presents two two-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 the shareholders' lagged total net assets (see Eq. (1)). "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 Eq. (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 columns 4–6. Statistical significance at 10%, 5%, and 1% is denoted with 1, 2, and 3 asterisks, respectively.

Most of the estimated coefficients for partial liquidations are positive. Although none is statistically significant, their sign does raise the possibility of portfolio rebalancing: perhaps shareholders periodically make a partial liquidation of their winning funds in order to reinvest in their losing funds.

4. 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 additional 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.

4.1. Taxes

If taxes were an important impediment, 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 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 house 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 the estimated coefficients for taxable and tax-deferred shareholders are not statistically different for any of the return lags. (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 economically strong, suggesting that other factors must be important. In particular, taxes do not explain most of the buy-sell asymmetry.

4.2. 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 et al. (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. In order to understand their trades, it is essential that both sides of the transaction be linked together for analysis; otherwise, an analysis of the sells by themselves would misleadingly indicate that shareholders ignore returns when trading.

The performance-flow literature has been unable to link sells to the following buys because the literature is grounded in fund-level shareholder flow, usually netted monthly. The present trade-by-trade database allows for a first look at the flow of individual dollars between individual funds. Although this is the best available database to test the proposed model, it is certainly not perfect. One potential shortcoming of the database is that most sells leave the mutual fund family. This contrasts with the ideal database that tracks sells that flow into financial assets outside the family as well

Table 8
Taxable households and tax-deferred households.

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

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 Eq. (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 intermediaries, including mutual fund supermarkets. 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 respective partition. Statistical significance at 10%, 5%, and 1% is denoted with 1, 2, and 3 asterisks, respectively.

as those sells that leave the financial markets. Second, this midsize fund family lacks the breadth of offerings generally available in the largest fund families and discount brokerages. A shareholder who anticipated actively switching among mutual funds—for example, using macroeconomic forecasts to time sector funds—would probably choose one of the industry's largest fund families or mutual fund supermarkets. The ideal database would also include shareholders from multiple fund families to test whether family-level product offerings are related to the shareholder's exchanges.

This subsection explores the performance-flow relationship for the set of exchange transactions. Specifically, it compares the returns of the origination fund (i.e., the old fund that is exited) with the returns of the destination fund (i.e., the new fund that is entered). This comparison is made using fund returns in the period prior to the exchange to test which fund's returns motivate the exchange—do shareholders flee poorly performing funds, or do they chase top performing funds? The comparison is also made after the exchange to show whether shareholders successfully select funds that have superior *future* performance (see Gruber (1996) and Zheng (1999)). In both cases, results over fixed horizons are reported as well as results based on the timing of shareholders' actual transactions. Using realized returns over individual holding periods is an important contribution to the "smart money" literature because data limitations have prevented that literature from saying how long shareholders hold funds. This is important because smart money has to know not only when to buy a fund but also when to sell it. Thus, inferences about shareholders' trading skills may be incorrect if they are based on evaluation periods that are different from the shareholders' actual holding periods (Frazzini and Lamont (2008) discuss this issue in more detail).

The following results show that exchanges are related more strongly to prior destination returns than to prior origination returns. This is consistent with the hypothesis that sells are motivated more by a desire to buy a new fund than they are by a desire to sell an old fund. The results also show that shareholder exchanges are neither smart nor dumb when evaluated between the time of the exchange and the time of the destination fund's closure: the destination fund outperforms the origination fund

in one-half of the exchanges, and the destination fund underperforms the origination fund in the other half of the exchanges. However, there is evidence of smart money among supermarket shareholders.

4.2.1. All exchanges

Each shareholder exchange in non-omnibus house accounts between distinct equity funds in the family is extracted from the database of shareholder transactions. Returns of the origination and destination funds are compared to test for return-chasing behavior (past returns) and the smart money effect (future returns). One set of tests is based on fixed evaluation periods ranging from one day to one year; the second set is based on account opening and closing dates; and the third set uses the account's last trade and next trade to define evaluation periods. Table 9 presents the results, reporting the proportion of exchange transactions for which the destination returns exceed the origination returns over the specified evaluation period. The first column reports results for all exchanges, and the other columns reports results for four subsets of the exchanges. The results in Panels A and B are equally- and dollar-weighted, respectively.

The first column shows that shareholders exchange into funds that have better track records than the funds they previously held. For example, 59.8% of exchanges move dollars into funds that outperform the origination fund over the prior week, and 81.9% of exchanges move dollars into funds that outperform the origination fund over the prior year.

Table 9
Equity exchange transactions.

Evaluation period	All (%)	Focused (%)	Taxable (%)	Tax deferred (%)	Supermarket (%)
A. Equally-weighted exchanges					
Account opening					
Origination	85.7	88.6	86.4	85.8	75.0
Destination	81.6		84.1	76.6	64.2
Most recent	83.7		86.1	79.4	67.9
Last trade					
Origination	72.2	70.5	71.2	74.2	72.1
Destination	73.4		72.5	77.1	56.9
Most recent	71.9		70.5	74.2	71.4
Fixed periods					
Year −1	81.9	83.8	82.0	81.1	76.8
Half −1	81.1	83.9	81.2	81.0	78.7
Quarter −1	70.0	65.5	70.0	70.0	76.4
Month −1	65.7	62.0	65.4	66.7	71.9
Week −1	59.8	60.3	60.0	58.9	65.8
Day of exchange					
Day 0	50.8	50.3	51.1	50.8	50.0
Fixed periods					
Week +1	54.1	53.7	54.7	53.7	55.1
Month +1	55.9	53.7	56.2	55.3	63.3
Quarter +1	55.1	51.7	55.4	54.9	52.4
Half +1	50.9	45.3	52.1	50.9	53.1
Year +1	62.2	64.2	62.8	63.1	59.2
Next trade					
First	54.0		53.3	56.7	56.5
Destination	53.7	48.8	53.6	55.7	57.2
Origination	59.0		58.7	61.1	58.9
Account closing					
First	58.0		56.4	64.2	62.1
Destination	50.6	46.2	49.6	55.5	53.5
Origination	61.6		59.9	67.4	63.5

(*continued on next page*)

Table 9 (*continued*)

Evaluation period	All (%)	Focused (%)	Taxable (%)	Tax deferred (%)	Supermarket (%)
B. Dollar-weighted exchanges					
Account opening					
Origination	77.4	80.8	77.7	86.5	66.2
Destination	69.1		75.9	75.5	39.9
Most recent	69.5		77.3	85.4	24.1
Last trade					
Origination	71.2	65.5	71.5	73.1	64.1
Destination	62.6		69.8	82.8	17.7
Most recent	70.9		70.6	72.9	64.0
Fixed periods					
Year −1	71.9	74.8	75.2	80.1	52.5
Half −1	75.2	77.6	76.0	81.2	67.2
Quarter −1	72.8	68.8	71.6	73.2	71.3
Month −1	68.0	66.9	67.6	70.0	61.0
Week −1	62.5	58.5	60.0	57.8	73.4
Day of exchange					
Day 0	48.2	47.7	50.0	55.2	43.1
Fixed periods					
Week +1	50.2	50.7	55.5	52.7	52.4
Month +1	53.6	49.4	59.6	52.3	64.6
Quarter +1	50.8	48.0	55.2	57.2	45.3
Half +1	49.1	46.4	51.3	52.5	45.8
Year +1	57.3	63.6	55.3	65.1	44.9
Next trade					
First	53.0		52.7	55.9	63.7
Destination	52.8	47.9	52.5	54.9	64.7
Origination	57.0		56.0	59.1	65.1
Account closing					
first	54.7		50.8	62.6	75.1
Destination	50.1	42.9	47.6	54.7	70.8
Origination	54.6		54.8	66.7	50.3

This table presents the fraction of equity fund exchanges in which the destination fund returns exceed the origination fund returns 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, fund distributions, and omnibus exchanges. The first column uses all exchanges; the second column uses only exchanges that both close the origination fund and open the destination fund; and the remaining three columns use only exchanges from retail taxable, retail tax deferred, and supermarket shareholders. Evaluation periods are based on fixed periods as well as on the individual shareholder's actual trades. Panel A weights each exchange equally, and Panel B weights each exchange by its dollar value.

A key problem in the literature is the identification of the period over which shareholders evaluate fund performance. Rather than looking at the same fixed periods for all shareholders, it makes more sense to examine evaluation periods based on each shareholder's actual trading history. For example, the account opening date might identify a period over which the shareholder cares about fund performance because that is when he first invests dollars with that fund. Results show that between the time the shareholder opens the *origination* account and the time of the exchange, the destination fund outperforms the origination fund 85.7% of the time. Between the time the shareholder opens the *destination* account and the time of the exchange, the destination fund outperforms the origination fund 81.6% of the time.[16] This fraction is 83.7% for the period between the most recent opening (i.e., the latest of the origination and destination opening dates) and the exchange.

[16] Each comparison in this table includes only exchanges for which the specified return is defined. For example, exchanges can be made into previously-existing accounts or into newly-opened accounts. In the former case, there is a well-defined return between the time of the destination opening and the exchange. In the later case, the return is undefined because the two events happen on the same day.

Another date that might mark an important evaluation period is the individual shareholder's last trade (which may be the account's opening buy). Thus, the table also reports results based on comparisons of returns between the last trade in the account and the exchange. The results indicate that for 72.2% of the exchanges, the destination fund outperforms the origination fund since the last trade in the *origination* fund. Comparing instead to the time since the last trade in the *destination* fund, the results indicate that 73.4% of the exchanges are into better-performing funds.

Overall, these results are consistent with the hypothesis that movements of dollars between mutual funds are motivated by good returns of the new fund rather than bad returns of the old fund. This sheds new light on the performance-flow relationship, and it may explain part of the asymmetry that has been documented by Sirri and Tufano (1998) and others.

The next set of tests shifts the focus from backward-looking returns to forward-looking returns. Specifically, the analysis explores whether the shareholder is "smart" in the sense that the destination fund continues to outperform the origination fund after the exchange. The fixed-horizon evidence is consistent with some smart money because the new fund outperforms the old fund in 51%–62% of the exchanges over periods ranging from one week to one year. Of course, in the other 38%–49% of the exchanges, the shareholders appear to be worse off because the fund they bought underperforms the fund they sold.

A more sophisticated test for smart money should control for the shareholder's actual holding period. After all, a shareholder probably would not care whether or not the destination fund outperforms the origination fund over a one-year period if his holding period were only six weeks. Thus, returns on six evaluation periods defined by the shareholder's actual post-exchange transactions are also reported in the first column. The evidence suggests that there is some smart money: 61.6% of the exchanges are into funds that outperform the prior fund between the exchange and the closure of the *origination* fund. However, during the more appropriate evaluation period based on the closure of the *destination* fund, only 50.6% of the exchanges were successful. So even though each of the six shareholder-specific holding periods resulted in a proportion of successful exchanges larger than 50%, it is probably better to conclude that shareholders, as a group, do not exchange into funds that subsequently outperform their prior funds. In other words, shareholders are neither smart nor dumb in their exchange transactions over their future individual holding periods.

If there are fixed costs in determining whether a particular fund is likely to outperform another, shareholders with large account balances should be more willing to bear it because they can spread the costs over a larger asset base by placing large exchanges. To test this hypothesis, each exchange in the database is weighted by its size, and all of the prior tests are reestimated. The results are tabulated in Panel B.

The evidence shows that large exchanges are generally *less* sensitive to lagged return over both fixed and shareholder-specific evaluation periods. There is also less support for the smart money hypothesis. For example, 62.2% of the equally-weighted exchanges outperform during the first year after the exchange while only 57.3% of the dollar-weighted exchange outperform during the same period. Nevertheless, the fraction of successful exchanges over the key evaluation period—between the exchange and the closure of the destination account—is essentially unchanged. It is 50.1% for dollar-weighted exchanges while it is 50.6% for equally-weighted exchanges. Thus, large exchanges are not smarter than the other exchanges.

4.2.2. Focused exchanges

Shareholders may make focused investments if they are confident in their ability to identify funds that will have good future performance. Specifically, they may be more likely to concentrate all of their dollars in their preferred fund rather than having them spread among multiple funds at the same time. To test this hypothesis, the analysis is repeated on the subset of exchanges that both close the origination fund and open the destination fund. These exchanges comprise approximately 36% of the total number of exchanges.

The second column of Table 9 reports the results. Contrary to expectations, these exchanges are not associated with higher future returns during the key evaluation period. In fact, they have the worst success rate reported in the table. On an equally-weighted basis, they underperform all exchanges by over 4% (46.2% versus 50.6%). On a dollar-weighted basis, they underperform by more than 7%

(42.9% versus 50.1%). These results are puzzling because they suggest that shareholders who move only part of their money when they exchange are more likely than the other shareholders to forecast future fund performance correctly.

4.2.3. Exchanges by shareholder type

The propensity to chase past returns and the ability to forecast future returns may vary across the types of shareholders that are in the fund. Columns 3–5 report results for exchanges from three shareholder types: retail taxable, retail tax-deferred, and supermarket (omitted categories includes trusts and institutions).[17] Although there is no compelling reason to expect that tax status should be related to future returns in this database, it might be related to lagged returns insofar as embedded capital gains and losses may affect how a shareholder trades his taxable account. Additionally, shareholders who choose to invest through mutual fund supermarkets may have very different trading strategies from the non-supermarket shareholders. The fraction of exchanges that are placed by taxable, tax-deferred, and supermarket shareholders is 54%, 30%, and 4%, respectively.

The largest differences in the results between the two types of retail shareholders come from Panel B. The dollar-weighted tax-deferred exchanges are more strongly related to prior returns over every horizon save two. Contrary to expectations, these exchanges are also smarter. In the key evaluation period between the exchange and the destination account closure, the destination fund beat the origination fund in 47.6% of the taxable exchanges and in 54.7% of the tax-deferred exchanges, a difference of 7.1%.

As expected, the supermarket shareholders differ from the non-supermarket shareholders. First, these shareholders are less likely to chase past returns. For example, while 77.4% of all dollar-weighted exchanges are into funds that outperform the origination fund since the opening of the origination fund, the corresponding fraction for supermarket shareholders is only 66.2%. Second, these shareholders are better able to forecast future returns. The key result here is that while 50.1% of the dollar-weighted exchange transactions are into funds with higher future returns over the exchange-closure window, this proportion is 70.8% for the supermarket shareholders. In other words, over two-thirds of supermarket dollar-weighted exchanges are successful over the individual shareholder's future holding period.

An important difficulty in the smart money literature lies with the identification of an appropriate evaluation period. For example, Frazzini and Lamont (2008) show that although mutual fund shareholders have some ability to trade into funds that perform well over a one-quarter horizon, they choose funds that perform poorly over longer horizons. Moreover, the magnitude of their "dumb money" effect is sensitive to the evaluation period—the three-year effect is more than double the one-year effect. The current database shows a similar feature: the fraction of successful dollar-weighted supermarket exchanges over a one-month period is 64.6% while it is only 44.9% over a one-year period. The fact that both of these proportions are smaller than the 70.8% documented through the shareholder's actual holding period further emphasizes the value of account-level data when making inferences about the trading skill of individual mutual fund shareholders.

4.3. 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,

[17] Some supermarkets establish omnibus house accounts with the fund. Other supermarkets break out shareholder accounts on the books of the fund's transfer agent. Because the database for this study was provided by the transfer agent, only exchanges made by shareholders in supermarkets that break out shareholder accounts are included in this analysis.

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 partly driven by the particular type of shareholders and transactions 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 included in Table 2 but excluded from Table 10). The second specification, a proper subset of the first, consists of just retail households (the 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, p. 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 any 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

Table 10
Shareholder sampling issues.

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

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 (see Eq. (2)). 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 so 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. Statistical significance at 10%, 5%, and 1% is denoted with 1, 2, and 3 asterisks, respectively.

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).

5. 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' trades incentivize the fund manager, 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' transactions 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 efficiently 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 incentive-alignment 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 incentive-alignment 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 fund distribution and other shareholder services.[18]

The typical shareholder exchange between two funds in the mutual fund family is not wealth enhancing. The fund that is bought is as likely to underperform as outperform the fund that is sold. However, shareholders who invest through mutual fund supermarkets are successful in over two-thirds of their exchanges. Because these findings are based on the individual shareholder's actual holding period of fund shares, they make a contribution to the "smart money" literature that examines whether mutual fund shareholders select funds that have good future performance (see Gruber (1996), Zheng (1999), and Frazzini and Lamont (2008)).

This paper studies the return-flow relationship in mutual funds, shedding new light on how closely fund shareholders watch the 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, p. 410), Del Guercio and Tkac (2008, Table 1), Greene and Hodges (2002, Table 3), and Sirri and Tufano (1998, Fig. 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 (2008)). Thus, the main results of this paper might actually be understated.

References

Agrawal, Anup, Knoeber, Charles R., 1996. Firm performance and mechanisms to control agency problems between managers and shareholders. J. Finan. Quant. Anal. 31, 377–397.

Almazan, Andres, Brown, Keith C., Carlson, Murray, Chapman, David C., 2004. Why constrain your mutual fund manager? J. Finan. Econ. 73, 289–321.

Barber, Brad, Odean, Terrance, 2000. Trading is hazardous to your wealth: the common stock investment performance of individual investors. J. Finan. 55, 773–806.

Barclay, Michael, Pearson, Neil, Weisbach, Michael, 1998. Open-end mutual funds and capital-gains taxes. J. Finan. Econ. 49, 3–43.

Bergstresser, Daniel, Poterba, James, 2002. Do after-tax returns affect mutual fund inflows? J. Finan. Econ. 63, 381–414.

Berk, Jonathan B., Green, Richard C., 2004. Mutual fund flows and performance in rational markets. J. Political Econ. 112, 1269–1295.

Bollen, Nicolas P.B., Busse, Jeffery A., 2005. Short-term persistence in mutual fund performance. Rev. Finan. Stud. 18, 569–597.

Brown, Keith, Harlow, Van, Starks, Laura, 1996. Of tournaments and temptations: an analysis of managerial incentives in the mutual fund industry. J. Finan. 51, 85–110.

Chevalier, Judith, Ellison, Glenn, 1999. Career concerns of mutual fund managers. Quart. J. Econ. 114, 389–432.

Christofferson, Susan E.K., Musto, David K., 2002. Demand curves and the pricing of money management. Rev. Finan. Stud. 15, 1499–1524.

Del Guercio, Diane, Tkac, Paula A., 2002. The determinants of the flow of funds of managed portfolios: mutual funds vs. pension funds. J. Finan. Quant. Anal. 37, 523–557.

Del Guercio, Diane, Tkac, Paula A., 2008. Star power: the effect of morningstar ratings on mutual fund flows. J. Finan. Quant. Anal. 43, 907–936.

Edelen, Roger M., 1999. Investor flows and the assessed performance of open-end mutual funds. J. Finan. Econ. 53, 439–466.

Fama, Eugene F., Jensen, Michael C., 1983. Separation of ownership and control. J. Law Econ. 26, 301–325.

Farnsworth, Heber, 2003. Reputational effects in portfolio management, Working paper, Washington University.

Frazzini, Andrea, Lamont, Owen A., 2008. Dumb money: mutual fund flows and the cross-section of stock returns. J. Finan. Econ. 88, 299–322.

Gervais, Simon, Lynch, Anthony W., Musto, David K., 2005. Fund families as delegated monitors of money managers. Rev. Finan. Stud. 18, 1139–1169.

Goetzmann, William N., Massa, Massimo, 2002. Daily momentum and contrarian behavior of index fund investors. J. Finan. Quant. Anal. 37, 375–390.

Goetzmann, William N., Massa, Massimo, 2003. Index funds and stock market growth. J. Business 37, 1–28.

Greene, Jason T., Hodges, Charles W., 2002. The dilution impact of daily fund flows on open-end mutual funds. J. Finan. Econ. 65, 131–158.

Gruber, Martin J., 1996. Another puzzle: the growth in actively managed mutual funds. J. Finan. 51, 783–810.

[18] I thank Sean Collins for making this point.

Hendricks, Darryll, Patel, Jayendu, Zeckhauser, Richard, 1993. Hot hands in mutual funds: short run persistence of relative performance, 1974–1988. J. Finan. 48, 93–130.
Ippolito, Roger, 1992. Consumer reaction to measures of poor quality: evidence from the mutual fund industry. J. Law Econ. 35, 45–70.
Ivković, Zoran, Weisbenner, Scott, 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. J. Finan. 59, 1979–2012.
Khorana, Ajay, Tufano, Peter, Wedge, Lei, 2007. Board structure, mergers, and shareholder wealth: a study of the mutual fund industry. J. Finan. Econ. 85, 571–598.
Lynch, Anthony W., Musto, David K., 2003. How investors interpret past returns. J. Finan. 58, 2033–2058.
Molton, Brent R., 1986. Random group effects and the precision of regression estimates. J. Economet. 32, 385–397.
Niehaus, Greg, Shrider, David G., 2006. Redemption behavior by load mutual fund investors: eidence on simplifying heuristics and framing, Working paper.
O'Neal, Edward S., 2004. Purchase and redemption patterns of US equity mutual funds. Finan. Manage. 33, 63–90.
Patel, Jayendu, Zeckhauser, Richard J., Hendricks, Darryll, 1994. Investment flows and performance: evidence from mutual funds, cross-border investments, and new issues. In: Sato, Ryuzo, Levich, Richard M., Ramachandran, Rama V. (Eds.), The International Financial Markets: Analytical and Empirical Perspectives, Japan, Europe. Cambridge University Press, New York, NY, pp. 51–72 (Chapter 4).
Sirri, Erik R., Tufano, Peter, 1998. Costly search and mutual fund flows. J. Finan. 53, 1589–1622.
Zheng, Lu, 1999. Is money smart? A study of mutual fund investors' fund selection ability. J. Finan. 54, 901–933.