

Individual investor mutual fund flows*

Zoran Ivković^{a,*}, Scott Weisbenner^{a,b}

^a*Michigan State University, East Lansing, MI 48824, USA*

^b*National Bureau of Economic Research, Cambridge, MA 02138, USA*

Abstract

This paper studies the relation between individuals' mutual fund flows and fund characteristics, establishing three key results. First, consistent with tax motivations, individual investors are reluctant to sell mutual funds that have appreciated in value and are willing to sell losing funds. Second, individuals pay attention to investment costs as redemption decisions are sensitive to both expense ratios and loads. Third, individuals' fund-level inflows and outflows are sensitive to performance, but in different ways. Inflows are related only to "relative" performance, suggesting that new money chases the best performers in an objective. Outflows are related only to "absolute" fund performance, the relevant benchmark for taxes.

JEL classification: G11; C41; D14; H20

Keywords: Mutual fund flows; Individual investor portfolio choice; Tax-motivated trading

*We thank an anonymous discount broker for providing data on individual investors' trades and Terry Odean for his help in obtaining and understanding the data set. Special thanks go to Joshua Pollet, Clemens Sialm, and Jay Wang for many insightful suggestions. We also thank seminar participants at the 2005 FEA Meetings at UNC (especially the discussant Jason Karceski), the 2006 EFA Meetings in Zürich (especially the discussant Daniel Bergstresser), the 2007 WFA Meetings in Big Sky, Montana (especially the discussant Sunil Wahal), Arizona State University, Florida State University, Hong Kong University of Science and Technology, Michigan State University, Nanyang Technological University, Purdue University, Queens University, Singapore Management University, the University of Illinois, the University of Indiana, the University of Münster, the University of Texas, and the University of Wisconsin for useful comments. Both authors acknowledge the financial support from the College Research Board at the University of Illinois at Urbana-Champaign.

*Corresponding author.

E-mail addresses: ivkovich@bus.msu.edu (Z. Ivković), weisbenn@bus.msu.edu (S. Weisbenner).

Individual investor mutual fund flows

Abstract

This paper studies the relation between individuals' mutual fund flows and fund characteristics, establishing three key results. First, consistent with tax motivations, individual investors are reluctant to sell mutual funds that have appreciated in value and are willing to sell losing funds. Second, individuals pay attention to investment costs as redemption decisions are sensitive to both expense ratios and loads. Third, individuals' fund-level inflows and outflows are sensitive to performance, but in different ways. Inflows are related only to "relative" performance, suggesting that new money chases the best performers in an objective. Outflows are related only to "absolute" fund performance, the relevant benchmark for taxes.

JEL Classification: G11; C41; D14; H20

Keywords: Mutual fund flows; Individual investor portfolio choice; Tax-motivated trading

1. Introduction

The mutual fund literature has long recognized that investors respond to mutual fund performance and has documented a robust, positive relation between net fund flows and past fund performance (e.g., Ippolito, 1992; Chevalier and Ellison, 1997; Sirri and Tufano, 1998). Net flows, however, are differences of two nearly equally large components—new purchases (inflows) and redemptions (outflows)¹—that might well follow different patterns, and aggregation into net flows may preclude the development of more detailed insights.

Although the existing studies largely rely upon net flows, as dictated by data availability,² a conventional wisdom developed that the net flow-performance relation stems from the strong performance-chasing exhibited by new buys, with little or no contribution to the relation from the

¹During the period from 1984 to 2002, redemptions were almost as large as new purchases, accounting for 48.5% of the sum of dollar amounts of new purchases and redemptions. This figure is based on authors' calculations from the data reported in the 2003 Mutual Fund Factbook (Investment Company Institute, 2004).

²The data collection of aggregate fund-level inflows and outflows, available from the Securities and Exchange Commission (SEC) in electronic form since the mid-1990s, is onerous and very few studies have pursued it (e.g., Edelen, 1999; Bergstresser and Poterba, 2002; O'Neal, 2004; Cashman, Deli, Nardari, and Villupuram, 2006), and then only, with the exception of Cashman, Deli, Nardari, and Villupuram (2006), for a relatively small number of funds. The only other data set that features the possibility of effective separation of net flows into inflows and outflows in the domain of U.S. mutual fund investments of which we are aware, used in Johnson (2007), are transactions of shareholders in one small, no-load mutual fund family. Although the number of investors covered by that data set is substantial (well over 50,000), its limitation is a narrow representation of funds (up to ten funds) and, thus, limited cross-sectional variation of their characteristics. Finally, data on inflows and outflows are available for U.K. mutual fund investors (Keswani and Stolin, 2008), although their use to date has been limited to assessing whether "money is smart," with separate consideration of inflows and outflows.

redemption side. That notion, however, is not grounded in a direct inquiry into the patterns of inflows and outflows, and it might well be that redemptions are related to past performance. Also, inflows and outflows might be related to other fund characteristics very differently.

This paper studies the relation between individual investors' mutual fund flows and a range of fund characteristics (past performance, proxies associated with potential future fund-driven tax liabilities, and investment costs). We use detailed brokerage data featuring investments a large sample of individual investors made in the period from 1991 to 1996. We first study individuals' fund share redemption decisions in both taxable and tax-deferred settings, and then proceed to study individual investor fund-level inflows and outflows by aggregating the purchases and redemptions from the brokerage data by month and by mutual fund and thus decomposing individuals' fund-level net flows into inflows and outflows (as well as, in additional analyses, decomposing the latter into outflows in taxable accounts and outflows in tax-deferred accounts).

Individual investors' mutual fund share redemption decisions might be related to fund performance since purchase for several (perhaps countervailing) reasons. For one, there is the tax motivation. In the U.S., capital gains are taxed on a realization basis, which provides investors with an incentive to hold on to mutual fund shares whose net asset value per share (NAV) has appreciated since purchase (thus delaying the payment of taxes) and redeem mutual fund shares whose NAV has fallen in value since purchase (thus using those losses to reduce taxes right away).

A belief in performance persistence could also lead to holding funds that have appreciated in value since purchase and selling those that have fallen since purchase. Both tax motivations and belief in performance persistence predict a negative relation between propensity to sell and past fund performance, although tax considerations matter only in taxable accounts.

On the other hand, psychological considerations seem to play an important role in individuals' trading decisions. For example, the disposition effect—the propensity to cash in gains and aversion to realize losses (Kahneman and Tversky, 1979; Shefrin and Statman, 1985)—would lead to a pronounced positive relation between redemption decisions and fund performance since purchase. Although the disposition effect appears to be a dominant determinant of individual investors' decisions to sell common stock shares (e.g., Odean, 1998; Grinblatt and Keloharju, 2001), there is little research inquiring whether such findings carry over to mutual funds.

To disentangle these competing hypotheses, we focus on the change of fund NAV since purchase. Naturally, that performance measure, as important as it is for taxation purposes and for psychological explanations, is unlikely to be the only determinant of individuals' redemption decisions. Tax-sensitive investors might focus not only on the direct tax consequences related to the change of the fund NAV since purchase, but also on fund characteristics that could provide information regarding *future* fund distribution policy (such as turnover, past distribution behavior,

and capital gains overhang), information relevant for taxable investors because the tax rate on distributions received generally exceeds the tax rate on capital gains realized in the future upon sale.

Individuals' mutual fund share redemption decisions in both their taxable and tax-deferred accounts might also be sensitive to investment costs. The literature to date has not focused on these sensitivities at the individual-investor level. Rather, as in Barber, Odean, and Zheng (2005), analyses focus on the effects that investment costs such as expenses and loads might have on mutual fund *net* flows. However, as with fund performance, useful insights may be gleaned by breaking net flows into their two components. For example, expense ratios might have only a modest relation with net flows, yet could have strong positive effects both on “new” money flowing into the fund (e.g., high expenses could partly be used for advertising to help attract new investors or could be interpreted as a signal of quality of fund management or services provided by the fund family) and “old” money leaving the fund (e.g., in response to the higher ongoing costs of maintaining the investment).

We establish three key findings. First, in stark contrast with individual investor behavior in regard to common stocks (Odean, 1998; Grinblatt and Keloharju, 2001), there is a *negative* relation between the likelihood of sale and past mutual fund performance for mutual funds held in taxable accounts. That is, investors holding mutual funds in taxable accounts are reluctant to sell funds that appreciated in value and willing to sell funds that have fallen in price.

A comparison of trades in taxable and tax-deferred accounts suggests that the negative relation can be explained by tax-motivated trading (i.e., capital gains lock-in and tax-loss selling) because there is no statistically significant relation between redemption probability and fund performance since purchase in tax-deferred accounts. Thus, on net, psychological motivations appear to play much less of a role in the domain of individuals' mutual fund investments than they do in the domain of investment in individual stocks.

In further support for tax-motivated trading, we find that taxable investors' redemption decisions are sensitive to proxies for future distribution behavior of the fund. In taxable accounts, the fund turnover ratio, the historical share of total fund returns distributed to the fund investors over the preceding five years, and the fund capital gains overhang—three proxies for future fund distribution behavior—are all positively related to redemption probability. By contrast, in tax-deferred accounts turnover does not play a role in redemption decisions and the relation between redemption probability and past fund distribution policy is significantly weaker than it is in taxable accounts.³

Our second key result is that mutual fund investors are sensitive to both front-end loads (an “in-your-face” cost of mutual fund investments) and expense ratios (a more subtle, ongoing cost). The latter yields a finding that differs from the conclusions reported in Barber, Odean, and Zheng (2005),

³Somewhat curiously, the redemption decision is sensitive to capital gains overhang not only in taxable accounts, but also in tax-deferred accounts (where tax motivations are absent).

though the differential could be attributed to the fact that we are looking into individuals' redemption decisions, whereas Barber, Odean, and Zheng (2005) focus on quarterly net fund flows.

Our third key result stems from aggregation of individual investors' buys and sells into monthly measures of fund-level inflows and outflows. Consistent with the notion that new money chases the best performers, we find that inflows are related only to funds' relative performance measures, that is, funds' one-year performance relative to other funds pursuing the same objective. On the other hand, as expected in light of the transaction-level results, outflows are related only to funds' one-year "absolute" returns. The latter finding is consistent with tax motivations, as it is the absolute change in NAV that is relevant for taxable investors' tax liability following a sale. Thus, both new money and old money are sensitive to past fund performance, but in very different ways.

Finally, we also consider the role that the investment costs play in the context of fund-level flows. The most striking finding is that individuals' fund-level inflows and outflows are *each* positively related to expense ratios—whereas higher expenses may attract more new money through advertising or a perception that higher expenses reflect better managerial talent or fund family services, they also prompt old money to leave the fund sooner than it otherwise would. We also find sensitivity to loads, particularly for outflows.

These rich characterizations are obscured when inflows and outflows are combined and only net flows are studied. Our results stress that the absence of a relation between net flows and a fund characteristic does not imply that the characteristic is unimportant for *all* investors. Indeed, given that new or potential investors in the fund and the incumbent investors may have different considerations, disentangling net flows into its components is important. The most glaring examples are considerations of absolute performance, expenses, and loads, all of which are of substantial importance to the incumbent investors considering a sale of their fund shares.

The remainder of this paper is organized as follows. In Section 2 we review the data and present some summary statistics. Section 3 presents the results of analyses that relate probability of sale of individuals' mutual fund investments with a range of fund characteristics, including past performance, determinants of future potential tax liabilities, and investment costs. In Section 4 we aggregate investors' buys and sells of mutual funds into monthly measures of inflows and outflows, and analyze the determinants of those flows. Section 5 concludes.

2. Data description and summary statistics

2.1. Data description

Our primary data set, trades that 78,000 households made in the period from January of 1991 to November of 1996, comes from a large discount broker. Mutual funds are the second most frequently used investment vehicle in the data set, accounting for 18% of the overall value of all the

trades investors in the sample made over the six-year period. They are second only to common stocks (which account for around two-thirds of the overall value of the investments in the sample). A number of households have multiple accounts (such as one taxable and one tax-deferred account); the median number of accounts per household is two.

Around 32,300 households made at least one mutual fund purchase during the sample period either in taxable or tax-deferred accounts (IRAs and Keogh plans; retirement plan accounts provided through employment such as 401(k)-type plans are not part of the data set). For a detailed description of the brokerage data set see Barber and Odean (2000).

Mutual fund returns, rankings, and fund characteristics come from the Center for Research in Security Prices (CRSP) Open-End Mutual Fund Database and Morningstar. We extract the relevant information regarding sample funds' investment objectives from the CRSP mutual fund database fields "Objective" and "ICDI Objective." Our brokerage sample contains transactions covering more than 1,100 different mutual funds across 200 different mutual fund families that span more than 40 different investment objective categories. We will control for heterogeneity both on the individual-investor level, as well as the mutual-fund-type level, by allowing sale decisions to vary by the mutual fund family, by the objective of the mutual fund, by whether the fund is actively or passively managed (i.e., is it an index fund), as well as by whether the transaction is in a taxable or a tax-deferred account.

Consistent with Ivković, Poterba, and Weisbenner (2005), we include in our sample all mutual fund share purchases (and follow the purchase to see whether there is a subsequent sale), with one exception: in the instances in which multiple buys are followed by a sale, it is not possible to match unambiguously which purchased fund shares actually have been sold without making assumptions such as FIFO (first share bought, first share sold) or LIFO (last share bought, first share sold), which by itself could drive the results.⁴ The exclusion of multiple buys preceding an ambiguous sale reduces the number of purchases in the sample by around 20%.

Also, in the instances in which multiple sales follow a single purchase, only the first sale is admitted into the sample, which means that our analyses may slightly understate the actual holding periods for these mutual fund investments. However, that bias is negligible because the vast majority of mutual fund sales in the sample (89%) are complete liquidations of the respective positions.

⁴For example, in an upward market (as generally was the case during much of the sample period) fund shares purchased first in a string of purchases would have a larger appreciation since purchase than the last share in a string of purchases. Therefore, assuming FIFO would induce a positive relation between redemption probability and past fund performance.

2.2. Summary statistics

Table 1 presents summary statistics on mutual fund purchases and subsequent sales in the sample. Applying the criteria outlined above results in 325,185 buys made over the sample period, representing 32,259 households that had at least one mutual fund purchase during the sample period. The numbers of mutual fund purchases in taxable accounts and tax-deferred accounts, as well as median dollar amounts of those purchases, are very similar. Approximately one-third of the purchases were followed by a sale during the sample period.

TABLE 1 ABOUT HERE

2.3. Graphical summary of hazard rates and past NAV change

Fig. 1 presents monthly hazard rates (i.e., the likelihood of sale during a given month after purchase conditional on having not sold up to that month) of individuals' sales of mutual fund shares held in their taxable accounts. The two solid lines depict hazard rates conditional upon the fund NAV having increased since purchase (gray solid line) and hazard rates conditional upon the fund NAV having decreased since purchase (black solid line) for each of the first 36 months following the purchase. For the purposes of this figure we restrict our attention to all mutual fund purchases in taxable accounts in January. This strategy allows for identification of end-of-year effects and other patterns potentially related to the calendar month. We obtain the confidence intervals presented in Fig. 1 by calculating standard errors that allow for heteroskedasticity as well as correlation across observations associated with the same individual.

The figure identifies two very pronounced empirical facts that differentiate sales of mutual fund shares from sales of common stocks (Odean, 1998; Grinblatt and Keloharju, 2001). First, in stark contrast with common stock investments, hazard rates conditional upon losses *exceed* those conditional upon gains. Thus, on net, psychological motivations such as the disposition effect appear to play much less of a role in the domain of individuals' mutual fund investments than they do in the domain of investment in individual stocks.

Second, hazard rates of selling mutual fund shares in taxable accounts, although declining like the hazard rates for common stocks, are significantly smaller than those for stocks. For example, in the first few months the unconditional hazard rates of selling mutual fund shares are around 3-4%, whereas the comparable hazard rates of selling common stocks start as high as 15% after one month, 10% after two months, and 8% after three months (Ivković, Poterba, and Weisbenner, 2005). This discrepancy suggests that high-frequency traders are not nearly as present in the arena of mutual fund investments.

FIGURE 1 ABOUT HERE

3. Analysis of redemption decisions

We proceed with analyses of the relation between propensity to sell fund shares and a range of fund characteristics, contrasting redemption behavior in taxable and tax-deferred settings. We begin by estimating a Cox proportional hazards model of mutual fund share sales over the sample of all trades in taxable accounts. We estimate the baseline hazard rate non-parametrically (Han and Hausman, 1990; Meyer, 1990). The proportional hazards specification assumes that the hazard function $h_i(t)$ for the sale of mutual fund investment i , t months after the purchase, takes the following form:

$$h_i(t) = \lambda(t) * e^{X_{i,t} \beta}, \quad (1)$$

where λ is the baseline hazard rate and X are the covariates that shift the baseline proportionately.

One possibility would be to set the baseline probability of sale over time to be the same for all investments i . However, this would impose the constraint that all trading decisions must conform to one general hazard function. An improvement over the use of a common baseline would be to allow investor-level heterogeneity by enabling each individual to have a personal, investor-specific baseline hazard function. This strategy would take into account that, regardless of past performance and other fund characteristics, some individuals simply are more likely to trade than others. Loosely speaking, allowing for investor-specific baselines in the present context is similar to the inclusion of individual fixed-effects in a linear regression model.

We address these heterogeneity issues by using an even more encompassing approach—we allow for separate non-parametric baseline hazard rates for each investor-mutual fund type combination. The fund type is defined by the interaction of a funds' objective with its degree of active management (index funds versus actively managed funds) and its fund family membership.⁵ Thus, the regression results concerning past performance, for example, will be identified by how a given individual trades two funds with the same objective, fund family, and degree of active management that have different performance since purchase. Investor-mutual fund type baselines, therefore, are a significant step toward alleviating the concern that investor and mutual fund heterogeneity and unobserved characteristics may be driving the results.

⁵There are 44 objectives and slightly more than 200 mutual fund families represented in the brokerage sample, leading to slightly more than 72,000 investor-mutual fund type combinations in the taxable account sample. Thus, the hazard model includes 72,000 separate non-parametric baselines (loosely speaking, 72,000 fixed effects). On average, an investor purchases two mutual funds in a particular objective-index-family combination.

The explanatory variables included in the specification that we estimate for observations in taxable accounts are:

$$\begin{aligned}
 X_{i,t} \beta = & \beta_1 * NAV_RETURN_{i,t-1} + \beta_2 * NAV_RETURN_{i,t-1} * December_{i,t} + & (2) \\
 & \beta_3 * December_{i,t} + \beta_4 * Relative\ Performance\ Controls_{i,t-1} + \\
 & \beta_5 * Future\ Distribution\ Controls_{i,t-1} + \\
 & \beta_6 * Cost\ Controls_{i,t-1} + \varepsilon_{i,t} ,
 \end{aligned}$$

where NAV_RETURN denotes the relative change in fund NAV since purchase, defined as $NAV_RETURN_{i,t} = NAV_{i,t} / NAV_{i,p} - 1$. $NAV_{i,p}$ denotes the NAV of investment i at the time of purchase and $NAV_{i,t}$ denotes its NAV at the end of month t since purchase. All NAV values are adjusted for splits. The percentage change in NAV since purchase captures capital appreciation/depreciation of the fund since purchase; it is the performance benchmark of direct relevance for tax considerations. The indicator variable December controls for end-of-year effects.

Relative Performance Controls include the percentile ranking of one-year total returns within the fund's investment objective and the fund's Morningstar 5-star rating. Future Distribution Controls are predictors of future distribution policy (turnover, past fund distribution policy, and fund overhang) that capture indirect tax motivations. Finally, Cost Controls capture investment costs (expense ratios, front-end loads, and back-end loads).

Investors covered by the data set can have both taxable and tax-deferred accounts (i.e., IRA and Keough plans; investments in 401(k) plans are not part of the brokerage sample). Under the assumption that the disposition effect and the belief in fund performance persistence do not differ across investments in taxable and tax-deferred accounts, comparing the propensities to sell across mutual fund holdings in the two types of accounts provides a direct way of identifying the impact of taxation because tax considerations should not affect trading decisions in tax-deferred accounts.⁶

Accordingly, we also estimate regressions over the full sample of taxable and tax-deferred accounts. The model allows for separate non-parametric baseline hazard rates for each investor-mutual fund type combination, introduced separately for an investor's holdings in taxable and tax-deferred accounts. We introduce an indicator variable TAX_i that denotes whether the mutual fund investment i is held in a taxable account and interact TAX_i with all of the preceding variables:

⁶This strategy is used in Ivković, Poterba, and Weisbenner (2005) to study individual investors' tax-motivated trading of common stocks. A stronger disposition effect in taxable accounts than in tax-deferred accounts would bias against, whereas a stronger belief in fund performance persistence in taxable accounts would bias in favor of, finding evidence of tax-motivated trading.

$$\begin{aligned}
X_{i,t} \alpha = & \alpha_1 * NAV_RETURN_{i,t-1} + \alpha_2 * NAV_RETURN_{i,t-1} * December_{i,t} + \\
& \alpha_3 * December_{i,t} + \mathbf{Other\ Controls}_{i,t-1} * \alpha_4 + \\
& \alpha_5 * NAV_RETURN_{i,t-1} * TAX_i + \alpha_6 * NAV_RETURN_{i,t-1} * December_{i,t} * TAX_i + \\
& \alpha_7 * December_{i,t} * TAX_i + \mathbf{Other\ Controls}_{i,t-1} * \alpha_8 * TAX_i + \varepsilon_{i,t}.
\end{aligned} \tag{3}$$

In this regression, estimated over the pooled sample of mutual fund investments in taxable and tax-deferred accounts, the coefficient α_1 represents the sensitivity of the sale decision to NAV performance since purchase in tax-deferred accounts, the coefficient α_5 represents the differential behavior in taxable accounts relative to tax-deferred accounts, and $\alpha_1 + \alpha_5$ equals the sensitivity of the sale decision to NAV performance since purchase in taxable accounts (which corresponds to the coefficient β_1 from Eq. (2)).

We present the results of the regressions based on Eq. (2) and Eq. (3) in Table 2.⁷ For presentational convenience, we group our discussion of covariates into those related to fund performance since purchase, relative performance rankings, indirect tax motivations, and investment costs (Sections 3.1 through 3.4, respectively).

3.1. Motivations for negative relation between sale propensity and performance since purchase

The negative relation between the likelihood of sale and past fund performance shown in Fig. 1 is consistent with tax-related motivations. In taxable environments (but not tax-deferred ones), a realization-based capital gains tax system provides incentives to sell investments that have fallen in price (“tax-loss selling”) and keep investments that have risen in price (the “lock-in” effect).

Another plausible explanation for the negative relation between the propensity to sell and performance since purchase is investors’ potential belief in fund performance persistence. If investors believe that funds’ past fund performance is indicative of their future performance, on the margin, they would be more likely to sell past losers and hold on to past winners. This should be equally true in taxable and tax-deferred (i.e., IRA and Keough) accounts.

In this section, we employ a hazard-model framework in a setting with numerous controls for fund characteristics to perform a detailed analysis of the negative relation found in the parsimonious estimates displayed in Fig. 1. We seek to differentiate between the alternative hypotheses for the

⁷A potential selection issue might arise because the sample consists of mutual fund trades placed by households that need not have both taxable and tax-deferred accounts. To address this concern, we run these analyses on a more restrictive sample of all mutual fund trades placed by the more than 17,000 households that have both types of accounts and find that the results are very similar to those reported in Table 2, a finding that is not surprising given all of the controls for heterogeneity of investors already included in the model.

negative relation by considering the results in taxable and tax-deferred settings. Finally, we also analyze the effects of other fund characteristics.

The results of estimating the model from Eq. (2) across all taxable mutual fund investments are presented in the first column of Table 2. Consistent with Fig. 1, evidence in support of a negative relation is very pronounced throughout the calendar year. The coefficient for NAV_RETURN is -0.78 , suggesting that, in months other than December, a 25% increase in NAV since purchase is associated with an 18% decrease in probability of sale (calculated as $e^{-0.78*0.25} - 1 = -0.18$). The coefficient for NAV_RETURN * December is large and negative (-1.21), indicating that tax-motivated trading is the most intense at the end of the year.

To gauge the economic significance of the redemption-performance relation, we perform a simple, back-of-the-envelope calculation. According to the Investment Company Institute (2004), over the past two decades the ratio of aggregate sales to aggregate assets across all equity mutual funds is roughly 30% on an annual basis. The estimate that a 25% increase in NAV since purchase reduces the probability of sale by 18% suggests that, in case of such a NAV increase, total assets under management would increase—very loosely—by about 5% ($30*0.18$) on an annual basis through the decreased redemptions, which constitutes a sizeable fraction of total assets under management.

The change in NAV since purchase, NAV_RETURN, by definition equals total returns since purchase minus total distributions since purchase. In an alternative specification, unreported for brevity, instead of including the NAV_RETURN, we include both total returns since purchase and distributions since purchase to check whether distinguishing between the two makes a difference. The coefficient on total returns since purchase is -0.78 (S.E. = 0.23, significant at the 1% level), and the coefficient on total distributions since purchase is 0.85 (S.E. = 0.50, significant at the 10% level). Their sum is statistically indistinguishable from zero (p -value = 0.90), implying that if, for example, the fund earns a 1% return, yet distributes it all to the investors (thus resulting in no change in the NAV), there is no net effect on the likelihood of sale. In other words, what matters to investors when deciding to sell mutual fund shares indeed is the change in “price” (NAV_RETURN), and not whether that price stems from higher/lower total returns since purchase or from lower/higher distributions since purchase.

However, this is not to say that past distributions do not affect sale decisions. Indeed, the results also imply that, *holding total returns constant*, the higher the distributions since the fund was bought, the higher the probability that the fund is sold (reflecting the mechanical reduction in NAV). Moreover, as we discuss in Section 3.3, long-term measures of past distribution policy (over the past five-year horizon), which likely signal future distribution policy, are also positively related to sale decisions.

Although the results we present up to this point establish a very robust negative relation between propensity to sell and performance since purchase in investors' taxable accounts, considering only taxable accounts does not enable us to disentangle the contributions of tax motivations from other factors potentially related to redemption decisions. Accordingly, we estimate a regression based on Eq. (3) that encapsulates trades in both taxable and tax-deferred accounts.

Redemption decisions in tax-deferred accounts are not related to the NAV change since purchase—the regression coefficient associated with NAV_RETURN in tax-deferred accounts for non-December months is small and statistically insignificant. Moreover, the differential between the coefficients associated with the two types of accounts is large and highly statistically significant. This suggests that the negative relation between the likelihood of redemption and past performance in taxable accounts is explained by tax-motivated trading. The lack of a relation in tax-deferred accounts suggests that, to whatever extent investor belief in fund performance persistence and the disposition effect are present in the domain of mutual fund investments, their net effect is such that they entirely offset each other (the negative relation resulting from belief in fund performance persistence cancels out the positive relation resulting from the disposition effect).

3.2. Effects of relative performance measures

Aside from the absolute performance of a fund since purchase (which is germane for tax purposes in case of redemption), also relevant for an investor may be the performance of that fund relative to funds pursuing the same investment objective. Indeed, investors are supplied routinely with the information regarding fund performance over certain investment horizons, as well as with ratings based on such performance, and may incorporate this information into their decision-making. The performance measures that we consider are the percentile ranking of recent one-year total returns within the investment objective (normalized to be between zero and one) and 5-star Morningstar rating, both of which are commonly used in the literature (see, e.g., Chevalier and Ellison, 1997; Sirri and Tufano, 1998; Bergstresser and Poterba, 2002).

Relative performance measures have no effects on the propensity to sell in either taxable or tax-deferred accounts (after all, absolute and not relative benchmarks are relevant for tax purposes). Thus, the sensitivity of *net* fund flows to *relative* performance rankings (e.g., Ippolito, 1992; Chevalier and Ellison, 1997; Sirri and Tufano, 1998) seems to be driven by inflows, rather than by redemption behavior (which we confirm in aggregated-flow analyses presented in Section 4).

3.3. Indirect tax considerations

Tax-sensitive investors might focus not only on the direct tax consequences related to the change of the fund NAV since purchase, but also on fund characteristics that could provide information

regarding future fund distribution policy because the tax rate on distributions received generally exceeds the tax rate on capital gains realized in the future upon sale. Having discussed the tax implications of NAV changes since purchase in Section 3.1, in this section we discuss the somewhat more subtle tax implications that may stem from fund managers' future distribution behavior.

In our empirical analyses we employ three proxies related to future distribution behavior. First, our specifications include fund turnover. According to Frazzini (2006), mutual fund managers exhibit behavior consistent with the disposition effect, that is, they are likely to sell the winners in their portfolios. Thus, turnover should be positively related to future distributions because there will be capital gains realized by such selling of the winners from the fund portfolio. Accordingly, on the margin, there should be a positive relation between the propensity to sell taxable mutual fund investments (and thereby avoid future distributions) and fund turnover.

Second, fund distribution policy might be highly persistent, in which case past distribution behavior might be indicative of future distributions. Thus, we construct our second proxy to reflect the fraction of total returns in the form of distributions over the past five years (months $t-61$ through $t-1$) preceding the month of potential sale t .⁸ Finally, the fund's capital gains overhang represents potential capital gains realizations that might be realized, depending on the fund manager's strategy and liquidity needs, in which case they would lead to future distributions and thus trigger a tax liability for the current taxable fund investments. Therefore, the relation between the propensity to sell taxable mutual fund investments and fund capital gains overhang should be positive.

Our results further reinforce the importance of tax-motivated behavior. In taxable accounts, the fund turnover ratio, the historical share of total fund returns distributed to the fund investors over the preceding five years, and the fund capital gains overhang are all positively related to the sale probability. By contrast, turnover does not play a role in sale decisions in tax-deferred accounts. Moreover, the relation between sale probability and historical distributions is also weaker in tax-deferred accounts. Somewhat curiously, the sensitivity of sale decisions to overhang is virtually identical across the two types of accounts. On net, the evidence suggests that direct (has the fund price gone up or down since purchase?) and more subtle, indirect (is the fund likely to pay out high future distributions?) tax motivations both play important roles in individuals' sale decisions.

Another way of interpreting these results is that both *past* and *future* distributions matter to investors. Both increase the probability that old money leaves the fund. Past distributions increase the probability of sale by mechanically reducing the NAV_RETURN, whereas proxies for future distributions increase the probability of sale because of likely higher associated future tax liabilities (with both of these effects much stronger in taxable accounts).

⁸It is defined as the ratio $(1 + \text{TOTAL_RETURN}_{t-61, t-1} - \text{NAV_RETURN}_{t-61, t-1}) / (1 + \text{TOTAL_RETURN}_{t-61, t-1})$.

3.4. *Effects of costs of mutual fund investment*

A priori, one might expect no relation between the propensity to sell and front-end charges (once fund shares are purchased, front-end charges are a sunk cost). On the other hand, expense ratios (costs that investors incur on a regular basis for as long as they hold the fund shares) and back-end loads are costs still ahead for mutual fund investors and they might alter the probability of sale. Higher expense ratios imply a stream of higher investment costs for as long as the investor owns the fund and thus, *ceteris paribus*, could be related positively with the probability of sale. By contrast, back-end loads can readily be conceived as deterrents to sale.

Barber, Odean, and Zheng (2005) consider the impact of front-end loads and expense ratios on individual investors' mutual fund investment decisions, but they limit their attention to the relation between *net* fund flows aggregated across a large number of individuals and lagged values of expense ratios and front-end loads, rather than on individuals' decisions to sell the mutual fund shares once they had acquired them. Barber, Odean, and Zheng (2005) report that net fund flows are sensitive to in-your-face costs such as front-end loads, yet are not sensitive to the more subtle, ongoing costs such as expense ratios.

To explore the impact of investment costs, we consider expense ratios charged by the funds at the time of potential sale, front-end charges that investors incurred at the time of purchase (expressed as the ratio between the fee charged to the investor and the size of the purchase), and an indicator variable denoting the presence of back-end loads for the respective fund investments at the time of potential sale.

First, the level of the expense ratio at the time of potential sale increases the likelihood that the investor will sell the mutual fund (effects are very similar across taxable and tax-deferred accounts). For example, compared to a fund with annual expenses of 50 basis points, a fund with annual expenses of 100 basis points is 11% more likely to be sold in taxable accounts ($e^{20.6*0.01}/e^{20.6*0.005} - 1 = 0.11$), suggesting that individual investors are sensitive to the ongoing, subtle costs of investments. This sensitivity is economically significant. Resorting once again to the data provided by the Investment Company Institute (2004) and performing another back-of-the-envelope calculation, such a 50-basis point increase in expense ratios loosely translates into a 3% decline ($30*0.11$) of total assets under management on an annual basis.

We also ran a specification in which we break the current expense ratio into the expense ratio at the time of purchase and the *change* in the expense ratio since purchase. In unreported results, we find that both are positively related to the propensity of sale and are statistically significant, with the coefficient on the change in the expense ratio of 54.6 (S.E. = 16.3) being significantly greater in magnitude than the coefficient for the expense ratio at the time of purchase of 19.2 (S.E. = 9.4).

Thus, investors who originally purchased a high-expense fund are more likely to sell that fund at any point in the future than are investors who purchased a low-expense fund, with investors responding particularly strongly if there was a change in the expense ratio since they made the purchase.

Second, investors appear to view front-end charges as an impediment to sale, potentially because they misperceive the front-end charge as a marginal rather than a sunk cost. The effect is more pronounced in taxable accounts, but both types of accounts feature a large and negative coefficient on the front-end load variable, suggesting that a front-end load of 5% reduces the monthly likelihood of sale by 75% in taxable accounts and 33% in tax-deferred accounts.

One might conjecture that this large effect simply reflects investor heterogeneity—households that invest in funds with front-end loads tend to have longer holding periods. However, the specification controls for considerable heterogeneity through investor-mutual fund type effects (non-parametric baselines). Thus, the correlation between front-end loads and sale decisions cannot simply be attributed, for example, to buy-and-hold investors purchasing funds with front-end loads. Rather, the regression results are identified by how a given individual trades two funds with the same objective, fund family, and degree of active management that have different front-end loads.

As a result, the front-end load effect likely does not merely reflect investor heterogeneity and may instead reflect investors' sunk cost fallacy (i.e., a confusion of sunk and marginal costs). Supporting this interpretation is survey evidence we obtained for a separate, ongoing project: nearly three-quarters of a random sample of 276 mutual fund investors that own funds with front-end loads report the need to hold the fund long enough to justify the front-end load; only one-quarter of the surveyed investors report that, *after* the fund had been purchased, the front-end load does not affect how long they hold on to the fund.

Finally, there is some evidence that investors are sensitive to back-end loads as well, but the results regarding the presence of a back-end load are not quite as strong as those for the front-end load. Whereas the negative and statistically significant coefficient for taxable accounts (-0.40 ; S.E. = 0.18) suggests that back-end loads serve as a deterrent to sales, this does not carry over to tax-deferred accounts, for which the coefficient is virtually zero (the difference between the coefficients across taxable and tax-deferred coefficients, though, is not statistically significant).⁹

TABLE 2 ABOUT HERE

⁹The results concerning loads are very robust. In addition to the specifications reported in Table 2, we explored alternative ways of capturing sensitivity to loads. First, we ran a specification in which both front-end loads and back-end loads are captured with indicator variables; we find investor sensitivity to both types of loads in this regression as well. Second, we ran a specification in which, in addition to front-end loads and back-end loads, we introduced an indicator variable that captures the absence of loads altogether (that is, it captures no-load funds). This additional no-load indicator variable was insignificant in its own right; moreover, it did not alter any of the results presented in Table 2 (the coefficients on the original load variables were essentially unchanged).

4. Determinants of monthly fund inflows and outflows

The preceding section reveals a very rich characterization of determinants of individuals' transaction-level mutual fund sale decisions. It shows that relative performance is unimportant for redemption decisions; the performance measure that does matter—consistent with tax motivations and present only in taxable accounts—is the absolute performance since purchase. These findings are delivered with a lot of precision because they are based upon very detailed, transaction-level data.

In this section, we turn our attention to individual investors' *fund-level* flows. We do so by aggregating all purchases and all redemptions of a fund in a month to compute dollar amounts of inflows and outflows, respectively. We seek to better understand how inflows and outflows, each possibly in a very different way, contribute to the overall net flow-performance relation previously studied in the literature (e.g., Ippolito, 1992; Chevalier and Ellison, 1997; Sirri and Tufano, 1998). We also analyze the effects of loads and expenses on old and new money.

Our specifications will encapsulate both absolute and relative measures of performance, as well as a range of other covariates (the same ones we employed in Section 3). Moreover, we also include three sets of fixed effects into the specifications to further alleviate concerns that omitted variables, endogeneity issues, or selection issues may be driving the results: time (monthly) fixed effects, objective/index fund fixed effects,¹⁰ and fund-specific fixed effects.

These steps advance the study of flows along two dimensions. First, we consider inflows and outflows separately, an approach pursued by very few studies to date. To our knowledge, the only studies that have done so in the domain of U.S. investments¹¹ are Edelen (1999), Bergstresser and Poterba (2002), O'Neal (2004), Cashman, Deli, Nardari, and Villupuram (2006), and Johnson (2007). Other than Johnson (2007), these studies rely upon SEC data regarding monthly fund-level purchases and redemptions, available in electronic form since the mid-1990s. The collection of these data is onerous and the researchers, with the exception of Cashman, Deli, Nardari, and Villupuram (2006), have limited their inquiry to a relatively small number of equity mutual funds. Johnson (2007) uses a different data set altogether—transactions of well over 50,000 shareholders in one small, no-load mutual fund family. The limitation of that data set is a very narrow representation of funds (up to ten funds) and thus limited cross-sectional variation of their characteristics.

Second, and more important, is the fact that, relative to previous work, we create superior opportunities for identification of competing explanations through the inclusion of a broader set of covariates (particularly those related to performance) and controls. Indeed, none of the

¹⁰Each investment objective is associated with up to two fixed effects—one for all actively managed funds with that investment objective, and the other for all index funds with that investment objective.

¹¹Data on inflows and outflows are available for U.K. mutual fund investors (Keswani and Stolin, 2008). Their use to date has been limited to assessing whether “money is smart,” with separate consideration of inflows and outflows.

aforementioned studies of inflow and outflow executes the natural step of considering both absolute and relative performance benchmarks, nor do they break out fund-level outflows into those generated by taxable investors and those generated by tax-deferred investors (a final step in our analyses that we employ to buttress our results further).

To obtain measures of flows, we compute the aggregate holdings of mutual funds in the sample at the end of each month and use them to scale the dollar inflows and outflows (both are non-negative by convention) over the next month as follows:

$$\text{Inflow}_{i,m+1} = \text{Buys}_{i,m+1} / \text{Positions}_{i,m}, \quad (4)$$

$$\text{Outflow}_{i,m+1} = \text{Sells}_{i,m+1} / \text{Positions}_{i,m}, \quad \text{and}$$

$$\text{Net Flow}_{i,m+1} = \text{Inflow}_{i,m+1} - \text{Outflow}_{i,m+1},$$

where $\text{Positions}_{i,m}$ is the total sum of all households' holdings of fund i at the end of month m , $\text{Buys}_{i,m+1}$ and $\text{Sells}_{i,m+1}$ are total sums of all sample households' purchases and sales, respectively, of fund i in month $m+1$. Finally, $\text{Inflow}_{i,m+1}$, $\text{Outflow}_{i,m+1}$, and $\text{Net Flow}_{i,m+1}$ are inflows, outflows, and net flows for fund i in month $m+1$, respectively, which we will often refer to as normalized flows.

A fund-month observation is admitted into the sample if at least five households held the fund at the end of the preceding month. In total, there are 18,038 fund-month observations for which we have both complete sample brokerage data and variables describing mutual fund characteristics from CRSP and Morningstar. The median number of households that hold a mutual fund at the end of a month is 32 (with an average of 99). The total number of funds appearing in the brokerage sample that we employ to compute flows is 529. Thus, our analysis is based on a fairly wide cross-section of mutual funds in existence at the time.

Our brokerage-level data provide an estimate of the aggregate inflows and outflows of individual investors. Given that individual investors hold about three-quarters of U.S. mutual fund assets (Investment Company Institute, 2004), their behavior in large part determines total mutual fund flows. However, because the normalized flows calculated from the brokerage sample (computed from investors' buys, sells, and positions) are only estimates of *actual* aggregate flows, this will add some imprecision, but not bias, to our regression coefficient estimates because we use these flow estimates as dependent variables in our regressions.

4.1. Flow-performance regressions: Inflows, outflows, and net flows

We relate monthly fund-flow variables to a range of covariates that correspond to those we employed in the previous section: one-year NAV returns preceding the month,¹² relative performance measures, proxies capturing future distribution behavior, and investment costs, as well as indicators for the date (month), indicators for the fund's objective and whether the fund is an index fund, and fixed effects for the individual funds themselves.

The results of relating fund-level flows to a range of covariates are presented in Table 3. The table consists of six columns; the first two columns are devoted to inflows, the next two to outflows, and the last two to net flows. For each flow measure, the first of the two columns features the specification in which one-year relative performance is represented with percentile rankings of one-year total returns within funds' investment objectives (normalized to be between zero and one). The second of the two columns, in light of the non-linearity of the well-known net flow-(relative) performance relation (e.g., Chevalier and Ellison, 1997; Sirri and Tufano, 1998), represents one-year relative performance through indicator variables, each denoting the decile of the percentile ranking of one-year total returns within the fund's investment objective (the omitted group is the combined 5th and 6th deciles). By construction, each coefficient in a column describing net flows is equal, modulo rounding errors, to the difference between the coefficients from the corresponding columns describing inflows and outflows.

The top portion of Table 3 focuses on performance measures. Consistent with the results from Section 3, fund-level outflows are sensitive only to absolute returns. As shown in the third and fourth columns, and as predicted by tax motivations, the relation is negative (the respective coefficients, both highly statistically significant, are -0.054 and -0.058). These coefficients are large in magnitude and are economically significant: a 10% increase in NAV over the past year decreases outflows over the next month by 0.5% of assets.

We find no relation between fund-level outflows and relative performance (in accordance with our results for individual sale decisions discussed in Section 3). This result is also in agreement with Bergstresser and Poterba (2002) and Johnson (2007), who find that outflows are not related to relative short-term performance. Our overall conclusions regarding fund outflows, however, differ from those outlined in Johnson (2007), who concludes that mutual fund share sales are idiosyncratic and are based on investors' liquidity needs. We, on the other hand, consider both absolute and relative performance since purchase and, consistent with tax motivations, find a strong and robust

¹²In the context of the individual-level transaction analyses from Section 3, it was natural to focus on holding-period returns. In the present analysis of fund-level inflows and outflows, there simply is no direct equivalent of holding-period returns for aggregated purchases. Therefore, we instead focus on the past one-year returns and thereby facilitate direct comparison across net flows, inflows, and outflows.

negative relation between the sale decision and absolute fund performance in investors' taxable accounts both at the level of individuals' mutual fund share selling decisions (Section 3) and at the monthly fund-level aggregation of outflows.

By contrast, inflows are sensitive only to a relative measure of performance—one-year objective rank—with no relation to absolute performance over the past year. The corresponding coefficient associated with the one-year objective rank in the first column is 0.037. The second column shows the pronounced non-linearity of that relation—the effect appears to be concentrated in the top decile of relative performance, wherein being in the top decile of performers in an objective is associated with an inflow that is 3.3% higher than the inflow into an otherwise identical fund with relative performance in the 9th decile. These results provide direct evidence of the notion that new money chases past fund relative performance.

Net flows aggregate these tendencies in ways consistent with the findings reported in the extant literature. They are highly sensitive to relative performance. Moreover, the non-linearity of the relation with relative performance, displayed in the sixth column, is very pronounced. As for absolute performance, whereas the point estimate of the relation between net flows and past one-year returns is positive, it is not statistically significant. What is obfuscated by focusing just on net flows, however, is that for a major subset of investors—the existing investors considering whether to sell or continue to hold the fund—there is a strong sensitivity between flows and absolute returns.

The coefficients on variables characterizing future distribution behavior of the fund (turnover, fraction of total returns distributed over the past five years, and capital gains overhang) are largely insignificant. On the outflow side, only overhang is statistically significant—as in the transaction-level analyses from Section 3, its relation with outflows is positive.

Finally, breaking net flows into their two components yields insights regarding the effect of investment costs as well. In contrast with Barber, Odean, and Zheng (2005), our specifications uncover sensitivity of net flows to expense ratios; moreover, both inflows *and* outflows are positively related with expense ratios, with the relation being more pronounced for inflows. It appears that there are two countervailing phenomena at play that, when differenced for net flows, still yield an overall positive relation. Specifically, the effects of marketing (or potentially a perceived signal of managerial quality and/or the service provided by the mutual fund family embedded in high expense ratios) appear to be strong when attracting new buys and, at the same time, consistent with the results from Section 3: existing investors are more likely to sell mutual funds with high expense ratios.

The economic impact of the expense ratio is fairly strong. A shift in the expense ratio equal to the interquartile range of 0.6% increases the monthly normalized outflows by 6.9% ($11.5 \times 0.006 = 0.069$), whereas it increases the inflows about twice as dramatically—by 13.6% ($22.7 \times 0.006 = 0.0136$).

One potential explanation for the positive relation between inflows and the expense ratio is that, consistent with the Berk and Green (2004) framework, the expense ratio (particularly the management fee component) reflects the underlying quality of the fund manager and, thus, the positive inflow-expense relation really reflects an underlying positive relation between inflows and past fund performance/quality (that is also picked up by the expense ratio coefficient). To differentiate between this hypothesis and the “advertising” explanation for the inflow-expense relation, in unreported analyses we break the expense ratio into the 12(b)-1 fee (which, to a large extent, represents fund advertising) and the non-12(b)-1 or “other” component (which, to a large extent, represents the fund management fee). If advertising attracts new money into the fund, the coefficient on the 12(b)-1 fee should be positive. If the quality of the manager attracts new money into the fund and, according to the Berk and Green (2004) framework, quality funds with good past performance have raised management fees to extract rents, the coefficient on the “other” component of the expense ratio should be positive. In the regression with the expense ratio separated into its two components, the coefficient on the 12(b)-1 fee is 20.5 and the coefficient on the “other” component is 22.9 (both coefficients are highly statistically significant), providing support for both hypotheses.

For both front-end loads and back-end loads, the results for the outflows are consistent with the transaction-level results from Section 3. For front-end loads, as discussed earlier, the negative relation between the outflows and the front-end load could reflect a sunk cost fallacy. The insignificant relation between net flows and back-end loads clouds the offsetting, large, and statistically significant sensitivities across new investors (inflows) and old investors (outflows). Back-end loads are associated with less new money entering the fund; at the same time, they are also associated with the existing money being less prone to leave.

In sum, many of the rich characterizations of investor sensitivities to various fund characteristics discussed above are obscured when inflows and outflows are combined and only net flows are studied. Our results stress that the absence of a relation between net flows and a fund characteristic does not imply that the characteristic is unimportant for *all* investors. Indeed, given that new or potential investors in the fund and the incumbent investors may have different considerations, disentangling net flows into its components is important. Thus, when assessing the overall effect of a fund-policy change, it is relevant to understand the differing sensitivities of new or potential investors and existing investors who already own shares.

TABLE 3 ABOUT HERE

4.2. Robustness check: Separating outflows into outflows in taxable and tax-deferred accounts

This section features the last robustness check. Results concerning the relation between outflow and performance, presented in detail in Section 4.1, suggest that the relation between outflows and performance is best characterized by a pronounced role of tax motivations for trade. Because our data enable us to disentangle outflows from taxable and tax-deferred accounts, we can assess this hypothesis even more directly—if tax motivations are driving the result for outflows, then absolute performance should matter only in the context of outflows from taxable accounts. We carry out this robustness check and report the results in Table 4. For brevity, we report only the regression coefficients related to absolute and relative one-year performance.

Table 4 is divided into two panels. Panel A employs a simpler, linear model of relative performance, and Panel B employs a non-linear representation of relative performance, as was done in Table 3. The first column of each panel features the regression results estimated over the fund outflows aggregated over all accounts. Thus, the first column in Panel A (Panel B) corresponds to the third (fourth) column in Table 3. The next two columns in each panel feature regression results estimated for the outflows computed from taxable accounts only and tax-deferred accounts only, respectively, and the last column in each panel features the difference between the two.

This robustness test of the hypothesis that tax motivations drive the relation between outflows and absolute performance affirms that interpretation very strongly. As in Table 3, there is no “action” in the domain of relative performance in either taxable or tax-deferred accounts. This is true for both the linear specifications (Panel A) and the nonlinear specifications (Panel B).¹³ Absolute performance, the central theme of this section, displays the predicted pattern: it is substantially larger in magnitude for outflows from taxable accounts (-0.293 or -0.295 , depending upon the specification) than for outflows from tax-deferred accounts (0.020 or 0.018 , depending upon the specification), and is statistically significant only for the taxable accounts.

Although not reported in the table for brevity, we also test whether the other variables in the outflow regressions (the proxies for future fund distributions and the mutual fund costs) have statistically different effects for taxable outflows than they do for tax-deferred outflows. There are no statistically significant differences across the coefficients in the taxable and tax-deferred regressions, with the exception of three variables—the past one-year absolute return (as discussed above), turnover, and overhang. Both the relation between outflows and turnover and the relation between outflows and overhang are stronger for taxable outflows (with both differences statistically

¹³Among the results of the non-linear specifications, one relative-performance coefficient—associated with the bottom decile in tax-deferred accounts—does attain statistical significance, but the p -value of the joint significance of all of the relative rank deciles is 0.27.

significant at the 5% level). This is consistent with these proxies for future fund distributions affecting fund outflows because they may carry with them higher tax liabilities.

TABLE 4 ABOUT HERE

5. Conclusion

This paper studies the determinants of mutual fund flows, with particular attention to individual investors' mutual fund selling decisions. In stark contrast with investor behavior regarding common stocks, there is a strong negative relation between the probability of sale and past mutual fund performance. Individuals hold on to mutual fund shares that have appreciated since purchase and are willing to sell those that have incurred losses. By comparing trading patterns in taxable and tax-deferred accounts, we confirm that the negative relation can be explained by tax-motivated trading.

Among a mutual fund's choice variables are how returns should be distributed to investors, with implications for taxable investments, and what kinds of fees—expenses and loads—should be charged. Both have implications for outflows from the fund. The evidence of tax-motivated trading in taxable accounts encompasses not only the negative relation between propensity to sell and NAV change since purchase, but also increased propensity to sell in response to three other variables that signal potential future tax liabilities—fund turnover, past distribution behavior of the fund, and capital gains overhang. Thus, individuals seem to be sensitive not only to *direct* tax considerations such as what taxes would be incurred in the event of selling the mutual fund, but to *indirect* tax considerations such as future distributions.

In regard to investment costs, somewhat surprisingly, investors from our sample seem to view front-end charges as impediments to sale. On the other hand, high expense ratios (an ongoing, not-in-your-face cost) are associated with a substantial increase in the probability of redemption, particularly if the expense ratio has increased since the investor purchased the fund.

Finally, we also use the brokerage data to aggregate individuals' purchases and redemptions into fund-level monthly measures of inflows and outflows and study the determinants of flows for inflows, outflows, and net flows. We find that, broadly speaking, inflows are driven by funds' relative performance measures, that is, funds' one-year performance relative to other funds pursuing the same objective. This suggests that new money indeed chases the best performers in an objective, with no relation to the fund's "absolute" performance. On the other hand, consistent with tax motivation and pronounced only in taxable accounts, outflows are driven by funds' one-year absolute returns. Thus, breaking net flows into their two components yields important insights on the behavior of old and new money.

References

- Barber, B., Odean, T., 2000. Trading is hazardous to your wealth: The common stock investment performance of individual investors. *Journal of Finance* 55, 773-806.
- Barber, B., Odean, T., Zheng, L., 2005. Out of sight, out of mind: The effects of expenses on mutual fund flows. *Journal of Business* 78, 2095-2119.
- Bergstresser, D., Poterba, J., 2002. Do after-tax returns affect mutual fund inflows. *Journal of Financial Economics* 63, 381-414.
- Berk, J., Green, R., 2004. Mutual fund flows and performance in rational markets. *Journal of Political Economy* 112, 1269-1295.
- Cashman, G., Deli, D., Nardari, F., Villupuram, S., 2006. Investors do respond to poor mutual fund performance: Evidence from inflows and outflows. Working paper. Arizona State University.
- Chevalier, J., Ellison, G., 1997. Risk taking by mutual funds as a response to incentives. *Journal of Political Economy* 105, 1167-1200.
- Edelen, R., 1999. Investor flows and the assessed performance of open-end mutual funds. *Journal of Financial Economics* 53, 439-466.
- Frazzini, A., 2006. The disposition effect and under-reaction to news. *Journal of Finance* 61, 2017-2046.
- Grinblatt, M., Keloharju, M., 2001. What makes investors trade? *Journal of Finance* 56, 589-616.
- Han, A., Hausman, J., 1990. Flexible parametric estimation of duration and competing risk models. *Journal of Applied Econometrics* 5, 1-28.
- Investment Company Institute, 2004. *Mutual fund factbook: A guide to trends and statistics in the mutual fund industry*, 43rd ed.
- Ippolito, R., 1992. Consumer reaction to measures of poor quality: Evidence from the mutual fund industry. *Journal of Law and Economics* 35, 45-70.
- Ivković, Z., Poterba, J., Weisbenner, S., 2005. Tax-motivated trading by individual investors. *American Economic Review* 95, 1605-1630.
- Johnson, W., 2007. Who monitors the mutual fund manager, new or old shareholders? Working paper. University of Oregon.
- Kahneman, D., Tversky, A., 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47, 263-291.
- Keswani, A., Stolin, D., 2008. Which money is smart? Mutual fund buys and sells of individual and institutional investors. *Journal of Finance* 63, 85-118.
- Meyer, B., 1990. Unemployment insurance and unemployment spells. *Econometrica* 58, 757-782.
- Odean, T., 1998. Are investors reluctant to realize their losses? *Journal of Finance* 53, 1775-1798.

- O'Neal, E., 2004. Purchase and redemption patterns of US equity mutual funds. *Financial Management* 33, 63-90.
- Shefrin, H., Statman, M., 1985. The disposition to sell winners too early and ride losers too long: Theory and evidence. *Journal of Finance* 40, 777-792.
- Sirri, E., Tufano, P., 1998. Costly search and mutual fund flows. *Journal of Finance* 53, 1589-1622.

Table 1

Summary statistics of mutual fund purchases and sales

The sample consists of 32,259 households that made at least one mutual fund purchase through the large discount broker during the period from January 1991 to November 1996. This table presents basic summary statistics (median dollar amount of purchase and number of purchases are reported in parentheses).

	Number of purchases	Average \$ amount of purchases (Median)	Average # of purchases per household, conditional on purchase in that type of account (Median)	Percentage of purchases sold during the sample period
All accounts	325,185	8,394 (3,000)	10.1 (4.0)	34
Taxable accounts	180,564	9,376 (3,000)	8.5 (3.0)	33
Tax-deferred accounts	144,621	7,169 (3,000)	7.2 (3.0)	35

Table 2

Relation between redemption decisions and fund characteristics

The Cox proportional hazards model employs a non-parametric estimate of the baseline hazard (i.e., the probability of selling the mutual fund in month t after the buy conditional on no prior sale). The model features separate non-parametric baseline hazard rates for each investor-mutual fund type combination, introduced separately for an investor's holdings in taxable and tax-deferred accounts. The fund type is defined by fund objective, degree of active management (index funds versus actively managed funds), and fund family membership. NAV_RETURN is defined as the relative change in NAV since purchase. December is an indicator variable. The model also incorporates relative performance measures (funds' percentile ranking of recent one-year total returns within the investment objective, normalized to be between zero and one, and Morningstar rating), as well as proxies of future fund distribution behavior (fund turnover, the fraction of total fund return over the past five years distributed to the investors, and fund capital gains overhang). Finally, the model also includes the expense ratio charged by the fund at the time of potential sale, the front-end charge incurred at the time of purchase (the ratio between the fee charged to the investor and the size of the purchase), and an indicator variable denoting the presence of back-end loads for the respective fund investments at the time of potential sale. The estimation is based on 325,185 buys made over the sample period from 1991 to 1996, representing 32,259 households that made at least one mutual fund purchase through the large discount broker. Standard errors (shown in parentheses) allow for heteroskedasticity as well as correlation across observations associated with the same transaction.

	Taxable accounts	All accounts	
		<i>Tax-deferred accounts</i>	Interaction w/ taxable
NAV_RETURN	-0.78 ^{***} (0.22)	0.21 (0.20)	-0.99 ^{***} (0.29)
NAV_RETURN*December	-1.21 [*] (0.70)	-0.95 (0.73)	-0.26 (1.00)
December	-0.11 [*] (0.06)	-0.01 (0.07)	-0.10 (0.09)
One-year objective rank	-0.10 (0.08)	-0.11 (0.08)	0.01 (0.12)
Morningstar rating	0.03 (0.03)	-0.01 (0.03)	0.05 (0.04)
Turnover	0.14 ^{**} (0.04)	-0.02 (0.03)	0.17 ^{***} (0.05)
Fraction of total returns distributed over past five years	0.66 ^{***} (0.12)	0.29 ^{**} (0.13)	0.37 ^{**} (0.17)
Overhang	0.36 ^{**} (0.17)	0.40 ^{**} (0.16)	-0.04 (0.24)
Expense ratio at time of potential sale	20.6 ^{**} (9.4)	14.2 ^{**} (7.1)	6.4 (12.4)
Front-end load (Fee charged normalized by purchase amount)	-27.4 ^{***} (4.4)	-7.9 ^{***} (2.7)	-19.5 ^{***} (5.2)
Back-end load?	-0.40 ^{**} (0.18)	0.01 (0.19)	-0.41 (0.26)
Number of observations	1,529,740	3,038,204	

***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 3

Aggregated flow-performance regression

The table presents regression results of relating monthly fund inflows, outflows, and net flows from the brokerage sample to absolute fund performance (expressed as the NAV return of the fund over the past year), its interaction with the December indicator variable, relative performance measures (percentile ranking of performance over the past year within the investment objective, normalized to be between zero and one, and Morningstar rating), proxies for future fund distribution behavior (turnover, fraction of total returns distributed over the past five years, and overhang), and investment costs (expense ratio, front-end load, and back-end load), as well as indicator variables that capture time (month) effects, indicator variables for the fund's objective and whether it is an index fund, and fixed effects for the individual funds. Inflows, Outflows, and Net Flows are computed according to Eq. (4) from Section 4—the aggregated flows of individual investors for a given month (in the sample provided by a large discount broker in the period from 1991 to 1996) are normalized by total holdings in the fund in the brokerage sample at the end of the prior month. For each of the flows, the first column features, among other controls, the percentile ranking of performance over the past year within the investment objective, whereas the second column features a non-linear representation of that ranking, expressed through indicator variables capturing its deciles (with the 5th and 6th deciles omitted). Standard errors are shown in parentheses.

Table 3: Continued

	Inflow		Outflow		Net Flow	
One-year NAV return	-0.019 (0.035)	-0.027 (0.037)	-0.054** (0.026)	-0.058** (0.028)	0.036 (0.025)	0.031 (0.025)
One-year NAV ret. x Dec.	-0.041 (0.105)	-0.043 (0.100)	-0.027 (0.038)	-0.026 (0.038)	-0.014 (0.099)	-0.017 (0.099)
One-year objective rank	0.037*** (0.012)		-0.014 (0.009)		0.051*** (0.009)	
One-year objective rank bottom decile ^a		-0.008 (0.011)		0.009 (0.008)		-0.017** (0.008)
One-year objective rank decile 2 ^a		-0.015 (0.009)		-0.007 (0.006)		-0.008 (0.008)
One-year objective rank decile 3 ^a		0.014 (0.017)		0.003 (0.008)		0.011 (0.014)
One-year objective rank decile 4 ^a		-0.009 (0.007)		-0.001 (0.005)		-0.008 (0.006)
One-year objective rank decile 7 ^a		0.009 (0.011)		-0.006 (0.005)		0.015 (0.010)
One-year objective rank decile 8 ^a		0.000 (0.008)		-0.005 (0.006)		0.006 (0.006)
One-year objective rank decile 9 ^a		0.007 (0.007)		-0.008* (0.005)		0.016*** (0.006)
One-year objective rank top decile ^a		0.040*** (0.009)		-0.004 (0.005)		0.044*** (0.009)
<i>p</i> -value for joint sig. of rank deciles		0.00***		0.37		0.00***
Morningstar rating	0.007 (0.005)	0.006 (0.004)	-0.000 (0.003)	-0.001 (0.003)	0.007** (0.003)	0.007** (0.003)
Turnover	0.007 (0.006)	0.006 (0.006)	-0.004 (0.004)	-0.004 (0.004)	0.011** (0.005)	0.010** (0.005)
Fraction of total returns distributed over past 5 years	0.033 (0.064)	0.029 (0.064)	-0.022 (0.043)	-0.024 (0.043)	0.055 (0.37)	0.053 (0.37)
Overhang	0.010 (0.010)	0.011 (0.010)	0.012** (0.006)	0.012** (0.006)	-0.002 (0.007)	-0.000 (0.007)
Expense ratio	22.7*** (4.9)	22.9*** (4.9)	11.5*** (3.3)	11.6*** (3.3)	11.2*** (3.0)	11.3*** (3.0)
Front-end load	-0.21 (0.61)	-0.25 (0.60)	-0.53** (0.27)	-0.52** (0.26)	0.32 (0.55)	0.27 (0.54)
Back-end load	-4.44* (2.34)	-4.28* (2.34)	-2.77** (1.16)	-2.74** (1.16)	-1.67 (1.98)	-1.55 (1.98)
Time (month) fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes
Objective & index fund f. e.?	Yes	Yes	Yes	Yes	Yes	Yes
Fund-specific fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.091	0.091	0.166	0.166	0.020	0.021
Number of observations	18,038	18,038	18,038	18,038	18,038	18,038

***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

^aThe omitted categories are deciles 5 and 6.

Table 4

Differences by tax status of outflows

The table presents regression results of relating monthly fund outflows from the brokerage sample to absolute fund performance (expressed as the NAV return of the fund over the past year), its interaction with the December indicator variable, relative performance measures (percentile ranking of performance over the past year within the investment objective normalized to be between zero and one, and Morningstar rating), proxies for future fund distribution behavior (turnover, fraction of total returns distributed over the past five years, and overhang), and investment costs (expense ratio, front-end load, and back-end load), as well as indicator variables that capture time (month) effects, indicator variables for the fund's objective and whether it is an index fund, and fixed effects for the individual funds. Outflows are computed according to Eq. (4) from Section 4—the aggregated flows of individual investors for a given month (in the sample provided by a large discount broker in the period from 1991 to 1996) are normalized by total holdings in the fund in the brokerage sample at the end of the prior month. Aside from the outflows computed across all accounts, we also compute outflows from taxable and tax-deferred accounts. Panel A features the percentile ranking of performance over the past year within the investment objective, whereas Panel B features a non-linear representation of that ranking, expressed through indicator variables capturing its deciles (with the 5th and 6th deciles omitted). Within each panel, the first column features outflows from all accounts (and is thus a restatement of the results from Table 3), the next two columns focus on outflows from taxable and tax-deferred accounts, respectively, and the last column reports the differences between the two. Standard errors are shown in parentheses.

Table 4: Continued

	<i>Panel A: Linear model of relative performance</i>				<i>Panel B: Non-linear model of relative performance</i>			
	All accounts combined	Separate consideration of taxable and tax-deferred accounts			All accounts combined	Separate consideration of taxable and tax-deferred accounts		
		Taxable accounts	Tax-deferred accounts	Difference		Taxable accounts	Tax-deferred accounts	Difference
One-year NAV return	-0.054** (0.026)	-0.293*** (0.089)	0.020 (0.014)	-0.313*** (0.092)	-0.058** (0.028)	-0.295*** (0.088)	0.018 (0.014)	-0.313*** (0.090)
One-year NAV ret. x Dec.	-0.027 (0.038)	-0.048 (0.119)	-0.009 (0.032)	-0.039 (0.123)	-0.026 (0.038)	-0.047 (0.119)	-0.009 (0.033)	-0.038 (0.124)
One-year objective rank	-0.014 (0.009)	-0.029 (0.019)	-0.003 (0.007)	-0.026 (0.019)				
One-year objective rank bottom decile ^a					0.009 (0.008)	0.019 (0.021)	0.016** (0.007)	0.002 (0.022)
One-year objective rank decile 2 ^a					-0.007 (0.006)	0.009 (0.013)	0.006 (0.005)	0.003 (0.014)
One-year objective rank decile 3 ^a					0.003 (0.008)	-0.001 (0.010)	0.006 (0.008)	-0.006 (0.013)
One-year objective rank decile 4 ^a					-0.001 (0.005)	-0.001 (0.010)	0.010 (0.007)	-0.011 (0.011)
One-year objective rank decile 7 ^a					-0.006 (0.005)	-0.003 (0.011)	-0.004 (0.004)	0.007 (0.012)
One-year objective rank decile 8 ^a					-0.005 (0.006)	-0.007 (0.007)	0.004 (0.004)	-0.010 (0.008)
One-year objective rank decile 9 ^a					-0.008* (0.005)	-0.012 (0.009)	0.006 (0.004)	-0.019 (0.012)
One-year objective rank top decile ^a					-0.004 (0.005)	-0.012 (0.010)	0.010 (0.007)	-0.022 (0.016)
<i>p</i> -value for joint sig. of rank deciles					0.37	0.48	0.27	0.32
Other controls from Table 3?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted <i>R</i> ²	0.166	0.218	0.080	0.088	0.166	0.080	0.088	0.190
Number of observations	18,038	15,136	14,044	29,180	18,038	15,136	14,044	29,180

***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

^aThe omitted categories are deciles 5 and 6.




Fig. 1. Hazard rates and the associated 95% confidence intervals of selling mutual funds in taxable accounts. This figure displays the average hazard rate for mutual fund share purchases conditional on whether the investors' fund has an accrued capital gain since purchase (gray line) or loss since purchase (black line) entering the month. The figure restricts attention to the observations initiated by January fund purchases in taxable accounts, a subset of the sample of 325,185 purchases that 32,359 households made through the large discount broker in the period from 1991 to 1996 that were .