



The effect of taxes on shareholder inflows around mutual fund distribution dates



Woodrow T. Johnson^a, James M. Poterba^{b,*}

^a U.S. Securities and Exchange Commission, United States

^b MIT and NBER, United States

ARTICLE INFO

Article history:

Received 8 September 2015

Accepted 24 October 2015

Available online 14 November 2015

Keywords:

Mutual funds

After-tax returns

Capital gains taxation

Taxable distributions

ABSTRACT

Taxable investors who are considering purchasing mutual fund shares around the dates when a mutual fund is planning a taxable distribution can reduce the present discounted value of their tax liability by delaying their purchase until after the distribution date. Non-taxable shareholders, such as those who invest through IRAs and other tax-deferred accounts, face no such incentive for delaying a purchase of the fund. This paper compares daily shareholder transactions by taxable and non-taxable investors in the mutual funds of a single no-load fund complex around distribution dates. Gross inflows to taxable accounts are significantly lower in the weeks preceding distribution dates than in the weeks following them, but gross inflows to tax-deferred accounts do not change around these dates. This finding suggests that some taxable shareholders time their purchase of mutual fund shares to avoid the tax acceleration associated with distributions. Taxable shareholders who purchase shares just before distribution dates also have shorter holding periods, on average, than those who buy just after a distribution. Since the cost of the distribution-related tax acceleration for pre-distribution buyers is related to the expected holding period of the shares, this finding provides some evidence of clientele formation among the buyers of mutual fund shares.

© 2015 University of Venice. Published by Elsevier Ltd. All rights reserved.

A voluminous literature has examined the impact of dividend taxes and capital gains taxes on the behavior of common stock investors, with particular emphasis on how these taxes affect portfolio composition and asset trading decisions. However, even though data from the [Investment Company Institute \(2015\)](#) suggest that 30.3 million (24 percent of) U.S. households own mutual funds outside tax-deferred retirement accounts, relatively little is known about how taxes affect the behavior of investors in intermediated investment vehicles such as mutual funds. The taxes associated with mutual fund ownership outside retirement accounts can be substantial. One study nearly a decade ago, [Lipper Associates \(2007\)](#), estimates that in 2006, open end mutual funds distributed \$234 billion in long-term capital gains, \$31 billion in short-term gains, and \$154 billion in taxable dividends, thereby generating at least \$24 billion in investor tax liability. Taxes on capital gain distributions accounted for nearly \$14 billion of this total. Total capital gain distributions vary from year to year, but these statistics indicate that they can be substantial and that they could have a non-trivial impact on investors' after-tax returns.

Mutual fund investors are taxed under a specialized set of tax rules. All fund shareholders as of a given date share in a fund distribution in proportion to their fund ownership, regardless of their holding period or whether the fund's share price

* Corresponding author at: MIT, 77 Massachusetts Avenue E17-214, Cambridge MA 02139.

E-mail address: poterba@mit.edu (J.M. Poterba).

has increased or decreased since their date of purchase. This tax provision raises the possibility that a fund investor who purchases fund shares just before a taxable distribution will be liable for capital gains taxes even though the value of her fund shares may have fallen since she purchased them. Financial advisors caution taxable investors against buying mutual fund shares just before distribution dates, particularly as the year-end distribution season approaches.

Several legislative proposals, for example [Saxton \(2004\)](#), have called for modifying the tax rules governing mutual fund distributions to allow deferral of the tax liability on distributed gains until the investor sells her fund shares. Assessing how these proposals would affect investor behavior requires information on several aspects of the mutual fund market, including the current effect of distribution taxes on the purchase and sale decisions of fund investors.

Tax return data leave little doubt that taxable investors report and pay taxes on billions of dollars in mutual fund distributions, but there have been relatively few studies of how taxes affect the behavior of mutual fund investors. [Barclay et al. \(1998\)](#) were among the first to focus on capital gains “overhang” in open end mutual funds. They found that funds with substantial embedded unrealized capital gains, which place investors at greatest risk for future capital gain distributions, attract smaller net inflows than comparable funds without a capital gains overhang. [Bergstresser and Poterba \(2002\)](#) examined gross inflows and gross outflows from equity mutual funds in an effort to understand whether investors consider the tax burden on fund returns in making their investment decisions. They found that open-end funds with high historical tax burdens on taxable shareholders attract smaller inflows than funds with comparable before-tax return experiences but lower shareholder tax burdens.

Several studies have suggested that fund managers consider the implications of their asset management decisions for taxable investors. [Planchich \(2003\)](#) and [Chen et al. \(2011\)](#) found that the payment of fund-level distributions changes around major tax law changes such as the Taxpayer Relief Act of 1997. [Khorana and Servaes \(1999\)](#) found that fund management companies responded to the accumulation of capital gains “overhang” in a fund by creating new funds that mimic existing funds, but lack the embedded capital gains that may discourage taxable investors from buying the fund's shares. [Christoffersen et al. \(2005\)](#) observed that the investing public displays substantial tax rate heterogeneity, which complicates the challenge facing managers who try to consider their investors' tax burdens. [Sialm and Starks \(2012\)](#) examined mutual funds that participate in the 401(k) market, and they found that funds with a greater share of their holdings through 401(k) accounts were less tax-efficient than funds that relied more heavily on taxable investors.

This paper investigates whether taxes affect the timing of mutual fund purchases around the dates when funds make taxable distributions. We exploit a proprietary six-year panel database that includes all trades within and across all funds in one actively managed no-load mutual fund family. We compare trades made through tax-deferred accounts—primarily traditional and Roth IRAs—with trades made through taxable accounts. While we cannot be certain of the marginal tax rate that applies to distributions received by the funds' taxable shareholders, we know that distributions paid to tax-deferred accounts are untaxed. We find that taxable shareholder inflow is 27 percent lower in the four weeks before a distribution date than in the four weeks following the distribution. By comparison, we cannot reject the null hypothesis that inflows from investors who hold fund shares in tax-deferred accounts are equal before and after a distribution date. These findings suggest that at least some taxable shareholders are aware of impending distributions and time their purchases to reduce their tax impact. These findings offer new evidence on how taxes affect the behavior of mutual fund shareholders.

Although we observe a modest decline in fund inflows just before the distribution date, we still find many taxable investors purchasing fund shares before distributions. These investors may regard the potential near-term change in the fund's net asset value (NAV) as large enough to warrant purchasing the fund, and the associated tax payment acceleration, before the distribution. We cannot observe investors' expectations of future returns, but we can measure investors' ex-post holding periods. We find that taxable fund investors who purchase shares just prior to distribution dates on average liquidate their accounts slightly sooner than those who purchase just after distribution dates. Because the cost of tax acceleration is increasing in the taxable investor's holding period, this finding suggests that pre-distribution buyers may have a smaller incentive than post-distribution buyers to defer their purchase.

Our study is similar in spirit to [Graham and Kumar's \(2006\)](#) analysis of how individual investor characteristics are correlated with trading decisions around ex-dividend days for common stocks. They found that older and lower income investors are more likely to purchase stocks before ex-dividend days than afterward, and they suggested that this could potentially be explained by tax considerations. They could not test this proposition directly, since their data set did not include any information on investors' tax circumstances. Our comparison of trading in taxable and tax-deferred accounts offers a more powerful test of how tax status affects trading around the dates of record that determine tax liability.

Our paper is divided into five sections. The first describes the tax treatment of mutual fund shareholders and explains how a mutual fund capital gain distribution accelerates capital gains tax liability. Section two describes the data set that we analyze and presents summary statistics. The third section presents our findings on the differences between gross inflows before and after fund distributions. It reports results for both taxable and tax-deferred households, and it contrasts patterns before and after the enactment of the Taxpayer Relief Act of 1997, which altered the marginal tax rates on capital gains. Section three also presents some tests for turn-of-year effects. The fourth section fits proportional hazard models for the risk that a shareholder decides to close her account, contrasting shareholders who open accounts just prior to distribution dates with those who open accounts just afterward. The last section summarizes our findings and outlines directions for future work.

1. Tax treatment of capital gains on mutual fund holdings

Investments in open-end mutual funds are taxed under a specialized set of income tax rules. One of the most important is the requirement that in order to avoid taxation as corporations, funds must “pass through” dividends and capital gains to their shareholders. Dickson et al. (2000) describe how the regulations affecting pass-through of realized capital gains from funds to their shareholders can accelerate capital gains tax liabilities for buy-and-hold fund shareholders. They also illustrate how the current tax rules create “externalities” by making the capital gains tax liability of each fund shareholder dependent on the redemption decisions of all other shareholders. When a fund manager realizes capital gains by selling portfolio holdings, the fund must make a taxable capital gain distribution which creates an immediate capital gains tax liability even for buy-and-hold investors. Such distributions do not increase the shareholders’ nominal capital gains tax liability over their entire holding period, but they increase the present value of this liability by accelerating the tax payments relative to what they would have been if the fund investment had been taxed like a common stock. Investors in mutual funds more generally lose the opportunity to apply tax-timing strategies of the type described by Constantinides (1984) in his analysis of portfolios consisting of directly-held common stock.

To quantify the tax consequences of deferring a mutual fund investment until after a taxable capital gain distribution, consider a shareholder who is planning to hold mutual fund shares for exactly 20 years. Assume that the nominal annual 20-year interest rate is $100 \cdot r_{20}$ percent, that the impending distribution equals $100 \cdot d$ percent of the fund’s net assets, and that the current and future tax rate on realized capital gains is τ_{cg} . Set the fund’s pre-distribution share price to unity. If the shareholder invests one dollar immediately prior to the capital gain distribution, she purchases one share of the fund, receives a capital gain distribution of d , and pays tax of $\tau_{cg} \cdot d$.

On the distribution date, the fund’s share price drops from 1 to $1 - d$. If the shareholder reinvests the distribution of d and draws on other resources to pay her capital gains tax bill, her tax basis in the fund will be $1 + d$. She will own $1 + d / (1 - d) = 1 / (1 - d)$ shares of the fund. Since shares cost $1 - d$ dollars, the shareholder once again has one dollar invested in the fund. If the fund’s assets appreciate at an annual rate of g , after 20 years the shareholder’s position will be worth $(1 + g)^{20}$. Her capital gains tax due at the time of sale will equal $\tau_{cg} \cdot [(1 + g)^{20} - (1 + d)]$. The present value of the tax due on the mutual fund purchase and sale, combining the tax due on the initial distribution and that due when the shares are sold, and computed using an appropriate after-tax discount rate r , is:

$$TAX_{buy \text{ before distribution}} = \tau_{cg} \cdot d + (1 + r_{20})^{-20} \cdot \tau_{cg} \cdot [(1 + g)^{20} - (1 + d)]. \tag{1}$$

In contrast, if the shareholder invests one dollar in the same mutual fund immediately after the distribution, she would also have a fund position in 20 years that would be worth $(1 + g)^{20}$, but her tax basis would be 1, so her capital gains tax liability at the time of sale would be $\tau_{cg} \cdot [(1 + g)^{20} - 1]$. The present value of her tax payment is:

$$TAX_{buy \text{ after distribution}} = (1 + r_{20})^{-20} \cdot \tau_{cg} \cdot [(1 + g)^{20} - 1]. \tag{2}$$

The difference between these two present value tax measures, $\tau_{cg} \cdot d \cdot [1 - (1 + r_{20})^{-20}]$, is the cost of accelerating capital gain realizations. A shareholder who purchases a fund just prior to a distribution date extends an interest-free loan of $\tau_{cg} \cdot d$ to the government. If a fund distributes eight percent of its net assets, the nominal 20-year interest rate is five percent, and the capital gains tax rate is 20 percent, the cost of acceleration equals 0.0099, or just under one percent of the fund investment. Buying the fund before the distribution is equivalent to paying a 100 basis point load. In some fund-years, the capital gains distribution can be as large as 20 percent of NAV. The median fund distribution for all equity mutual funds in our 1994–2000 sample period was approximately 3.5 percent of NAV. Some funds made much larger distributions in some years, thereby generating larger investor tax liabilities.

The expected tax cost of purchasing fund shares prior to a taxable distribution is an increasing function of the shareholder’s expected holding period. If a fund shareholder who has not yet sold her shares has a constant probability p of selling her shares each year, the expected present discounted value of the difference between the tax burden from buying just before and just after a distribution d is:

$$EPDV(p, r, d, \tau_{cg})_{tax \text{ differential from distribution}} = \tau_{cg} \cdot d - \sum_{j=1}^{\infty} p \cdot (1 - p)^{j-1} \cdot \tau_{cg} \cdot d \cdot (1 + r)^{-j} = \tau_{cg} \cdot d \cdot [r / (r + p)]. \tag{3}$$

In Eq. (3), we assume that the one-period interest rate, r , is the same in all periods, so there is no need to use different interest rates for different holding periods as in the reference to r_{20} above. Note that when $p = 1$, so that the mutual fund position will be sold with certainty one year after it is purchased, this expression reduces to $\tau_{cg} \cdot d \cdot [r / (r + 1)]$. This is the cost of loaning the government $\tau_{cg} \cdot d$ for one year without interest.

Table 1 reports the value of $EPDV/d$, the increased tax burden from buying before a distribution as a fraction of the distribution. The table considers a variety of discount rates, turnover probabilities, and capital gains tax rates. The results illustrate that the incremental tax burden as a fraction of the distribution can be substantial when the probability of sale is low and the nominal discount rate is high. When the sale probability is two percent per year, for example, the nominal interest rate is nine percent, and the tax rate on long-term gains is 15 percent, the present value of the tax-acceleration cost is 12.3 percent of the current tax burden associated with buying a fund share. With lower discount rates, lower tax rates, or higher redemption probabilities, the acceleration of the tax burden is smaller. The effective cost of buying mutual fund

Table 1

Present discounted value of tax saving from trading after vs. before taxable distribution. Source: Authors' calculations as described in the text.

Nominal discount rate	Sale probability=0.02/year		Sale probability=0.05/year		Sale probability=0.25/year	
	$\tau_{cg}=0.15$	$\tau_{cg}=0.28$	$\tau_{cg}=0.15$	$\tau_{cg}=0.28$	$\tau_{cg}=0.15$	$\tau_{cg}=0.28$
0.03/year	0.090	0.168	0.056	0.105	0.016	0.030
0.06/year	0.113	0.210	0.082	0.153	0.029	0.054
0.09/year	0.123	0.229	0.096	0.180	0.040	0.074

Table 2

Summary statistics on daily gross inflows. Source: Authors' calculations as described in the text. Daily shareholder transactions are aggregated to the fund level and scaled by the fund's total net assets on the prior day. The resulting ratio is multiplied by 10,000. Thus, an entry of 53.02 implies inflow of 0.0053, or just over one half of one percent of assets.

	Mean	Median	Standard deviation
All accounts	53.02	14.17	295.02
Taxable direct household accounts	8.23	0.70	89.52
Tax-deferred direct household accounts	4.78	0.21	119.19

shares before the distribution date is likely to vary across taxpayers, with the expected holding period playing an important role in generating this variation.

Because funds not only tend to make distributions at similar times each year but also generally announce their planned distributions several weeks before each record date, shareholders who make the effort to obtain this information can respond to impending distributions. To avoid accelerating their capital gains tax liabilities, they might defer their purchase in a fund that is about to make a distribution until after the distribution date or shift their intended purchase to a different fund. Our empirical work explores the extent of these responses by studying fund inflows before and after distribution dates.

2. Account-level data on mutual fund shareholders

We examine data on the daily trading behavior of all mutual fund shareholders at a no-load open-end mutual fund complex. The anonymous fund family that provided the data operates with fees and policies that are comparable to industry standards, and it is larger than the median fund family in terms of total assets under management. It includes approximately 10 funds, including both equity and fixed-income portfolios, but we focus our analysis on the equity funds. Data from this mid-size fund family are the basis for [Johnson's \(2004, 2010\)](#) studies of heterogeneity in shareholder trading patterns. The transfer agent for the mutual fund family provided a database consisting of over 50,000 shareholders and nearly one million transactions in the family of funds between late 1994 and mid-year 2000. The database includes information on the registration information for each account, as well as on daily net asset values (NAVs) and distributions for each fund.

Shareholders have traditionally purchased no-load mutual fund shares directly from the fund. These are "direct" shareholders. However, many shareholders now purchase shares through an intermediary, such as a mutual fund super-market, that collects the transactions of its customers and passes them through to the fund. These are "indirect" shareholders. The distinction between direct and indirect shareholders is important because the tax status of the former but not the latter can be determined reliably at the account level in our database. Our strategy for identifying the effects of taxes, therefore, requires that we exclude the indirect shareholders from some of the analysis. Moreover, we generally further restrict our analysis to households in order to avoid confounding tax effects with other clientele effects than might arise in a broader sample of account types. "Non-household" shareholders include trusts, college endowments, and corporate treasury accounts.

We investigate whether shareholders reduce their purchases of fund shares just before the fund pays a distribution. Although distributions could also affect shareholder redemptions under certain assumptions, the effect of taxes on redemption decisions are less clear than those on inflows; therefore, we focus on this margin of investor behavior. In an earlier version of this paper, we studied gross outflows and found no unusual patterns around distribution dates.

We aggregate individual purchases to create daily gross inflows for each fund. Following previous studies such as [Bergstresser and Poterba \(2002\)](#), we scale daily gross inflows by the fund's lagged net assets. In all cases, we exclude fund distributions from the inflows whether or not the shareholder reinvests them. Most shareholders reinvest dividends. In our sample, 97.1 percent of accounts, and 95.9 percent of the dollar-weighted accounts, reinvest dividends. These reinvestment rates are similar to those found in the mutual fund industry as a whole.

[Table 2](#) presents sample summary statistics on gross inflows. The first row presents results for aggregate inflows, while the last two focus only on gross inflows from households who are direct shareholders. The average daily gross inflow is 0.0053 times fund assets, or just over one-half of one percent. There is substantial variation in daily gross inflows, which is

indicated by the disparity between the mean and median as well as the large standard deviation. Table 2 also shows that gross inflows to taxable direct household accounts represent roughly one-seventh of the aggregate gross inflows, while inflows to tax-deferred direct household accounts represent about one-eleventh of all gross inflows. Some of our analysis focuses on inflows from subsets of accounts, and in such cases, the total dollar value impact of proximity to a distribution date is smaller than when we consider total inflows.

The mean return on the equity funds in our sample is 0.084 percent per day. Multiplying by an average of 252 trading days per year, this translates to an average annual return of 21.2 percent. The standard deviation of returns is 1.24 percent per day, or 19.7 percent per year. The corresponding benchmark returns for the equity funds in our sample average 16.6 percent, and the annual standard deviation is also 16.6 percent. Thus, the average return on the sample funds exceeded the average benchmark returns, but the sample funds were also substantially more volatile than the benchmark.

Each of the funds in our sample usually makes distributions in December although there are several years in which some funds made two distributions, the first in the second half of the year but before December, and the second in December. The sample fund family is somewhat unusual in not making a distribution in the first half of the calendar year. For the entire universe of U.S. equity mutual funds making distributions in 2000, 38.3 percent of the distributions occurred in December, 11.4 percent in June, 10.2 percent in March, and 9.8 percent in September. The median fund distribution in our sample is about four percent of NAV while the mean is approximately five percent. The average annual distribution at the funds in our sample is very close to the average of all funds for our sample period.

3. Gross inflows before and after distribution dates

Many previous studies of mutual fund flows, including Warther and Warther (1995) and Sirri and Tufano (1998), suggest that fund flows are sensitive to lagged returns. Johnson (2010) confirms this relationship for our sample funds. We develop a simple empirical model for gross inflows and use it to consider inflow patterns around distribution dates. We refine our analysis by examining differences in distribution-related effects before and after the Taxpayer Relief Act of 1997, which changed relevant shareholder tax rates.

3.1. Specification issues

We focus on empirical models that relate gross inflows at fund i to the lagged return on fund i ($R_{i,t}$) as well as lagged returns on either the aggregate equity market ($R_{m,t}$) or on the benchmark index identified in the fund's prospectus ($R_{b,t}$). With respect to the lag structure for returns, even with daily flows data, a long and unrestricted lag structure would absorb many degrees of freedom. To allow lagged returns to affect fund inflows for periods as long as a year, but to reduce the number of parameters that we estimate, we allow separate coefficients for the inflow effect of returns on each of the previous five days, but we constrain all returns in the previous three weeks to have an identical effect by including the average return over that period (trading days $t-6$ to $t-21$). We also include the average return in each of the previous five months, and the average return over the period between seven and 12 months earlier. In this way, returns for the entire previous year enter the specification but we do not lose many degrees of freedom. Thus, we would include 12 lagged return variables when explaining the inflow to a fund on the first trading day in January. These variables would be five distinct daily returns for the last five trading days in December, the average daily return over days $t-21$ through $t-6$, five variables corresponding to lagged monthly returns between July and November, and one variable measuring the average return over the period between the roughly six month interval that is $t-252$ through $t-127$. Our findings were not substantively affected when we included longer lags.

Our basic estimating equation is

$$\text{INFLOW}_{i,t} = \alpha_1 + \sum_j \beta_{1j} * R_{i,t-j} + \sum_j \gamma_{1j} * R_{b,t-j} + \sum_{k=-1,1} \lambda_{1k} * \text{WINDOW30}_{i,k,t} + \epsilon_{1,t}. \quad (4)$$

The variable capturing patterns around distribution dates, $\text{WINDOW30}_{i,k,t}$, is set equal to 1 if day t occurs in the 30-day window before or after (k equals -1 or 1 , respectively) a distribution date for fund i . We define a similar variable (WINDOW15) for a 15-day window. We are primarily concerned with the $\{\lambda_k\}$ coefficients. We often estimate two inflow equations, such as one for taxable accounts and one for non-taxable accounts, and we are interested in testing cross-equation coefficient restrictions. To simplify such testing, we estimate the two equations using a Seemingly Unrelated Regression (SUR) procedure, and we calculate test statistics using the resulting variance-covariance matrix for all estimated parameters.

Table 3 reports return and WINDOW30 and WINDOW15 coefficients from estimates of (4) for all equity funds in our sample. The table shows estimates with robust and heteroscedastic-consistent standard errors in parentheses. The first model uses a 30-calendar-day window on each side of the distribution while the second uses a 15-calendar-day window. The first two rows present evidence on gross inflow patterns in the periods just before and just after fund distributions, and the third row reports p -values for the hypothesis test of equality between the estimates in the first two rows. The estimates imply that for each of the 30 days before a distribution, gross inflows average 24 basis points lower per day than the average at other times. In the 30 days after the distribution date, inflow averages 11 basis points per day below their average other times.

Table 3

Gross inflows around fund distribution dates.

Independent variables	30-Day window		15-Day window	
Indicator for pre-distribution period	–23.68 (4.33)		–16.05 (5.68)	
Indicator for post-distribution period	–10.97 (4.42)		–0.27 (6.03)	
<i>p</i> -Value (equality test)	0.020		0.040	
	Lagged returns ($\times 10^{-2}$)		Lagged returns ($\times 10^{-2}$)	
	Fund return	Benchmark	Fund return	Benchmark
Day t-1	6.59 (1.53)	–3.20 (1.74)	6.57 (1.53)	–3.21 (1.74)
Day t-2	2.93 (1.53)	–1.71 (1.75)	2.86 (1.53)	–1.64 (1.75)
Day t-3	4.17 (1.53)	–1.81 (1.75)	4.09 (1.53)	–1.72 (1.75)
Day t-4	4.90 (1.53)	–3.84 (1.75)	4.84 (1.54)	–3.82 (1.75)
Day t-5	3.57 (1.53)	–1.82 (1.74)	3.58 (1.54)	–1.77 (1.75)
Days t-6 through t-21	2.90 (0.37)	–1.99 (0.43)	2.99 (0.37)	–1.91 (0.43)
Month t-2	2.95 (0.32)	–2.30 (0.36)	3.03 (0.32)	–2.20 (0.36)
Month t-3	2.74 (0.32)	–3.09 (0.37)	2.62 (0.32)	–2.82 (0.36)
Month t-4	3.68 (0.34)	–3.76 (0.40)	3.58 (0.34)	–3.48 (0.39)
Month t-5	3.23 (0.36)	–3.80 (0.40)	3.02 (0.36)	–3.36 (0.43)
Month t-6	3.64 (0.39)	–3.50 (0.48)	3.51 (0.39)	–3.27 (0.47)
Months t-7 through t-12	2.05 (0.18)	–2.28 (0.22)	2.04 (0.18)	–3.34 (0.22)
R ²	0.112		0.109	

Notes: The dependent variable is the sum of daily shareholder transactions, divided by the fund's total net assets on the prior day and multiplied by 10,000. The first coefficient therefore translates into a 24 basis point decline in fund inflow in each of the 30 days immediately prior to a fund distribution. Robust standard errors are shown in parentheses. See text for further discussion.

Table 4

Robustness of estimates to return specification.

Independent variables	Fund returns	Fund and benchmark returns	Fund and market returns	Fund–benchmark returns	Fund–market returns
Indicator for 30 days prior to distribution	–9.84 (4.20)	–23.68 (4.33)	–16.70 (4.37)	–25.01 (4.07)	–16.05 (4.08)
Indicator for 30 days following distribution	2.96 (4.28)	–10.97 (4.42)	–9.92 (4.44)	–9.93 (4.15)	–3.08 (4.16)
<i>p</i> -Value (equality test)	0.020	0.020	0.213	0.005	0.016
R ²	0.061	0.112	0.120	0.103	0.089

Notes: Equations are estimated by SUR using daily observations from late 1994 through the middle of 2000. All equations include fund-specific intercepts as well as the indicated returns for the previous five days, the next 15 days, and then months t-2, t-3, t-4, t-5, and t-6, and the previous six months. Robust standard errors are shown in parentheses. See text for further discussion.

Thus, we estimate a “distribution effect” of –13 basis points per day, which cumulates over 21 trading days in the 30 calendar days before a distribution to a reduction in net fund inflows of 2.73 percent (21×0.0013) of fund net assets. We reject the null hypothesis of equal coefficients in the pre- and post-distribution periods at the 95 percent confidence levels. The difference between the before-distribution and after-distribution inflows is even more pronounced when we use a 15-day window. In that case the difference is 16 basis points per day, and we once again reject the null hypothesis that gross inflows before and after the distribution date are the same.

The coefficient estimates below the first three rows of Table 3 show that lagged fund returns have a positive effect on shareholder inflows, with the largest estimated coefficients in the previous few trading days. There is relatively little decay in the effect of returns as the lag length increases. In the 30-day window specification, a one percentage point increase in the fund return on a given day increases the gross inflows to the fund by 22 basis points in the next five days.

The coefficients on the benchmark portfolio returns are always of the opposite sign of the lagged fund returns. In many cases, the individual coefficients on daily or monthly benchmark returns are not statistically significantly different from zero, but we reject the null hypothesis that all of the benchmark return coefficients are jointly zero. The coefficient patterns suggest that some shareholders base their fund purchase decisions on relative performance.

Before comparing shareholder behavior in taxable and non-taxable accounts, we consider the sensitivity of the gross inflow model to the choice of alternative specifications of lagged returns. Table 4 presents the coefficients on the 30-day indicator variables $\{\lambda_k\}$ from five different specifications of (4) that differ only in the included set of lagged return measures. The specification with the highest explanatory power is the one that includes both fund returns and market returns, measured as the CRSP value-weighted index return. The specification that includes both fund and fund benchmark index

Table 5

Direct household accounts vs. all other accounts.

Independent variables	Direct household accounts		All other accounts	
	30-Day window	15-Day window	30-Day window	15-Day window
Indicator for pre-distribution period	−4.37 (0.72)	−4.46 (0.94)	−19.31 (4.16)	−11.60 (5.45)
Indicator for post-distribution period	−2.15 (0.73)	−1.15 (1.00)	−8.81 (4.24)	0.88 (5.79)
p-Value (equality test)	0.015	0.009	0.046	0.090
R ²	0.153	0.150	0.085	0.082

Notes: Equations are estimated by SUR using daily observations from late 1994 through the middle of 2000. All equations include fund-specific intercepts as well as fund and benchmark returns for the previous five days, the next 15 days, and then months t-2, t-3, t-4, t-5, and t-6, and the previous six months. Robust standard errors are shown in parentheses. See text for further discussion.

Table 6

Taxable vs. tax-deferred direct household accounts.

Independent variables	Taxable direct household accounts		Tax-deferred direct household accounts	
	30-Day window	15-Day window	30-Day window	15-Day window
Indicator for pre-distribution period	−3.60 (0.60)	−3.50 (0.79)	−0.76 (0.27)	−0.94 (0.35)
Indicator for post-distribution period	−1.38 (0.61)	−0.85 (0.84)	−0.78 (0.28)	−0.38 (0.38)
p-Value (equality test)	0.004	0.013	0.960	0.244
R ²	0.133	0.130	0.068	0.067

Notes: Equations are estimated by SUR using daily observations from late 1994 through mid 2000. All equations include fund-specific intercepts as well as fund and benchmark returns for the previous five days, the next 15 days, and then months t-2, t-3, t-4, t-5, and t-6, and the previous six months. Robust standard errors are shown in parentheses. See text for further discussion.

returns has the second highest explanatory power. Constraining the coefficients by including either (fund return – market return) or (fund return – benchmark return) reduces explanatory power, as does excluding both market and benchmark returns. While the fit of the various models differs, the estimated differences in both inflows in the 30-day window before and after the distribution date are remarkably stable across specifications. The change is approximately 13 basis points in every specification except for the third one, in which it is 7 basis points. Our conclusions about flow patterns are robust to alternative choices with regard to return specification choices.

3.2. Taxable and non-taxable accounts

We follow in the tradition of Barber and Odean (2004) and Ivković et al. (2005) by using differences between the behavior of household shareholders in taxable and tax-deferred settings to estimate the impact of taxes on fund purchases. We focus on households to avoid confounding tax effects with other clientele effects than might arise in a broader sample of account types. Data limitations prevent us from reliably identifying households in the indirect distribution channel. Thus, our household sample is drawn solely from the direct distribution channel. To understand how this selection rule might affect our findings, we compare gross inflows around distribution dates from direct household accounts with gross inflows around distribution dates from all other accounts. The second group includes all shareholders in the indirect channel, as well as all non-households in the direct channel.

Table 5 presents the results, which suggest that gross inflows increase after distribution dates for both direct household accounts and for all other accounts. The change in gross inflows is statistically significant at the 10 percent confidence level for both the 30-day window and the 15-day window. Focusing on the 30-day inflow results for the direct household sample, the pre- and post-distribution coefficients are −4.37 and −2.15, respectively. For the other sample, the corresponding coefficients are −19.31 and −8.81. Note that the absolute values of the WINDOW coefficients in Table 5 are smaller than those in Tables 3 and 4, mechanically reflecting our focus on only a subset of inflows.

Table 6 presents our central results on the differences in gross inflow patterns around distribution dates between taxable and tax-deferred direct household accounts. The results suggest that some taxable shareholders alter their fund trades around distribution dates while tax-deferred shareholders ignore distribution dates. For taxable shareholders, gross inflows increase after distribution dates by a statistically significant 2.22 basis points in the case of the 30-day window and a statistically significant 2.65 basis points in the case of the 15-day window. For tax-deferred shareholders, we cannot reject the null hypothesis that inflows are the same before and after distribution dates. This is true for both the 30-day and 15-day windows. To place the results in perspective, recall that in Table 2 the average daily gross inflow from taxable households is 8.23 basis points. Thus, taxable household inflows decline about 27 percent (2.22/8.23) in the period prior to fund distributions.

3.3. The Taxpayer Relief Act of 1997 and the composition of fund distributions

The tax code changed during our sample period. The Taxpayer Relief Act of 1997 (TRA97) reduced the top marginal tax rate on realized long-term capital gains from 28 percent to 20 percent for taxpayers in the highest income tax brackets. The transition from one tax regime to another was complex because the legislation changed both the definition of the long-term holding period and the rates that applied to long-term gains. Prior to TRA97, gains were defined as long-term if the underlying asset had been held for at least 12 months. TRA97 redefined long-term gains as gains on assets held for at least 18 months, effective July 29, 1997.

TRA97 also retroactively reduced long-term capital gains tax rates from 28 percent to 20 percent with an effective date of May 6, 1997. Gains on assets held for at least 18 months and sold after that date faced a maximum tax rate of 20 percent. Gains on assets held for less than 12 months were taxed as ordinary income throughout the year. Gains on assets held for between 12 and 18 months, however, faced a top marginal tax rate of 28 percent prior to May 6, when they were long-term gains taxed under the old top marginal tax rate, 20 percent between May 6 and July 28, when they were long-term gains taxed under the new but retroactive rate, and 28 percent after July 29, when they were defined as a new category of “mid-term” capital gains that were subject to a new rate. The Internal Revenue Service Restructuring and Reform Act of 1998 eliminated mid-term gains, and effective January 1, 1998, gains on assets held for more than 12 months qualified as long-term and were taxed at a maximum tax rate of 20 percent.

TRA97 also changed other tax rules that may affect inflow patterns to tax-deferred accounts. The legislation introduced Roth IRAs, Education IRAs, and 529 plans. Contributions to Education IRAs and 529 plans must be made by the end of the calendar year, rather than by April 15 of the following year as with traditional and Roth IRAs. While this could affect gross inflows, neither Education IRAs nor 529 plans account for substantial gross fund inflows in our database.

The 1997 and 1998 tax reforms reduced mutual fund shareholders' tax penalty from the acceleration of capital gains tax liability. We test whether the inflow difference around distribution dates was larger before the TRA97 than afterward by interacting an indicator variable for observations in our sample after July 29, 1997 (POSTTRA97) with the WINDOW variables. We do not attempt to capture the effects of the brief mid-term holding period because it applies for only five months. We also control for other potential changes in gross inflows before and after TRA97 by also adding an indicator variable that equals one in the post-TRA97 regime to our specifications. The new inflow specification is

$$\text{INFLOW}_{i,t} = \alpha_1 + \alpha_2 * \text{POSTTRA97}_{i,t} + \sum_j \beta_{1j} * R_{i,t-j} + \sum_j \gamma_{1j} * R_{b,t-j} + \sum_{k=-1,1} (\lambda_{1k} + \lambda_{2k} * \text{POSTTRA97}_{i,t}) * \text{WINDOW30}_{i,k,t} + \epsilon_{1i,t}. \quad (5)$$

Table 7 presents estimates using the 30-day window. Findings from the unreported 15-day specification are similar. The first specification in Table 7 shows that the decline in taxable inflows before a distribution date was larger for taxable direct household accounts before TRA97, at 4.89 (0.14 – (–4.75)) basis points per day, than afterwards, when it was 0.67 basis points per day (–0.66 – (–1.33)). Additionally, the pre-1997 decline is statistically significantly different from zero while the post-1997 decline is not. The post-1997 indicator variable, included without an interaction with distribution date variables, shows that gross inflows were lower in the later part of our sample than they were earlier. The differences between gross inflows in tax-deferred accounts before and after distribution dates are never statistically significantly different from zero. We note that this model has higher explanatory power than the previous one: for example, the R2 for the taxable shareholders increases from 13.3% (Table 6) to 19.1% (Table 7).

Table 7
The Taxpayer Relief Act of 1997 and the composition of the fund distribution.

Independent variables	Taxable direct household accounts		Tax-deferred Direct household accounts	
	TRA97	Tax impact	TRA97	Tax impact
Pre-1997	–4.75	–3.52	–0.02	0.12
Pre-distribution	(1.00)	(0.82)	(0.46)	(0.37)
Post-1997	0.14	0.55	–0.48	–0.42
Post-distribution	(0.96)	(0.78)	(0.44)	(0.36)
p-Value (equality test)	0.000	0.000	0.448	0.271
Post-1997	–1.33	–1.01	–0.48	–0.41
Pre-distribution	(0.72)	(0.55)	(0.33)	(0.25)
Post-1997	–0.66	–0.31	–0.39	–0.31
Post-distribution	(0.76)	(0.57)	(0.35)	(0.26)
p-Value (equality test)	0.459	0.311	0.833	0.764
Post-1997 indicator variable	–7.59	–7.56	–2.68	–2.67
	(0.42)	(0.41)	(0.19)	(0.19)
R2	0.191	0.191	0.110	0.110

Notes: Equations are estimated by SUR using daily observations from late 1994 through the middle of 2000. All equations include fund-specific intercepts as well as fund and benchmark returns for the previous five days, the next 15 days, and then months t-2, t-3, t-4, t-5, and t-6, and the previous six months. Robust standard errors are shown in parentheses. See text for further discussion.

We have not yet considered differences in the tax burden on different distributions for taxable investors, even though Plancich (2003) explains that the mix of long-term gains and ordinary income in fund distributions varies across distributions. The acceleration of capital gains tax liability is greater when a distribution contains short-term gains than when it contains only long-term gains. If shareholders can estimate the composition of distributions prior to the distribution date, there may be larger changes in trading patterns around distributions that generate larger tax penalties.

We measure each distribution's relative tax burden, IMPACT, as:

$$\text{IMPACT}_{i,k,j} = \left[\tau_{\text{ordinary}} * (1 - \mu_{l_{t,i,k,j}}) + \tau_{\text{long-term cg}} * \mu_{l_{t,i,k,j}} \right] / \tau_{\text{long-term cg}}. \quad (6)$$

In this expression, $\mu_{l_{t,i,k,j}}$ is the fraction of the distribution that is taxed as a long-term gain. If $\mu_{l_{t,i,k,j}}$ equals one, IMPACT for a distribution equals one. If $\mu_{l_{t,i,k,j}}$ equals zero, so that the entire distribution is taxed as ordinary income, the tax impact factor equals $\tau_{\text{ordinary}} / \tau_{\text{long-term cg}}$, which is greater than one. In the pre-1997 period, this ratio could be as large as $39.6 / 28 = 1.41$. In the post-1997 period, when the long-term capital gains tax rate declined, the ratio could be as large as 1.98 ($= 39.6 / 20$). Thus, it is potentially more important to distinguish long-term and short-term gains after the TRA97 than before. In our sample, the average value of $\mu_{l_{t,i,k,j}}$ is roughly two-thirds. We aggregate distributions of ordinary income with short-term capital gains when constructing this variable.

We test for the importance of tax status variation across distributions by interacting IMPACT with WINDOW30 and estimating the following model

$$\begin{aligned} \text{INFLOW}_{i,t} = & \alpha_1 + \alpha_2 * \text{POSTTRA97}_{i,t} + \sum_j \beta_{1j} * R_{i,t-j} + \sum_j \gamma_{1j} * R_{b,t-j} \\ & + \sum_{k=-1,1} (\lambda_{1k} + \lambda_{2k} * \text{POSTTRA97}_{i,t}) * \text{WINDOW30}_{i,k,t} * \text{IMPACT}_{i,k,t} + \epsilon_{1i,t}. \end{aligned} \quad (7)$$

The second and fourth columns of Table 7 present estimates of this specification. The coefficient patterns are very similar to those in equations that do not allow for variation in the tax burden across distributions. One potential explanation for the similarity of the results is that shareholders may not be aware of the subtle tax consequences that come from differences in the composition of different distributions. Alternatively, the composition may not have been announced by the funds.

The cost of tax acceleration is proportional to the size of the distribution, so tax-sensitive shareholders should respond more strongly to large distributions than to small ones. In an unreported specification patterned after Eq. (7), we consider this possibility by interacting WINDOW30 with the size of the distribution, measured as a percentage of NAV. The results are qualitatively similar to the first specification in Table 7. We were not able to reject the null hypothesis that the change in inflows around a distribution is unrelated to the size of the distribution, but we note that our small sample size makes strong conclusions difficult.

3.4. Controlling for turn-of-the-year effects

The funds often have distributions in late December. Therefore, the pre-distribution period for the late-December distributions will overlap with the last few weeks of the calendar year and the post-distribution period will overlap with the first few weeks of the New Year. Choi (2015) studies fund-level flow data collected from N-SAR filings and finds that net shareholder flows are higher in January than in December.

We test for the possibility that a turn-on-the-year seasonality drives our results by adding indicator variables for each of the first four and last four weeks in the calendar year to Eq. (7) and then re-estimating the gross inflow equation for both taxable and tax-deferred direct household accounts. The results, shown in Table 8, suggest that allowing for turn-of-the-year effects does not substantially alter our results on trading patterns. For taxable shareholders, we continue to find a statistically significantly lower rate of gross inflows in the 30 days before a distribution than in the 30 days afterward, with a more pronounced effect prior to the enactment of TRA97 than afterward. Tax-deferred accounts remain insensitive to distribution dates. For both types of accounts, we cannot reject the null hypothesis that the eight indicator variables for weeks around the turn of the year are jointly equal to zero.

4. Holding periods before and after distribution dates

Our findings suggest that at least some taxable shareholders reduce their mutual fund purchases in the weeks before distributions. Yet it is still evident that many taxable investors purchase fund shares in the weeks and days prior to distribution dates. These buyers may not be aware that a distribution is about to be paid, or they may have specialized tax circumstances, such as capital loss carryforwards, that eliminate their tax liability on the distribution. It is also possible that these buyers believe that the fund's NAV will rise enough by the distribution date to offset the cost of capital gains tax acceleration. Buyers with such beliefs would expect a higher after-tax return from purchasing fund shares before, rather than after, the distribution date.

Our data set does not contain any information that enables us to measure a fund buyer's expected NAV appreciation, or that identifies the buyer's detailed tax positions. We do, however, observe one investor attribute that is correlated with the cost of tax acceleration: the investor's realized holding period of fund shares. The cost of buying shares before rather than after a distribution is smaller for an investor with a short holding period than for one who expects to hold the fund shares

Table 8
Turn-of-the-year effects.

Independent variables	Taxable direct	Tax-deferred direct
	Household accounts	Household accounts
Pre-1997	–2.20	–0.13
Pre-distribution	(1.07)	(0.49)
Post-1997	1.48	–0.63
Post-distribution	(1.10)	(0.50)
<i>p</i> -Value (equality test)	0.009	0.433
Post-1997	–0.44	–0.54
Pre-distribution	(0.64)	(0.29)
Post-1997	0.22	–0.50
Post-distribution	(0.71)	(0.32)
<i>p</i> -Value (equality test)	0.421	0.914
Post-1997 indicator variable	–7.52	–2.66
	(0.42)	(0.19)
Week 49	–1.45	1.27
	(1.44)	(0.66)
Week 50	–2.12 (1.46)	–0.12
		(0.66)
Week 51	–1.70	–0.22
	(1.43)	(0.65)
Week 52	–2.49	0.45
	(1.36)	(0.62)
Week 1	–1.36	1.58
	(1.52)	(0.70)
Week 2	–0.74	–0.11
	(1.46)	(0.67)
Week 3	–0.23	–0.02
	(1.50)	(0.68)
Week 4	–1.84	–0.43
	(1.27)	(0.58)
<i>p</i> -Value (eight weeks jointly equal to zero)	0.541	0.125
<i>R</i> ²	0.192	0.112

Notes: Equations are estimated by SUR using daily observations from late 1994 through the middle of 2000. All equations include fund-specific intercepts, fund and benchmark returns, and controls for “tax impact”. Robust standard errors are shown in parentheses. See text for further discussion.

for many years. We can study whether investors who buy fund shares before distribution dates sell them sooner, on average, than post-distribution buyers.

We test this possibility by fitting a Cox (1975) proportional hazards model to the shareholder's risk of account closure. This is an appropriate measure of how long shareholders hold fund shares because, as Johnson (2004) finds, shareholders do not generally trade between account opening and account closing. We measure time in calendar days since the account opening, and we model the probability of closure as a function of the same explanatory variables that we used in the last section.

Table 9 reports our estimates of hazard ratios, which are easier to interpret than the underlying coefficients from the proportional hazards model, along with robust standard errors. Estimated hazard ratios that are smaller than one are associated with low closure risk and long holding periods while those that are larger than one are associated with high closure risk and short holding periods. The results suggest some differences between pre- and post-distribution buyers. The hazard ratios indicate that taxable shareholders who open accounts after a distribution are approximately 25 percent less likely each day to close their accounts than those who open accounts before a distribution. Tax-deferred shareholders have a similar tendency, but it is economically smaller and statistically insignificant.

The hazard ratio differences translate into substantial differences in expected holding periods for taxable shareholders. In the pre-TRA97 period, shareholders who invested in the 30 days before a distribution had a 27 percent chance of liquidating their position within one year while those who invested in the 30 days after a distribution had a 19 percent chance of liquidation within one year. In the post-TRA97 period, the corresponding numbers are 28 percent and 15 percent. Thus, the evidence indicates that the differences in one-year closing rates between those who buy before distributions versus those who buy after distributions are 8 percent and 13 percent. For tax-deferred shareholders, expected holding periods do vary around distribution dates; however, the differences are much smaller. In both the pre- and post-TRA97 regimes, the difference is four percent.

This evidence suggests that there are differences in holding periods between investors who open accounts before distributions and those who open after distributions. We reject the null hypothesis of equal account-closure rates for taxable accounts, but we cannot reject the corresponding null hypothesis for tax-deferred accounts. This is suggestive evidence that

Table 9
Account closure risk.

Independent variables	Taxable direct household accounts	Tax-deferred direct household accounts
Pre-1997	0.995	1.408 (0.179)
Pre-Distribution	(0.095)	
Post-1997	0.749	1.248
Post-distribution	(0.074)	(0.196)
p-Value (equality test)	0.012	0.479
Post-1997	1.008	1.274
Pre-distribution	(0.131)	(0.176)
Post-1997	0.720	1.165
Post-distribution	(0.077)	(0.180)
p-Value (equality test)	0.027	0.614
Post-1997 indicator variable	1.001	1.087
	(0.050)	(0.081)
Week 49	1.033	1.012
	(0.146)	(0.178)
Week 50	0.701	0.750
	(0.109)	(0.146)
Week 51	0.689	0.686
	(0.111)	(0.131)
Week 52	1.338	0.756
	(0.171)	(0.129)
Week 1	1.296	0.915
	(0.171)	(0.194)
Week 2	1.207	0.830
	(0.158)	(0.177)
Week 3	1.266	0.886
	(0.161)	(0.185)
Week 4	1.235	0.904
	(0.140)	(0.162)

Notes: Entries are hazard ratios from Cox proportional hazards models estimated by maximum likelihood. The equations are estimated on daily data from late 1994 through the middle of 2000 and are adjusted for “tax impact”. The equations include fund-specific intercepts as well as fund and benchmark returns. Robust standard errors for the hazard ratios are shown in parentheses. See text for further discussion.

investors may consider their own investment and tax circumstances when they make trading decisions: the anticipated tax-acceleration penalty for opening an account prior to the distribution is smaller than that for opening an account after the distribution for those shareholders who expect to have short holding periods.

5. Conclusion

Our analysis of account-level mutual fund purchases suggests that a substantial group of shareholders heed the advice of financial advisors who warn about the tax penalty associated with buying mutual fund shares prior to a distribution date. At the fund family we study, gross inflows from taxable investors decline by approximately one quarter in the four weeks before a distribution, and on average a fund’s net assets are about three percent lower by the distribution date than they would have been in the absence of a distribution. We attribute this decline in inflows to the tax-avoidance behavior of individual investors whose capital gains tax liability would be accelerated, and increased in present value, if they purchased fund shares before rather than after the distribution date. The decline in pre-distribution fund inflows is focused in the period prior to the enactment of the Taxpayer Relief Act of 1997, which reduced the income tax penalty for receiving a capital gains distribution by cutting the top long-term capital gains tax rate from 28 percent to 20 percent. We also find that the likelihood of closing an account within a year of opening is higher for taxable shareholders who open accounts just before fund distribution dates than it is for taxable shareholders who open accounts just after fund distribution dates. We find no such difference for investors in tax-deferred accounts. This is consistent with investors strategically timing their investments based on the tax consequences of expected future trades.

One important question about our findings is the extent to which they generalize to the broader universe of mutual fund shareholders. Since the results are based on data from only one mid-size mutual fund family, they may not be representative; nevertheless, Johnson (2010) provides evidence that the pattern of flows at the fund family we study is similar to that at other fund families. One additional risk is that because our tests focus on the behavior of households who invest through the direct distribution channel, they may not generalize to other shareholders who trade through intermediaries such as brokers or fund supermarkets. It is impossible to resolve this question definitively using our data, but it warrants study with data from other fund families.

A second issue, which we have not been able to address with any precision in our database, concerns investor heterogeneity. We found, but do not report, some evidence that investors who make large transactions are more sensitive to

distribution dates than their smaller-transaction counterparts, but the pattern varied within our sample period. Data collected from mutual fund accounts offer very limited information on investor attributes, but future work could usefully explore ways to either collect additional information, perhaps with a follow-on survey of shareholders, or to impute additional data, perhaps using geographical identifiers.

Our results relate to a number of issues that are broader than the behavior of mutual fund investors. One is the ongoing public policy discussion about the tax treatment of mutual fund distributions and mutual fund shareholders more generally. Our evidence on the difference between fund inflows before and after distribution dates suggests that changing the tax law to allow deferral of tax liability on realized gains within funds, as suggested by Saxton (2004) and others, would change fund inflow patterns. Our results do not address the important question of whether such a policy change would increase the assets held in open-end mutual funds relative to other financial products. The potential substitution between traditional open-end mutual funds and exchange-traded funds (ETFs) is particularly intriguing. Most ETFs avoid taxation of realized capital gains by using redemption-in-kind distributions as described in Poterba and Shoven (2002). Fama and French (2006) suggest that this difference between ETFs and traditional mutual funds “gives ETFs a substantial advantage over traditional open-end funds in the competition for customers”. Our results support Bergstresser and Poterba’s (2002) finding that mutual fund investors consider after-tax returns in their investment decisions. This body of work broadly suggests the potential importance of taxation in affecting investment in intermediated financial products.

Second, our findings also bear on a broad issue involving asset market equilibrium in the presence of taxes. They underscore the possibility, documented in Graham and Kumar’s (2006) work as well, that the clientele of shareholders holding a security may vary around significant events, such as a fund distribution date or an ex-dividend date. They provide some evidence that investors view the involuntary realization of capital gains as costly, as Chay et al. (2006) suggest based on security price movements. The unresolved but central empirical question in this field is whether clientele variation is large enough to induce changes in asset prices and the associated signals that asset markets send to managers.

Finally, while our findings suggest that shareholder behavior is affected by taxes, they do not provide any guidance on the efficiency costs of tax-induced distortions. If shareholders simply delay the purchase of a particular mutual fund by a few days or weeks, the efficiency cost can be evaluated by the cost of holding an alternative portfolio that is not the desired one for a modest period of time. If shareholders respond to the unusual tax treatment of mutual funds by choosing to hold other asset classes, the efficiency calculation must recognize the cost of such portfolio changes. Analysis of the efficiency costs of these behavioral changes is an important direction for future work.

Acknowledgments

We are grateful to Raj Chetty, Joel Dickson, Bill Gentry, Dan Halperin, Jiekun Huang, Lenore Robbins, several anonymous referees, and seminar participants at the American Economic Association annual meetings, Boston College, Hamilton College, NBER, the Securities and Exchange Commission, and the University of Chicago for comments. We thank the fund family that provided data to make this project possible, and the National Science Foundation (Poterba) and a Keller Research Award (Johnson) for research support. Poterba is a trustee of the College Retirement Equity Fund (CREF) and the TIAA-CREF mutual funds. 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 authors and do not necessarily reflect the views of the Commission or of its staff, or of the NBER, TIAA-CREF, or any other organization.

References

- Barber, Brad M., Odean, Terrance, 2004. Are individual investors tax savvy? Evidence from retail and discount brokerage accounts. *J. Public Econ.* 88, 419–442.
- Barclay, Michael J., Pearson, Neal D., Weisbach, Michael S., 1998. Open end mutual funds and capital gains taxes. *J. Financ. Econ.* 49, 3–43.
- Bergstresser, Daniel, Poterba, James, 2002. Do after-tax returns affect mutual fund inflows? *J. Financ. Econ.* 63, 381–414.
- Chay, J.B., Choi, Dosoung, Pontiff, Jeffrey, 2006. Market valuation of tax-timing options: evidence from capital gains distributions. *J. Financ.* 61, 837–865.
- Chen, Feng, Kraft, Arthur, Weiss, Ira, 2011. Tax planning by mutual funds: evidence from changes in the capital gains tax rate. *Natl. Tax J.* 64 (1), 105–134.
- Choi, Hyung-Suk, 2015. Seasonality in mutual fund flows. *J. Appl. Bus. Res.* 31 (2), 715–726.
- Christoffersen, Susan, Geczy, Christopher, Musto, David, Reed, Adam, 2005. Cross-border dividend taxation and the preferences of taxable and non-taxable investors: evidence from Canada. *J. Financ. Econ.* 78, 121–144.
- Constantinides, George, 1984. Optimal stock trading with personal taxes: implications for prices and the abnormal January returns. *J. Financ. Econ.* 13, 65–89.
- Cox, David R., 1975. Partial likelihood. *Biometrika* 62, 269–276.
- Dickson, Joel M., Shoven, John B., Sialm, Clemens, 2000. Tax externalities of equity mutual funds. *Natl. Tax J.* 53, 607–628.
- Fama, Eugene F., French, Kenneth R., 2006. Keep it simple. *Wall Str. J.* 25, A10.
- Graham, John R., Kumar, Alok, 2006. Do dividend clienteles exist? Evidence on dividend preferences of retail investors. *J. Financ.* 61, 1305–1336.
- Investment Company Institute, 2015. *Investment Company Fact Book*. Investment Company Institute, Washington.
- Ivković, Zoran, Poterba, James, Weisbenner, Scott, 2005. Tax-motivated trading by individual investors. *Am. Econ. Rev.* 95, 1605–1630.
- Johnson, Woodrow T., 2004. Predictable investment horizons and wealth transfers among mutual fund shareholders. *J. Financ.* 59, 1979–2012.
- Johnson, Woodrow T., 2010. Who incentivizes the mutual fund manager, new or old shareholders? *J. Financ. Intermed.* 19, 143–168.
- Khorana, Ajay, Servaes, Henri, 1999. The determinants of mutual fund starts. *Rev. Financ. Stud.* 12, 1043–1074.
- Lipper Associates, 2007. *Taxes in the Mutual Fund Industry – 2007: Assessing the Impact of Taxes on Shareholder Returns*. Lipper Associates, New York.
- Planchich, Stephanie, 2003. Mutual fund capital gain distributions and the Tax Reform Act of 1997. *Natl. Tax J.* 56, 271–296.
- Poterba, James, Shoven, John, 2002. Exchange traded funds: a new investment option for taxable investors. *Am. Econ. Rev.* 92, 422–427.

- Saxton, Representative Jim, 2004. Providing Tax Equity for Mutual Fund Investors: Changing the Tax Treatment of Capital Gain Distributions. U.S. Congress, Joint Economic Committee, Washington.
- Sialm, Clemens, Starks, Laura, 2012. Mutual fund tax clienteles. *J. Financ.* 67 (4), 1397–1422.
- Sirri, Erik R., Tufano, Peter, 1998. Costly search and mutual fund flows. *J. Financ.* 53, 1589–1622.
- Warther, Vincent, Warther, Vincent A., 1995. Aggregate mutual fund flows and security returns. *J. Financ. Econ.* 39, 209–235.