

The Effect of the Frequency of Holdings Data on Conclusions About Mutual Fund Management Behavior

by

Edwin J. Elton*

eelton@stern.nyu.edu

Martin J. Gruber*

mgruber@stern.nyu.edu

Christopher R. Blake**

cblake@fordham.edu

Yoel Krasny***

jkrasny@stern.nyu.edu

Sadi Ozelge***

sozelge@stern.nyu.edu

June 29, 2009

* Nomura Professors of Finance, Stern School of Business, New York University
** Joseph Keating, S.J., Distinguished Professor, Fordham University
*** Doctoral Students, Stern School of Business, New York University

Abstract

A number of articles in financial economics have used quarterly or semi-annual mutual fund holdings data to test hypotheses about investment manager behavior. This article reexamines four well-known hypotheses in finance to determine whether the results of prior tests of these hypotheses remain valid when higher frequency (monthly) holdings data are employed. The areas examined are: momentum trading, tax-motivated trading, window dressing, and tournament behavior. We find that the use of monthly holdings data rather than quarterly holdings data or, in the case of tournament behavior, holdings data rather than monthly return data, change, and in some cases reverse, previous results. This occurs because monthly holdings data capture a large number of trades missed by quarterly data (18.5% of the trades) and permit a more precise estimation of the timing of trades.

The empirical issues and topics examined in financial economics and possibly the conclusions reached are influenced by the type of data available. Past studies of mutual fund behavior have been constrained by the lack of availability of funds' holdings data at less than a quarterly interval. In this paper we use a new source of data which reports monthly holdings data to reexamine the conclusions reached in classic articles in each of several areas of financial economics dealing with mutual fund manager behavior. We do not attempt to examine all articles in any area, but rather have selected for each area a classic article to see if the conclusions are affected by a new type of data. We find in many cases the conclusions are affected by the frequency of available data.

Among the questions we examine are: 1) whether mutual funds trade on momentum (Grinblatt, Titman and Werners (1995)); 2) whether mutual funds trade for tax purposes (Gibson, Safieddine and Titman (2000)); and 3) whether mutual funds engage in window dressing (Sias (2006), O'Neal (2001) and Meier and Schaumburg (2004)).¹ All of these studies use mutual fund holdings data measured at quarterly or semi-annual intervals. There are three problems with the use of quarterly data (which are accentuated with the use of semi-annual data). First, quarterly holdings data miss within-quarter round-trip transactions (either a purchase followed by sale or a sale followed by a repurchase). This is a serious problem; we show the use of quarterly data misses 18.5% of a typical fund's trades revealed using monthly data. Second, the use of quarterly data means there is great imprecision in estimating the timing of a trade since a trade could occur anytime in the quarter. The use of monthly data allows a more precise estimate of timing and a better test of many hypotheses. Third, if an event might occur each quarter, such as window

¹ We also examine tournament models because monthly holdings data allows a more precise estimate of risk during partial years than monthly return data.

dressings, then the only way to test if the event has impact is to use intra-quarter data and test whether the impact occurs at the appropriate time in the quarter.

The purpose of this article is to reexamine some well-known hypotheses in financial economics, and the classic article in which each was tested, to see whether the results of prior tests of the hypotheses remain valid when higher frequency holdings data are used. We will find that (in many cases) the results change when monthly holdings data are used.

Section I describes our sample. Section II explores the amount of missed trades when alternative reporting periods are employed. In Section III we examine the results of Grinblatt, Titman and Wermers (1995), who, using quarterly data, showed that mutual funds on average trade based on momentum. We reexamine their results using monthly data and find that on average the effect of momentum essentially disappears and about as many funds trade against past price increases (contrarians) as trade based on these increases (momentum traders). In Sections IV and V we examine whether mutual funds engage in extra trading for tax purposes and whether they window dress. We find evidence of both tax trading and window dressing. The use of monthly data allows us to better analyze the existence, timing and the magnitude of tax trading and window dressing. Finally, Section VI examines tournament behavior, that is, whether underperforming mutual funds increase risk in the latter part of the year. This phenomenon has been studied using return data, and different studies have reached different conclusions. The use of a new methodology based on monthly holdings data allows a more precise estimation of portfolio characteristics each month and thus a more precise measurement of tournament behavior. We find that mutual funds that had high return in the first part of the year increase risk while low return funds decrease risk.

I. SAMPLE

Data on the holdings of individual mutual funds were obtained from Morningstar. Morningstar supplied us with all of its holdings data for all domestic (U.S.) stock mutual funds they followed during the period 1994 to 2005. The data are free of survivorship bias; for once a fund enters the database for the first time it remains in the database until it ceases to exist. Note that until 2008 Morningstar was the source for the CRSP data.

Our initial sample consisted of all common stock funds that reported at least two consecutive years of monthly holdings. From this sample we eliminated all funds for which the aggregate asset value reported in Morningstar or CRSP did not closely match the aggregate asset value computed using holdings data.² Note that the Morningstar data include all holdings of mutual funds, including traded equity, bonds, preferred stock, options, futures, non-traded equity and cash.³ We eliminated all funds which held less than 93% of their assets in U.S. stock and cash in any month. Next, we eliminated all index funds and specialty funds. Finally, we eliminated one fund that had a stated policy of holding a fixed portfolio over the year and rebalancing only once a year. This left us with a sample of 215 funds and 6,432 fund months.

Table 1 shows the dispersion across ICDI objectives. The most common objective is long-term growth with almost half of our sample in that category. Aggressive growth represents about one-third of our sample, and growth and income about 19% of the sample.

In analyzing window dressing it is important to know each fund's fiscal year to detect if there is excessive turnover prior to the date on which holdings are computed for reporting

² "Closely matched" is defined as aggregate asset value from holdings data within 3% of that reported by CRSP or Morningstar. Note that the average difference between Morningstar and CRSP in the reported aggregate asset value was about 3%, with neither being consistently higher or lower.

³ Previous studies using holding-period data have employed the Thomson database. Thomson only reports holdings for traded equity. The use of this data means that conclusions about management behavior are based solely on the equity holdings of a set of funds, many of which hold a significant portion of their assets in other securities. Most funds hold other investments. The presence of these other securities can significantly change the characteristics of the portfolio.

purposes. Table 2 shows the distribution of fiscal year ends.⁴ Although December is the most common month for fiscal year end, only 27% of the funds choose a fiscal year ending in December. September and October are the next most commonly chosen months for fiscal year ends. Almost 45% of the funds choose a fiscal year ending in one of these two months.

In the paper we sometimes find results that vary quite significantly from those reported in prior studies. Before presenting these results, we need to discuss whether the differences we find are likely to be caused by differences in the data used. There are two issues with the data that need to be addressed. First, since most prior studies use Thomson data and this study uses Morningstar data, how do these databases compare? Second, since we use a sample of the Morningstar data (those funds that report monthly data), are there any differences between the sample and the population that are likely to impact our results?

First we examine the differences between Morningstar and Thomson. We compare the number of stock funds reporting quarterly data in the Morningstar database to the number of stock funds reporting to Thomson. For the first quarter of 2003 Morningstar reported quarterly data for 3,273 funds that they classified as common stock funds. This includes funds that are commonly classified as growth, aggressive growth, balanced, growth and income, international equity, metals, etc. When we examined the data provided by Thomson for this same period of time, we found quarterly holdings data for 3,600 common stock funds in these categories.⁵ The Thomson database includes closed-end funds, whereas Morningstar segregates closed-end funds into a separate database. Adjusting for this difference accounts for most of the difference in the number of funds that report quarterly data in each database. Examining the holdings data that are

⁴ Some funds have fiscal years ending in different months in different years. In reporting fiscal years, we include all funds in our sample whether or not we have monthly holdings data over the period. Thus we include roughly the same funds each year, and our results are not due to different numbers of years for different funds.

⁵ This is an estimate, for over half the funds in the Thomson database are not classified. We sampled these funds to estimate what percentage were bond funds and what percentage were equity funds.

reported in the two data sets shows that the Thomson database does not include holdings of preferred stock, convertibles, warrants, bonds, options, futures, non-traded equity and cash. Thus the Morningstar database is much more complete. To determine what difference this makes, we examined a random sample of 87 funds that were included in both databases. For these funds Thomson does not include holdings for 10% of the assets held. Also, in 10% of the cases Thomson misses holdings for 30% of the assets in a fund, and in one case (a fund that makes extensive use of futures and cash to replicate indexes) misses over 90%. The principal reason for the differences in the two databases is the inclusion by Morningstar of bonds, futures, non-traded equity, and options holdings.

The question remains as to whether funds in the Morningstar database that report at monthly intervals are fundamentally different from funds that report only quarterly. To explore this question, we need to examine why funds might not report monthly. Three reasons have been proposed for why funds do not report more frequently: cost, front running and free riding (see Ge and Zheng (2006) and Wermers (2001)). The cost of reporting monthly is trivial, since mutual funds have data on their holdings at the end of each day.

Front running can be explained as follows. Assume a mutual fund takes days to purchase a position. Assume further that investors, by examining changes in fund holdings, can infer what securities the fund will be purchasing. An investor can then purchase the stock before the fund does, driving up the price and causing the fund to pay more. The potential for front running is easily mitigated by reporting holdings with a delay. The average delay for the reports in our sample is two months, with a range from zero to five months. Thus, unless funds take months to purchase or sell a position, front running is not likely to be an issue.

Free riding or mimicking involves copying the holdings of a mutual fund. Mimicking is

advantageous if a fund can benefit from another fund's security selection without bearing the cost of doing research. Frank, Poterba, Shackelford and Shoven (2005) show that copycat funds can earn a substantial portion of the return of the fund being copied. The ability to successfully copycat depends on how quickly the private information of the fund manager is incorporated in price. If it is incorporated within a couple of months, then the normal delays funds employ in reporting should eliminate most of the benefit to copycat funds. In addition, since the length of delay in monthly reporting is largely controlled by the funds, funds have the ability to determine the amount of delay.⁶ They should choose a time for release that does not help the copycat fund.

Finally, consider why a fund might be concerned with investors or other funds copying its portfolio. The concern is that funds or investors can duplicate the portfolio at a lower cost (since they don't incur analyst costs) and therefore capital will flow out of the fund. However, Ge and Zheng (2006) show that funds that report more frequently (in their case quarterly rather than semi-annually) have greater inflows, not smaller inflows. They speculate that this occurs because investors can better evaluate the strategy of the fund (growth, value, midcap, etc.) and hence can better estimate asset allocation. An alternative explanation is that there are very few copycat funds because of the delay in obtaining data and that reporting more frequently is viewed as good corporate governance. Thus the costs of more frequent reporting should be small.

If the costs of reporting monthly are minimal, why don't all funds report monthly? Perhaps the biggest reason is inertia. However, the impact of inertia is changing, as seen from the fact that the fraction of funds that report monthly continues to increase. From 1998 to 2005 the fraction reporting monthly increased from 3.5% to 18%. We can explore possible types of differences in funds that report monthly rather than quarterly by examining Ge and Zheng (2006). They found four differences in funds that voluntarily reported quarterly rather than semi-

⁶ The SEC has set a maximum 90-day period for reporting of mandated holdings disclosure.

annually. They found these funds had 0.04% lower expenses, 10% less turnover, were less likely to engage in fraud, and differed in performance. To determine whether some of these same characteristics were present when we examined funds that engage in voluntary monthly reporting rather than mandatory quarterly reporting for each of the funds in our sample, ten funds that only reported quarterly but had the same objective code were drawn at random. This gave us a comparison sample of 2,150 funds. We ran the standard Fama-French model and the Fama-French model with a momentum variable added. We then tested differences in alpha across the two groups. There was no statistical or economic difference between the two groups. Like Ge and Zheng, we did find slightly lower expenses and turnover for funds that reported more frequently.⁷ In later sections of the paper we will return to the question of bias when interpreting our results.

II. MISSING TRADES AND TURNOVER

One of the principal disadvantages of using quarterly (or semi-annual) holdings data to examine the decisions of mutual fund managers is that quarterly or semi-annual data miss round-trip trades (purchase and sale or sale and repurchase) that take place within the quarter or half year. In this section we will examine how much trading is missed by not using monthly data. As we will show in latter sections of this paper, missing intra-quarter trades changes the results for several of the hypotheses tested in the literature. To examine the amount of missed trades we use a commonly employed metric in the industry: turnover.

⁷ Looking at names in the two groups suggests more institutional funds report monthly, and this may account for the difference. However, we were unable to identify funds by institutional holdings in any meaningful way, so this could not be formally tested. Bias could also be generated if poorly performing funds stopped reporting monthly and started to report quarterly, and thus would be excluded from the sample. To measure this, we examined all funds that had one year of monthly data but did not have two years, and examined the reporting interval the subsequent year. In only four cases out of over 600 observations was it quarterly. In the vast majority of the cases a fund was not included in a year because it reported 10 or 11 months of holdings data rather than a full 12 months.

Turnover is defined by industry practice (see, for example, Morningstar) as the lesser of purchases or sales (excluding all securities with maturities less than one year) divided by the average monthly net asset value. These numbers are calculated by the individual funds and supplied to data purveyors.⁸

To be more precise we calculate turnover in a period as follows:

$$C_t^+ = \sum_i (N_{it} - N_{it-1}) \bar{P}_{it} \text{ for all } i \text{ where } (N_{it} - N_{it-1}) \geq 0 \quad (1)$$

$$C_t^- = \sum_i (N_{it-1} - N_{it}) \bar{P}_{it} \text{ for all } i \text{ where } (N_{it} - N_{it-1}) < 0 \quad (2)$$

where

1. N_{it} = the number of shares of stock i held at the end of month t ;
2. \bar{P}_{it} = the average of the prices of stock i at the beginning and end of month t .

For any year

$$C^+ = \sum_{t=1}^{12} C_t^+ \quad (3)$$

$$C^- = \sum_{t=1}^{12} C_t^- \quad (4)$$

Turnover equals the smaller of C^+ and C^- divided by the average of the net asset value of the portfolio over the 13 end-of-month values in the year (including December 31 of the prior year).⁹

⁸ Turnover, as computed by mutual funds, uses the exact timing and prices at which purchases and sales take place. There has to be some loss of precision in estimating turnover even when monthly data are used because the transaction price is not known and we use the average price over the month to approximate it. However, deviations of average price from actual price are somewhat random. Thus, our final estimates using mutual fund holdings data closely match reported turnover.

⁹ Using the average of 13 net asset values is necessary to compare our results to reported numbers, since this is industry practice.

To estimate turnover based on quarterly, semi-annual, or annual data, the same formulas are used, but equations (1), (2), (3) and (4) are estimated over the appropriate intervals. For example, to obtain a quarterly estimate equations (1) and (2) are estimated using three-month changes in holdings, and equations (3) and (4) are summed over the four quarterly values (still dividing by average monthly net asset value over the year).

Table 3 presents the distribution of the turnover estimated from monthly, quarterly, semi-annual and annual data. Differences in these estimates result from missing round-trip transactions in individual stocks over the period being examined. Note that as expected due to missing trades, the larger the interval between observations of holdings the lower the estimate of turnover. Monthly holdings data are the most frequent holdings data available to us, and leads to estimates of turnover very close to those reported by the mutual funds themselves.¹⁰ This suggests that management does not perform many round-trip transactions (purchase followed by sale or sale followed by purchase) within any given month. Table 3 allows us to estimate how much turnover we miss by having less frequent data. Use of quarterly data misses 18.5% of the trades estimated using monthly data. Similarly, using semi-annual holdings or yearly holdings to measure turnover misses respectively 34.2% and 55.4% of the trades estimated using monthly data. The differences in means for turnover measured at different intervals are all statistically significant at the 1% level. Clearly, studies that employ semi-annual or quarterly data to study phenomena such as tax effects or momentum are overlooking a large proportion of trades.

III. MOMENTUM

¹⁰ The numbers from monthly data are actually slightly higher than the turnover numbers reported by mutual funds. The difference is caused by a small difference in the price used.

Grinblatt, Titman and Wermers (1995) (hereafter GTW) examined the issue of whether mutual funds tended to buy stocks that went up in the past and sell stocks that went down in the past. This is related to the more general issue of herding (see Lakonishok, Shleifer and Vishny (1992)); namely, do institutional investors tend to trade in the same stocks at the same time (herd)?

If institutions herd, a concern is that there will be a price impact, leading to a much more volatile stock market. If institutions herd by buying what other funds bought in prior periods, then the purchases will be spread over time and the price impact is likely to be relatively small. However, if they all follow the same mechanical trading rules, they are likely to purchase or sell at the same time and the potential impact will be large. The mechanical rule suggested, and the one that researchers have argued produces excess returns, is momentum (see Jegadeesh and Titman, (1993)). Momentum involves buying stocks that have significantly increased in price in the past and selling those that have significantly decreased in price. The most widely used measure of momentum for mutual funds, and the one we will explore in this paper, is that proposed by GTW. We present evidence in this section that employing monthly rather than quarterly holdings data with the GTW momentum measure yields different results. Before we measure momentum with monthly holdings data, we repeat the GTW measure with quarterly data. The GTW measure of momentum for a single fund is

$$m = \frac{1}{T} \sum_{t=1}^T \sum_{j=1}^{N_t} \left[\frac{S_{jt} \bar{P}_{jt}}{A_t} - \frac{S_{j,t-1} \bar{P}_{jt}}{A_{t-1}} \right] \times R_j \quad (5)$$

where

$$A_t = \sum_{j=1}^{N_t} S_{jt} \bar{P}_{jt}$$

$$A_{t-1} = \sum_{j=1}^{N_t} S_{jt-1} \bar{P}_{jt}$$

and

1. m is momentum.
2. S_{jt} is the number of shares held in asset j at time t .
3. \bar{P}_{jt} is the average of the price of security j at time t and time $t-1$.
4. N_t is total number of individual assets held by the fund at time t and time $t-1$.
5. T is number of periods.
6. R_j is return on asset j over alternative time (contemporaneous or lagged) periods as explained below.

GTW employ quarterly data to estimate equation (5). The use of the same price at the beginning and end of any quarter prevents having a weight increase over the quarter simply because the return on the security was above the average return in the portfolio. GTW measure return in equation (5) in two ways: as three-month contemporaneous return and as the three-month prior return. The three-month contemporaneous return is measured as the return over the quarter in which the change in portfolio investment is calculated. Which measure is correct depends on when the trade took place. Contemporaneous measurement implicitly assumes that the purchase or sale took place on the last day of the quarter. If the purchase took place earlier in the quarter, then part of the return that is assumed to occur before purchase or sale actually took place after the trade. If management has the ability to select superior securities, then using data after the purchase will cause the contemporaneous GTW measure to show that management was buying on momentum when actually they were simply good forecasters. Narrowing the

uncertainty of when the security was purchased or sold from a three-month period to a one-month period greatly reduces the amount of post-trade data that might be included.

Alternatively, GTW computed the return using data for the three months prior to the start of the quarter over which the weight change is calculated. The disadvantage of this measure is that, if the trade takes place at the end of the quarter, the data on return can be very old and therefore not the data a manager using momentum is likely to be examining in determining the buy and sell decisions. Using monthly holdings data narrows the time frame over which a trade could have taken place and means that such data are likely to be used by the manager following a momentum strategy.

We will compute both of the GTW measures using quarterly holdings data, exactly as they did, and we will also compute them using monthly holdings data. When we use monthly data to compute changes in holdings, we define a contemporaneous period as that three-month period ending at the end of the month over which holding changes are measured; we define the prior period as the return for the three months ending at the beginning of the month in which we measure the change in holdings.¹¹ The problem associated with using contemporaneous return should be partially mitigated, but not eliminated, by using monthly data.

Using quarterly estimates of momentum, we compute equation (5) using only holdings data four times per year (every quarter) as in GTW. While applying equation (5) will correctly

¹¹ There is a second question that could be raised with the GTW measure. Mutual funds have inflows and outflows over the quarter. Assume that over the quarter there were inflows equal to 4% of total net assets. Then any security held in the portfolio where the shares were not increased by this amount would enter the GTW measure as sold. Likewise, if there is an outflow over the quarter, all existing stocks will appear to have involved additional purchases where there is no change in the shares held. We also examined an alternative to the GTW measure that does not have this property, but rather measures momentum using as purchase and sales actual changes in shares.

The alternative measure is $m = \frac{1}{T} \sum_{t=1}^T \sum_{j=1}^{N_t} \left[\frac{(S_{jt} - S_{jt-1}) \bar{P}_{jt}}{(A_t + A_{t-1})/2} \right] \times R_j$. We repeated all our tests using this measure.

The results were very similar in their pattern, and therefore are not reported here.

measure momentum, over both quarterly and monthly intervals the momentum measures from these two intervals are not directly comparable. If using quarterly and monthly data captured the same trades, then the momentum measure using holdings computed 12 times a year would be one-third the size of the momentum metric computed four times per year. Therefore, the results from applying equation (5) to monthly holdings data must be multiplied by 3 to be directly comparable to those computed using quarterly intervals.

Table 4 presents the momentum measure based on examining changes in holdings using both quarterly data (holdings changes measured over a three-month period) computed as in GTW, and monthly holdings data. In Panels A and B three-month returns are used, where return is measured either for the period ending in the last month of the period over which change in holdings is computed (labeled “contemporaneous”) or the three-month period prior to the period over which change is computed (labeled “prior”). Panel A is based on quarterly holdings changes, and Panel B is based on monthly changes.¹² In Panel C we compute momentum measures using the one-month change in holdings, with return defined contemporaneously and in each of the prior three months. This makes apparent which month is contributing to the momentum factor.¹³

Examining Panel A for quarterly data and contemporaneous returns, we find general support for the momentum hypothesis. The momentum measure is equal to 0.273 and statistically significantly different from zero at the 1% level. The second column of Table A shows that when we measure returns over the prior three months the momentum measure is

¹² The mean numbers in Table 4 are the average of the individual fund numbers, each weighted by the ratio of one over the variance divided by the sum of one over the variances for all funds. This puts more weight on the more precisely measured funds. The weighting on all funds adds to one.

¹³ In a later section we show that funds window dress in the month before the end of the fiscal year by selling losers and buying winners. This might affect the momentum measures and might make it appear that funds follow momentum when they are window dressing. While this does affect the results in the hypothesized direction, the changes in results are quite small.

much smaller (0.043) and no longer statistically significant. When quarterly holdings data are used there is much more uncertainty as to when the trade occurs, and many trades are missed. When returns are measured prior to the change, the strength of the momentum factor is weaker. Thus, using quarterly data, one would reach different conclusions depending on the belief of when the trade took place (beginning or end of quarter) and whether returns are lagged or not. Examining momentum on a fund-by-fund basis shows that for 152 of the 215 funds using contemporaneous return and for 137 funds using prior return, the momentum measure is positive.

Examining Part B, which reports the momentum measure using the change in monthly holdings, shows that the overall contemporaneous momentum effect is no longer statistically significant, and its magnitude is decreased by a third.¹⁴ The difference between monthly and quarterly momentum measures is due to missed trades and more precise measurement of the time of the trades. There is an extremely large decrease in the momentum factor when prior rather than contemporaneous return is employed; in fact, it is close to zero. .

Monthly data and contemporaneous returns show that 31% of funds appear to trade against momentum, while with lagged return 41% of funds appear to trade against momentum. Thus a large percentage of fund managers seem to be contrarians (trading against momentum) rather than momentum traders.¹⁵

Panel C in Table 4 separately presents measures of momentum using monthly holdings and returns calculated contemporaneously and in each of the prior three months. Changing the period over which returns used to compute momentum are calculated from three months to one month allows identification of which months are contributing to the overall measure. Table 4C

¹⁴ Note that the numbers in Parts B and C are reported on the monthly momentum measures multiplied by three (put on a quarterly measurement unit) so that they can be directly compared with the quarterly momentum measures.

¹⁵ We also examined the momentum measure in months after the purchase or sale. The momentum measure was positive, consistent with forecasting, but it was very small and not statistically significant.

shows that the positive values for the overall contemporaneous measure using monthly data are the result of returns occurring in the month over which change in holdings is measured.¹⁶ These returns could well take place after the trade. This casts further doubt on whether a contemporaneous measure is, in fact, measuring momentum at all.

In summary, the presence of momentum trading uncovered using quarterly changes in holdings all but disappears when monthly changes in holdings are used instead. A plausible explanation for this is that quarterly holdings imprecisely measure both the amount and time of trades, incorporates too much post-trade return, and misses intra-period trades.

IV. TAX-MOTIVATED TRADES

Another issue which has been examined, using holdings data, is whether mutual funds engage in trades motivated by taxes. This topic is interesting by itself, and carries further significance because tax trades could be construed as window dressing for those funds with an October or December fiscal year.¹⁷ Why might mutual funds want to trade for tax reasons in a particular month? Mutual funds meeting certain criteria, including the paying out of at least 98% of all dividends and capital gains to their investors, don't pay taxes. However, investors who receive these payments are taxed on them. Given this, why should a mutual fund care about capital gains and losses? One reason is a desire by the fund to maximize assets under management. A second reason may be a desire to postpone capital gains reported to shareholders so that shareholders can exercise their tax timing option. Mutual fund management company compensation is generally in the form of a percentage of assets under management. A typical fee

¹⁶ When the results are repeated excluding the month in which the fiscal year ends, the results are very similar. The contemporaneous month is economically large and statistically significant. No other month is economically or statistically significant at conventional levels.

¹⁷ The most relevant antecedent is Gibson, Safieddine and Titman (2000), which examines this issue directly. For related articles, see Chan (1986), Lee (1998), Huddart and Narayanan (2002), and Ritter and Chopra (1989).

would be a payment to the management company of 1.2% of assets under management. Thus management companies have an incentive to keep assets. Consider a mutual fund that has incurred a capital gain. If they pay it out they lose fees on all assets for shareholders not members of an automatic dividend reinvestment plan. The fund can reduce the capital gains to be paid out by taking capital losses and incurring some transaction costs. Since transaction costs are likely to be small, there should be little impact on returns and hence on future inflows to the fund. In addition, taking capital losses to offset gains minimizes realized capital gains, and since investors are taxed on realized gains, this minimizes the present value of any tax investors must pay and thereby helps the fund's rankings among rating services that rank on after-tax returns. Thus, rationally we would expect an increase in trading activity in key (tax computation) months for funds that have capital gains.¹⁸

What if the fund enters a key month with a capital loss? Capital losses are not distributed to investors. However, they can be carried forward by the fund, reducing future gains and future disbursements. Since losses can be carried forward, there is no advantage in taking gains in order to immediately offset losses since this involves transaction costs and changing the portfolio. We would expect that most funds should have no tax-motivated trading in loss years.¹⁹

What are the key months in which tax laws encourage trading activity? Mutual funds are subject to two sets of tax regulations that could lead to tax-motivated trades. The first, under Section 52 of the Internal Revenue Code, states that to qualify as a registered investment company a fund must pay out 90% of its investment income computed as of the end of the fund's

¹⁸ There is one influence that could potentially act in the opposite direction. Barclay, Pearson and Weisbock (1998) made the case that potential investors are concerned about tax overhang and thus prefer to see capital gains realized early rather than later. This would suggest that mutual funds might choose not to realize losses in order to pay out more capital gains. While this could potentially mitigate the amount of tax-loss selling at the end of the tax year, the data soon to be presented clearly indicate that reducing current taxes by tax-loss selling plus raising current after-tax returns are in aggregate the more important influences.

¹⁹ There is a limit to how far a fund can carry the loss into the future. We would expect funds that have expiring capital losses to offset these by taking capital gains.

fiscal year. This was the operative section of the Internal Revenue Code until Section 4982 of a revised code was passed as part of the Tax Reform Act of 1986. Under the revised code (subject to a short adjustment period) an excise tax of 4% was placed on all mutual funds that did not distribute by the end of the calendar year 98% of ordinary income received in the calendar year and 98% of capital gains realized in the 12-month period ending on October 31. This new rule was passed to stop funds from realizing gains in one year but postponing the tax liability to the following fiscal year. The new tax rule had an additional complication: funds that had a fiscal year ending in November or December could elect to use their fiscal year rather than October 31 as the basis for the excise tax computation. Discussions with practitioners indicate that few funds take this alternative.

The vast majority of mutual funds calculate their capital gains tax on October 31, though some will have tax computations as of December 31. (Few have fiscal years ending in November.) This suggests that to the extent a fund wishes to modify its tax position, we should see tax-motivated transactions in September and October.

The use of monthly holdings data allows us to examine the extent to which funds which have accumulated realized capital gains take capital losses to offset these gains. We first form a variable which represents the gains against which a fund might want to take losses in September or October. This variable is the sum of all realized gains for the tax year November through October minus all realized capital losses taken in November through August.

The purchase price was calculated as the average of the beginning and ending price in the month the stock was purchased. Gains or losses are calculated by comparing the price computed as the average of the beginning and ending price in the month of sale to the purchase price.

The question we examine is whether the accumulated net realized capital gains affect the size of losses taken in September and October.²⁰ To address this question, we divide our sample into two groups: those funds that had an accumulated net realized capital gain and those funds that had a net realized capital loss excluding any losses occurring in September and October. As shown in Table 5, funds with net gains take unusually large losses in September and October. The average monthly loss taken in September and October for funds with net realized gains is 3.95 times the average losses taken in the other ten months. This is over three times larger than the ratio for the group with net realized losses. This difference is statistically significant at the 0.01 level. When the losses are examined only for the month of October, even stronger results are obtained, with the average loss in October over five times the average loss in the other months.

These numbers clearly indicate that funds that would have to disburse taxable capital gains to their investors without tax-loss selling in September and October increase their tax-loss selling in these months much more than funds that do not have accumulated gains.

Next we examine, for funds that have accumulated capital gains, whether the magnitude of tax-loss selling is related to the potential tax liability of the fund. To examine this we regress the average of the tax losses taken in September and October divided by the average loss in the prior 10 months against the gain that would be subject to taxes had sales in September and October not taken place divided by the average loss in the prior 10 months.²¹ Define normal loss

²⁰ A strong indication that tax-loss selling might be taking place is the large increase in turnover in October. Excluding October fiscal year funds (to avoid picking up window dressing), October turnover is significantly higher than turnover in the average month. The turnover is not caused by shareholder cash flows. In our sample, October has lower net cash flows than the average month, while December has the highest cash flows. Furthermore, the net cash flows were positive in October which would lead to net purchases, not sales.

²¹ We also ran the analysis with only October losses. The results are similar with the t value associated with the independent variable dropping to 6.16. In addition, to ensure that we have appropriately removed the impact of size on the relationship, both regressions were re-run with total asset value as a second independent variable. The results are virtually unchanged. The t value associated with the size variable is less than 0.15.

as the average monthly loss in the 10 months prior to September. Then the dependent variable (RL) is the number of months of normal loss taken in September and October. The independent variable (C_1) is the number of months of normal loss that would be needed to eliminate all gains.²² The results are:

$$RL = -0.2155 + 0.0993C_1$$

where $R^2 = 0.518$ and the t values for the intercept and coefficient are -0.27 and 13.51 , respectively.

The regression results indicate a large and statistically significant positive relationship between tax-loss selling and capital gain overhang.

Our results differ from those reported by Gibson, Safieddine and Titman (2000). There are two reasons for this difference. We measure gains and losses much more accurately, and we examine tax-loss using monthly rather than quarterly holdings data. GST assume that stocks that were in the lowest quarter of all stocks based on the prior year's rate of return are the ones with capital losses and hence the candidates to sell to take a tax loss. We define loss from actual purchase prices of securities rather than assuming all securities were purchased as of a specific calendar date as GST do. They show that funds in aggregate sell more of these stocks in the second and third quarter of the year than in the fourth quarter. This is somewhat surprising if tax-loss selling is the motivation for sale, since the key tax month, October, is in the fourth quarter. They hypothesize that the drop in the fourth quarter sales of stocks with no losses is due to funds repurchasing in the months of November and December stock they had previously sold. When we compute the loss realized each quarter from the sale of securities that had accumulated losses using our definition, we find slightly greater dollar losses realized in the third and fourth quarter.

²² This formulation has the added benefit of the regression not being affected by the size of the fund.

However, losses in all quarters are very similar and not statistically different. Thus, using quarterly data, even with a better measure of loss, a researcher would find no evidence of tax-loss selling. Our study, by using monthly holdings data, allows us to isolate the September and October period (months which fall into two different quarters) and to show how significant tax-loss selling occurs in these two months. Furthermore, using monthly data we are able to measure accumulated gains and losses on each stock in a fund's portfolio and thus the size of potential tax losses and gains which can be taken in any portfolio at the end of the tax year. Finally, our methodology allows us to examine, cross-sectionally, tax-loss selling as a function of capital gains accumulated by each fund. We find strong evidence of tax-loss selling in the two months preceding the end of the tax year.

V. WINDOW DRESSING

A number of researchers examine whether mutual funds window dress (see, for example, Musto (1997, 1999); He, Ng and Wang (2004); Lakonishok, Shleifer, Thaler and Vishny (1991), Meier and Schaumburg (2004), Sias and Starks (1997), Sias (2006), and O'Neal (2001)). Most studies of window dressing do not examine equity mutual funds. Window dressing involves selling securities that have done poorly before holdings are reported to investors and adding securities that have done well. This gives the impression to fund investors that the fund is holding winners and allows investors to infer that the fund held the stocks during the time the stock had superior performance. Window dressing implies excessive trading prior to report dates. Until recently mutual funds were required to report their holdings only twice per year in their annual and semi-annual reports.²³ This has recently been changed to mandatory quarterly

²³ All but one of the prior studies of stock mutual funds (O'Neal (2005)) have assumed funds report their holdings as of December 31. As shown in Table 2, only 37.9% of the sample has a December fiscal year. About 69% of our

reporting. Mutual funds are required to report their holdings as of the end of their fiscal year. Thus a fund with an August fiscal year would issue its annual report stating its annual holdings as of the end of August and a semi-annual report listing holdings as of the end of February.

Monthly data are essential to the study of window dressing for two reasons.²⁴ First, if funds window dress in the month before each quarterly report, then observations on trades within the quarter (between reports) are necessary to study whether funds window dress. Second, if window dressing is present (or much greater) only at the time of the annual report, it is still difficult to detect with quarterly data, since the month of window dressing is mixed with two months where it is not present.²⁵ Monthly holdings data allow a more precise measure of trading near the end of the quarter.

In this section we test the tendency to window dress in three ways. First we examine whether funds add and delete securities more actively in months before holdings are reported. Second, we examine whether funds tend to buy winners and sell losers in the month before holdings are reported. Third, we examine whether funds avoid adding illiquid assets in the month prior to the fiscal year end.

Table 6 presents evidence on window dressing. Panel A of Table 6 shows the number of new stocks added to the portfolio plus the number of stocks eliminated from the portfolio in the month directly preceding each quarter and year-end dates. We refer to this as new turnover. Given that we find evidence of tax-motivated trading, in computing Panel A of Table 6 we eliminate funds where the fiscal year end coincides with a key tax month. In Table 6 we report

sample will report holdings in December, but almost half of these reports will contain quarterly or semi-annual holdings.

²⁴ The funds in our sample report holdings to Morningstar, but not to their shareholders, on a monthly basis. The holdings are reported with no delay up to five months' delay. Thus an analyst getting monthly data from Morningstar could get holdings at a common point in time by merging data from Morningstar over six months. However, the holdings data shareholders see are tied to the fiscal year.

²⁵ The difficulty of measuring window dressing with quarterly holdings data is probably why more investigators have chosen to study window dressing for types of non-stock funds for which monthly data are available.

the new turnover at the end of the first, second, third and fourth (year-end) quarter as well as the average monthly new turnover. The year-end new turnover is the highest of any month during the year and differs from the average monthly turnover at the 0.01 significance level. Two of the other quarter ends have results below the average, and in the remaining quarter (first) the results are less than one standard deviation above the mean.

While there is evidence of annual window dressing, there is no evidence of window dressing in any other quarter.

The second way we examine window dressing is to examine the past returns on securities bought minus the past return on securities sold in each month prior to the fiscal year-end month (t). The period for measuring return was the three-month period prior to the month under study. The presence of window dressing would be indicated by a much higher value in the final month than those found in most other months.

Panel B of Table 6 presents the ratio of the measure described above to its average value in each month of the fiscal year. The difference in return between securities bought and those sold in the last month of the fiscal year is nearly twice the average of all months.²⁶ It is higher than all other months except for month $t-7$, which has essentially the same value. The results for quarters 1, 2 and 3 show no pattern of window dressing. We have found results consistent with window dressing at year-end, but not at earlier dates.

We found no evidence of quarterly or semi-annual window dressing. However, it is possible that it is obscured because only a few funds window dress and quarterly or semi-annual window dressing is not apparent when these funds are analyzed as a part of the larger sample. To further examine quarterly and semi-annual window dressing, we constructed a sample of funds which have a greater than average tendency to window dress on an annual basis. The sample

²⁶ The difference is statistically significant at the 10% level.

consisted of all funds where the return on purchases minus sales in December was above the average for other months. The results are shown in Panel C of Table 6. The return on purchases minus sales is not abnormally high at the end of any quarter except for the year end. In fact, all the quarterly differential returns are below the mean of the 11 months prior to the end-of-year report.

Final evidence of fiscal year window dressing can be seen by examining the quantity of illiquid assets (e.g., restricted stock) purchased or sold in the month before fiscal year end. Illiquid assets are likely to be less well known, to involve higher transaction costs when sold, and to be perceived as more risky than liquid assets. The percentage of a portfolio that is added in illiquid assets in the final month is 0.274%, where in other months the percentage added is 0.371%. The difference between these is highly statistically significant, although of small size. Likewise, the percentage of illiquid assets sold in the final month of the fiscal year is higher than average, although the difference is not statistically significant. Mutual funds take actions that reduce the quantity of illiquid assets prior to issuance of the annual report.

Trading is highest in the final month of the fiscal year when funds have a greater tendency to dump losers and buy winners and to take actions that reduce their illiquid holdings. All of this is indicative of fiscal year window dressing. We find no evidence of semi-annual or quarterly window dressing. This is consistent with the belief that investors pay more attention to annual reports than to interim reports.

Would funds that voluntarily report on a monthly basis be more or less likely to window dress? Window dressing can be detected with monthly data and is difficult to detect with quarterly reporting. Thus funds that are concerned about having window dressing detected

should not report monthly. Thus, if there is a bias in our sample it should be against finding window dressing.

VI. TOURNAMENT MODEL AND MUTUAL FUND BEHAVIOR

One of the more interesting hypotheses about mutual fund behavior is that in order to be viewed as a top-performing fund (win the tournament), managers will change their behavior during the year. Because of the large payoff from being a top-performing fund versus the small penalty from falling further in the rankings, it is argued that mutual fund managers will increase the risk they take during a year if they are doing poorly or decrease the risk if they are doing well. Brown, Harlow and Starks (1996) (hereafter BHS), in a classic paper in finance, tested this hypothesis and presented evidence supporting it.²⁷

Several authors have questioned the theory, methodology and empirical results of BHS. In terms of theoretical underpinnings, Taylor (2000) constructed a game-theoretic model in which winners may increase risk more than losers. The different implications of the model arise from the assumption that funds that are doing well will take into consideration that poorly performing funds will increase their risk.

Perhaps the principal challenge to BHS comes from the research of Chevalier and Ellison (1997). They study the relationship between new cash flows and returns, and find that it is non-linear. They conclude that, given this relationship, the incentive to change risk is more complex than that hypothesized by BHS. They construct and test a model that hypothesizes that a fund that slightly underperforms the market will increase risk in the hope of beating the market, while a fund that slightly outperforms the market will decrease risk to lock in gains. However, this risk

²⁷ This phenomenon should be even stronger for funds that have incentive fees, and it has been shown to exist for these funds (Elton, Gruber and Blake (2003). While Elton, Gruber and Blake can use a three-year window because of the way incentive fees are computed, BHS had to employ a one-year window.

relationship is reversed for more extreme fund behavior. They hypothesize and test the notion that funds with returns well above the market will increase risk, while those with returns well below the market will decrease risk. The former change occurs because of the extreme payoff from winning the tournament as opposed to being one of the last knights killed in the jousting competition. The latter may be due to funds not wanting to risk being one of the worst funds. While they hypothesize and find evidence for this non-linear relationship, when they perform a time series analysis they find that risk-taking is positively related to past return.

Busse (1996) questions the empirical results reported by BHS, and in particular the negative relationship between first period returns and subsequent risk changes. Busse finds that these conclusions are functions of the data selected for the study, the return data used to calculate standard deviation (daily returns versus monthly), and the definition of starting and ending dates within a month. As an example, when Busse uses daily rather than monthly returns to compute standard deviation, the relationship between return and subsequent risk-shifting found by BHS reverses itself.

Though the logic behind the BHS hypothesis is appealing, there is both theoretical and empirical evidence that suggests more complex hypotheses, and empirical research results opposite to that found by BHS.

Monthly holdings data allow a much more precise measurement of the BHS risk-shifting hypothesis because we can observe every month exactly how the managers are changing their portfolios. We will follow BHS in looking for evidence of risk-shifting in the last five months of the year versus the first seven months.²⁸

²⁸ As in BHS, we tried leaving out the last month. This did not affect the results. We use the seven/five split because this is what BHS found worked best.

Using holdings data, not only can we more precisely measure whether risk has shifted, but we can observe how a fund shifts risk. Funds can shift risk in several ways. First, the manager could shift risk by changing the proportion invested in cash rather than risky assets. Second, the manager could alter risk characteristics by changing the riskiness of the equity portfolio. Both of these will be examined in turn. In all of the analysis which follows we divide for each year all funds into a high-return group and a low-return group where the designation depends on whether the return in the first seven months of the year was above or below the median value.

1. Change in Asset Properties

In Table 7 we present the proportion of assets held in cash and stock for the last five months of the year compared to the first seven months of the year. Note that for the low-return group the percentage invested in the risky asset goes up while the percentage in the safe asset goes down. The same is true for the high-return funds. However, the change in asset weighting is so small that it has no perceptible impact on risk. If management is not risk-shifting through the change in asset mix, perhaps it is doing so by changing the risk characteristics of the non-cash portfolio.

2. Change in the Risk of the Stock Portfolio

Management could change the risk of the portfolio by adding and subtracting securities to change the standard deviation of the portfolio or the sensitivity of the portfolio to market returns. Because we know the holdings in a given portfolio at the beginning of each month, we can estimate the sensitivity of the portfolio to the market or its standard deviation using the actual securities held.

We follow BHS by dividing funds into four groups based on return and risk to examine risk-shifting. In particular we first divide funds into two groups based on return: those above median return in the first seven months of the year and those below median return. For each fund in each of these return groups we form two new fund groups based on whether a fund has increased or decreased risk more than the median fund in that year. More specifically, we define the shift in risk for any fund as the ratio of risk in the last five months of the calendar year to the first seven months of the year. Funds are placed in a group for increased (decreased) risk if this ratio is higher (lower) than the median across all funds in that year.²⁹ We use two different measures of risk: beta and standard deviation.

For each month we calculate the average portfolio beta with respect to the market for a portfolio by weighting the beta on each individual security by the percentage of the portfolio held in that security. We calculate beta using the prior 36 months of returns. If the security has less than 36 months of returns, we use returns that are available if there are at least 12 months of data or assume a beta of one (1.2% of the cases) for those with fewer than 12 months of data.³⁰ We assume cash has a beta of zero. To calculate variance for any particular month we compute portfolio returns over the prior 36 months using security returns combined with security weights for the month for which variance is being computed. Since the weights are known and, in fact, unchanged, we can compute variances from portfolio returns without having to estimate the

²⁹ BHS also examine risk-shifting excluding the month of December (because of the possibility of window dressing). We repeated our tests excluding the month of December. The results are substantially unchanged and we do not report them. In addition, we compared estimates using weekly returns rather than monthly returns to estimate beta and found results similar to those reported, with the same significance.

³⁰ To get the beta and return on securities other than traded equity, the following methods were used. For bonds and preferred stock, returns and betas were computed based on indexes of bonds and preferred stocks. Betas and returns of options and futures were computed using standard formulas and underlying instruments. Non-traded equity was assumed to have the same beta and return as the traded equity in the portfolio under consideration

variance/covariance matrix of security returns.³¹ By constructing the variance of the portfolio, the variance of the individual securities plus the covariance between them is taken into account. We then repeat the BHS tests using the risk definitions above where change in risk is defined in two ways: change in beta and change in variance.

The results of changes in beta from the first seven months of the calendar year to the last five months are shown in Table 8. Note that high-return funds increase beta, while low-return funds decrease beta, and the difference is statistically significant at the 1% level. We also examine change in betas for the top 10% and bottom 10% of return with respect to the first seven months. For the bottom 10% the difference in beta is -0.003 , while for the top 10% of the funds with the highest returns it is $+0.041$. Clearly the action of funds in the tails as well as funds in general support the idea that high-return funds increase risk while low-return funds decrease risk.

The other estimate of the change in risk which we calculated was the ratio of standard deviations where, as described above, standard deviations are estimated from the actual security holdings. As shown in Table 9, high-return funds increase standard deviation while low-return funds decrease it, and the difference is statistically significant at the 1% level. Looking at the tails, we find that the ratio of the standard deviations for the last five months of the year versus the first seven months for the 10% of the funds with the worst return is 0.95, while for the 10% with the highest return it is 1.03.³² Clearly, the risk relationship continues in the tails.

Examination of shifts in the cash-stock mix shows almost no change from the first seven months to the last five months. When we measure risk-shifting by examining changes in beta or

³¹ Some securities have less than 36 months of returns. For these securities we assume inclusion would not change portfolio risk by calculating monthly return using the remaining securities where the weights are recomputed to add to one. This has *de minimis* effect on our results.

³² When the percent of increases and decreases are examined for the tails, a larger percent of high-return funds increase risk and a larger percent of low-return funds decrease risk than those in the population, and the results are significant at the 1% level.

changes in standard deviation (both of which incorporate the effects of stock and cash changes), high-return funds increase risk while low-return funds decrease risk. Using more precise measures of risk, we find evidence inconsistent with the standard tournament model.

VII CONCLUSIONS

In this paper we use monthly holdings data to examine some accepted hypotheses about mutual fund behavior. Existing empirical tests of these hypotheses have analyzed return data on funds or have used holdings data observed at quarterly or semi-annual intervals. The conclusions change when monthly holdings data are used. There are two reasons for this. The first is that the timing of events can be more precisely measured using more frequent data. The second is that using quarterly or semi-annual holdings data misses a significant proportion of trades – the very trades that might well be initiated by the phenomenon under investigation. In fact, we have shown that employing quarterly data misses 18.5% of the trades found using monthly data, while employing semi-annual data misses 34.2% of such trades. Furthermore, the fact that the estimate of turnover based on monthly holdings data is so close to that reported by Morningstar means that few trades are missed with monthly data.

The first hypothesis we examine concerns momentum trading. Past studies using quarterly data have found that funds on average follow a momentum strategy, buying winners and selling losers. But these studies suffer from the very problems described above, and when fund actions are viewed on a monthly basis momentum disappears and as many funds seem to buy past losers (contrarians) as buy past winners (momentum buyers).

The second hypothesis we examine is tax selling. We find that funds increase their turnover in the month preceding the calculation of taxes and that funds that have realized capital

gains during the year increase their realization of tax loss in the two months before taxes are computed. Using losses to offset realized capital gains is in the best interests of both managers and investors, and seems to take place.

The third hypothesis we examine is window dressing: do funds trade in the period immediately preceding quarterly, semi-annual and annual reports to show more attractive holdings? We find that funds trade more in the months preceding the annual report, that they tend to lower the illiquid securities in their portfolio in that month, and that they tend to buy past winners and sell past losers. While we find strong evidence of window dressing with respect to the annual report, we find no such evidence for quarterly or semi-annual reports.

The final subject we examine is tournament models of mutual fund behavior. The existing literature has used return data to examine whether mutual funds that are doing well increase risk in the second part of the reporting year. Theoretical results and empirical results are mixed. Using monthly holdings data we observe what funds are doing over time, yielding more precise estimates of risk changes. We find evidence that high-return funds increase risk and low-return funds decrease risk.

The use of monthly holdings data leads to new conclusions regarding well-accepted beliefs. This paper presents important evidence on the usefulness of monthly data and suggests that correct public policy should encourage the availability of such data on a universal basis.

Table 1
Number of Funds with a Particular ICDI Classification

Aggressive Growth	71
Long-Term Growth	100
Growth and Income	40
Income	4
Total	215

The table shows the Investment Company Institute classification for the funds in our sample. Two funds changed from aggressive growth to long-term growth. They are classified by their predominant classification.

Table 2
Distribution of Fiscal Year Ends
(in percent)

JAN	FEB	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC
0.56%	1.12%	3.36%	1.49%	1.68%	9.51%	3.36%	9.14%	15.86%	25.56%	0.75%	27.61%

The table shows the percentages of funds in our sample with fiscal years ending in the indicated months.

Table 3
Estimated Turnover at Different Intervals
(in percent)

	Monthly Observations	Quarterly Observations	Semi-annual Observations	Annual Observations
Mean	88.12	71.89	57.96	39.35
Median	74.90	65.05	54.90	38.30
10%	23.36	20.19	18.73	14.60
25%	40.51	33.76	29.53	21.28
75%	123.02	101.07	79.97	55.05
90%	164.67	130.00	101.94	66.80

The table shows turnover estimated using holdings data at various intervals. Turnover is the lesser of yearly sales or purchases divided by average monthly net asset value. Turnover is precisely defined in equations 1 and 2.

Table 4
Momentum Measures

Panel A
Quarterly Change in Holdings

	Contemporaneous	Prior
Overall Momentum	0.273 (2.40)	0.043 (0.80)
Num Positive	152	137
Sig Positive	89	51
Num Negative	63	78
Sig Negative	21	23

Panel B
Monthly Change in Holdings

	Contemporaneous	Prior
Overall Momentum	0.181 (1.35)	0.003 (0.04)
Num Positive	149	126
Sig Positive	99	73
Num Negative	66	89
Sig Negative	37	42

Panel C
Monthly Value of Momentum Measure with Different Lags

Lag	Momentum Measure
Contemporaneous	0.185 (2.32)
1	-0.009 (-0.23)
2	0.015 (0.67)
3	0.044 (2.07)

Panels A and B of the table show three-month return measured contemporaneously and lagged on net portfolio changes. The entry called “Overall” is the three-month return on an arbitrage portfolio of buys minus sells. In addition, the table shows the number of funds for which the three-month return on net portfolio changes is positive or negative. Panel A uses quarterly holding data to compute portfolio changes, and Panel B uses monthly holding data. Panel C uses monthly holding data and monthly returns to compute momentum values for the contemporaneous month and each of the three months preceding it. For all panels, the numbers in parentheses are *t* values.

Table 5
Tax Trading: Gains and Losses

	September and October Relative Gains	September and October Relative Losses
Funds with net realized capital gain	0.92	3.95
Funds with net realized capital loss	1.20	1.23

The table shows separately the average monthly gains (and losses) taken in September and October divided by the average monthly gains (and losses) in the prior ten months in two categories according to whether the fund has gains or losses for the year. Gains and losses are computed by comparing the current price with the original purchase price. Price is computed as the average of the beginning and ending price in the appropriate month.

Table 6
Evidence on Window Dressing

Panel A

Number of Stocks Newly Added or Completely Eliminated in the Ending Month of Each Quarter

Quarter				Average of All Months
1	2	3	4	
703	659	669	761	670

Panel B

Return on Purchases minus Sales Relative to The Average Monthly Value

Quarter				Average of All Months
1	2	3	4	
1.402	.615	.901	1.845	1.00

Panel C

Return on Purchases minus Sales for Funds with Above-Average Fourth-Quarter Values

Quarter				Average of All Months
1	2	3	4	
.921	.798	.474	3.694	1.00

Panel A reports the number of stocks added for the first time or completely eliminated for all mutual funds in our sample in the ending month of the quarter.

Panel B shows the return on all purchases minus the return on all sales for the three months prior to the month shown divided by the average of this metric across all months.

Panel C shows the return on all purchases minus the return on all sales for the three months prior to the month shown, aggregated only for the funds that were above the median in month 12 and divided by the average value in all months.

Table 7
Change in Stock/Cash Mix from First Seven Months to Last Five Months

	% in stock first seven months	% in stock last five months
Low Return Funds	93.96	94.88
High Return Funds	93.42	93.86

The percentage a fund has invested in stock was computed each month. Funds were divided into two groups by their relationship to the median return over the first seven months. Then the average investment in stock is shown for each return category for both the first seven months and the last five months.

Table 8
Percentage of High- or Low-Return Funds
That Increase or Decrease Beta

	Increased Beta	Decreased Beta
High Return	28.1	21.9
Low Return	21.7	28.3

The table shows the percentage of funds in each cell. The cells were formed by dividing return relative to the median return in the first seven months and change in beta for August to December compared to January through July. Each year is calculated separately, the results divided by the medium ratio for that year, and the results are aggregated. Betas are calculated each month from fund holdings and then averaged over the period in question.

Table 9
Percentage of High- or Low-Return Funds
That Increase or Decrease Standard Deviation

	Increased Standard Deviation	Decreased Standard Deviation
High Return	30.9	19.1
Low Return	19.0	31.0

This table shows the percentage of funds in each cell where cells are formed on the basis of median returns in the first seven months and changes in standard deviation for the last five months divided by the standard deviations in the first seven months. Every year is calculated separately, the results divided by the median for that year, and the results are aggregated. Standard deviations are estimated each month using the prior three years of returns on a portfolio of the funds' actual security holdings.

BIBLIOGRAPHY

- Barclay, Pearson and Weisbach (1998). Open-end mutual funds and capital gains taxes. *Journal of Financial Economics* 4, 3-43.
- Busse, Jeff (2001). Another look at mutual fund tournaments. *Journal of Financial and Quantitative Analysis* 26, 53-73.
- Brown, Keith; Harlow, W.V.; Starks; Keith Laura (1996). Of tournaments and temptations: an analysis of managerial incentives in the mutual fund industry. *Journal of Finance* 51(1), 85-110.
- Chan, K.C. (1986). Can tax-loss selling explain the January seasonal in stock returns? *Journal of Finance* XLI(5), 1115-1128.
- Chevalier, Judith and Ellison, Glenn (1997). Risk taking by mutual funds as a response to incentives. *Journal of Political Economy* 105(6), 1167-1206.
- Elton, Edwin J.; Gruber, Martin J.; Blake, Christopher (2007). Monthly holdings data and the selection of superior mutual funds. NYU working paper.
- Elton, Edwin J.; Gruber, Martin J; Blake, Christopher (2003). Incentive fees and mutual funds. *Journal of Finance* 58(2), 779-804.
- Ge, Weili and Zheng, Le (2005). The frequency of mutual fund portfolio disclosure. University of Michigan working paper.
- Gibson, Scott; Safieddine, Assem; and Titman, Sheridan (2000). Tax-motivated trading and price pressure: an analysis of mutual fund holdings. *Journal of Financial and Quantitative Analysis* 35(3), 386-398.
- Grinblatt, Mark; Titman, Sheridan; and Wermers, Russ (1995). Momentum investment strategies, portfolio performance, and herding: A study of mutual fund behavior. *American Economic Review* 85(5):1088-1105.
- He, Jia; Ng, Lilian; and Wang, Qinghai (2004). Quarterly trading patterns of financial institutions. *Journal of Business* 77(3) 493-510.
- Huddart, Steven J. and Narayanan, V. G. An Empirical Evaluation of Tax Factors and Mutual Fund Stock Sales Decisions. *Review of Accounting Studies* 7 (2-3), 319-341.
- Jegadeesh, Narasimhan and Titman, Sheridan (1993). Returns to buying winners and selling losers: implications for stock market efficiency. *Journal of Finance*, 48 (1), 65-91.
- Lakonishok, J., Schleifer, A., Thaler, R. and Vishny, R. (1991). Window dressing by pension fund manager. *American Economic Review* 81(2), 227-231.

- Lakonishok, Josef; Shliefier, Andrei; and Vishny, Robert (1992). The impact of institutional trading on stock prices. *Journal of Financial Economics* 32:23-43.
- Lee, Cheng-Fen (1998). Indirect tests of the Haugen-Lakonishok small-firm/January effect hypothesis: window dressing versus performance hedging. *Financial Review* 2, pp.70-81.
- Meier, Iwan and Schaumburg, Ernst (2004). Do funds window dress? Evidence for U.S. equity mutual funds. Unpublished manuscript, Northwestern University.
- Morey, Matthew and O'Neal, Edward (2005), Window dressing in bond mutual funds. Unpublished manuscript, Pace University.
- Musto, David (1997). Portfolio disclosures and year-end price shifts. *Journal of Finance* LII:4, 1563-1588.
- Musto, David (1999). Investment decisions depend on portfolio disclosures. *Journal of Finance* LIV(3), 935-952.
- O'Neal, Edward (2001). Window dressing and equity mutual funds. Unpublished manuscript, Wake Forest University.
- Ritter, Joy and Chopra, Navin (1989). Portfolio rebalancing and the turn-of-the-effect. *Journal of Finance* XLIV (3), 149-165.
- Sias, Richard (2006). Window dressing, tax-loss selling and momentum profit seasonality. Unpublished manuscript, Washington State University.
- Sias, Richard and Starks, Laura (1997). Institutions and Individuals at the Turn-of-the-Year. *Journal of Finance* LII (4), 1543-1562.
- Taylor, Jonathan (2000). A role for the theory of tournaments in studies of mutual fund behavior. Working paper, Sloan School of Management, Operations Research Center, Massachusetts Institute of Technology.
- Wermers, Russ (2001). The potential effects of more frequent portfolio disclosure on mutual fund performance. *Investment Company Institute Perspective* 7(3), 1-11.