Do Redemption Fees Hurt Long-term U.S. Mutual Fund Investors? *

Anthony W. Lynch†
New York University and NBER

Sinan Tan‡
Fordham University

This Version: 18 July 2008

Comments welcome.

*The authors would like to thank Ned Elton, Steve Figlewski, Marti Gruber, Joel Hasbrouck, Hong Liu, Lasse Pedersen, Matt Richardson, and seminar participants at NYU for their helpful comments. All remaining errors are of course the authors' responsibility.

†Stern School of Business, New York University, 44 West Fourth Street, Suite 9-190, New York, NY 10012-1126, alynch@stern.nyu.edu, (212) 998-0350.

‡Fordham University, GBA 113 West 60th Street, Suite 403B, New York, NY 10023, tan@fordham.edu, (212) 636-6118.

Do Redemption Fees Hurt Long-term U.S. Mutual Fund Investors?

Abstract

Redemption fees have been proposed as a way to curb trading on "stale" prices by short-horizon investors to make profits at the expense of long-horizon investors who only trade to rebalance back to their optimal allocations. For redemption fees to be viable, they must have a negligible impact on the utility of the multiperiod agents who only trade to rebalance. In this paper, we examine how the imposition of redemption fees for selling mutual funds purchased within a prespecified period affects the utility of long-horizon agents and we focus on 3 types of funds: domestic funds holding large cap stocks; domestic funds holding value stocks; and, domestic funds holding small cap stocks We find that redemption fees chosen to be large enough to make the "stale" price trading strategies in these funds examined by Chalmers, Edelan and Kadlec (2001) unprofitable, do not materially affect the utility of long-horizon investors. The trade triggers employed in these "stale" price trading strategies imply an average round trip trade length of 2 months but the utility reduction for a 6-month redemption period can be bounded above by 0.53% of initial wealth or per period consumption; and based on exact cost numbers when the agent does not have access to a second fund, the utility cost is likely never more than 0.27% of initial wealth or per period consumption even when a second fund is available. And for a 4-month period, the exact utility cost is even smaller, never more than 0.16% of initial wealth or per period consumption, irrespective of whether a second fund is available or not. The implication is that redemption fees may be a viable device to curb trading on "stale" prices by short-horizon investors at least for the 3 fund types we consider.

1 Introduction

Redemption fees have been proposed as a way to curb trading on "stale" prices by short-horizon investors to make profits at the expense of long-horizon investors who only trade to rebalance back to their optimal allocations (see for example, Chalmers, Edelan and Kadlec, 2001, and Boudoukh, Subrahmanyam, Richardson, and Whitelaw, 2002) and their use has been shown to reduce the volatility of daily fund flows (see Greene, Hodges, Rakowski, 2007). For redemption fees to be viable, they must have a negligible impact on the utility of the long-horizon agents who only trade to rebalance. In this paper, we examine how the imposition of redemption fees for selling mutual funds purchased within a prespecified period affects the utility of long-horizon agents and we focus on 3 types of funds: domestic funds holding large cap stocks; domestic funds holding value stocks; and, domestic funds holding small cap stocks. We find that redemption fees chosen to be large enough to make the "stale" price trading strategies in these funds examined by Chalmers, Edelan and Kadlec unprofitable, do not materially affect the utility of long-horizon investors. The trade triggers employed in these "stale" price trading strategies imply an average round trip trade length of 2 months but the utility reduction for a 6-month redemption period can be bounded above by 0.53% of initial wealth or per-period consumption; and based on exact cost numbers when the agent does not have access to a second fund, the utility cost is likely never more than 0.27% of initial wealth or per-period consumption, even when a second fund is available. And for a 4-month period, the exact utility cost is even smaller, never more than 0.16% of initial wealth or per-period consumption, irrespective of whether a second fund is available or not. The implication is that redemption fees may be a viable device to curb trading on "stale" prices by short-horizon investors at least for the 3 fund types we consider.

Our paper considers a long-horizon individual with power utility who has access to a fund which charges redemption fees, possibly a second fund, and the riskless asset. Redemption fee rates are chosen to make "stale" price trading unprofitable, based on empirical results about the profitability of such trading contained in Chalmers, Edelan and Kadlec (2001). We allow for the possibility that the long-horizon investors are rebalancing to take advantage of return predictability.[1] Fund returns and expense ratios are calibrated to U.S. data. Buying on margin and margin requirements are also incorporated. We report the utility costs of redemption fees. Specifically, we calculate the fraction of wealth (or the fraction of per-period consumption) that an investor would be prepared to give

[1]Campbell (1987) and Fama and French (1989), among others, find that stock return variation can be explained by the one-month Treasury bill rate, by a contemporaneous and a lagged measure of the term premium, and by the dividend yield.

up to be allowed to not pay the redemption fees.

For each of the three funds, we choose a redemption fee for the fund which is larger than the average round trip return reported by Chalmers, Edelan and Kadlec for 'stale" price trading of funds that hold similar stocks to those in the fund using a trigger that implies 6 round-trip trades per year or an average round trip length of 2 months. Averaging across all domestic funds in their sample, this trade trigger delivers an average round trip return of 0.66%, and so the redemption fee we use for the high book-to-market (B-M) fund is 0.75%. For funds that hold large firms, this trade trigger delivers an average round trip return of 0.48%, and so the redemption fee we use for the market fund is 0.5%. Lastly, for funds that hold small firms, this trade trigger delivers an average round trip return of 0.92%, and so the redemption fee we use for the small firm fund is 1.0%. Based on Chalmers, Edelan and Kadlec's results, so long as most of the round trips implied by the trigger are less than k months, making these fees payable for sales within k months of buys would make the "stale" price trading in these funds unprofitable. Since the average round trip length given the trade trigger used by Chalmers, Edelan and Kadlec is 2 months, it seems likely that the probability of one of these round-trip trades taking more than 6 months is extremely small, which means that a 6-month redemption period is likely long enough to deter "stale" price trading. However a redemption period shorter than 6 months but longer than 2 months, say 4 months, may also be sufficient to deter such trading, depending on the distribution of round trip lengths.

To examine the utility cost of redemption fees for rebalancing long horizon agents, we consider the following cases: 1) The investor has access to the riskless asset and one fund that charges a redemption fee; 2) The investor has access to the riskless asset and two funds, one that charges a redemption fee and another that allows costless rebalancing; and, 3) The investor has access to the riskless asset and two funds, one that charges a redemption fee and another that charges a proportional cost to sell the fund equal to the redemption fee rate for that fund. For shorter redemption periods, we are able to calculate exactly the utility cost while for longer periods, computational considerations mean we can only place an upper bound on the utility cost. We do this by calculating the utility cost of the n-month redemption fee when the investor is only allowed to sell every n months but must pay the redemption fee on any fund shares sold that had been bought since the last sell date.

When the agent has access to a single fund and the riskless asset, the single fund is either the market fund, the high B-M fund, or the small firm fund. When returns are predictable, the utility cost of a 6-month redemption fee is never more than 0.27% when risk aversion is 6, and never more than 0.25% when risk aversion is 2, irrespective of which fund is available. The utility costs are much lower when returns are i.i.d. rather than predictable. Whether returns are predictable or i.i.d., the cost is always the highest for the small firm fund, which is not surprising since the redemption fee rate needed to deter "stale" price trading is higher for the small firm fund than the other 2 funds. If a 4 month redemption period is sufficient to deter "stale" price trading, the utility costs of redemption fees are even smaller, never more than 0.15% when risk aversion is 6, and never more than 0.16% when risk aversion is 2, irrespective of which fund is available. These are extremely small cost numbers.

We are also interested in how the availability of a second domestic stock fund affects the utility cost of a redemption fee on a fund. Potentially, the presence of a second fund would reduce the utility cost of the redemption fee on the first asset, since positive return correlation across the 2 funds means that larger sales or smaller purchases of the second fund could substitute for sales of the first fund that would otherwise incur redemption fees. We consider 3 cases: 1) the high B-M fund charges a redemption fee and the other fund is the market fund; 2) the small firm fund charges a redemption fee and the other fund is the market fund; and, 3) the high B-M fund charges a redemption fee and the other fund is the small firm fund. For all three cases, we find that, consistent with the intuition above, the utility cost of the redemption fee on a given fund always remains the same or is lowered once the investor is given access to a second fund. Thus, for a 4-month redemption period, the utility cost of the redemption fee is always the same as or less than the cost when only that fund is available. The implication is that the utility costs of redemption fees when only one fund is available are likely to represent tight upper bounds on the costs when a second fund is made available for the redemption periods longer than 4 months that can't be solved in a manageable time frame when there are 2 funds. As mentioned above, another upper bound for the utility cost of an n-month redemption fee is the utility cost of this fee when the investor can only sell the fund every n months. With predictable returns and the agent having access to a second fund (so it's not feasible to calculate the actual utility cost), this upper bound on the utility cost of a 6-month redemption fee is never more than 0.53% when risk aversion is 6 and never more than 0.33% when risk aversion is 2, irrespective of which two funds are available. These upper bounds are very small for a 6-month redemption period and are even smaller for shorter redemption periods.

To summarize the evidence we present in this paper, it appears that redemption fees chosen to make "stale" price trading of the funds we've examined unprofitable have a negligible impact on the utility of multiperiod agents who only trade to rebalance. This finding suggests that redemption fees may be a viable measure to curb trading on "stale" prices by short-horizon investors to make profits at the expense of long-horizon investors, at least for the three types of domestic fund we've considered.

The paper is organized as follows. Section 2 describes the investor's dynamic optimization problem with i.i.d. or predictable returns, redemption fees and transaction costs, and possibly buying on margin. Section 3 calibrates the state variables and asset returns to the U.S. economy while section 4 discusses the parameter calibration. Section 5 discusses the utility cost results and Section 6 concludes.

2 The investor problem

This section lays out the preferences of and constraints faced by the investor. We characterize the optimization problem for a dynamic investor who does or does not take into account asset-return predictability. We describe the problem when redemption fees are imposed on one asset and proportional selling costs are imposed on the rest. The specification allows margin requirements to be imposed on the agent when buying on margin. We also describe the nature of the utility cost calculations to be performed later in the paper, and the solution technique for numerically solving the investor's problem.

2.1 *Constraints*

In this subsection, we consider an agent's portfolio allocation between N risky assets and a riskless asset. We describe the constraints facing the agent when: 1) margin requirements are imposed for borrowing on margin and the borrowing rate is greater than the riskless rate; 2) one of the risky assets imposes a redemption fee; and 3) the remaining $N-1$ risky assets impose selling costs that are proportional to value of the asset sold for a given choice of portfolio weights. No short-selling of the risky assets is allowed and there is no cost associated with buying any of the risky assets. The redemption fee on the first asset is incurred whenever the agent sells an amount of the risky asset within n months of purchasing at least that amount of the asset.

The presence of such a redemption fee on the first risky asset means that the agent's problem at t depends on the inherited purchases of this asset for the last n periods. Let risky asset 1 be the asset with the redemption fee and let $Y_{j,t}$ be the dollar holding at t of asset 1 acquired j periods ago, $j = 1, ..., n-1$. We let $X_{i,t}$ be the holding of the ith risky asset at t and $X_{0,t}$ be the holding of the riskless asset at t. So letting $X_{b,t}$ be the amount borrowed at t at the margin borrowing rate, $X_{b,t} = \min(X_{0,t}, 0)$, and letting $X_{f,t}$ be the amount invested at t at the riskless rate, $X_{f,t} = \max(X_{0,t}, 0)$. We let $\hat{X}_{i,t}$ be the dollar allocation to the ith risky asset inherited from the previous period. Its evolution equation is given by:

$$\hat{X}_{i,t+1} = X_{i,t} R^i_{t+1}. \tag{1}$$

We need to keep track of this variable for risky asset 1 so we can determine the redemption fee, and for each $i = 2, ..., N$ so we can determine the proportional selling cost. The law of motion for wealth becomes:

$$W_{t+1} = \sum_{i=1}^{N} X_{i,t} R^i_{t+1} + X_{f,t} R^f + X_{b,t} R^b, \tag{2}$$

$$\begin{aligned} X_{0,t} = W_t - C_t - \sum_{i=1}^{N} X_{i,t} - \sum_{i=2}^{N} \phi_s^i \max(\hat{X}_{i,t} \frac{W_t - C_t}{W_t} - X_{i,t}, 0) \\ -\phi_r g(\hat{X}_{1,t} \frac{W_t - C_t}{W_t} - X_{1,t}, \sum_{j=1}^{n} Y_{j,t} \frac{W_t - C_t}{W_t}) \end{aligned} \tag{3}$$

where C_t is consumption at t, R^i_{t+1} is the ith element of $\mathbf{R_{t+1}}$ an Nx1 vector of returns on the N risky assets from t to $t+1$, R^f is the risk-free rate, R^b is the rate charged on margin loans, and $g(x, y) = \min(x, y)$ if $x > 0$, $g(x, y) = 0$ if $x \leq 0$, ϕ_s^i is the proportional cost rate associated with selling the ith risky asset, and ϕ_r is the redemption fee rate on risky asset 1. For simplicity, the riskless rate and the margin loan rate are both assumed to be constant.

Let $\alpha_{i,t}$ be the portfolio fraction invested in the ith asset at t, κ_t be the fraction of wealth consumed at t, $\alpha_{b,t} \equiv \min(\alpha_{0,t}, 0)$, $\alpha_{f,t} \equiv \max(\alpha_{0,t}, 0)$, and $y_{j,t} \equiv \frac{Y_{j,t}}{W_t}$. We let $\hat{\alpha}_{i,t}$ be the allocation to the ith risky asset inherited from the previous period. Its evolution equation is given by:

$$\hat{\alpha}_{i,t+1} = \frac{\alpha_{i,t} R^i_{t+1}}{R^W_{t+1}}. \tag{4}$$

We need to keep track of this variable for risky asset 1 so we can determine the redemption fee, and for each $i = 2, ..., N$ so we can determine the proportional selling cost. This law of motion for W can be rewritten in terms of the chosen portfolio weights $\alpha_{i,t}$ and inherited weights $\hat{\alpha}_{i,t}$ as follows:

$$W_{t+1} = (W_t - C_t) \left(\sum_{i=1}^{N} \alpha_{i,t} R^i_{t+1} + \alpha_{f,t} R^f + \alpha_{b,t} R^b \right) = W_t (1 - \kappa_t) R^W_{t+1}, \tag{5}$$

$$\alpha_{0,t} = 1 - \sum_{i=1}^{N} \alpha_{i,t} - \sum_{i=2}^{N} \phi_s^i \max(\hat{\alpha}_{i,t} - \alpha_{i,t}, 0) - \phi_r g(\hat{\alpha}_{1,t} - \alpha_{1,t}, \sum_{j=1}^{n} y_{j,t}) \tag{6}$$

where R_{t+1}^W is the portfolio return, net of the transaction and redemption costs from t to $t+1$.

The margin requirements imply the following additional constraint on the portfolio weights $\alpha_{i,t}$:

$$1 - \phi_s^i \max(\hat{\alpha}_{i,t} - \alpha_{i,t}, 0) - \phi_r g(\hat{\alpha}_{1,t} - \alpha_{1,t}, \sum_{j=1}^{n} y_{j,t}) \geq \sum_{i=1}^{N} m_b \alpha_{i,t}, \tag{7}$$

where m_b is the margin requirement on margin buying (equity stake as a fraction of the stock holding). The margin requirement m_b each gives the minimum possible equity stake as a fraction of the absolute value of the stock holding. The m_b parameter can take values from 0 to 1, with 0 allowing unlimited buying on margin and 1 ruling it out,

To determine the redemption fee on the first asset, we need to keep track of $\hat{\alpha}_t^1$, but we also need to keep track of the $y_{j,t}$s, and we adopt the FIFO or "first in, first out" method to do so. This method assumes that when any of asset 1 is sold, the oldest shares still purchased within the n-month window are the ones sold. Their evolution equations are given by

$$y_{j+1,t+1} = \frac{\max(y_{j,t} - \max(\hat{\alpha}_{1,t} - \alpha_{1,t} - \sum_{k=j+1}^{n-1} y_{k,t}, 0), 0) R_{t+1}^1}{R_{t+1}^W} \tag{8}$$

if $\alpha_{1,t} < \hat{\alpha}_{1,t}$; and

$$y_{j+1,t+1} = \frac{y_{j,t} R_{t+1}^1}{R_{t+1}^W} \tag{9}$$

if $\alpha_{1,t} \geq \hat{\alpha}_{1,t}$, for $j = 1, ..., n-2$. And

$$y_{1,t+1} = 0 \tag{10}$$

if $\alpha_{1,t} < \hat{\alpha}_{1,t}$; and

$$y_{1,t+1} = \frac{(\alpha_{1,t} - \hat{\alpha}_{1,t}) R_{t+1}^1}{R_{t+1}^W} \tag{11}$$

if $\alpha_{1,t} \geq \hat{\alpha}_{1,t}$.

The law of motion for wealth in eqs. (2) and (5) implicitly assumes that consumption at time t is obtained by liquidating costlessly the ith risky asset and the riskless asset in the proportions $\hat{\alpha}_t^i$ and $(1 - \hat{\alpha}_{\mathbf{t}}' \mathbf{i_N})$, where $\hat{\alpha}_t^i$ is the allocation to the ith risky asset inherited from the previous period. This assumption is not so onerous given the availability of money-market bank accounts and given that equities pay dividends. To the extent that the sum of the risky assets' dividends exceeds the consumption out of the risky asset, $\kappa(\hat{\alpha}'\mathbf{i_N})W$, a dividend reinvestment plan can be used to costlessly reinvest the excess dividend in the risky asset. Dollar transaction costs at t are $W_t(1-\kappa_t)[\sum_{i=2}^{N}\phi_s^i \max(\hat{\alpha}_{i,t}-\alpha_{i,t},0)+\phi_r g(\hat{\alpha}_{1,t}-\alpha_{1,t},\sum_{j=1}^{n} y_{j,t})]$, and are paid by liquidating the riskless asset.

2.2 *Preferences and the optimization problem*

We consider the optimal portfolio problem of a investor with a finite life of T periods and utility over intermediate consumption. Preferences are time separable and exhibit constant relative risk aversion (CRRA):

$$E\left[\sum_{t=1}^{T}\delta^t \frac{c_t^{1-\gamma}}{1-\gamma}|W_1, d_1, \mathbf{S_1}\right], \tag{12}$$

where γ is the relative-risk-aversion coefficient and δ is the time-discount parameter, d_t is a variable that summarizes the future opportunity sets available to the agent by determining the conditional means of future returns, and $\mathbf{S_t}$ is the vector of state variables other than d_t for the agent's problem at time t. Now d_t is a state variable for the agent's problem at t, when returns are predictable. Further, the inherited portfolio weight for the ith risky asset $\hat{\alpha}_{i,t}$, $i=2,...,N$, is a state variable at t (and so an element of $\mathbf{S_t}$) when ϕ_s^i is greater than zero, since the value of this inherited portfolio weight determines the transaction costs to be paid at time t. Since the first asset incurs a redemption fee when sold, the state vector at t, $\mathbf{S_t}$, is also augmented by $\hat{\alpha}_{1,t}$ and the $n-1$ holdings $y_{1,t},...y_{n-1,t}$. These preferences have been extensively used in empirical work by Grossman and Shiller (1981), Hansen and Singleton (1982), and many others.

Facing a given return generating process, the investor can make unconditional (U) or conditional (C) portfolio choices. When making unconditional choices, the investor uses the steady state distribution and ignores any predictability of returns. In other words, the investor assumes returns are i.i.d. when making unconditional choices, which means the investor's unconditional choices would be optimal if returns were in fact i.i.d. with the unconditional distribution. In contrast, the investor exploits return predictability when making conditional choices. Thus, we can evaluate the impact of return predictability on the utility costs of redemption fees.

We now present the Bellman equations associated with these two problems. For the conditional problem (C), the fraction of portfolio value allocated to the risky assets at time t, $\alpha_{\mathbf{t}}$, is denoted by $\alpha(d_t, \mathbf{S_t}, \mathbf{t})$, which is time dependent since the time horizon T is finite. In contrast, if the investor ignores predictability (U), the optimal allocation to the risky asset no longer depends on the d_t state and so $\alpha(\mathbf{S_t}, t)$ can denote α_t.

2.2.1 *Conditional (C) portfolio choices*

Given our parametric assumptions, the Bellman equation faced by the investor is given by

$$\begin{aligned} \frac{a(d_t, \mathbf{S_t}, t) W_t^{1-\gamma}}{1-\gamma} &= \max_{\kappa_t, \alpha_t} \left\{ \frac{\kappa_t^{1-\gamma} W_t^{1-\gamma}}{1-\gamma} \right. \\ & \left. + \frac{(1-\kappa_t)^{1-\gamma} W_t^{1-\gamma}}{1-\gamma} E\left[a(d_{t+1}, \mathbf{S_{t+1}}, t+1) R_{W,t+1}^{1-\gamma} | d_t, \mathbf{S_t} \right] \right\}, \\ \text{for } t = 1, \ldots, T-1, & \end{aligned} \tag{13}$$

where $E[.|d, \mathbf{S}]$ denotes the expectation taken using the conditional distribution given d. This form of the value function derives from the CRRA utility specification in eqs. (12), and from the linearity in W of the budget constraint, eq.(5).

2.2.2 *Unconditional (U) portfolio choices*

The Bellman equation faced by the investor is given by

$$\begin{aligned} \frac{a(\mathbf{S_t}, t) W_t^{1-\gamma}}{1-\gamma} &= \max_{\kappa_t, \alpha_t} \left\{ \frac{\kappa_t^{1-\gamma} W_t^{1-\gamma}}{1-\gamma} \right. \\ & \left. + \frac{(1-\kappa_t)^{1-\gamma} W_t^{1-\gamma}}{1-\gamma} E^U\left[a(\mathbf{S_{t+1}}, t+1) R_{W,t+1}^{1-\gamma} | \mathbf{S_t} \right] \right\}, \\ \text{for } t = 1, \ldots, T-1. & \end{aligned} \tag{14}$$

where $E^U[.|\mathbf{S_t}]$ denotes expectations taken assuming returns are i.i.d. with the same unconditional distribution as in the conditional case. This expectation does not depend on d_t, and neither does α_t. Note that the Bellman eqs. (13) and (14) are solved by backward iteration, starting with $t = T-1$ and either $a(d_T, \mathbf{S_T}, T) = 1$ or $a(\mathbf{S_T}, T) = 1$.

2.3 *Utility cost calculations*

Each of the investor problems described above imply a policy function that, in turn, yields a particular level of expected lifetime utility. Specifically, the policy functions $\{\alpha(d_t, \mathbf{S_t}, t)\}_{t=1}^{T-1}$ and $\{\kappa(d_t, \mathbf{S_t}, t)\}_{t=1}^{T-1}$ can be substituted into the actual law of motion for investor's wealth eq. (5) to obtain the consumption sequence $\{c_t = \kappa(d_t, \mathbf{S_t}, t) W_t\}_{t=1}^{T}$. This consumption sequence is then substituted into eq. (12) to obtain the investor's expected lifetime utility. In our utility cost calculations, we examine how the imposition of redemption fees for selling mutual fund shares purchased within a prespecified period affects the utility of agents with a multiperiod horizon.

As mentioned above, the expected lifetime utility at time 1 depends on the initial value of the predictive variable, d_1, the initial value of the inherited portfolio allocation for risky asset 1, $\hat{\alpha}_{1,1}$, the $n-1$ lagged purchases of asset 1, $y_{1,1}, ..., y_{n-1,1}$, and when the remaining $N-1$ risky assets impose a proportional selling cost, the initial values of the inherited portfolio allocations for the remaining assets, $\hat{\alpha}_{2,1}, ..., \hat{\alpha}_{N,1}$. For any investor problem, whenever $\hat{\alpha}_i$ is a state variable, $i = 1, 2, ..., N$, it is always set equal to the optimal $\alpha_{i,1}$ for the analogous problem with zero transaction costs.

2.4 *Solution technique*

The dynamic programming problems are solved by backward recursion. With two risky assets and a redemption fee on asset 1, the state variables are the inherited allocations, $\hat{\alpha}_{1,t}$, $\hat{\alpha}_{2,t}$, and the $n-1$ lagged purchases of asset 1, $y_{1,t}, ..., y_{n-1,t}$. These state variables are discretized and the value function is linearly interpolated between $(\hat{\alpha}_{1,t}, \hat{\alpha}_{2,t}, y_{1,t}, ..., y_{n-1,t})$ points. This technique yields an approximate solution that converges to the actual solution as the $(\hat{\alpha}_{1,t}, \hat{\alpha}_{2,t}, y_{1,t}, ..., y_{n-1,t})$ grid becomes finer. We use a 1% grid for the inherited allocations and a 4% grid for the lagged purchases.

3 Return calibration

We use the one-month Treasury-bill rate as a proxy for the risk-free rate R^f, and the 12-month dividend yield on the value-weighted NYSE index as a proxy for the predictive variable D. We use three stock portfolios to calibrate the three types of domestic funds that we consider. The monthly rate of return on the value-weighted NYSE index is used to calibrate the monthly return on the market fund while a high B-M portfolio and a small firm portfolio are used to calibrate the value fund and the small-cap fund respectively. These two portfolios are formed from the six value-weighted portfolios SL, SM, SH, BL, BM, and BH from Fama and French (1993) and Davis, Fama and French (2000). The notation S (B) indicates that the firms in the portfolio are smaller (larger) than 50 percent of NYSE stocks. The notation L indicates that the firms in the portfolio have B-M ratios that place them in the bottom three deciles for all stocks; analogously, M indicates the middle four deciles and H indicates the top three deciles. The high B-M portfolio is an equally-weighted portfolio of SH and BH while the small firm portfolio is an equally-weighted portfolio of SL, SM and SH. The three stock return series and the T-bill rate series are deflated using monthly

CPI inflation. The dividend yield series is from CRSP, the CPI series is from CITIBASE and the T-bill return and the three stock portfolios are constructed from data on Ken French's website. The data period used is from 1927:1 to 2003:12. The continuously compounded riskfree rate is estimated to be the mean of the continuously compounded one-month Treasury-bill rate over this period, which gives a value for R^f of 0.056 percent.

Assume that $\mathbf{R}$ is an Nx1 return vector and let $\mathbf{r} \equiv \ln(\mathbf{1}+\mathbf{R})$ and $d \equiv \ln(1+D)$. We assume that $[\mathbf{r}' \; d]'$ follows the vector autoregressive model (VAR):

$$\mathbf{r_{t+1}} = \mathbf{a_r} + \mathbf{b_r} d_t + \mathbf{e_{t+1}}, \tag{15}$$

$$d_{t+1} = a_d + b_d d_t + v_{t+1}, \tag{16}$$

where $\mathbf{a_r}$, Nx1, and a_d are intercepts, $\mathbf{b_r}$, Nx1, and b_d are coefficients and $[\mathbf{e}' \; v]'$ is an i.i.d. vector of mean-zero, disturbances, with *constant* covariance matrix $\mathbf{\Sigma_{ev,ev}}$; the covariance matrix of v is $\mathbf{\Sigma_{v,v}}$ and the variance of e is σ_e^2. Similarly, the unconditional covariance matrix for $[\mathbf{r}' \; d]'$ is $\mathbf{\Sigma_{rd,rd}}$; the unconditional variance matrices for $\mathbf{r}$ and d are $\mathbf{\Sigma_{r,r}}$ and σ_d^2 respectively. Without loss of generality, we normalize the mean of d, μ_d, to be zero and its variance, σ_d^2, to be 1. The disturbance vector $[\mathbf{e_{t+1}}', v_{t+1}]'$ is assumed to be multivariate normally distributed, but with truncation for extreme realizations. Truncation is assumed so that short selling is not ruled out by extreme realizations of $\mathbf{e_{t+1}}$ that have positive probability under the normal distribution, but which are, in fact, implausible. Specification (15)-(16) assumes that d_t is the only state variable needed to forecast $\mathbf{r_{t+1}}$ which is in line with other papers on optimal portfolio selection (e.g., Barberis (2000) and Campbell and Viceira (1999)).

The data VAR is estimated using ordinary least squares (OLS) and discretized using a variation of Tauchen and Hussey's (1991) Gaussian quadrature method; the variation is designed to ensure that d is the only state variable (see Balduzzi and Lynch (1999) for details). However, following Lynch (2001), this study implements the discretization in a manner that produces exact matches for important moments for portfolio choice. In particular, the procedure matches both the conditional mean vector and the covariance matrix for log returns at all grid points of the predictive variables, as well as the unconditional volatilities of the predictive variables and the correlations of log returns with the predictive variables. We choose 19 quadrature points for the dividend yield and 3 points for the stock-return innovations since Balduzzi and Lynch (1999) find that the resulting approximation is able to capture important dimensions of the return predictability in the data. Finally, the data values for $\mathbf{\Sigma_{ev,ev}}$ are taken to be the covariance matrix for the associated untruncated normal distributions when performing the quadrature approximation. Because the truncation typically uses extreme cutoffs, the misstatement of $\mathbf{\Sigma_{ev,ev}}$ by the approximation is likely to be small.

Table 1 presents VAR parameter values for both the data and the various quadrature approximations used. Panel A reports the slope coefficients $\mathbf{b_r}$ and b_d as well as unconditional means for $\mathbf{r}$ and d. Panel B reports the unconditional covariance matrix for $[\mathbf{r}'\ d]$ and and the cross-correlations. Panel C reports the unconditional covariance matrix for $[\mathbf{e}'\ v]$ and the cross-correlations. Throughout the table, the quadrature values almost always replicate the data values, which suggests that the discretization is capturing the important features of the data. Further evidence on this point can be found in Balduzzi and Lynch (1999).

With risk aversion of 2, we find that increasing the number of grid points for stock returns from 3 to 5 has virtually no effect on the optimal portfolio weights chosen by the investor when given access to any 2 of the funds with no transaction costs, no wedge between the borrowing and lending rates, and no margin requirements for borrowing on margin. Lynch (2001) describes how the investor's optimal portfolio is largely unaffected by the severity of a symmetric truncation that is sufficient to ensure that the possibility of negative wealth does not drive the investor's portfolio choice. It appears that here three grid points for returns implies a truncation that is sufficiently severe to ensure this. Moreover, this number of return grid points in conjunction with 19 grid points for dividend yield as the predictive variable implies a less severe truncation than that implied by the data, for all three stock portfolios: 1) For the market portfolio, a one-month return that is less than -54% and a one-month return that is greater than 122% both have positive probability in the quadrature, while the smallest one-month return in the data for this portfolio is -28%, and the largest is only 38%. 2) For the high B-M portfolio, a one-month return that is less than -64%; and a one-month return that is greater than 184% both have positive probability, while the smallest one-month return in the data for this portfolio is -33%, and the largest is only 71%. 3) For the small firm portfolio, a one-month return that is less than -63%; and a one-month return that is greater than 174% both have positive probability, while the smallest one-month return in the data for this portfolio is -30%, and the largest is only 67%.

4 Parameter choices

The investor's risk aversion parameter, γ, is set to either 2 or 6. These γ choices are motivated by the Mehra and Prescott (1985) argument that the existing evidence from macro and micro studies constitutes an *a priori* justification for restricting the value of γ to be less than ten. The horizon

T of the young investor is 240 periods or 20 years, since the return processes are calibrated to monthly returns. A 20-year horizon is a realistic investment horizon for an investor who retires at time 1. The time preference parameter, δ is set equal to the inverse of the riskfree return.

Holding mutual funds is subject to back-end redemption fees and total annual operating expenses. Redemption fees are incurred whenever assets are liquidated within a prespecified time period and are paid as a fraction of liquidation value. Total annual expenses include management fees, distribution fees, service (12b-1) fees,and other expenses and are paid as a fraction of the total portfolio value on a periodic basis. Annual fees are calibrated using mutual fund prospectus information as of May 18, 2008, and subtracted from the quadrature return at all grid points. Redemption fees are calibrated to be sufficiently large to make "stale" price trading unprofitable. We use empirical results in Chambers, Kadlec and Edlin (2001) about the profitability of "stale" price trading of domestic equity mutual funds to perform the calibration. These no-profit cutoff redemption fee rates are also compared to redemption fee rates charged by mutual funds as of May 18, 2008.

We first consider operating expenses. Fidelity Brokerage Services LLC. (Fidelity, henceforth) offers a Spartan Total Market Index fund and the total annual operating expense for this fund is reported to be 0.10%, while the Vanguard Group, Inc. (Vanguard, henceforth) offers an S&P 500 index fund and reports a per annum total annual operating expense of 0.15%. Both funds are designed to replicate S&P 500 returns and therefore we use them to calibrate the operating expense for the market fund. We fix the total annual operating expense ratio for this fund to be 0.18% which is in the ballpark of the ratios quoted for these 2 funds.[2] Turning to the high B-M fund, Fidelity's Midcap Value Fund charges a per-annum total annual operating expense of 0.83% on the total account balance, while Vanguard reports much lower per-annum total annual expense ratios of 0.20% for its Value Index fund and 0.22% for its Small-cap Value Fund. The Vanguard cost figures are taken as representative of costs of trading the high B-M portfolio by holding these funds and so the annual expense ratio for the high B-M fund is set equal to 0.25%. Turning lastly to the small firm fund, Fidelity's Small Cap Enhanced Index Fund charges a per-annum total annual operating expense of 0.67% on the total account balance, while Vanguard again reports a much lower per-annum total annual expense ratios of 0.22% for its Small-cap Index fund and 0.22% for its Small-cap Value Fund. Again, the Vanguard cost figure is taken as representative of costs of trading the small firm portfolio by holding these funds and so the annual expense ratio for the

[2] Using mutual fund prospectus information as of May 7, 2004., Vanguard's S&P 500 index fund charged a total annual operating expense of 0.18% at that time.

small firm fund is set equal to 0.25%.

We next consider redemption fees. Chalmers, Edelan and Kadlec (2001) consider "stale" price trading that involves using a relevant benchmark's futures return in the last two hours of the day to decide whether to exercise a wildcard call option and buy the fund (which will be exercised if the futures return is sufficiently high) or a wildcard put option and sell the fund (which will be exercised if the futures return is sufficiently low). Chalmers, Edelan and Kadlec use a trigger that implies 6 round-trip trades per year or an average trade length of 2 months, and we choose the redemption fee for each of our three funds to be larger than the average round trip return reported by them when they use this trigger to trade funds that hold similar stocks to those in our fund. Based on their results, so long as most of the round trips implied by the trigger are less than k months, making this fee payable for sales within k months of buys for the typical large-firm fund would again make the "stale" price trading unprofitable. Since the average round trip length given the trade trigger used by Chalmers, Edelan and Kadlec is 2 months, a redemption fee period of at least 2 months is needed make sure that almost all the round trips are shorter in length than the redemption fee period. Unfortunately, the paper doesn't discuss the distribution of the round trip length associated with using this trigger for any of its fund samples, and conceivably the distribution could vary depending on the sample. For this reason, we report utility costs for redemption fee periods that range from 1 month up to 12 months for the three funds that we consider.

When Chalmers, Edelan and Kadlec (2001) consider only funds that hold large firms, the trade trigger that implies 6 round-trip trades per year delivers an average hedged next-day return of 0.24%. This implies an average round trip return of 0.48%. We want to use a redemption fee for the market fund that is larger than this average round trip return and so the redemption fee we use for the market fund is 0.50%. Averaging across all domestic equity funds in their sample and using a trade trigger that implies 6 round-trip trades per year, they find an average hedged next-day return of 0.33% and an average round trip return of 0.66%. Since the high B-M portfolio that we calibrate has both large and small firms, we want to use a redemption fee for the high B-M fund that is larger than this average round trip return, and so the redemption fee we use for the market is 0.75%. When they consider only funds that hold small firms, the trade trigger that implies 6 round-trip trades per year delivers an average hedged next-day return of 0.46% and an average round trip return of 0.92%. We want to use a redemption fee for the small firm fund that is larger than this average round trip return and so the redemption fee we use for the small-fund fund is 1.00%.

Comparing these cutoff redemption fee rates to rates used in the industry, Fidelity charges a redemption fee of 0.50% on balances of its Spartan 500 index fund redeemed within 3 month, which is equal to the cutoff no-profit rate for market fund, calculated using results in Chalmers, Edelan and Kadlac. For its Midcap Value Fund, Fidelity charges a redemption fee of 0.75% on balances redeemed within 1 month, which again is equal to the cutoff no-profit rate we calculated for the high B-M fund. Finally, For its Small Cap Enhanced Index Fund, Fidelity charges a redemption fee of 1.50% on balances redeemed within 3 month, which is a little higher than the 1.00% cutoff no-profit rate we calculated for the small firm fund. However, there are reasons other than 'stale" price trading to impose redemption fees on funds containing small firms. These funds want to discourage redemptions because small firms tend to be illiquid and have high trading costs. It appears that the redemption fees being charged by Fidelity are in the ballpark of those needed to make the "stale" price trading described in Chalmers, Edelan and Kadlac unprofitable. On the flipside, Vanguard does not charge redemption fees on any of its four funds that we discussed above.

We now characterize some essential aspects of margin borrowing. Margin loans allow investors to borrow against their securities with the brokerage firm and are subject to initiation and maintenance requirements. The initiation requirement is set by the Federal Reserve Board and currently dictates the borrower's equity share in the account to be at least 50%. The maintenance requirements are set by the NYSE, NASD, and/or the brokerage firm. As of May 18, 2008, Vanguard and Fidelity require that an investor's equity share in an account which is using mutual funds as collateral must always remain at or above 30% as a fraction of the collateral value. But only mutual funds that have been held for 30 days can be used as collateral for margin borrowing. We take the interest on the margin loan to be the Prime Bank Loan Rate, which we obtain from the Federal Reserve Bank of St. Louis web-site. Since the period for this data set (1947-12 to 2003-6) is shorter than our return and riskless rate calibration (1927-1 to 2003-12), to obtain a rate estimate for our longer data period, we add the mean difference between the Prime Bank Loan Rate and the riskless rate (the 1 month T-bill return series from Kenneth French's web-site) to the mean riskless rate for the longer data period. This procedure produces 0.4977 % per month as the nominal margin loan rate. For simplicity, we set the margin requirement to 30%, the same as the maintenance requirement for Fidelity and Vanguard. Since an investor using either of these fund families can only buy mutual funds on margin after holding them for at least a month, our margin is looser than the one such an investor faces. But since looser margin requirements encourage more trading and so typically increase the cost of redemption fees and other trading costs, the utility cost numbers we present in the next section likely represent an upper bound on the utility cost borne by investors who can buy mutual funds on margin.

5 Utility cost of redemption fees for long-horizon rebalancing agents

Tables 2 through 5 report utility costs of redemption fees. The redemption fee is a proportional transaction cost imposed only on sale transactions when the sale transaction is within a prespecified period of the purchase transaction. The utility cost of a redemption fee is defined to be the fraction of wealth (or per-period consumption) that an agent subject to the redemption fee is prepared to give up to be able to trade without the redemption fee in an otherwise identical economic environment. The liquidation policy is FIFO (first in first out), and is described in detail in section 2 above. The tables also report the utility cost of an n-month redemption fee when the investor is only allowed to sell every n months but must pay the redemption fee on any fund shares sold that had been bought since the last sell date. This cost number represents an upper bound on the the actual utility cost of an n-month redemption fee since the set of possible trades available to this agent is a subset of the set available to an agent facing an n-month redemption fee. This problem is also easier to solve, since the redemption fee paid every n periods only depends on the total amount of the fund bought since the last sell date. Consequently, it can be solved in a manageable time frame for large n values for which the full n-month redemption fee problem cannot be solved.

In each table, Panels A and B report results for agents with risk aversions of 2 and 6 respectively. In each panel, the first seven rows report results with no buying on margin while the last 7 reports results with buying on margin allowed and the borrowing rate calibrated to the Prime Bank Loan Rate as described in section 4. The redemption fee numbers are calibrated as described in section 4 to make "stale" price fund trading unprofitable, and expenses are calibrated to those charged by Vanguard and Fidelity as described in section 4. Shorting is never allowed. All cost numbers are reported in percent.

Table 2 reports utility costs for an agent who has access to a single fund and the riskless asset. The single fund is either the market fund, the high B-M fund, or the small firm fund. Utility costs for the i.i.d. return cases are reported in the first 6 columns of Table 2 while those for the predictable return cases are reported in the last 6. The table shows that the utility cost of redemption fees is much higher when returns are predictable rather than i.i.d., irrespective of the fund, risk aversion or whether buying on margin is allowed. But the utility costs of redemption fees are similar irrespective of risk aversion for predictable returns but are typically larger when risk aversion is 6 rather than 2 for i.i.d. returns. And being able to buy on margin has a negligible effect on this cost. When returns are predictable, the utility cost of a 6-month redemption fee is never more than 0.27% when risk aversion is 6, and never more than 0.25% when risk aversion is 2, irrespective of which fund is available. When returns are i.i.d., the utility cost of a 6-month redemption fee is even smaller, never more than 0.12% when risk aversion is 6, and never more than 0.01% when risk aversion is 2, irrespective of which fund is available. Whether returns are predictable or i.i.d., the cost is always the highest for the small firm fund, which is not surprising since the redemption fee rate needed to deter "stale" price trading is higher for the small firm fund than the other 2 funds.

The "stale" price trading strategies examined by Chalmers, Edelan and Kadlec use a trade trigger that implies 6 round-trip trades per year, and an average round trip length of 2 months. It seems likely that the probability of one of these round-trip trades taking more than 6 months is extremely small, which means that a 6-month redemption period is likely long enough to deter "stale" price trading. But the same may also be true for the probability of such round-trip trades taking more than 4 or 5 months, in which case the utility cost for the long-horizon rebalancing investor of a redemption fee period sufficiently long to deter "stale" price trading is even smaller than the numbers quoted above for a 6-month period. As an illustration, the utility cost of a 4-month redemption fee when returns are predictable and a second fund is not available is never more than 0.15% when risk aversion is 6 and never more than 0.16% when risk aversion is 2, irrespective of which fund is available. These are extremely small cost numbers.

We are also interested in how the availability of a second domestic stock fund affects the utility cost of a redemption fee on a fund. Potentially, the presence of the second fund would reduce the utility cost of the redemption fee on the first asset, since, if the two funds have positively correlated returns, larger sales or smaller purchases of the second fund could substitute for sales of the first fund that would otherwise incur redemption fees. To address this question, Tables 2 through 5 report utility costs for an agent who has access to two funds and the riskless asset. One fund charges a redemption fee and the other can either be traded costlessly or only be sold by paying a proportional trading cost set equal to the redemption fee rate for that fund. We would like to charge an m-month redemption fee on the second asset but this problem with redemption fees on both funds cannot be solved in a manageable time frame. Instead we impose a proportional cost to sell the other fund and we set the proportional cost rate equal to the redemption fee rate for that fund. Intuition suggests that making it costly to sell the other fund increases the utility cost associated with the redemption fee on the first asset. And since a proportional cost for selling the other fund is more severe than an n-month redemption fee (of the same rate), the utility cost reported for the case with a proportional cost on sales of the other fund likely represents an upper bound on the utility cost for the case with an m-month redemption fee on the other fund.

Utility costs for the i.i.d. return cases are reported in the first 4 columns of Tables 3 through 5 while those for the predictable return cases are reported in the last 4. In Table 3, the high B-M fund charges the n-month redemption fee and the other fund is the market fund, while in Table 4, the small firm fund charges the n-month redemption fee and the other fund is the market fund. In Table 5, the high B-M fund charges the n-month redemption fee and the other fund is the small firm fund. As is the case when only one fund is available, the tables show that the utility costs of redemption fees are much higher when returns are predictable rather than i.i.d.. Again, the utility costs of redemption fees are similar irrespective of risk aversion with predictable returns but are typically larger when risk aversion is 6 rather than 2 for i.i.d. returns. Once again, being able to buy on margin has a negligible effect on this cost. Consistent with the intuition developed above, the utility cost of the redemption fee on the first fund is always higher (or the same) when selling the second fund incurs a proportional cost as compared to when the second fund can be traded costlessly, but the difference is typically small and never more than 0.01%. When the redemption fee is on the high B-M fund and the other fund is the market fund, the difference is always 0 (at least to 2 decimal places when the utility cost expressed as a percentage).

Comparing Tables 2 and 3, we see that the utility cost of the redemption fee on the high B-M fund is always the same, irrespective of whether the investor has access to the market fund as a second risky asset or not. This finding indicates that the investor never holds more than a trivial amount of the market fund when it is available. This is because the high B-M fund is so attractive relative to the market fund that the investor almost always wants to short the market fund, irrespective of whether her risk aversion is 2 or 6. This intuition also explains why, when the redemption fee is on the high B-M fund and the other fund is the market fund, the utility cost is always the same irrespective of whether the market fund can be traded costlessly or a proportional cost must be paid to sell it. Comparing Table 4 to Table 2, we see that the utility cost of the redemption fee on the small firm fund always remains the same or is lowered once the investor is given access to the market fund as a second risky asset. This suggests that the investor sometimes substitutes larger sales or smaller purchases of the market fund for sales of the small firm fund that would have incurred redemption fees. Similarly, comparing Table 5 to Table 2, we see that the utility cost of the redemption fee on the high B-M fund always remains the same or is lowered once the investor is given access to the small firm fund as a second risky asset. Again, this suggests that the investor sometimes substitutes larger sales or smaller purchases of the small firm fund for sales of the high B-M that would have incurred redemption fees. For a 4-month redemption period, the utility cost of the redemption fee is never more than 0.16%, regardless of which of the 3 pairs of funds are available, which is the same as the maximum cost of 0.16% when only 1 fund is available. This maximum utility cost is extremely small.

Unfortunately, computational constraints mean we do not have utility costs for redemption periods greater than 4 months when the other fund charges a proportional cost to sell the fund, and we do not have utility costs for redemption periods greater than 5 months when the other fund can be traded costlessly. However, the results we discussed above for the shorter redemption fee periods suggest that, for periods of 5 and 6 months, the utility costs when a second fund is available are likely to be the same or smaller than those when a second fund is not available. The implication is that the utility costs for redemption fees that we discussed above for the cases when only one fund is available are likely to represent tight upper bounds on the costs for those cases when a second fund is made available as well. Tables 3 through 5 also report the utility costs of n-month redemption fees when the investor can only sell the fund every n months. For the reasons given above, these costs represent upper bounds on the actual utility costs of the redemption fees. When returns are predictable and the agent has access to a second fund (so it's not feasible to calculate the actual utility cost), this upper bound on the utility cost of a 6-month redemption fee is never more than 0.53% when risk aversion is 6 and never more than 0.33% when risk aversion is 2, irrespective of which fund is available. When returns are i.i.d., this upper bound on the utility cost of a 6-month redemption fee is even smaller, never more than 0.17% when risk aversion is 6 and never more than 0.01% when risk aversion is 2, irrespective of which two funds are available. These upper bounds are very small for a 6-month redemption period and are even smaller for shorter redemption periods.

To sum up, it appears that redemption fees chosen to make "stale" price trading of the funds we've examined unprofitable have a negligible impact on the utility of long-horizon agents who only trade to rebalance. This finding suggests that redemption fees may be a viable measure to curb trading on "stale" prices by short-horizon investors to make profits at the expense of long-horizon investors, at least for the three types of domestic fund we've considered.

6 Conclusions

We consider a long-horizon individual with power utility who has access to a fund which charges redemption fees, possibly a second fund, and the riskless asset. Redemption fee rates are chosen to make "stale" price trading unprofitable, based on empirical results about the profitability of such trading contained in Chalmers, Edelan and Kadlec (2001). We focus on 3 types of funds: domestic funds holding large-cap stocks; domestic funds holding value stocks; and, domestic funds holding small-cap All funds are allowed to have predictable returns. We examine the cost of redemption fees for a long-horizon mutual fund investor for whom rebalancing is the only motive for buying or selling a fund. We numerically solve the individual's multiperiod problem calibrating return predictability to the U.S. stock market and expense ratios to those charged by Vanguard and Fidality.

We find that redemption fees chosen to be large enough to make the "stale" price trading strategies in these funds examined by Chalmers, Edelan and Kadlec unprofitable, do not materially affect the utility of long-horizon investors. The trade triggers employed in these "stale" price trading strategies imply an average round trip trade length of 2 months but the utility reduction for a 6-month redemption period can be bounded above by 0.53% of initial wealth or per-period consumption; and based on exact cost numbers when the agent does not have access to a second fund, the utility cost is likely never more than 0.27% of initial wealth or per-period consumption even when a second fund is available. And for a 4-month period, the exact utility cost is even smaller, never more than 0.16% of initial wealth or per-period consumption, irrespective of whether a second fund is available or not. The implication is that redemption fees may be a viable device to curb trading on "stale" prices by short-horizon investors at least for the 3 fund types we consider.

On the flip side, Chalmers, Edelan and Kadlec do report considerably higher per trade next-day returns for funds that hold high-beta firms and for foreign funds which suggests that the fees we consider may not be high enough to deter "stale" price wildcard option trading in these funds. Consequently, our findings may not generalize to funds that hold high-beta firms or to foreign funds.

References

Balduzzi, P., Lynch, A., 1999. Transaction costs and predictability: some utility cost calculations. Journal of Financial Economics 52, 47-78.

Boudoukh, K., Subrahmanyam, M., Richardson, M. and Whitelaw, R., 2002. Stale Prices and Strategies for Trading Mutual Funds. Financial Analysts Journal 58, 53-71.

Campbell, J., 1987. Stock returns and the term structure. Journal of Financial Economics 18, 373-99.

Chalmers, J.M.R., Edelen, R.M., Kadlec, G.B., 2001. On the Perils of Security Pricing by Financial Intermediaries: The Wildcard Option in Transacting Mutual Fund Shares. Journal of Finance 56, 2209-2236.

Constantinides, G., 1986. Capital market equilibrium with transaction costs. Journal of Political Economy 94, 842-62.

Davis,.J.J., Fama, E.F., French, K.R., 2000. Characteristics, covariances, and average returns: 1927-1997. Journal of Finance 55, 389-406.

Fama, E., French, K., 1989. Business conditions and expected returns on stocks and bonds. Journal of Financial Economics 25, 23-49.

Fama, E., French, K., 1993. Common risk factors in the returns on stocks and bonds. Journal of Financial Economics 33, 3-56.

Greene, J. T, Hodges, C.W., Rakowski, D. A., 2007. Daily mutual fund flows and redemption policies. Journal of Banking and Finance, 31, 3822-3842.

Grossman, S., Shiller, R., 1981. The determinants of the variability of stock market prices. American Economic Review 71, 222-227.

Hansen, L., Singleton, K., 1982. Generalized instrumental variables estimation of nonlinear rational expectations models. Econometrica 50, 1269-1286.

Lynch, A. W., 2001. Portfolio Choice and Equity Characteristics: Characterizing the Hedging Demands Induced by Return Predictability. Journal of Financial Economics 62, 67-130.

Mehra, R., Prescott, E., 1985. The equity premium: a puzzle. Journal of Monetary Economics 15, 145-161.

Tauchen, G., Hussey, R., 1991. Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models. Econometrica 59, 317-96.

Table 1. Sample Statistics, VAR Coefficients and Quadrature Approximation: High Book-to-Market, Small Firm and Market Portfolios. The table reports moments and parameters for a high book-to-market portfolio, a small firm portfolio and a market portfolio of stocks estimated from US data and calculated for the quadrature approximation of Tauchen and Hussey (1991) as described in section 3. The high book-to-market and small firm portfolios are denoted High B-M and Small respectively and are are described in section 3. The Market portfolio is also described in section 3. The data period is from January 1927 to December 2003. The discretization is based on a VAR that uses log dividend yield as the only state variable. Panel A reports unconditional means,VAR slopes and VAR R^2s for the data and the quadrature approximation: b is the vector of VAR slopes and R^2 denotes the regression R^2. Panel B reports the unconditional covariance matrix for the data and for the quadrature approximation. Panels C reports the conditional covariance matrices for the data VAR and the quadrature VAR. All results are for continuously compounded returns. Returns are expressed per month and in percent.

	Data				Quadrature			
Panel A: Unconditional sample moments and VAR coefficients								
Asset/Variable		Uncond. Mean	b	R^2		Uncond. Mean	b	R^2
High B-M		0.83	0.47	0.40		0.83	0.47	0.40
Small		0.75	0.43	0.34		0.75	0.43	0.34
Market		0.54	0.30	0.30		0.54	0.30	0.30
Dividend Yield		0.00	0.98	95.98		0.00	0.96	92.58
Asset/Variable	High B-M	Small Firm	Market	Dividend Yield	High B-M	Small Firm	Market	Dividend Yield
Panel B: Unconditional standard deviations, covariances (above diagonal), and correlations (below)								
High B-M	7.37	52.17	36.91	-0.83	7.37	52.17	36.91	-1.30
Small Firm	0.96	7.41	37.38	-0.81	0.96	7.41	37.38	-1.17
Market	0.91	0.91	5.51	-0.70	0.91	0.91	5.51	-1.06
Dividend Yield	-0.11	-0.11	-0.13	1.00	-0.18	-0.16	-0.19	1.00
Panel C: Conditional standard deviations, covariances (above diagonal), and correlations (below) for the VAR with Dividend Yield as Predictor								
High B-M	7.35	51.97	36.77	-1.29	7.35	51.97	36.77	-1.75
Small	0.96	7.40	37.25	-1.24	0.96	7.40	37.25	-1.59
Market	0.91	0.91	5.51	-1.00	0.91	0.91	5.51	-1.35
Dividend Yield	-0.88	-0.83	-0.90	0.20	-0.87	-0.83	-0.90	0.27

Table 2. Utility Costs of Redemption Fees on a Fund when No Other Fund is Available. The table reports utility costs of redemption fees. This cost is defined to be the fraction of wealth an agent subject to redemption fees is prepared to give up to be able to trade without the redemption fees in an otherwise identical economic environment. The numbers are given in percent. The redemption fee is a proportional cost imposed only on sale transactions when the sale transaction is within a prespecified period of the purchase transaction. Some columns report the actual utility cost of an n-month redemption fee while others report the utility cost of an n-month redemption fee that is coupled with a restriction that only allows selling every n periods. The latter utility cost represents an upper bound on the the actual utility cost of an n-month redemption fee. The agent has access to a single domestic stock fund, which is subject to redemption fees and fund expenses, and to the riskless asset. The fund is a market fund, a high book-to-market(B-M) fund, or a small firm fund. Definitions of the market portfolio, the small firm portfolio and the high B-M portfolio used to calibrate the fund returns are given in section 3. The liquidation policy is FIFO (first in first out), and is described in detail in section 2. Redemption fee rates and expense ratios are described in section 4. Redemption fee rates are calibrated to make the "stale" price trading in Chalmers, Edelan and Kadlec, (2001) unprofitable and the expense ratios are calibrated to U.S. mutual fund expense ratios. Shorting is never allowed. Buying on margin is subject to collateral requirements as described in section 4 of the text and the borrowing rate is calibrated to the Prime Bank Loan Rate as described in section 4. Returns can be i.i.d. or predictable and in each case the dynamics are discretized by the procedure described in section 3. The return processes are calibrated to data from January 1927 to December 2003. Panel A and B report results when risk aversion is 2 and 6 respectively. The cases marked with a "—" could not be solved in a manageable time frame.

	iid						predictable					
	Market		High B-M		Small Firm		Market		High B-M		Small Firm	
periods	redemption	sell only every n	redemption	sell only every n	redemption	sell only every n	redemption	sell only every n	redemption	sell only every n	redemption	sell only every n
Panel A: Risk Aversion = 2												
No buying on margin												
2	0.00	0.00	0.00	0.00	0.00	0.00	0.03	0.05	0.07	0.13	0.08	0.14
3	0.00	0.00	0.00	0.00	0.00	0.00	0.07	0.07	0.09	0.16	0.12	0.17
4	0.00	0.00	0.00	0.00	0.00	0.00	0.10	0.12	0.11	0.18	0.14	0.19
5	0.00	0.00	0.00	0.00	0.00	0.00	0.14	0.19	0.17	0.19	0.18	0.23
6	0.00	0.00	0.01	0.01	0.01	0.01	0.19	0.21	0.20	0.24	0.21	0.29
12	—	0.01	—	0.02	—	0.02	—	0.27	—	0.47	—	0.49
∞	0.00		0.02		0.03		0.31		0.43		0.48	
Buying on margin with borrowing-lending wedge												
2	0.00	0.00	0.00	0.00	0.00	0.00	0.04	0.06	0.09	0.14	0.10	0.16
3	0.00	0.00	0.00	0.00	0.00	0.00	0.09	0.08	0.12	0.18	0.13	0.17
4	0.00	0.00	0.00	0.00	0.00	0.00	0.12	0.14	0.13	0.20	0.16	0.22
5	0.00	0.00	0.00	0.00	0.00	0.00	0.17	0.22	0.19	0.23	0.19	0.24
6	0.00	0.00	0.01	0.01	0.01	0.01	0.24	0.24	0.24	0.28	0.25	0.33
12	—	0.01	—	0.02	—	0.02	—	0.30	—	0.50	—	0.52
∞	0.00		0.02		0.03		0.34		0.48		0.53	

Table 2 cont.

	iid						predictable					
	Market		High B-M		Small Firm		Market		High B-M		Small Firm	
periods	redemption	sell only every n	redemption	sell only every n	redemption	sell only every n	redemption	sell only every n	redemption	sell only every n	redemption	sell only every n
Panel B: Risk Aversion = 6												
No buying on margin												
2	0.00	0.01	0.00	0.02	0.01	0.05	0.04	0.07	0.09	0.15	0.09	0.16
3	0.00	0.01	0.00	0.03	0.02	0.06	0.08	0.12	0.12	0.21	0.13	0.24
4	0.01	0.02	0.02	0.05	0.04	0.09	0.13	0.15	0.14	0.30	0.15	0.31
5	0.01	0.02	0.03	0.08	0.08	0.13	0.18	0.23	0.17	0.36	0.19	0.39
6	0.02	0.04	0.05	0.10	0.12	0.17	0.21	0.30	0.25	0.48	0.27	0.54
12	—	0.08	—	0.17	—	0.19	—	0.48	—	0.91	—	1.05
∞	0.09		0.20		0.28		0.51		0.81		0.89	
Buying on margin with borrowing-lending wedge												
2	0.00	0.01	0.00	0.02	0.01	0.05	0.04	0.07	0.09	0.15	0.09	0.16
3	0.00	0.01	0.00	0.03	0.02	0.06	0.08	0.12	0.12	0.21	0.13	0.24
4	0.01	0.02	0.02	0.05	0.04	0.09	0.13	0.15	0.14	0.30	0.15	0.31
5	0.01	0.02	0.03	0.08	0.08	0.13	0.18	0.23	0.17	0.36	0.19	0.39
6	0.02	0.04	0.05	0.10	0.12	0.17	0.21	0.30	0.25	0.48	0.27	0.54
12	—	0.08	—	0.17	—	0.19	—	0.48	—	0.91	—	1.05
∞	0.09		0.20		0.28		0.51		0.81		0.89	

Table 3. Utility Cost of Redemption Fees on a High B-M Fund when a Market Fund is Available as Well. The table reports utility costs of redemption fees. This cost is defined to be the fraction of wealth an agent subject to redemption fees is prepared to give up to be able to trade without the redemption fees in an otherwise identical economic environment. The numbers are given in percent. The redemption fee is a proportional cost imposed only on sale transactions when the sale transaction is within a prespecified period of the purchase transaction. Some columns report the actual utility cost of an n-month redemption fee while others report the utility cost of an n-month redemption fee that is coupled with a restriction that only allows selling every n periods. The latter utility cost represents an upper bound on the the actual utility cost of an n-month redemption fee. The agent has access to two domestic stock funds, a high B-M fund and a market fund, in conjunction with the risk free asset. The high B-M fund trades with redemption fees, while the market fund trades costlessly or trades with proportional transaction costs on sales, where the proportional cost rate is set equal to the redemption fee rate for the market fund. Definitions of the high B-M portfolio and the market portfolio used to calibrate the funds are given in section 3. The liquidation policy is FIFO (first in first out), and is described in detail in section 2. Redemption fee rates and expense ratios are described in section 4. Redemption fee rates are calibrated to make the "stale" price trading in Chalmers, Edelan and Kadlec, (2001) unprofitable and the expense ratios are calibrated to U.S. mutual fund expense ratios. Shorting is never allowed. Buying on margin is subject to collateral requirements as described in section 4 of the text and the borrowing rate is calibrated to the Prime Bank Loan Rate as described in section 4. Returns can be i.i.d. or predictable and in each case the dynamics are discretized by the procedure described in section 3. The return processes are calibrated to data from January 1927 to December 2003. Panel A and B report results when risk aversion is 2 and 6 respectively. The cases marked with a "—" could not be solved in a manageable time frame.

	iid				predictable			
Market	Costless		Prop. when sale		Costless		Prop. when sale	
periods	redemption	sell only every n	redemption	sell only every n	redemption	sell only every n	redemption	sell only every n
Panel A: Risk Aversion = 2								
No buying on margin								
2	0.00	0.00	0.00	0.00	0.07	0.13	0.07	0.13
3	0.00	0.00	0.00	0.00	0.09	0.16	0.09	0.16
4	0.00	0.00	0.00	0.00	0.11	0.18	0.11	0.18
5	0.00	0.00	—	0.00	0.17	0.19	—	0.19
6	—	0.01	—	0.01	—	0.24	—	0.24
12	—	0.02	—	0.02	—	0.47	—	0.47
∞	0.02		0.02		0.43		0.43	
Buying on margin with borrowing-lending wedge								
2	0.00	0.00	0.00	0.00	0.09	0.14	0.09	0.14
3	0.00	0.00	0.00	0.00	0.12	0.18	0.12	0.18
4	0.00	0.00	0.00	0.00	0.13	0.20	0.13	0.20
5	0.00	0.00	—	0.00	0.19	0.23	—	0.23
6	—	0.01	—	0.01	—	0.28	—	0.28
12	—	0.02	—	0.02	—	0.50	—	0.50
∞	0.02		0.02		0.48		0.48	
Panel B: Risk Aversion = 6								
No buying on margin								
2	0.00	0.02	0.00	0.02	0.09	0.15	0.09	0.15
3	0.00	0.03	0.00	0.03	0.12	0.21	0.12	0.21
4	0.02	0.05	0.02	0.05	0.14	0.30	0.14	0.30
5	0.03	0.08	—	0.08	0.17	0.36	—	0.36
6	—	0.10	—	0.10	—	0.48	—	0.48
12	—	0.17	—	0.17	—	0.91	—	0.91
∞	0.20		0.20		0.81		0.81	
Buying on margin with borrowing-lending wedge								
2	0.00	0.02	0.00	0.02	0.09	0.15	0.09	0.15
3	0.00	0.03	0.00	0.03	0.12	0.21	0.12	0.21
4	0.02	0.05	0.02	0.05	0.14	0.30	0.14	0.30
5	0.03	0.08	—	0.08	0.17	0.36	—	0.36
6	—	0.10	—	0.10	—	0.48	—	0.48
12	—	0.17	—	0.17	—	0.91	—	0.91
∞	0.20		0.20		0.81		0.81	

Table 4. Utility Cost of Redemption Fees on a Small Firm Fund when a Market Fund is Available as Well. The table reports utility costs of redemption fees. This cost is defined to be the fraction of wealth an agent subject to redemption fees is prepared to give up to be able to trade without the redemption fees in an otherwise identical economic environment. The numbers are given in percent. The redemption fee is a proportional cost imposed only on sale transactions when the sale transaction is within a prespecified period of the purchase transaction. Some columns report the actual utility cost of an n-month redemption fee while others report the utility cost of an n-month redemption fee that is coupled with a restriction that only allows selling every n periods. The latter utility cost represents an upper bound on the the actual utility cost of an n-month redemption fee. The agent has access to two domestic stock funds, a small firm fund and a market fund, in conjunction with the risk free asset. The small firm fund trades with redemption fees, while the market fund trades costlessly or trades with proportional transaction costs on sales, where the proportional cost rate is set equal to the redemption fee rate for the market fund. Definitions of the small firm portfolio and the market portfolio used to calibrate the funds are given in section 3. The liquidation policy is FIFO (first in first out), and is described in detail in section 2. Redemption fee rates and expense ratios are described in section 4. Redemption fee rates are calibrated to make the "stale" price trading in Chalmers, Edelan and Kadlec, (2001) unprofitable and the expense ratios are calibrated to U.S. mutual fund expense ratios. Shorting is never allowed. Buying on margin is subject to collateral requirements as described in section 4 of the text and the borrowing rate is calibrated to the Prime Bank Loan Rate as described in section 4. Returns can be i.i.d. or predictable and in each case the dynamics are discretized by the procedure described in section 3. The return processes are calibrated to data from January 1927 to December 2003. Panel A and B report results when risk aversion is 2 and 6 respectively. The cases marked with a "—" could not be solved in a manageable time frame.

	iid				predictable			
Market	Costless		Prop. when sale		Costless		Prop. when sale	
periods	redemption	sell only every n	redemption	sell only every n	redemption	sell only every n	redemption	sell only every n
Panel A: Risk Aversion = 2								
No buying on margin								
2	0.00	0.00	0.00	0.00	0.08	0.14	0.08	0.14
3	0.00	0.00	0.00	0.00	0.12	0.17	0.12	0.17
4	0.00	0.00	0.00	0.00	0.14	0.19	0.14	0.19
5	0.00	0.00	—	0.00	0.18	0.23	—	0.23
6	—	0.01	—	0.01	—	0.29	—	0.29
12	—	0.02	—	0.02	—	0.49	—	0.49
∞	0.03		0.03		0.48		0.48	
Buying on margin with borrowing-lending wedge								
2	0.00	0.00	0.00	0.00	0.10	0.16	0.10	0.16
3	0.00	0.00	0.00	0.00	0.13	0.17	0.13	0.17
4	0.00	0.00	0.00	0.00	0.16	0.22	0.16	0.22
5	0.00	0.00	—	0.00	0.19	0.24	—	0.24
6	—	0.01	—	0.01	—	0.33	—	0.33
12	—	0.02	—	0.02	—	0.52	—	0.52
∞	0.03		0.03		0.53		0.53	
Panel B: Risk Aversion = 6								
No buying on margin								
2	0.01	0.05	0.01	0.05	0.08	0.15	0.09	0.15
3	0.02	0.06	0.02	0.06	0.12	0.23	0.13	0.24
4	0.04	0.09	0.04	0.09	0.14	0.29	0.15	0.31
5	0.08	0.13	—	0.13	0.17	0.36	—	0.37
6	—	0.17	—	0.17	—	0.51	—	0.53
12	—	0.19	—	0.19	—	0.86	—	0.89
∞	0.28		0.28		0.84		0.85	
Buying on margin with borrowing-lending wedge								
2	0.01	0.05	0.01	0.05	0.08	0.15	0.09	0.15
3	0.02	0.06	0.02	0.06	0.12	0.23	0.13	0.24
4	0.04	0.09	0.04	0.09	0.14	0.29	0.15	0.31
5	0.08	0.13	—	0.13	0.17	0.36	—	0.37
6	—	0.17	—	0.17	—	0.51	—	0.53
12	—	0.19	—	0.19	—	0.86	—	0.89
∞	0.28		0.28		0.84		0.85	

Table 5. Utility Cost of Redemption Fees on a High B-M Fund when a Small Firm Fund is Available as Well. The table reports utility costs of redemption fees. This cost is defined to be the fraction of wealth an agent subject to redemption fees is prepared to give up to be able to trade without the redemption fees in an otherwise identical economic environment. The numbers are given in percent. The redemption fee is a proportional cost imposed only on sale transactions when the sale transaction is within a prespecified period of the purchase transaction. Some columns report the actual utility cost of an n-month redemption fee while others report the utility cost of an n-month redemption fee that is coupled with a restriction that only allows selling every n periods. The latter utility cost represents an upper bound on the the actual utility cost of an n-month redemption fee. The agent has access to two domestic stock funds, a high B-M fund and a small firm fund, in conjunction with the risk free asset. The high B-M fund trades with redemption fees, while the small firm fund trades costlessly or trades with proportional transaction costs on sales, where the proportional cost rate is set equal to the redemption fee rate for the small firm fund. Definitions of the high B-M portfolio and the small firm portfolio used to calibrate the funds are given in section 3. The liquidation policy is FIFO (first in first out), and is described in detail in section 2. Redemption fee rates and expense ratios are described in section 4. Redemption fee rates are calibrated to make the "stale" price trading in Chalmers, Edelan and Kadlec, (2001) unprofitable and the expense ratios are calibrated to U.S. mutual fund expense ratios. Shorting is never allowed. Buying on margin is subject to collateral requirements as described in section 4 of the text and the borrowing rate is calibrated to the Prime Bank Loan Rate as described in section 4. Returns can be i.i.d. or predictable and in each case the dynamics are discretized by the procedure described in section 3. The return processes are calibrated to data from January 1927 to December 2003. Panel A and B report results when risk aversion is 2 and 6 respectively. The cases marked with a "—" could not be solved in a manageable time frame.

	iid				predictable			
Small Firm	Costless		Prop. when sale		Costless		Prop. when sale	
periods	redemption	sell only every n	redemption	sell only every n	redemption	sell only every n	redemption	sell only every n
Panel A: Risk Aversion = 2								
No buying on margin								
2	0.00	0.00	0.00	0.00	0.07	0.13	0.07	0.13
3	0.00	0.00	0.00	0.00	0.09	0.15	0.09	0.15
4	0.00	0.00	0.00	0.00	0.10	0.17	0.11	0.18
5	0.00	0.00	—	0.00	0.16	0.18	—	0.19
6	—	0.01	—	0.01	—	0.22	—	0.23
12	—	0.02	—	0.02	—	0.43	—	0.44
∞	0.02		0.02		0.42		0.42	
Buying on margin with borrowing-lending wedge								
2	0.00	0.00	0.00	0.00	0.08	0.14	0.08	0.14
3	0.00	0.00	0.00	0.00	0.11	0.17	0.12	0.18
4	0.00	0.00	0.00	0.00	0.12	0.19	0.13	0.19
5	0.00	0.00	—	0.00	0.18	0.23	—	0.23
6	—	0.01	—	0.01	—	0.27	—	0.28
12	—	0.02	—	0.02	—	0.47	—	0.48
∞	0.02		0.02		0.46		0.47	
Panel B: Risk Aversion = 6								
No buying on margin								
2	0.00	0.02	0.00	0.02	0.08	0.14	0.09	0.15
3	0.00	0.03	0.00	0.03	0.11	0.19	0.12	0.20
4	0.02	0.05	0.02	0.05	0.13	0.28	0.14	0.30
5	0.03	0.08	—	0.08	0.15	0.34	—	0.35
6	—	0.10	—	0.10	—	0.46	—	0.47
12	—	0.17	—	0.17	—	0.89	—	0.90
∞	0.20		0.20		0.78		0.80	
Buying on margin with borrowing-lending wedge								
2	0.00	0.02	0.00	0.02	0.08	0.14	0.09	0.15
3	0.00	0.03	0.00	0.03	0.11	0.19	0.12	0.20
4	0.02	0.05	0.02	0.05	0.13	0.28	0.14	0.30
5	0.03	0.08	—	0.08	0.15	0.34	—	0.35
6	—	0.10	—	0.10	—	0.46	—	0.47
12	—	0.17	—	0.17	—	0.89	—	0.90
∞	0.20		0.20		0.78		0.80	