Career Concerns and the Active Fund Manager's Problem

David A. Chapman*†
Boston College

Richard B. Evans*
University of Virginia

Zhe Xu*
Boston College

February 2009

Abstract

Existing empirical evidence of career concerns in mutual fund management actually confounds concerns about future labor income risk with the fact that younger managers are also learning about the value of their private information. We examine the implications of career concerns, with and without learning, in a calibrated dynamic model of the active fund manager's investment problem. We find that career concerns in the absence of learning cannot account for the observed differences in the allocations of young and old managers. When managers learn about the speed with which the market incorporates their private information, young managers hold larger (smaller) positions than old managers in response to bad (good) news. The measured differences are economically significant to the manager.

*Mailing Address: Chapman and Xu: 140 Commonwealth Avenue, Chestnut Hill, MA 02467, e-mail: david.chapman@bc.edu and xuze@bc.edu. Evans: Darden School of Business, University of Virginia, P.O. Box 6550, Charlottesville, VA 22906-6550 and evansr@darden.virginia.edu. We would like to thank Rui Albuquerque, Pierluigi Balduzzi, Jonathan Berk, Michael Brennan, Jerome Detemple, Evan Gatev, Jakub Jurek, Darren Kisgen, Martin Lettau, Neil Pearson, Josh Pollett, Jeff Pontiff, Jun Qian, Josh White, Moto Yogo, Josef Zechner, and seminar participants at Boston College, Boston University, and the 2007 UBC Summer Research Conference for helpful comments. The Boston College research services department provided support for this research in the form of access to the Linux cluster at Boston College. The current version of this paper is available online at http://www2.bc.edu/~chapmadb/cex.htm

†Corresponding author.

1 Introduction

There is evidence of systematic differences between the observable actions of young versus old mutual fund managers. For example, Chevalier and Ellison (1999) find that termination decisions differ by age, and this leads younger managers to assume less unsystematic risk in their managed portfolios. As a second example, Almazan, Brown, Carlson, and Chapman (2004) show that manager age is significantly related to the use of direct investment restrictions in the mutual fund management process. In a fully rational model, age differences reflect both standard life-cycle consumption/portfolio decisions and the effect of learning about how to trade on valuable private information. We examine the interaction between a manager's lifetime consumption-portfolio choice problem and his decision to trade on valuable private information.

The active manager's problem is a stylized one: he observes the true underlying value of the benchmark portfolio, without cost, at each possible trading date. He knows that the observed per share price and true per share value are related in the long run (i.e., these time series are cointegrated), and that the two series revert toward each other at a constant rate (i.e., the cointegrating residual has a constant rate of mean reversion). This means that the manager has a nontrivial forecast of the (stochastic) time until the market price and his private valuation converge. This is a pure "market timing" problem, although it can also be considered as a simplified version of the selection problem for a manager who is evaluated relative to a benchmark.

The manager is risk averse, and he can hold a personal portfolio that consists of the benchmark and a risk-free asset, with dynamic portfolio weights. The manager's labor income is a fraction of the assets under management, which increases (or decreases) with the performance of the managed portfolio relative to the benchmark return. In addition, the manager may be either "promoted" or "demoted." Promotion (demotion) is modelled as a large, discrete increase (decrease) in assets under management. These events are stochastic, and they are determined by a counting process whose arrival intensity depends on recent past performance relative to the benchmark.

There are two different forms of learning in the model. First, the manager learns about the speed of adjustment of market prices to his private valuation; i.e., his ability. The manager is assumed to start his career with a diffuse prior about this parameter of the model and to update his beliefs as a proper Bayesian. This will naturally induce a difference between younger and older managers based on their knowledge of their own ability. The second form of learning is by investors about managerial ability. Since we do not model the investor's problem explicitly, this type of learning is handled implicitly through the specification of the flow-for-performance function designed to match the empirical flow-for-performance relation. It is also captured, implicitly, through age effects in the exogenous promotion and demotion probabilities.

The solution to the manager's problem consists of a consumption policy, personal portfolio investment policy, and managed portfolio investment policy that maximizes the manager's lifetime utility from consumption, subject to short-sale constraints in both the personal and managed portfolio and nonnegativity constraints for consumption and the value of the manager's personal portfolio of financial assets. These optimal policies are functions of the state variables of the manager's problem, which include the asset wealth to labor income ratio, the value of the manager's private information, and a state that summarizes the recent performance of the portfolio relative to the benchmark. This is a problem with endogenous stochastic labor income that is correlated with the time-varying performance of asset markets. The optimal policies do not exist in closed-form, and they must be computed numerically, given values of the model parameters.

Our first finding is that labor income risk including the probability of promotions and demotions – in the absence of managerial learning – is not large enough to account for the observed differences in portfolio weights or fund policies found in the existing literature. We find that the differences in the optimal managed portfolio weights between a young manager (defined as a manager with an investment horizon of 30 years) and an old manager (defined as a manager with an investment horizon of 2 years) are quite small. In our baseline parameterization, the portfolio weight differences between a young and an old manager is less than one percent. This is true for all levels of wealth, prior performance, and private information that we examine.

When managerial learning is introduced into the problem, we find that there is some scope for observable differences in the choices of young and old managers. In response to extreme negative private information, the young manager takes a larger position in the risky asset than old manager. In response to extreme positive private information, the young manager takes a smaller position than the old manager in the risky asset. These effects are readily explained by the greater uncertainty facing the young manager about the speed of adjustment. These differences are insensitive to managerial wealth, and they decrease in the persistence of private information.

We measure the economic significance of these differences, to the manager, by computing the certainty equivalent wealth cost that the young manager would pay to be allowed to switch from the optimal choice of an old manager to his perceived optimal choice. This cost is increasing in the absolute value of the realized signal, decreasing in wealth, and varies from 4.8 percent of initial wealth to 12.9 percent of initial wealth in the baseline parameterization. We judge these amounts to be economically significant. When information is more persistent than the baseline case, the utility costs decrease to 1.6 percent of initial wealth to 4.3 percent of initial wealth, but when information is less persistent than the baseline case, costs range from 7.2 percent to 19.3 percent of initial wealth.

2 Literature Review

In choosing an active portfolio position, a manager must not only recognize current misvaluation, he must also have a forecast of the time until the market incorporates his private information. This expected period of misvaluation introduces the general notion of career risk, and this risk will have an impact on the investment decisions of a risk-averse manager. The notion of "career concerns", see Holmström (1999) for a general discussion of the labor market impact on managerial incentives, has been considered in the existing literature on fund management in the work of Chevalier and Ellison (1999). They find that the decision to fire a manager for poor performance is more sensitive to risk-adjusted performance for younger managers, and they also find that the probability of termination leads younger managers to deviate less from the allocations in the fund's objective group and to take on less unsystematic risk. Almazan, Brown, Carlson, and Chapman (2004) consider the impact of career concerns, measured as manager age, in resolving the incentive alignment problem inherent in delegated portfolio management.

These findings describe (conditional) correlations between manager age and performance, manager age and asset holdings, and manager age and governance mechanisms. These correlations reflect both concerns about future labor income and learning about market dynamics and the value of the manager's private information. Are the observed correlations between age and performance plausibly consistent with a complete specification of a manager's decision? The active manager's problem that we consider is a finite-horizon consumption/portfolio choice problem with stochastic labor income. Zeldes (1989) and Deaton (1991) are two early influential articles that formulate fully dynamic versions of this problem.[1] These authors focus on optimal consumption choice with exogenous stochastic labor income, with varying forms of liquidity constraints. Consequently, asset markets are extremely simple. In Zeldes (1989), the only asset is a one-period (real) risk-free bond, and in Deaton (1991), the only asset is a real risk-free bond with a constant rate of return. Despite this simplicity – and with significant parameter restrictions and additional restrictions on the functional form of per period utility – these papers find that it is impossible to characterize the optimal policy in closed-form. This is the first indication in the literature that numerical

[1] There is a literature that examines the consumption/portfolio problem with stochastic labor income in a continuous-time setting; see Merton (1971) for an early example and Duffie (2001) for a textbook review of this research.

methods will be required to understand this decision problem.

Koo (1999) extends this earlier work to examine a version of the stochastic income portfolio choice problem that includes a single risky asset (in addition to the constant return bond) and a liquidity constraint of the form of Deaton (1991).[2] The investment opportunity set is constant over time, and labor income can be characterized by purely transitory shocks, purely permanent shocks, or a mixture of transitory and permanent shocks. He is able to demonstrate, numerically, that the optimal consumption and optimal risky asset allocation are both lower in the presence of liquidity constraints. This finding is important because it is the first indication that career concerns reduce the allocation to the risky asset, even in this (comparatively) simple setting.

There is a large literature in financial economics that examines optimal portfolio choice in the presence of time-varying investment opportunities; see, for example, Balduzzi and Lynch (1999) and Campbell and Viceira (1999). The goal of these models is to understand the extent to which optimal asset holdings change with realistic levels of return predictability. They must be solved either approximately or numerically for parameter values that are calibrated to asset market data. However, these papers generally do not simultaneously consider labor income.

Viceira (2001) is an exception. Here, a long-lived investor faces exogenous stochastic labor income with purely permanent shocks. At a random date in the future, the investor faces retirement and the associated prospect of consuming only out of asset wealth. The conditional mean of the risky asset varies over time and asset variances are constant. The focus of the paper is on the optimal investment policy. The primary findings are that the allocation to the risky asset is larger in the employment state relative to the retirement state, savings increases in labor risk with a smaller relative allocation to the risky asset, and if labor income is positively correlated with stock returns, the allocation to the risky asset decreases

[2]Specifically, the liquidity constraint in Koo (1999) requires that asset wealth be non-negative at all dates and states.

further. This final result is important for the following analysis, since the active manager will have endogenous conditional correlation between labor income and asset returns.

Jurek and Yang (2007) examine the problem facing an arbitrageur, trading in continuous-time, who faces what they term as “horizon” and “divergence” risks.[3] They derive, in closed form, the optimal investment policies of a finite-lived arbitrageur who is investing in a relative value trading problem (e.g., Siamese twins or stubs trading). They consider two alternative situations in which the manager maximizes the utility from terminal wealth and one in which he maximizes the utility from intermediate consumption. The spread in the Jurek-Yang problem follows an exogenous Ornstein-Uhlenbeck process, so the timing and date of the convergence in the relative value trade is uncertain. This investment situation differs from the decision problem facing the active manager considered here in two important ways. First, the manager, below, is investing in an absolute valuation setting. Second, and more importantly, however, the active manager examined here has inflows and outflows of assets that affect his income.

Kondor (2006) develops a general equilibrium model of the matching problem between risk neutral arbitrageurs (managers) and risk neutral investors. Arbitrageurs with ability have the opportunity to invest in temporary mispricing of identical assets trading in distinct (local) markets. This elegant analysis shows that, in general equilibrium, the possibility of being fired leads managers to optimally shorten the time horizon of their investment, and there is an increased probability of liquidity crises in markets with long-horizon arbitrage opportunities. In contrast to Kondor (2006), we develop a structure that is more closely related to the problems of an equity mutual fund manager, but this complexity forces us to consider a partial equilibrium model.

The notion of time-horizon to realize mispricing is also common in the practitioner literature. Grinold and Kahn (2000), a popular portfolio management textbook, defines the concept of the “information half-life” as the exponential rate of decay in the value of the

[3]See Liu and Longstaff (2004) for a similar problem, with divergence risk but not horizon risk.

manager's private information. This notion is consistent with the idea of information decay embodied in the dynamics of the cointegrating relationship defined in Section 3. However, the practitioner notion of exponential decay does not explicitly examine the dynamics of the arrival of new information that may push the market price towards or further away from the manager's private valuation. The practitioner literature also does not emphasize the interaction between the manager's private information and his compensation function.

Berk and Green (2004) emphasize the importance of learning about managerial ability from past returns in a model with no separate role for career concerns. They develop a partial equilibrium model that endogenizes the flow-for-performance function as investors in a competitive market for the supply of funds learn about managerial ability. A parameterized version of their simple and elegant model is capable of rationalizing a convex flow-for-performance function that decreases in convexity with a manager's age. It also reconciles investors' focus on prior returns even though there is no return predictability in equilibrium. In a related continuous-time, infinite-horizon model, Dangl, Wu, and Zechner (2008) examine the role of learning in both the investor's capital allocation decision and the fund management company's firing decision.

3 Stating the Manager's Problem

In defining the manager's problem, we make the following assumptions:

A1: Time is discrete and finite, with the time index set $\{0, 1, 2, \ldots, T\}$.

A2: The manager is a pure price taker, and he trades, without cost, in a liquid and anonymous market.

A3: The manager trades in two assets: the benchmark portfolio, with a gross return from $t-1$ to t of $R_{b,t}$, and a single-period risk-free asset with a constant per period return of

R_f. The continuously compounded return on the two assets are denoted $r_{b,t} = \log R_{b,t}$ and $r_f = \log R_f$, respectively.

These assumptions define the general investment environment. **A2** implies that the manager does not trade in a strategic manner.[4] **A3** is made for tractability. Although it defines the manager's problem as a case of "market timing," it can also be thought of a simplified version of the selection process. A manager whose performance is judged relative to a benchmark portfolio is constantly evaluating his active underweighting or overweighting of different stocks and/or sectors. Any deviation from the benchmark contributes to both the portfolio's alpha and its active risk. As a result, the stock or sector selection problem is a (multidimensional) timing problem. If the stock or sector is going to be undervalued for an extended period of time, there is significant additional risk in holding a position that deviates from the benchmark.

A4: At each date, t, the manager can, without cost, observe the true underlying value of the benchmark asset. This true value is denoted $V_{b,t}$. We assume that

$$r_{v,t} \equiv \ln \left(V_{b,t} / V_{b,t-1} \right) \overset{iid}{\sim} \mathcal{N} \left(\mu_v, \sigma_v^2 \right). \tag{1}$$

A5: Let $P_{b,t}$ denote the market price of the benchmark portfolio on exit from time t. $v_{b,t} = \log V_{b,t}$ and $p_{b,t} = \log P_{b,t}$ are cointegrated with a known cointegrating vector $(-1, 1)$. The cointegrating residual is

$$\xi_t = v_{b,t} - p_{b,t}. \tag{2}$$

Furthermore

$$\xi_{t+1} = \rho \xi_t + \zeta_{t+1}, \tag{3}$$

where $\zeta_{t+1} \overset{iid}{\sim} \mathcal{N} \left(0, \sigma_\zeta^2 \right)$.

[4]This is not an innocuous assumption, since it eliminates any interaction of fund size with the manager's ability to outperform the benchmark; see Berk and Green (2004).

A4 defines the mechanism through which the manager might outperform the benchmark. It is particularly simple. It is certainly feasible to extend the problem to allow for noise in the signal about current underlying value, but it is much more challenging to introduce the decision to invest in information about the benchmark according to some cost function. **A5** is critical.[5] It defines the way in which the market for the benchmark is efficient in a dynamic sense. Clearly, it is not efficient with respect to the manager's private information unless $\sigma_\zeta = 0$, which implies that $v_{b,t} = p_{b,t}$ for all t. **A5** is a natural extension of the practitioner notion of an "information half-life", see Grinold and Kahn (2000), to a fully dynamic environment; i.e., the market price is not always right, but the length of time until prices are "virtually" right depends on both the size of the shock that drives value away from price, σ_ζ, and the persistence of that shock, ρ.

A6: The manager does not know ρ with certainty. At date 0, he starts with a flat prior for ρ on the half-open interval $[0, 1)$. He updates each period, as a proper Bayesian, to a new posterior distribution based on observations of ξ_t (and given the other features of the return process).

A6 is also a critical assumption for our analysis. It defines the form of managerial learning. In particular, the manager is assumed to have perfect information about the data-generating process for returns and about the flow-for-performance function (see **A9**, below). Our implicit assumption is that these market-wide processes are stationary and can be observed over a long prior period. The manager learns about how quickly his private information is incorporated into the market price. The initial prior distribution is partially informative; i.e., ρ equal to one is ruled out. This is partly a technical assumption that imposes stationarity on the manager's problem. Economically, it reflects the assumption that the manager's information is *eventually* impounded in the market price. The impounding may occur immediately, ρ equal to zero, in which case the information is of limited economic

[5] Jurek and Yang (2006) use a virtually identical structure for modelling a mean reverting spread pricing relation in studying a continuous-time version of an arbitrageur's problem.

value, or it may take years, ρ equal to 0.999. Along with life-cycle considerations (including the risk of termination, see **A9**, below), learning about the value of his private information is the fundamental way in which a young manager differs from an old manager.

The manager's expected (continuously compounded) abnormal return from t to $t+1$ denoted α_t, is equal to the expected change in the cointegrating residual – positive in the case of an undervalued asset and negative in the case of an overvalued asset. Given (3), this can be written as

$$\alpha_t = (1 - \overline{\rho}_t)\, \xi_t, \tag{4}$$

where $\overline{\rho}_t$ denotes the mean of the posterior distribution of ρ based on information through time t. The manager's overall parameter uncertainty about α_t is captured in the posterior distribution of ρ given the information available at time t, denoted $p\left(\rho \mid \sigma\left(\xi_t\right)\right)$ where $\sigma\left(\xi_t\right)$ is the σ-algebra formed from the history of observations of ξ through time t. The posterior is updated at $t+1$ in a standard Bayesian manner as the product of the prior density, now $p\left(\rho \mid \sigma\left(\xi_t\right)\right)$, and the likelihood function.

Given our earlier assumptions, the manager's dynamics for the continuously compounded returns to the benchmark asset are

$$r_{b,t} = \alpha_{t-1} + \mu_v + \varepsilon_{v,t} + \zeta_t + \varrho_t, \tag{5}$$

where ϱ_t is the uncertainty associated with α_{t-1}; i.e., a draw from the prior. Given the manager's private information, benchmark returns are not *iid*. However, investors do not have access to the manager's private information, and they cannot do better than to assume that alpha is equal to its unconditional mean of zero. Therefore, based on public information, the benchmark return appears to be *iid*.[6]

[6] Our analysis is explicitly partial equilibrium in nature. The nature of a general equilibrium that embeds privately informed fund managers remains an open question. One possible structure is to assume that there are a continuum of managers with total measure 1. A proportion I of the managers have access to information about the true underlying value, while $(1-I)$ are uninformed managers. The proportion of informed managers is large enough to ensure that when they trade on their information, they move price back to value, but they are individually small relative to the market. In each period, some managers retire,

The expected alpha on the managed portfolio depends on both the manager's private information and how that information is converted to an optimal portfolio holding, conditional on both the manager's private information *and* the manager's compensation function. The manager cares about how he uses his private information because it affects his compensation.

A7: The true portfolio alpha is unobservable to investors, and the manager cannot credibly convey his private information to the market.

The next two assumptions link the portfolio's realized returns to the manager's decisions. The returns to the managed portfolio are predictable because the manager's private information is persistent. Investors care about *past* returns only insofar as they help to forecast future returns. The timing conventions here are important. There is an inherent two-date lag between a manager's portfolio decisions and the consequences of those decisions for realized returns (and for future labor income). A portfolio decision made at time t affects the managed returns that investors observe at $t+1$. Their fee income consequences are realized at $t+2$, through their impact on both expected flows and (time-varying) promotion and demotion probabilities.

A8: The manager's income is a constant proportion of assets under management,

$$L_t = \phi F_t, \tag{6}$$

where $0 < \phi << 1$ and F_t denotes assets under management at time t.

A9: The (gross) growth rate of F (and by extension L) is conditionally lognormal. It has three components: internal growth through the return to the portfolio, (exogenous) *external* flow of assets into the fund, and the possibility of a discrete change in assets

and they are replaced by a new cohort of managers with the same proportion of informed and uninformed managers. Investors cannot observe which managers are informed and which are uninformed, but investors know the proportions of informed versus uninformed in the overall population. Whether or not this general structure can be formalized into an overall equilibrium is a subject for future research.

under management. The discrete change – a "promotion" or "demotion" event – occurs after investment and consumption decisions are made but before asset market uncertainty is resolved. The process for assets under management is

$$F_{t+1} = \theta_{t+1} F_t, \tag{7}$$

where

$$\theta_{t+1} \equiv \begin{cases} \exp\left(r_{m,t+1} + \delta_{t+1}\right) & \text{w/prob} \quad \left(1 - q^p_{t+1|t} - q^d_{t+1|t}\right) \\ \nu^d & \text{w/prob} \quad q^d_{t+1|t} \\ \upsilon^p & \text{w/prob} \quad q^p_{t+1|t}. \end{cases} \tag{8}$$

$r_{m,t+1}$ is the continuously compounded return to the managed portfolio, and δ_{t+1} is the external fund flow.

δ_{t+1} is defined as

$$\delta_{t+1} = \mathbb{S}_3\left(z_t\right) + \epsilon_{t+1}, \tag{9}$$

where $\epsilon_{t+1} \overset{iid}{\sim} \mathcal{N}\left(0, \sigma_\epsilon\right)$, and $\mathbb{S}_3\left(z_t\right)$ is a cubic spline with possibly K interior knot points. So, the conditional mean is modeled as a flexible nonlinear function of lagged relative performance, consistent with the empirical evidence in Chevalier and Ellison (1997) and Sirri and Tufano (1998). ϵ_{t+1} represents the cumulative effect of all determinants of fund flow other than prior fund performance. The measure of prior relative performance is

$$z_{t+1} = \psi z_t + \kappa\left(R_{m,t+1} - R_{b,t+1}\right), \tag{10}$$

where $z_0 = 0$ and and $0 \leq \kappa \ll 0$. This is an important state variable for promotion/demotion and fund flow dynamics.

The functional form for the promotion/demotion probabilities conform to a multinomial logit specification:

$$q^d_{t+1|t} = \frac{\exp\left(\boldsymbol{\varphi}^{d\prime}\mathbf{x}_t\right)}{\left[1 + \exp\left(\boldsymbol{\varphi}^{d\prime}\mathbf{x}_t\right) + \exp\left(\boldsymbol{\varphi}^{p\prime}\mathbf{x}_t\right)\right]}, \tag{11}$$

$$q^p_{t+1|t} = \frac{\exp\left(\boldsymbol{\varphi}^{p\prime}\mathbf{x}_t\right)}{\left[1 + \exp\left(\boldsymbol{\varphi}^{d\prime}\mathbf{x}_t\right) + \exp\left(\boldsymbol{\varphi}^{p\prime}\mathbf{x}_t\right)\right]}, \tag{12}$$

where $\mathbf{x}_t$ is the set of variables that determines a transition, and the probability of neither promotion nor demotion is

$$\left(1 - q^p_{t+1|t} - q^d_{t+1|t}\right) = \frac{1}{\left[1 + \exp\left(\boldsymbol{\varphi}^{d\prime}\mathbf{x}_t\right) + \exp\left(\boldsymbol{\varphi}^{p\prime}\mathbf{x}_t\right)\right]}. \tag{13}$$

$\mathbf{x}_t$ consists of managerial age and prior performance, z_t, consistent with the demotion evidence in Chevalier and Ellison (1999) and discussed further in Section 5, below. A promotion or demotion event at t is independent of a promotion or demotion event at date $\tau \neq t$. If a manager is promoted, we assume that

$$F_{t+1} = \nu^p F_t, \tag{14}$$

where $\upsilon^p >> 1$. If the manager is demoted at t, we assume that he is re-hired at $t+1$ to manage a smaller fund:

$$F_{t+1} = \nu^d F_t, \tag{15}$$

where $0 < \nu^d << 1$.[7] The gross return to the managed portfolio is

$$R_{m,t+1} = \omega_t \left(R_{b,t+1} - R_f\right) + R_f, \tag{16}$$

where ω_t defines the weight on the benchmark asset on exit from time t.

We are not claiming that **A8** and **A9** define an optimal managerial compensation contract. Rather, this form of the manager's problem is motivated by the empirical literature on fund manager compensation and performance. For example, Starks (1987) describes the basic symmetric compensation contract of the form of (6). Chevalier and Ellison (1997) and Sirri and Tufano (1998) provide evidence of a positive (and nonlinear) relationship between fund performance and new asset flows into the fund. Chevalier and Ellison (1999)

[7] Demotion can be interpreted as being fired and rehired at a smaller fund complex. It can also be interpreted as being fired permanently by setting $\nu^d = 0$. In addition, it is certainly possible to make both the rehiring date and the value loss stochastic, but we have not done so here, in the interest of simplicity.

and Evans (2006) document that manager promotions and demotions are related to risk-adjusted performance, and Brown, Harlow, and Starks (1996) provides additional evidence that managerial decisions reflect relative performance.

Since we are abstracting from the role of the fund family in monitoring, rewarding, punishing, and developing mutual fund products and managers, there is no explicit principal-agent problem to solve.[8] Rather, there is a relationship between the manager and the market as a whole. There is no explicit bargaining over the terms of employment and compensation. There is, instead, an aggregate response to the observable characteristics of the fund.

The manager's objective function is standard:

A10: The manager is risk averse with time-separable power utility over the real consumption of a single good:

$$U\left(C_t\right) = \frac{C_t^{1-\gamma}}{1-\gamma}, \tag{17}$$

where γ is the coefficient of relative risk aversion.[9]

The manager's budget constraint reflects variation over time in both labor income – derived from the managed portfolio – and the returns to the manager's personal portfolio. Let A_t denote the manager's personal financial asset wealth on exit from t. If $W_t \equiv A_t + L_t$ is the "cash-on-hand" (to follow the terminology in Koo, 1999), then the (real) budget constraint is

$$W_{t+1} = \left(W_t - C_t\right) R_{p,t+1} + \theta_{t+1} L_t, \tag{18}$$

and

$$R_{p,t+1} = \eta_t \left(R_{b,t+1} - R_f\right) + R_f. \tag{19}$$

(18) assumes that the manager can consume out of fee income in the period that it is earned.

[8]See Bhattacharya and Pfleiderer (1985) or Stoughton (1993) for examples of the optimal contracting approach to describing the manager's problem.

[9]In the interest of simplicity, we assume that the inflation rate is constant. So, we avoid introducing the additional complication of introducing price level notation and interpret all quantities as real.

A11: The manager faces the following additional constraints:

a. Short positions are not allowed in the manager's personal portfolio, $\eta_t \in [0,1]$ for all t. The benchmark asset cannot be sold short in the actively managed portfolio, but the manager is allowed to borrow at the risk-free rate to support investment in the risky asset in order to be able to exploit information that the benchmark is undervalued, $\omega_t \in [0, \overline{\omega}]$, where $\overline{\omega} \geq 1$.[10]

b. Consumption is restricted to be nonnegative, $C_t \geq 0$ for all t, and the manager is liquidity constrained, in the standard definition of that term, $A_t \geq -L_t$ for all t.

4 Solving the Manager's Problem

Under assumptions **A1** to **A11**, the *active manager's problem*, at time 0, is

$$\max_{\{C_t,\omega_t,\eta_t\}_{t=1}^{T}} E^{\rho}\left[\sum_{t=1}^{T}\beta^t U\left(C_t\right) \mid \mathbf{S}_0\right], \tag{20}$$

where $\beta < 1$ is a constant discount parameter, $\mathbf{S}_0 = (W_0, L_0, \xi_0, z_0)'$ denotes the state at time 0, $W_0, L_0 \geq 0$ and $E_t^{\rho}[\cdot]$ denotes the fact that the expectation is over both the objective uncertainty and the subjective posterior distribution for the manager at date t. The maximization is also subject to the budget constraint, income dynamics, return dynamics, and the constraints in **A11**.[11] We use ξ, the shock to α, rather than α as a a state variable due to the different effects of learning on managers of different ages.[12]

The value function at the terminal date is simply the terminal value of the manager's objective function

$$V_T\left(\mathbf{S}_T\right) = \frac{C_T^{1-\gamma}}{1-\gamma} = \frac{W_T^{1-\gamma}}{1-\gamma}, \tag{21}$$

[10] Almazan, Brown, Carlson, and Chapman (2004) contain a detailed discussion of the evidence on the incidence of short-sale (and other) mutual fund constraints in U.S. domestic equity funds.

[11] This form of the manager's objective function is consistent with the finite-horizon stochastic labor income literature; see, for example Zeldes (1989) and Koo (1999), although this literature does not include learning.

[12] In the case without learning, ξ and α are effectively the same because the α is just proportional to ξ.

where $\mathbf{S}_T = (W_T, L_T, \xi_T, z_T)'$. This result follows because utility is strictly increasing. Following Koo (1999), we exploit the fact that the value function is homogeneous of degree $(1-\gamma)$ in financial wealth and labor income.

Let $w_{T-1} \equiv W_{T-1}/L_{T-1}$. The normalized value function is

$$\mathcal{V}_{T-1}(w_{T-1}, \xi_{T-1}, z_{T-1}) \equiv V_{T-1}(w_{T-1}, 1, \xi_{T-1}, z_{T-1}) L_{T-1}^{1-\gamma}. \quad (22)$$

or

$$V_{T-1}(\mathbf{S}_{T-1}) = \mathcal{V}_{T-1}(\mathbf{s}_{T-1}) L_{T-1}^{1-\gamma}, \quad (23)$$

where $\mathbf{s}_{T-1} \equiv (w_{T-1}, \xi_{T-1}, z_{T-1})'$.

Since the backward recursion is standard, the Bellman equation in the scaled value function at any time t from 0 to $T-1$ is

$$\mathcal{V}_t(\mathbf{s}_t) = \max_{c_t, \omega_t, \eta_t} \left\{ \frac{c_t^{1-\gamma}}{1-\gamma} + \beta E_t^{\rho} \left[\mathcal{V}_{t+1}(\mathbf{s}_{t+1}) \theta_{t+1}^{1-\gamma} \right] \right\}, \quad (24)$$

where $c_t \equiv C_t/L_t$. The dynamics for scaled wealth, w_t, are

$$w_{t+1} = (w_t - c_t) \frac{R_{p,t+1}}{\theta_{t+1}} + 1. \quad (25)$$

The dynamics for the prior performance state variable, z_t, are defined in (10), and the dynamics for α_t are in (4). The optimization is also subject to the definition of asset returns, the growth of assets under management, equation (8), the demotion/promotion probabilities in equations (11) and (12), the fund flows in equation (9), and the short-sale and non-negativity restrictions defined in Section 3.

The first-order conditions for this problem are

$$\begin{aligned} c_t: \quad c_t^{-\gamma} + \beta E_t^{\rho} \left[\theta_{t+1}^{1-\gamma} \tfrac{\partial}{\partial c_t} \mathcal{V}_{t+1}(\mathbf{s}_{t+1}) + \mathcal{V}_{t+1}(\mathbf{s}_{t+1}) \tfrac{\partial}{\partial c_t} \theta_{t+1}^{1-\gamma} \right] &= 0, \\ \eta_t: \quad E_t^{\rho} \left[\theta_{t+1}^{1-\gamma} \tfrac{\partial}{\partial \eta_t} \mathcal{V}_{t+1}(\mathbf{s}_{t+1}) + \mathcal{V}_{t+1}(\mathbf{s}_{t+1}) \tfrac{\partial}{\partial \eta_t} \theta_{t+1}^{1-\gamma} \right] &= 0, \end{aligned} \quad (26)$$

and the first-order condition with respect to ω_t is

$$E_t^{\rho}\left[\sum_{j=0}^{J}\left\{\theta_{t+1+j}^{1-\gamma}\frac{\partial}{\partial\omega_t}\mathcal{V}_{t+1+j}\left(\mathbf{s}_{t+1+j}\right)+\mathcal{V}_{t+1+j}\left(\mathbf{s}_{t+1+j}\right)\frac{\partial}{\partial\omega_t}\theta_{t+1+j}^{1-\gamma}\right\}\right]=0 \tag{27}$$

The first-order condition with respect to the managed portfolio choice reflects the fact that allocations today have an effect on future maximized utility through the relative performance state variable. These conditions simplify because the growth of assets under management is unrelated to managerial consumption choice and the manager's holdings in his personal portfolio:

$$\begin{aligned} c_t: \quad c_t^{-\gamma}+\beta E_t^{\rho}\left[\theta_{t+1}^{1-\gamma}\tfrac{\partial}{\partial x_{t+1}}\mathcal{V}_{t+1}\left(\mathbf{s}_{t+1}\right)\tfrac{\partial}{\partial c_t}x_{t+1}\right] &= 0, \\ \eta_t: \quad E_t^{\rho}\left[\theta_{t+1}^{1-\gamma}\tfrac{\partial}{\partial x_{t+1}}\mathcal{V}_{t+1}\left(\mathbf{s}_{t+1}\right)\tfrac{\partial}{\partial \eta_t}x_{t+1}\right] &= 0, \end{aligned} \tag{28}$$

and the first-order condition with respect to ω_t is

$$\begin{aligned} E_t^{\rho}\Bigg[\sum_{j=0}^{J}\Bigg\{&\theta_{t+1+j}^{1-\gamma}\frac{\partial}{\partial z_{t+1+j}}\mathcal{V}_{t+1+j}\left(\mathbf{s}_{t+1+j}\right)\frac{\partial z_{t+1+j}}{\partial\omega_t} \\ &+\mathcal{V}_{t+1+j}\left(\mathbf{s}_{t+1+j}\right)\frac{\partial}{\partial\omega_t}\theta_{t+1+j}^{1-\gamma}\Bigg\}\Bigg]=0. \end{aligned} \tag{29}$$

Solving the active manager's problem is computationally demanding because of learning, the number of state variables, and the inclusion of an endogenous state variable. The first step in our solution is to apply a Gibbs sampling algorithm to construct the posterior distribution of α at each date. The second step uses a numerical procedure that is derived from the simulation based algorithm developed in den Haan and Marcet (1990) and Brandt, Goyal, Santa-Clara, and Stroud (2005). We follow the refinements of the basic simulation-based algorithm as developed in Barillas and Fernándes-Villaverde (2007) and Koijen, Nijman, and Werker (2008), where simulations are also done over the posterior distribution for α at each date. The method is described in detail in the online Appendix.[13] Before we simulate the optimal solution, however, we must choose parameter values for the manager's objectives, private information, flow-for-performance, and benchmark and risk-free rate dynamics.

[13] The appendix and the Matlab code used to solve the model are both available on the web at http://www2.bc.edu/˜chapmadb/cex.htm.

5 Model Calibration

In calibrating the model, we start by specifying a period as corresponding to one calendar quarter. This is not an innocuous choice, but it is a compromise between the manager's portfolio decisions, which are likely made at a much higher frequency, and investors' rebalancing decisions, which likely have a distribution that varies from daily to monthly to annual horizons (or beyond). The remaining parameters can (roughly) be divided into two groups: parameters that can be calibrated to observed returns and fund flows and those that are inherently unobservable to an econometrician.

We define the Russell 2000 Growth Index as the benchmark. The continuously compounded returns are measured quarterly from the first quarter of 1994 to the third quarter of 2006, and they are expressed in percent at a quarterly rate.[14] The average yield-to-maturity on the three-month Treasury bill available at the beginning of each quarter is the constant risk-free return. The unconditional moments of returns are reported in Table 1. This benchmark has an average return in excess of the average risk free rate of 0.48 percent per quarter, and it has excess return volatility of 12.988 percent per quarter.

In order to calibrate fund flows, we construct a database of manager promotions and demotions by combining CRSP and Morningstar data on managers and supplementing, where necessary, from the mutual fund prospectuses. The manager data consists of domestic equity managers from 1995 to 2002. Managers of fixed-income, precious metal, international and utility funds were excluded from the sample. Both CRSP and Morningstar list the start date and the name of each manager. The end date for a manager is assumed to be one month before the starting date of the next manager listed for the fund. In some cases there were discrepancies between the databases as to manager name or starting date. When CRSP and Morningstar disagree, we check the original prospectus to ensure accuracy. Also, if multiple managers were identified, a separate record was created for each manager.

[14] The quarterly total return values were downloaded from http://www.russell.com.

In order to account for the impact of age on promotion and demotion probabilities, we categorize managers as young (manager tenure of less than or equal to 3 years), middle-aged (manager tenure greater than 3 years but less than or equal to 7 years), or old (manager tenure greater than 7 years). Manager tenure is defined as the difference between the current date and the first date that the manager appeared in the database (measured in years). We classify managers as growth managers based on the coefficient on the book-to-market portfolio (HML) in a regression of the fund's return on the 4-factors of Fama and French (1993) and Carhart (1998). For each date, the median HML coefficient is calculated and funds with an HML coefficient below the median are classified as Growth.

Chevalier and Ellison (1997) and Sirri and Tufano (1998) establish that there is a nonlinear function relating future fund flows to past performance.[15] In fitting the flow-performance function, we first considered a generalized cubic spline approximant, but the data suggested that a simple third-order polynomial defined over the entire range of the performance variable provided an adequate fit. Figure 1 shows this function. z is measured in units of percent per month ($z = 0.1$ means 0.1% per month).[16] Fund flows are measured quarterly, and prior performance is defined as z in (10) with ψ and κ set to 0.59 and 0.15, respectively. The flow-performance function is steeper and convex at lower levels of relative performance. The coefficients that define the polynomials are shown in Table 2. The R-squared from the cubic regression on flow-for-performance is 7.41%, and the standard deviation of fund flows is 13.36%. This suggests a value for σ_ε of approximately 0.13. The upper bound on the managed portfolio weight, $\overline{w}$, is set to 1.33, consistent with Section 18 of the Investment Company Act of 1940 and by the reported use of leverage by funds in their semi-annual N-SAR filings.

[15]These papers use data on growth, aggressive growth, and growth and income funds, over the period from 1983 to 1993 for Chevalier and Ellison (1997) and 1971 to 1990 for Sirri and Tufano (1998), in estimating the flow-for-performance function.

[16]The total range of z from -1% to 1% corresponds to a range of excess return of -2% to 2% in monthly terms. The first percentile of the distribution of monthly Jensen's alphas for the entire CRSP equity mutual fund universe from 1963 to 2006 is -2.63% per month, and the 99th percentile of the alpha distribution is 2.68% per month. So a range of z from -1% to 1% seems reasonable.

The calibration of manager transitions – demotions and promotions – requires data on both the probability of a transition in any quarter and the average change in assets under management associated with a transition event. The database records managerial changes but whether specific managers are promoted or demoted is not stated directly. Therefore, we use the same decision rule as Chevalier and Ellison (1999), Hu, Hall, and Harvey (2000), and Baks (2006) for identifying promotions and demotions. A managerial change is a promotion if the total assets under management (adjusting for fund growth) for the 12 months after the change is greater than the assets under management for the 12 months before the change. The rule would classify a manager who moves from managing a small fund to a larger fund, or who adds an additional fund to his management responsibilities as being promoted. Identifying a demotion is done symmetrically. In the case of a management team of two or more managers, the assets are equally divided among the members of the team in defining a transition event. Managers who leave the sample completely after poor performance are treated as demotions, and managers who leave the sample after superior performance are treated as promotions.[17]

Consistent with the functional forms for promotion and demotion probabilities in (11) through (13), we fit a multinomial logit model to our managerial data where transition probabilities depend on prior performance (z) and managerial age, defined in the broad categories stated earlier. The coefficients from this logit regression, along with coefficient standard errors, are shown in Table 3. They seem reasonable in that promotion (demotion) probabilities are positively (negatively) and statistically significantly related to recent relative performance. Age effects appear to be significant, but the interaction terms between age and performance are generally not significant. The plots of age-related transition probabilities are presented in Figures 2 and 3. Demotion probabilities decrease with increasing performance. The overall (quarterly) probability of demotion ranges from zero (for high relative performance funds) to in excess of 4 percent for the worst performers. There

[17]This rule likely understates the wealth effects of firing, since a manager in our analysis does get another job with a smaller fund. It likely understates the wealth effect of promotions if managers who leave mutual funds go on to become *successful* hedge fund managers.

is virtually no difference in demotion probabilities for young and middle-aged managers, but older managers appear to be less likely to be demoted for the same level of prior performance. Promotion probabilities are shown in Figure 3. The quarterly probability of demotion is always less than 2.5 percent per quarter, substantially lower than the demotion probability. They increase in prior performance, with younger managers less likely to be promoted, for a given level of prior performance than either older or middle-aged managers. In promotion events, middle-aged managers behave much more like old managers instead of young managers. The average change in assets under management is not statistically significantly related to either age or prior performance. Therefore, we set υ^d equal to its sample average of 0.423 and υ^p equal to its sample average of 1.72.

The parameters that describe the manager's private information and his utility function are inherently unobservable. They are shown in Table 4. In the following simulations, ρ can assume a value in the set $\{0.1, 0.3, 0.5\}$, and given a choice of ρ, σ_ζ is chosen to ensure that the distribution of α is normally distributed with a mean of zero and a standard deviation of 2 percent, expressed at an annual rate. As ρ increases, the length of time until the market incorporates the manager's information increases. For example, the information half-life in the baseline case of ρ equal to 0.3, is 0.58 quarters, whereas when ρ is 0.5, the information half-life is one quarter. The time discount parameter, β, is set equal to 0.99, and the coefficient of relative risk aversion, γ, is set equal to 4.

6 An Example Without Managerial Learning

6.1 A Few Brief Observations

There are two non-standard features of our formulation of the active manager's problem, even in the absence of learning. The first is that the labor income growth rate, θ_t, is endogenous, persistent, and it is correlated with the return to the manager's personal portfolio. In particular, using the definitions above, the conditional correlation of labor income and the personal portfolio return can be written as

$$\frac{\text{cov}_t\left(R_{p,t+1},\theta_{t+1}\right)}{\sigma_t\left(R_{p,t+1}\right)\sigma_t\left(\theta_{t+1}\right)}=\frac{\text{cov}_t\left(\eta_t\left(R_{b,t+1}-R_f\right)+R_f,R_{m,t+1}\widetilde{\delta}_{t+1}\left(1-q^p_{t+1|t}-q^d_{t+1|t}\right)\right)}{\eta_t\left(1-q^p_{t+1|t}-q^d_{t+1|t}\right)\sigma_t\left(R_{b,t+1}\right)\sigma_t\left(R_{m,t+1}\widetilde{\delta}_{t+1}\right)},\tag{30}$$

where $\widetilde{\delta}_{t+1}\equiv\exp\left(\delta_{t+1}\right)$. (30) can be rewritten as

$$\frac{\text{cov}_t\left(R_{p,t+1},\theta_{t+1}\right)}{\sigma_t\left(R_{p,t+1}\right)\sigma_t\left(\theta_{t+1}\right)}=\frac{\text{cov}_t\left(R_{b,t+1},R_{m,t+1}\widetilde{\delta}_{t+1}\right)}{\sigma_t\left(R_{b,t+1}\right)\sigma_t\left(R_{m,t+1}\widetilde{\delta}_{t+1}\right)},\tag{31}$$

or (using the definition of $R_{m,t+1}$)

$$\frac{\text{cov}_t\left(R_{p,t+1},\theta_{t+1}\right)}{\sigma_t\left(R_{p,t+1}\right)\sigma_t\left(\theta_{t+1}\right)}=\frac{w_t\text{cov}_t\left(R_{b,t+1},R_{b,t+1}\widetilde{\delta}_{t+1}\right)+\left(1-w_t\right)R_f\text{cov}_t\left(R_{b,t+1},\widetilde{\delta}_{t+1}\right)}{\sigma_t\left(R_{b,t+1}\right)\left\{w_t\sigma_t\left(R_{b,t+1}\widetilde{\delta}_{t+1}\right)+\left(1-w_t\right)R_f\sigma_t\left(\widetilde{\delta}_{t+1}\right)\right\}}.\tag{32}$$

Since η_t drops out of (32), the only control that the manager has over the correlation between labor income and his personal portfolio is indirectly through the composition of the managed portfolio, w_t.

The choice of ω_t that hedges the manager's labor income risk perfectly is

$$\omega^H_t=\frac{-R_f\Xi_t}{\text{cov}_t\left(R_{b,t+1},R_{b,t+1}\widetilde{\delta}_{t+1}\right)+\sigma_t\left(R_{b,t+1}\right)\sigma_t\left(R_{b,t+1}\widetilde{\delta}_{t+1}\right)-R_f\Xi_t},\tag{33}$$

where

$$\Xi_t\equiv\text{cov}_t\left(R_{b,t+1},\widetilde{\delta}_{t+1}\right)+\sigma_t\left(R_{b,t+1}\right)\sigma_t\left(\widetilde{\delta}_{t+1}\right).$$

There are no simple conditions on the exogenous parameters of the model that will ensure that $\omega^H_t\in[0,\overline{\omega}]$.[18] In general the manager's labor income risk can only be partially hedged by trade in the personal portfolio. Intuitively, this follows from the exogenous component of flow-for-performance that is unrelated to the managed portfolio's prior return.

The second non-standard feature of the problem is that managed portfolio choices at time t have an impact on expected future labor income until the end of the manager's horizon.

[18] If $\Xi_t<0$, then the numerator of (33) positive. This occurs if the contemporaneous covariance between the benchmark return and the flow is negative and larger (in absolute value) than the product of the volatilities of benchmark return and flow. This condition, however, is insufficient to ensure that the denominator of (33) is also positive. Positivity of the denominator requires, additionally, that $\text{cov}_t\left(R_{b,t+1},R_{b,t+1}\widetilde{\delta}_{t+1}\right)$ not be too negative.

Intuitively, (10) implies that the impact of the realized return generated by ω_t^m declines geometrically over time. This covariance is nonnegative because of the short sale constraint in the managed portfolio. Although the structure of the problem is different, the results in Viceira (2001) suggest that a positive correlation between labor income and expected returns reduces the optimal allocation to the risky asset in the manager's personal portfolio.

6.2 Portfolio Rules

The optimal w for the benchmark parameterization (see Table 4) is shown in Table 5. Panel A of the table corresponds to a manager with a low level of asset wealth relative to labor income, whereas Panels B and C consider moderate and high relative wealth levels, respectively.[19] The wealth levels used in the table are derived from the support of the endogenous wealth distribution. Each panel is divided in half, with the left side of the panel reporting the weights for a young manager, defined as having 30 years to retirement, and the right side reporting the weights for an old manager, defined as having 2 years to retirement. For each manager type in each panel, the optimal weight is shown for different levels of the other state variables of the problem: relative past performance, z, and the manager's private information, ξ.

Since our sample period includes the run-up of the internet bubble and the bursting of the bubble, growth stock returns have average returns that are low relative to their total volatility; see Table 1. The annualized ex post unconditional Sharpe ratio for this period was only 0.04. This explains why the optimal portfolio weights range from 0.00 to 0.61. The first general observation – an observation that holds true for both young and old managers at all wealth levels – from Table 5 is that the optimal portfolio weight is a positive function of the manager's private information. This is intuitively reasonable. Given the level of ξ, the portfolio is not sensitive to the relative prior performance variable. We conjecture that this

[19] In Table 5 (and our subsequent discussion), wealth refers to wealth after consumption. We make this (slight) change in terminology to accommodate the endogenous state variables in the solution algorithm; see the online Appendix for a discussion.

finding follows because both the flow-for-performance and firing probabilities are smooth (nearly linear) functions of prior performance. In comparing the optimal managed portfolio holdings of young and old managers for a given level of wealth (comparing across the rows in each of the three panels of Table 5), we find that there is virtually no difference in the portfolio allocations of young versus old managers. This finding seems to be inconsistent with the Chevalier and Ellison (1999) finding where younger managers assume less risk than older managers.

The most relevant comparison in Table 5, for empirical purposes, is between the young manager in Panel A and the old manager in Panel C. This follows if older managers, in general, have accumulated more wealth than younger managers. This would certainly be true in a simple life cycle model of consumption and saving, but it also follows if relatively poor managers are demoted (or fired) earlier in their careers. Even accounting for the combined effects of age and wealth, the differences between the young and old managers are trivial.

Tables 6 and 7 repeat the calculations reported in Table 5 for alternative parameterizations in which private information is more or less persistent than in the baseline parameterization. The portfolio choices reported in this table are completely consistent – and very close to – the values reported in Table 5.

7 Incorporating Managerial Learning

7.1 Evolution of the Posterior Distribution

As the last section demonstrated, life-cycle considerations do not have a significant impact on the way in which young versus old managers use private information – when the rate at which the market incorporates that information is known. However, when managers learn about the persistence of their private information, young managers will differ from old managers along the learning dimension. The dynamics of this learning is captured in the evolution of the posterior distribution as managers age.

Given the uniform prior on $[0,1)$, $\pi_0(\rho)$, and the normally distributed pricing errors, $f(\xi_t, \cdots, \xi_1 \mid \rho, \xi_0)$, it follows that the posterior distribution at t is

$$\begin{aligned}
p_t(\rho \mid \xi_t, \cdots, \xi_0) &\propto f(\xi_t, \cdots, \xi_1 \mid \rho, \xi_0)\, \pi_0(\rho) \\
&= \exp\left[-\frac{1}{2\sigma^2}\sum_{s=1}^{t}(\xi_s - \rho\xi_{s-1})^2\right] 1_{\rho\in[0,1)} \\
&= \exp\left[-\frac{1}{2\sigma^2}\left(\sum_{s=1}^{t}\xi_{s-1}^2\right)\left(\rho^2 - 2\frac{\sum_{s=1}^{t}\xi_s\xi_{s-1}}{\sum_{s=1}^{t}\xi_{s-1}^2}\rho - \frac{\sum_{s=1}^{t}\xi_s^2}{\sum_{s=1}^{t}\xi_{s-1}^2}\right)\right] 1_{\rho\in[0,1)} \\
&\propto \exp\left[-\frac{\sum_{s=1}^{t}\xi_{s-1}^2}{2\sigma^2}\left(\rho - \frac{\sum_{s=1}^{t}\xi_s\xi_{s-1}}{\sum_{s=1}^{t}\xi_{s-1}^2}\right)^2\right] 1_{\rho\in[0,1)}.
\end{aligned}$$

If we define

$$\sigma_{p,t} = \frac{\sigma}{\sqrt{\sum_{s=1}^{t}\xi_{s-1}^2}}$$

and

$$\overline{\rho}_t = \frac{\sum_{s=1}^{t}\xi_s\xi_{s-1}}{\sum_{s=1}^{t}\xi_{s-1}^2},$$

then

$$p_t(\rho \mid \xi_t, \cdots, \xi_0) \propto \exp\left[-\frac{1}{2\sigma_{p,t}^2}(\rho - \overline{\rho}_t)^2\right] 1_{\rho\in[0,1)}, \tag{34}$$

which is a truncated normal.

Figures 4 through 6 show the posterior distribution for managers of different ages for the baseline parameterization and the more and less persistent information parameterizations. In Figure 4, the solid line shows the posterior of a manager with two quarters of experience. As expected, with this small amount of data, the posterior is dominated by the uniform prior, indicating that there is still substantial uncertainty about the speed of the market's adjustment to the manager's information. For a manager with just over 10 years of experience, shown in the dashed line, the posterior now places very small weight on the information being extremely persistent ($\rho > 0.8$) and the most weight on $0.2 \leq \rho \leq 0.4$

(recall that the true value of ρ is 0.3). The posterior for a manager with more than 37 years of experience is shown in the dotted line. Consistent with (34), the distribution with the long period of experience is converging to a truncated normal distribution. Intuitively, a manager with substantial experience has a much more precise idea of the speed at which the market responds to his valuable private information. This general finding is confirmed in Figures 5 and 6 for the alternate parameterizations.

7.2 Differences in Portfolio Choices Under Learning

The differences between young and old managers when learning is a part of the problem are shown in Tables 8 through 10. The results for the baseline parameterization are shown in Table 8. In comparison to the full information case in Table 5, the young manager takes larger positions when receiving an extreme negative position, and he takes smaller positive positions when receiving an extreme positive signal.[20] As in Table 5, allocations in the managed portfolios are not sensitive to the manager's wealth. In comparing the results for Tables 5 and 8 for the older manager, we see that the effects of learning about ρ are largely finished, since the portfolio weights in these tables are virtually identical. This indicates that the older manager has sufficient information in the posterior about the value of ρ to make decisions that are almost identical to the certainty case.

A comparison of the young and old managers' portfolio allocations under learning shows some larger differences than in the case without learning. The younger manager holds (roughly) 5 percent larger positions in the risky asset in response to a negative signal when compared to the older manager. At the other extreme, when the managers receive a strong positive signal, the younger manager holds a (roughly) 7 percent smaller risky asset allocations. These results are consistent with the simple intuition that the younger manager has greater uncertainty about ρ. Therefore, they interpret extreme news more

[20] As a technical aside, the model is solved over a grid for the manager's information of $[-0.8, 0.8]$, but we do not report results near the boundary of the state space out of concern for an increase in approximation error.

cautiously. As noted earlier, since managerial allocations are not sensitive to wealth, it does not generate additional differences when young but relatively poor managers are compared to old but relatively rich managers over and above the differences observed for the same level of managerial wealth.

Table 9 and 10 extend the comparison to the alternative information persistence parameterizations. In Table 9, the information is more persistent than in the baseline parameterization. This is equivalent, for inference purposes, to reducing the sample size. The older manager does not learn as much about the value of ρ in (roughly) thirty years to make a significant difference from the uncertainty experienced by the younger manager. As a result, the differences in their choices are smaller, and the young and old managers (again) behave with very similar risky portfolio allocations. In Table 10, when the information is less persistent, we observe that the differences between the young and old managers are exacerbated. They can be as large as 12 percent (or a little more than 1/3 to 1/4 of the total allocation). The issue becomes whether or not these differences – in any of the parameterizations – are economically significant. This is the topic that we consider in the next subsection.

7.3 Economic Significance of the Portfolio Differences

In order to assess economic significance, we consider the implications of the portfolio differences from the perspective of the younger manager. We do not present the utility costs for the older manager since, with a two-year horizon, the CEW estimates are generally small, as would be expected. The manager's measure of economic significance is a version of a standard certainty equivalent wealth (CEW) calculation. In particular, we compare the utility cost of a young manager following the optimal policy of an otherwise identical old manager. "Otherwise identical" means that the young manager has the same level of wealth and prior performance as the old manager.[21] It is impossible to examine the utility cost to a typical investor, in our framework, because we do not explicitly model the investor's decision problem.

The results of the utility cost calculations are presented in Table 11. We consider a range of values for the manager's private information and five different levels of relative wealth. The utility costs are decreasing in the wealth ratio and increasing in the absolute value of the current signal, regardless of the persistence of information. The costs appear to be symmetric about zero for all levels of persistence and manager wealth. In the baseline persistence parameterization, the utility costs are between almost 5 percent and nearly 13 percent of initial wealth, which we judge to be economically significant. In the low information persistence parameterization, these costs are significantly higher, ranging from just over 7 percent to as high as 19 percent of wealth. In the high information persistence case, since managerial learning is less, the utility cost to the young manager of following the old manager's policy is substantially lower, ranging from 1.6 percent to 4.3 percent of initial wealth.

8 Conclusions

How important are career concerns in determining the optimal portfolio choices of mutual fund managers? The existing empirical evidence suggests that younger managers make systematically different investment decisions than older managers (Chevalier and Ellison, 1999) and that they are monitored differently by fund investors (Almazan et al, 2004). However, these findings confound the risk of future labor income with learning about the nature of the manager's private information and how that information is incorporated into

[21] In an earlier version of the paper, we compared the utility of the closet-benchmarking policy with the optimal managed investment policy for a young manager with a current track record equal to the recent prior benchmark performance; i.e., we fix z at zero. This makes the comparison with the closet-benchmarking strategy more direct, since as we noted earlier z is identically zero in all dates and states when the manager closet-benchmarks. These results also suggest substantial utility costs to non-optimal choices, and they are available on request.

market prices.

We formulate a version of the active manager's decision problem as the lifetime consumption and portfolio choices of a market timer with access to costless valuable private information. We deliberately choose simple dynamics for asset returns and information and a standard form of preferences. The manager solves a finite-horizon, discrete-time consumption-portfolio choice problem with endogenous stochastic labor income. The endogeneity of the labor income follows from the manager's active portfolio choices and the flow-for-performance and promotion/demotion probabilities that follow from realized differences between the managed portfolio and the benchmark portfolio. Despite the economic simplicity of the problem, there is no closed-form solution to the manager's problem, and we approximate the value function and the optimal policies using numerical techniques that build on the work of den Haan and Marcet (1991) and Brandt et al (2005).

In the absence of learning and for all parameterizations of the model, we find that the differences in the choices between young managers (defined as a 30 year investment horizon) and old managers (defined as 2 year investment horizon) are small in absolute value. When we consider a version of the problem in which the manager learns about the speed with which the market learns about his private information, we find that there is a larger scope for differences in the responses of young and old managers to extreme information realizations. This difference is not sensitive to wealth, but it is sensitive to the persistence of the value of the private information. We measure the economic significance of the differences in managed portfolio choice in the certainty equivalent wealth loss to young managers from following the optimal policy of an otherwise identical old manager. These costs are significant, ranging from 3 percent of wealth to more than 19 percent of wealth, depending on the realized level of the manager's information, wealth level, and the persistence of the manager's information.

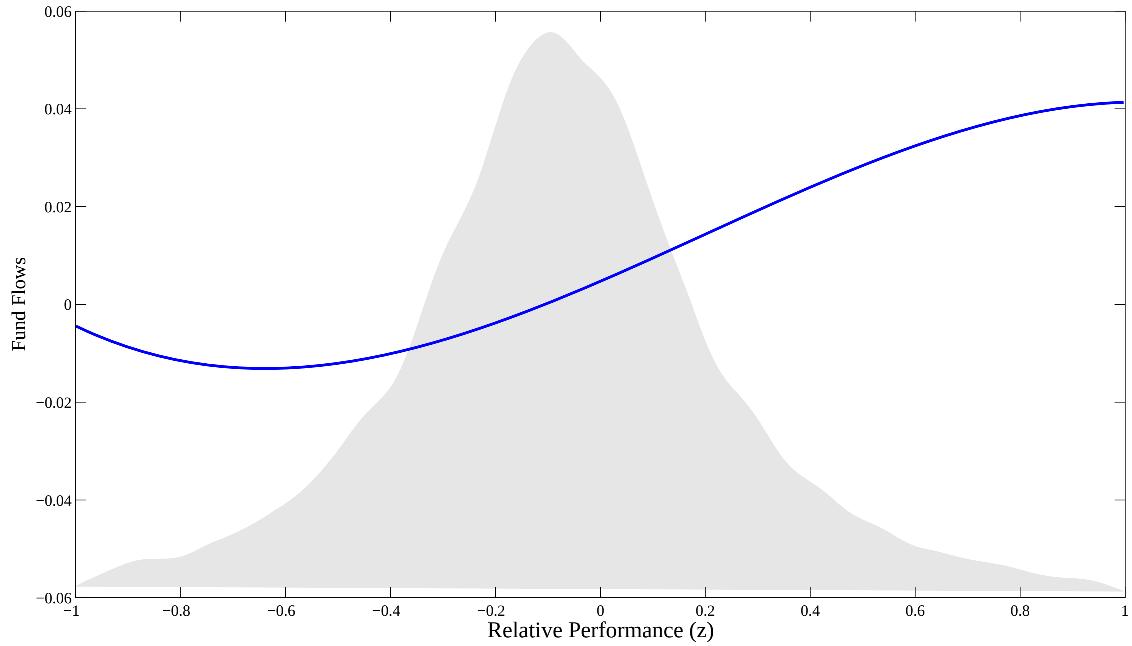

Figure 1: Fund flow as a function of prior relative performance. Prior performance is defined as z from (10), with $\psi = 0.59$. The flow data are quarterly at a quarterly rate. The shaded regions in the background show the empirical density for prior performance.

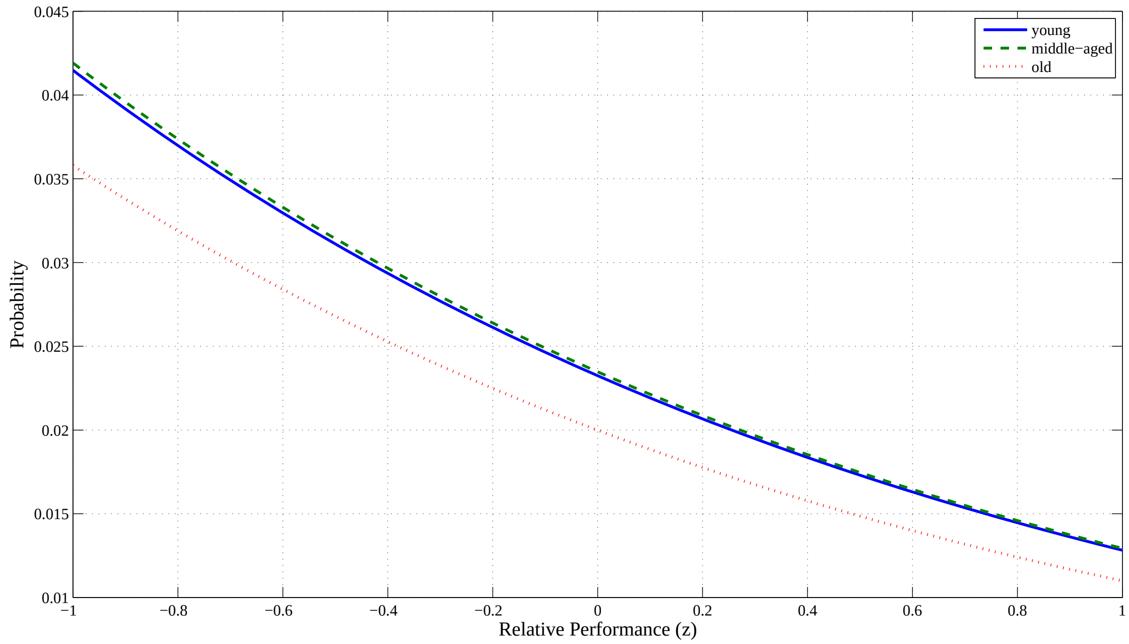

Figure 2: Demotion probabilities as a function of age and prior performance.

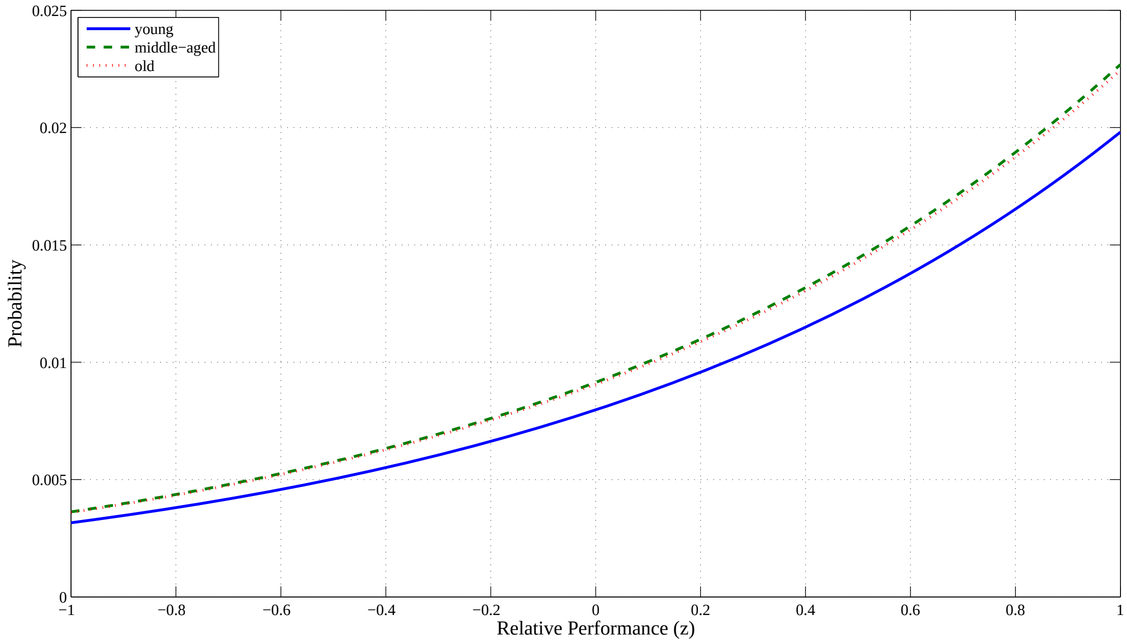

Figure 3: Promotion probabilities as a function of age and prior performance.

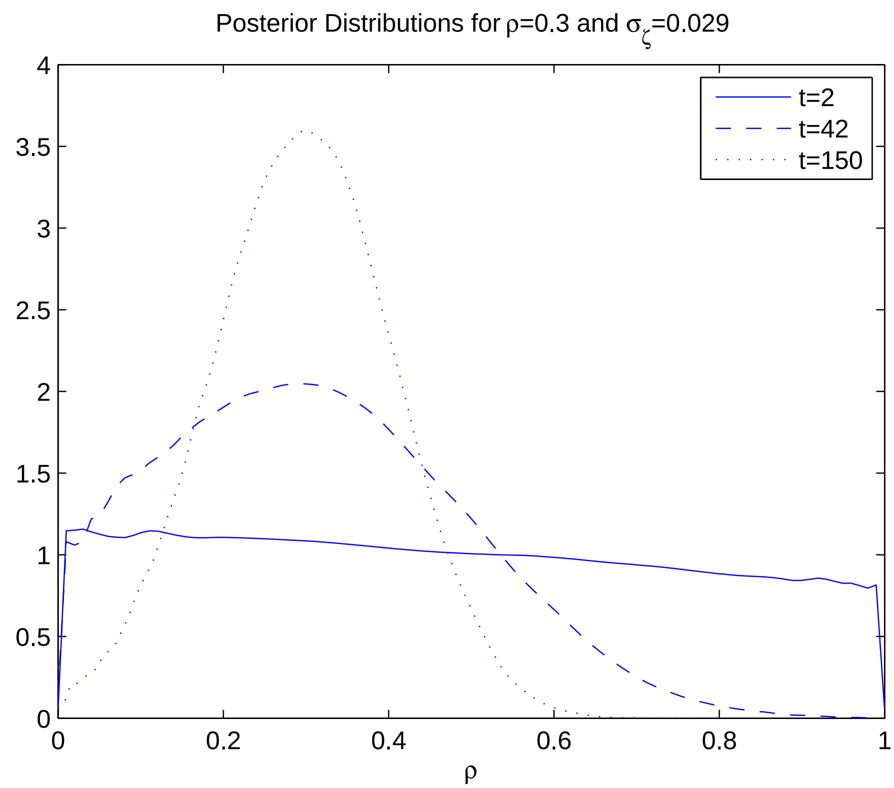

Figure 4: This plot shows the shape of the posterior distribution for ρ, the persistence of information parameter, for different periods of learning, in the baseline private information parameterization. The solid line, $t = 2$, is the posterior for a manager who has been learning for two quarters, the dashed line is the posterior for a manager who has been learning for 42 quarters, and the dotted line is the posterior for a manager who has been learning for 150 quarters.

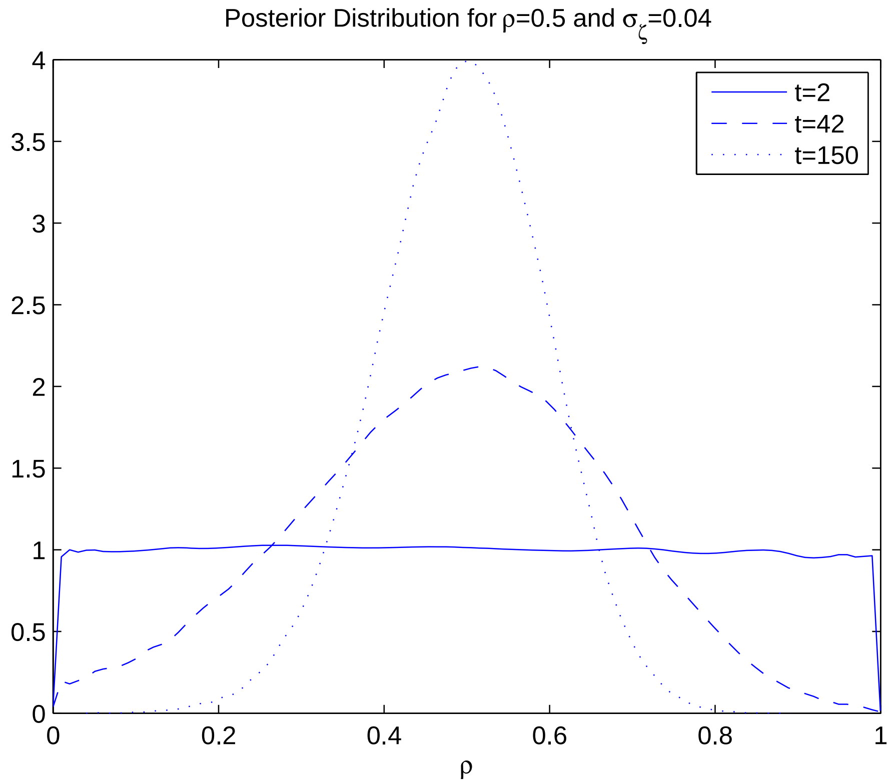

Figure 5: This plot shows the shape of the posterior distribution for ρ, the persistence of information parameter, for different periods of learning, in the private information parameterization with high persistence. The solid line, $t = 2$, is the posterior for a manager who has been learning for two quarters, the dashed line is the posterior for a manager who has been learning for 42 quarters, and the dotted line is the posterior for a manager who has been learning for 150 quarters.

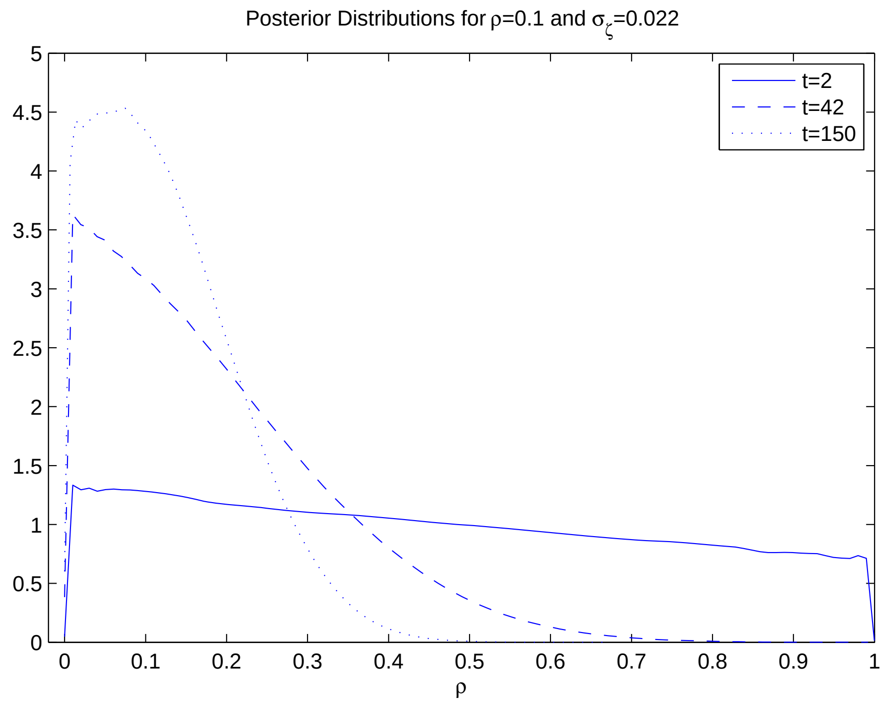

Figure 6: This plot shows the shape of the posterior distribution for ρ, the persistence of information parameter, for different periods of learning, in the private information parameterization with low persistence. The solid line, $t = 2$, is the posterior for a manager who has been learning for two quarters, the dashed line is the posterior for a manager who has been learning for 42 quarters, and the dotted line is the posterior for a manager who has been learning for 150 quarters.

Table 1: Asset Returns: 1994 Q1 to 2006 Q3

Benchmark Portfolio:

Series	Mean	Median	Std. Dev.	Skewness	Kurtosis
Russell 2000 Growth	1.460	1.469	12.988	−0.410	0.245

Risk-free Rate: 0.986

The benchmark returns are nominal, continuously compounded, quarterly, total return series downloaded from http://www.russell.com. All returns are expressed at a quarterly rate, in percent. The risk-free rate is the sample average of the nominal, continuously compounded, end-of-quarter yield-to-maturity on a three-month Treasury bill. The interest rate series is constructed from data in the Federal Reserve Board Statistical Release H15, available at http://www.federalreserve.gov/releases/h15/data.htm.

Table 2: Mutual Fund Flow and Compensation Parameters

Description	Value			
Flow Nonlinearity	δ_0	δ_1	δ_2	δ_3
	0.0142	0.1389	0.0411	-0.0703
Flow Persistence	$\psi = 0.59$			
Flow Volatility	$\sigma_\epsilon = 0.13$			
Promotion Asset Growth	$\upsilon^p = 1.72$			
Demotion Asset Growth	$\upsilon^d = 0.423$			
Fee Share	$\phi = 0.02$			

Table 3: Logit Estimates of Managerial Promotion and Demotion Probabilities

Variable	Demotion	Promotion
alpha	−0.5930 (0.1032)	0.9118 (0.1637)
Young	−3.7300 (0.0565)	−4.8004 (0.1081)
Middle-age	−3.7186 (0.0650)	−4.6615 (0.1144)
Old	−3.8829 (0.0554)	−4.6746 (0.0855)
alpha × *Young*	0.0331 (0.1217)	0.2188 (0.2177)
alpha × *Middle-age*	−0.1348 (0.1037)	0.0482 (0.1997)

These are the coefficients from multinomial logit regressions that fit the parameters φ^d and φ^p in equations (11) though (13). *Alpha* is prior relative performance, defined as z in equation (10), with ψ set to 0.35 and three annual lags of relative returns. *Young* is a dummy variable equal to 1 if a manager's experience is less than or equal to 3 years. *Middle-aged* is a dummy variable equal to 1 if a manager has more than 3 years experience but less than or equal to 7 years experience, and *Old* is a dummy variable equal to 1 if a manager has more than 7 years experience. Standard errors from the maximum likelihood estimation of the parameters are reported in parentheses. There is no intercept in the regression, so young, medium and old all have coefficients. There is no value reported for *alpha* × *Old* because, with young, medium and old all having intercepts, the *alpha* × *old* coefficient is redundant.

Table 4: Managerial Ability and Utility Parameters

Panel A: Managerial Ability

Description	Value		
	Low	Baseline	High
Shock Persistence (ρ)	0.1	0.3	0.5
	Low	Baseline	High
Shock Volatility (σ_ζ)	0.022	0.029	0.040

Panel B: Managerial Utility

Description	Value
Risk Aversion (γ)	4.00
Time Preference (β)	0.99

Table 5: Optimal Managed Portfolio Weights: Baseline Parameterization without Learning

Panel A: Wealth-to-Income Ratio of 1.413

	Young Manager					Old Manager					
	z						z				
ξ	−0.60	−0.20	0.00	0.20	0.60	ξ	−0.60	−0.20	0.00	0.20	0.60
−0.028	0.031	0.036	0.038	0.040	0.041	−0.028	0.023	0.029	0.032	0.034	0.036
−0.014	0.135	0.140	0.142	0.143	0.144	−0.014	0.134	0.139	0.140	0.142	0.143
0.000	0.250	0.254	0.256	0.257	0.258	0.000	0.248	0.253	0.254	0.255	0.257
0.014	0.378	0.382	0.384	0.385	0.386	0.014	0.376	0.381	0.382	0.383	0.384
0.028	0.517	0.521	0.522	0.523	0.524	0.028	0.516	0.520	0.521	0.522	0.523

Panel B: Wealth-to-Income Ratio of 3.110

	Young Manager					Old Manager					
	z						z				
ξ	−0.60	−0.20	0.00	0.20	0.60	ξ	−0.60	−0.20	0.00	0.20	0.60
−0.028	0.036	0.041	0.043	0.044	0.045	−0.028	0.033	0.038	0.040	0.041	0.003
−0.014	0.137	0.142	0.144	0.145	0.145	−0.014	0.138	0.143	0.144	0.145	0.137
0.000	0.250	0.256	0.258	0.259	0.260	0.000	0.252	0.260	0.264	0.265	0.244
0.014	0.378	0.384	0.386	0.387	0.387	0.014	0.378	0.388	0.393	0.399	0.368
0.028	0.516	0.522	0.524	0.525	0.525	0.028	0.514	0.526	0.532	0.542	0.505

Panel C: Wealth-to-Income Ratio of 13.527

	Young Manager					Old Manager					
	z						z				
ξ	−0.60	−0.20	0.00	0.20	0.60	ξ	−0.60	−0.20	0.00	0.20	0.60
−0.028	0.033	0.039	0.041	0.042	0.043	−0.028	0.026	0.031	0.032	0.034	0.035
−0.014	0.134	0.140	0.142	0.143	0.144	−0.014	0.128	0.133	0.135	0.137	0.137
0.000	0.249	0.255	0.257	0.258	0.259	0.000	0.245	0.252	0.254	0.255	0.256
0.014	0.377	0.384	0.385	0.387	0.387	0.014	0.376	0.384	0.386	0.387	0.388
0.028	0.517	0.523	0.525	0.526	0.527	0.028	0.518	0.526	0.529	0.530	0.531

The wealth-to-income ratio is the ratio of wealth after consumption to labor income. z is the relative performance state variable, and ξ is the shock to the manager's private information. A young manager is one with thirty years to the end of the investment horizon, whereas an old manager has only two years to the end of the horizon. The optimal managed portfolio weight is constrained to lie in the interval $[0, 1.33]$.

Table 6: Optimal Managed Portfolio Weights: High Persistence Parameterization without Learning

Panel A: Wealth-to-Income Ratio of 1.413

	Young Manager					Old Manager					
	z						z				
ξ	−0.60	−0.20	0.00	0.20	0.60	ξ	−0.60	−0.20	0.00	0.20	0.60
−0.028	0.091	0.096	0.098	0.099	0.095	−0.028	0.078	0.095	0.096	0.097	0.090
−0.014	0.167	0.172	0.173	0.175	0.175	−0.014	0.166	0.170	0.172	0.173	0.174
0.000	0.251	0.256	0.257	0.258	0.259	0.000	0.250	0.254	0.256	0.257	0.258
0.014	0.342	0.347	0.348	0.349	0.350	0.014	0.341	0.345	0.346	0.347	0.348
0.028	0.439	0.443	0.445	0.446	0.446	0.028	0.438	0.442	0.443	0.444	0.445

Panel B: Wealth-to-Income Ratio of 3.110

	Young Manager					Old Manager					
	z						z				
ξ	−0.60	−0.20	0.00	0.20	0.60	ξ	−0.60	−0.20	0.00	0.20	0.60
−0.028	0.092	0.097	0.098	0.100	0.100	−0.028	0.091	0.096	0.097	0.098	0.097
−0.014	0.168	0.174	0.175	0.176	0.177	−0.014	0.169	0.176	0.177	0.178	0.163
0.000	0.252	0.257	0.259	0.260	0.261	0.000	0.253	0.261	0.264	0.266	0.239
0.014	0.342	0.348	0.350	0.351	0.352	0.014	0.343	0.352	0.357	0.362	0.325
0.028	0.439	0.445	0.446	0.448	0.448	0.028	0.438	0.449	0.454	0.462	0.462

Panel C: Wealth-to-Income Ratio of 13.527

	Young Manager					Old Manager					
	z						z				
ξ	−0.60	−0.20	0.00	0.20	0.60	ξ	−0.60	−0.20	0.00	0.20	0.60
−0.028	0.089	0.095	0.097	0.098	0.099	−0.028	0.082	0.087	0.089	0.090	0.091
−0.014	0.166	0.171	0.173	0.175	0.175	−0.014	0.160	0.166	0.168	0.169	0.170
0.000	0.250	0.256	0.258	0.259	0.260	0.000	0.246	0.253	0.255	0.256	0.257
0.014	0.341	0.347	0.349	0.351	0.351	0.014	0.340	0.347	0.349	0.350	0.351
0.028	0.439	0.445	0.447	0.448	0.449	0.028	0.439	0.447	0.449	0.450	0.451

The wealth-to-income ratio is the ratio of wealth after consumption to labor income. z is the relative performance state variable, and ξ is the shock to the manager's private information. A young manager is one with thirty years to the end of the investment horizon, whereas an old manager has only two years to the end of the horizon. The optimal managed portfolio weight is constrained to lie in the interval $[0, 1.33]$.

Table 7: Optimal Managed Portfolio Weights: Low Persistence Parameterization without Learning

Panel A: Wealth-to-Income Ratio of 1.413

	Young Manager						Old Manager				
			z						z		
ξ	−0.60	−0.20	0.00	0.20	0.60	ξ	−0.60	−0.20	0.00	0.20	0.60
−0.028	0.000	0.000	0.000	0.000	0.000	−0.028	0.000	0.000	0.000	0.000	0.000
−0.014	0.105	0.110	0.112	0.113	0.114	−0.014	0.104	0.109	0.110	0.112	0.113
0.000	0.248	0.253	0.255	0.256	0.257	0.000	0.247	0.252	0.253	0.254	0.256
0.014	0.415	0.419	0.421	0.422	0.422	0.014	0.413	0.418	0.419	0.420	0.421
0.028	0.597	0.601	0.602	0.603	0.603	0.028	0.596	0.600	0.601	0.602	0.602

Panel B: Wealth-to-Income Ratio of 3.110

	Young Manager						Old Manager				
			z						z		
ξ	−0.60	−0.20	0.00	0.20	0.60	ξ	−0.60	−0.20	0.00	0.20	0.60
−0.028	0.000	0.000	0.000	0.000	0.000	−0.028	0.000	0.000	0.000	0.000	0.000
−0.014	0.106	0.111	0.113	0.114	0.115	−0.014	0.106	0.111	0.112	0.113	0.112
0.000	0.249	0.255	0.257	0.258	0.259	0.000	0.251	0.259	0.263	0.264	0.246
0.014	0.414	0.420	0.422	0.423	0.424	0.014	0.414	0.425	0.431	0.438	0.408
0.028	0.595	0.601	0.603	0.604	0.605	0.028	0.592	0.605	0.611	0.325	0.589

Panel C: Wealth-to-Income Ratio of 13.527

	Young Manager						Old Manager				
			z						z		
ξ	−0.60	−0.20	0.00	0.20	0.60	ξ	−0.60	−0.20	0.00	0.20	0.60
−0.028	0.000	0.000	0.000	0.000	0.000	−0.028	0.000	0.000	0.000	0.000	0.000
−0.014	0.104	0.109	0.111	0.112	0.113	−0.014	0.097	0.102	0.104	0.105	0.106
0.000	0.248	0.254	0.256	0.257	0.258	0.000	0.244	0.250	0.252	0.254	0.255
0.014	0.414	0.421	0.423	0.424	0.424	0.014	0.414	0.421	0.424	0.425	0.426
0.028	0.597	0.603	0.605	0.606	0.606	0.028	0.600	0.608	0.611	0.612	0.612

The wealth-to-income ratio is the ratio of wealth after consumption to labor income. z is the relative performance state variable, and ξ is the shock to the manager's private information. A young manager is one with thirty years to the end of the investment horizon, whereas an old manager has only two years to the end of the horizon. The optimal managed portfolio weight is constrained to lie in the interval $[0, 1.33]$.

Table 8: Optimal Managed Portfolio Weights: Baseline Parameterization with Learning

Panel A: Wealth-to-Income Ratio of 1.413

	Young Manager					Old Manager					
	z						z				
ξ	−0.60	−0.20	0.00	0.20	0.60	ξ	−0.60	−0.20	0.00	0.20	0.60
−0.028	0.075	0.086	0.088	0.089	0.085	−0.028	0.024	0.030	0.033	0.035	0.037
−0.014	0.160	0.165	0.167	0.168	0.169	−0.014	0.135	0.139	0.141	0.142	0.143
0.000	0.249	0.254	0.255	0.257	0.257	0.000	0.248	0.253	0.254	0.256	0.257
0.014	0.346	0.351	0.352	0.353	0.354	0.014	0.376	0.380	0.382	0.383	0.384
0.028	0.449	0.454	0.455	0.456	0.457	0.028	0.514	0.518	0.520	0.520	0.521

Panel B: Wealth-to-Income Ratio of 3.110

	Young Manager					Old Manager					
	z						z				
ξ	−0.60	−0.20	0.00	0.20	0.60	ξ	−0.60	−0.20	0.00	0.20	0.60
−0.028	0.081	0.087	0.088	0.089	0.090	−0.028	0.034	0.040	0.041	0.043	0.004
−0.014	0.161	0.167	0.169	0.170	0.170	−0.014	0.138	0.143	0.145	0.146	0.138
0.000	0.250	0.256	0.257	0.259	0.259	0.000	0.252	0.260	0.264	0.265	0.244
0.014	0.346	0.352	0.354	0.355	0.356	0.014	0.378	0.388	0.393	0.399	0.368
0.028	0.449	0.455	0.457	0.458	0.458	0.028	0.513	0.525	0.531	0.541	0.504

Panel C: Wealth-to-Income Ratio of 13.527

	Young Manager					Old Manager					
	z						z				
ξ	−0.60	−0.20	0.00	0.20	0.60	ξ	−0.60	−0.20	0.00	0.20	0.60
−0.028	0.079	0.085	0.087	0.088	0.089	−0.028	0.027	0.032	0.034	0.035	0.036
−0.014	0.159	0.165	0.167	0.168	0.169	−0.014	0.128	0.134	0.136	0.137	0.138
0.000	0.248	0.254	0.256	0.257	0.258	0.000	0.245	0.252	0.254	0.255	0.256
0.014	0.346	0.352	0.354	0.355	0.356	0.014	0.376	0.383	0.385	0.387	0.388
0.028	0.449	0.455	0.457	0.459	0.459	0.028	0.517	0.525	0.527	0.529	0.529

The baseline persistence parameter is $\rho = 0.3$. The wealth-to-income ratio is the ratio of wealth after consumption to labor income. z is the relative performance state variable, and ξ is the shock to the manager's private information. A young manager is one with thirty years to the end of the investment horizon, whereas an old manager has only two years to the end of the horizon. The optimal managed portfolio weight is constrained to lie in the interval $[0, 1.33]$.

Table 9: Optimal Managed Portfolio Weights: High Persistence Parameterization with Learning

Panel A: Wealth-to-Income Ratio of 1.413

	Young Manager					Old Manager					
	z						z				
ξ	−0.60	−0.20	0.00	0.20	0.60	ξ	−0.60	−0.20	0.00	0.20	0.60
−0.028	0.075	0.086	0.088	0.089	0.085	−0.028	0.072	0.089	0.091	0.092	0.085
−0.014	0.161	0.166	0.168	0.169	0.170	−0.014	0.163	0.167	0.169	0.170	0.171
0.000	0.251	0.256	0.257	0.258	0.259	0.000	0.250	0.254	0.256	0.257	0.258
0.014	0.349	0.353	0.355	0.356	0.357	0.014	0.344	0.349	0.350	0.351	0.356
0.028	0.453	0.457	0.459	0.460	0.460	0.028	0.445	0.450	0.451	0.452	0.453

Panel B: Wealth-to-Income Ratio of 3.110

	Young Manager					Old Manager					
	z						z				
ξ	−0.60	−0.20	0.00	0.20	0.60	ξ	−0.60	−0.20	0.00	0.20	0.60
−0.028	0.082	0.087	0.088	0.090	0.090	−0.028	0.085	0.090	0.092	0.093	0.093
−0.014	0.162	0.168	0.170	0.171	0.171	−0.014	0.166	0.173	0.174	0.175	0.160
0.000	0.252	0.257	0.259	0.260	0.261	0.000	0.253	0.261	0.265	0.266	0.240
0.014	0.349	0.355	0.357	0.358	0.359	0.014	0.346	0.356	0.360	0.366	0.330
0.028	0.452	0.458	0.460	0.461	0.462	0.028	0.446	0.457	0.462	0.470	0.470

Panel C: Wealth-to-Income Ratio of 13.527

	Young Manager					Old Manager					
	z						z				
ξ	−0.60	−0.20	0.00	0.20	0.60	ξ	−0.60	−0.20	0.00	0.20	0.60
−0.028	0.080	0.085	0.087	0.088	0.089	−0.028	0.076	0.082	0.084	0.085	0.086
−0.014	0.160	0.166	0.167	0.169	0.169	−0.014	0.157	0.163	0.165	0.166	0.167
0.000	0.250	0.256	0.258	0.259	0.260	0.000	0.246	0.253	0.255	0.256	0.257
0.014	0.348	0.354	0.356	0.358	0.358	0.014	0.344	0.351	0.353	0.354	0.355
0.028	0.453	0.459	0.461	0.462	0.463	0.028	0.447	0.454	0.457	0.458	0.459

The high persistence parameterization is $\rho = 0.5$. The wealth-to-income ratio is the ratio of wealth after consumption to labor income. z is the relative performance state variable, and ξ is the shock to the manager's private information. A young manager is one with thirty years to the end of the investment horizon, whereas an old manager has only two years to the end of the horizon. The optimal managed portfolio weight is constrained to lie in the interval $[0, 1.33]$.

Table 10: Optimal Managed Portfolio Weights: Low Persistence Parameterization with Learning

Panel A: Wealth-to-Income Ratio of 1.413

	Young Manager					Old Manager					
	z						z				
ξ	−0.60	−0.20	0.00	0.20	0.60	ξ	−0.60	−0.20	0.00	0.20	0.60
−0.028	0.075	0.086	0.088	0.089	0.085	−0.028	0.000	0.000	0.000	0.000	0.003
−0.014	0.160	0.165	0.166	0.167	0.168	−0.014	0.114	0.118	0.120	0.121	0.122
0.000	0.248	0.253	0.254	0.255	0.256	0.000	0.247	0.252	0.253	0.254	0.255
0.014	0.344	0.349	0.351	0.352	0.352	0.014	0.400	0.405	0.406	0.407	0.408
0.028	0.447	0.452	0.453	0.454	0.454	0.028	0.568	0.572	0.573	0.574	0.575

Panel B: Wealth-to-Income Ratio of 3.110

	Young Manager					Old Manager					
	z						z				
ξ	−0.60	−0.20	0.00	0.20	0.60	ξ	−0.60	−0.20	0.00	0.20	0.60
−0.028	0.081	0.086	0.088	0.089	0.090	−0.028	0.000	0.004	0.005	0.007	0.000
−0.014	0.161	0.166	0.168	0.169	0.170	−0.014	0.116	0.121	0.122	0.123	0.120
0.000	0.249	0.254	0.256	0.257	0.258	0.000	0.250	0.259	0.263	0.264	0.245
0.014	0.345	0.350	0.352	0.354	0.354	0.014	0.401	0.412	0.418	0.425	0.349
0.028	0.447	0.453	0.454	0.456	0.456	0.028	0.565	0.577	0.584	0.596	0.560

Panel C: Wealth-to-Income Ratio of 13.527

	Young Manager					Old Manager					
	z						z				
ξ	−0.60	−0.20	0.00	0.20	0.60	ξ	−0.60	−0.20	0.00	0.20	0.60
−0.028	0.079	0.085	0.086	0.088	0.088	−0.028	0.000	0.000	0.000	0.000	0.000
−0.014	0.158	0.164	0.166	0.167	0.168	−0.014	0.106	0.112	0.114	0.115	0.116
0.000	0.247	0.253	0.255	0.256	0.257	0.000	0.244	0.250	0.252	0.254	0.255
0.014	0.344	0.350	0.352	0.353	0.354	0.014	0.401	0.408	0.410	0.412	0.413
0.028	0.447	0.453	0.455	0.456	0.457	0.028	0.572	0.580	0.582	0.583	0.584

The low persistence parameterization is $\rho = 0.1$. The wealth-to-income ratio is the ratio of wealth after consumption to labor income. z is the relative performance state variable, and ξ is the shock to the manager's private information. A young manager is one with thirty years to the end of the investment horizon, whereas an old manager has only two years to the end of the horizon. The optimal managed portfolio weight is constrained to lie in the interval $[0, 1.33]$.

Table 11: Managerial Utility Costs with Learning

Panel A: Baseline Information Persistence

	Wealth				
ξ	1.413	1.887	3.110	5.148	13.527
−0.028	0.129	0.123	0.115	0.105	0.088
−0.014	0.108	0.103	0.095	0.085	0.068
0.000	0.088	0.083	0.075	0.065	0.048
0.014	0.108	0.103	0.095	0.085	0.068
0.028	0.128	0.123	0.115	0.105	0.089

Panel B: High Information Persisence

	Wealth				
ξ	1.413	1.887	3.110	5.148	13.527
−0.028	0.043	0.041	0.038	0.035	0.029
−0.014	0.036	0.034	0.032	0.028	0.023
0.000	0.030	0.028	0.025	0.022	0.016
0.014	0.036	0.034	0.032	0.028	0.023
0.028	0.043	0.041	0.038	0.035	0.029

Panel C: Low Information Persistence

	Wealth				
ξ	1.413	1.887	3.110	5.148	13.527
−0.028	0.193	0.184	0.172	0.157	0.132
−0.014	0.162	0.154	0.143	0.128	0.102
0.000	0.133	0.125	0.112	0.097	0.072
0.014	0.162	0.155	0.142	0.127	0.103
0.028	0.191	0.184	0.173	0.157	0.133

Wealth is defines as the ratio of wealth after consumption to labor income, and ξ is the shock to the manager's private information. A young manager is one with thirty years to the end of the investment horizon, whereas an old manager has only two years to the end of the horizon. The utility cost is the proportion of current wealth that the young manager would pay to be allowed to switch from the old manager's policy to his perceived optimal policy.

References

[1] Almazan, A., K. C. Brown, M. Carlson, and D. A. Chapman, 2004, "Why Constrain Your Mutual Fund Manager?" **Journal of Financial Economics, 73**, pages 289-321.

[2] Baks, K. P., 2006, "On the Performance of Mutual Fund Managers," forthcoming, **Journal of Finance**.

[3] Balduzzi, P. and A. W. Lynch, 1999, "Transactions Costs and Predictability: Some Utility Cost Calculations," **Journal of Financial Economics, 52**, pages 47-78.

[4] Barillas, F., and J. Fernández-Villaverde, 2007, "A Generalization of the Endogenous Grid Method," **Journal of Economic Dynamics and Control**, 31, pages 2698-2712.

[5] Berk, J. B., and R. C. Green, 2004, "Mutual Fund Flows and Performance in Rational Markets," **Journal of Political Economy, 112**, pages 1269-1295.

[6] Bhattacharya, S., and P. Pfleiderer, 1985, "Delegated Portfolio Management," **Journal of Economic Theory, 36**, pages 1-25.

[7] Brandt, M. W., A. Goyal, P. Santa-Clara, and J. R. Stroud, 2005, "A Simulation Approach to Dynamic Portfolio Choice with an Application to Learning About Return Predictability," **Review of Financial Studies, 18**, pages 831-873.

[8] Brown, K. C., W. V. Harlow, and L. T. Starks, 1996, "Of Tournaments and Temptations: An Analysis of Managerial Incentives in the Mutual Fund Industry," **Journal of Finance, 51**, pages 85-110.

[9] Campbell, J. Y. and L. M. Viceira, 1999. "Consumption and Portfolio Decisions When Expected Returns are Time Varying," **Quarterly Journal of Economics, 114**, pages 433-495.

[10] Carhart, M. M., 1997, "On Persistence in Mutual Fund Performance," **Journal of Finance, 52**, pages 57–82.

[11] Chevalier, J., and G. Ellison, 1997, "Risk Taking by Mutual Funds as a Response to Incentives," **Journal of Political Economy**, 105, pages 1167-1200.

[12] Chevalier, J., and G. Ellison, 1999, "Career Concerns of Mutual Fund Managers," **Quarterly Journal of Economics, 114**, pages 389-432.

[13] Dangl, T., Y. Wu, and J. Zechner, 2008, "Market Discipline and Internal Governance in the Mutual Fund Industry," **Review of Financial Studies**, 21, pages 2307-2343.

[14] Deaton, A., 1991, "Saving and Liquidity Constraints," **Econometrica, 59**, pages 1221-1248.

[15] den Haan, W. J., and A. Marcet, 1990, "Solving the Stochastic Growth Model by Parameterizing Expectations," **Journal of Business and Economic Statistics, 8**, pages 31-34.

[16] Duffie, D., 2001, **Dynamic Asset Pricing Theory, 3rd Ed.** Princeton: Princeton University Press.

[17] Evans, R. B., 2006, "Does Alpha Really Matter? Evidence from Mutual Fund Incubation, Termination, and Manager Change," manuscript, Boston College.

[18] Fama, E. F., and K. R. French, 1993, "Common Risk Factors in the Returns on Stocks and Bonds," **Journal of Financial Economics, 33**, pages 3–56.

[19] Grinold, R. C., and R. N. Kahn, 2000. **Active Portfolio Management, 2nd Ed.** New York: McGraw-Hill.

[20] Holmström, B, 1999, "Managerial Incentive Problems: A Dynamic Perspective," **Review of Economic Studies, 66**, pages 169-182.

[21] Hu, F., A. R. Hall, and C. R. Harvey, 2000, "Promotion or Demotion? An Empirical Investigation of the Determinants of Top Mutual Fund Manager Change," Manuscript, Duke University.

[22] Jurek, J. W., and H. Yang, 2007, "Dynamic Portfolio Selection in Arbitrage," manuscript, Princeton University and Harvard University.

[23] Koijen, R. S. J, T. E. Nijman, and B. J. M. Werker, 2008, "When Can Life-cycle Investors Benefit from Time-varying Bond Risk Premia?" (with Appendix), Manuscript, University of Chicago and Tilburg University.

[24] Kondor, P., 2006, "Career Concerns and Dynamic Arbitrage," manuscript, London School of Economics.

[25] Koo, H. K., 1999, "Consumption and Portfolio Choice with Labor Income: A Discrete-Time Approach," **Mathematical Methods of Operations Research, 50**, pages 219-243.

[26] Liu, J., and F. A. Longstaff, 2004, "Losing Money on Arbitrage: Optimal Dynamic Portfolio Choice in Markets with Arbitrage Opportunities," **Review of Financial Studies, 17**, pages 611-641.

[27] Merton, R. C., 1971, "Optimum Consumption and Portfolio Rules in a Continuous-Time Model," **Journal of Economic Theory, 3**, pages 373-413.

[28] Sirri, E., and P. Tufano, 1998, "Costly Search and Mutual Fund Flows," **Journal of Finance, 53**, pages 1589-1622.

[29] Starks, L. T., 1987, "Performance Incentive Fees: An Agency Theoretic Approach," **Journal of Financial and Quantitative Analysis, 22**, pages 17-32.

[30] Stoughton, N. M., 1993, "Moral Hazard and the Portfolio Management Problem," **Journal of Finance, 48**, pages 2009-2028.

[31] Viceira, L. M., 2001, "Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income," **Journal of Finance, 56**, pages 433-470.

[32] Xiong, W., 2001, "Convergence Trading with Wealth Effects: An Amplification Mechanism in Financial Markets," **Journal of Financial Economics, 62**, pages 247-292.

[33] Zeldes, S. P., 1989, "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivelance," **Quarterly Journal of Economics, 104**, pages 275-298.