Can Information Costs Explain the Equity Premium and Stock Market Participation Puzzles?

Hui Chen*

University of Chicago GSB

November 2006

Abstract

Unlikely. Rational delay models argue that costs of acquiring and processing information will result in delays in consumption adjustments, which could potentially explain the low contemporaneous covariation between stock returns and consumption growth. However, information costs are likely to be small. In order to explain the equity premium puzzle with such costs, it is important that small costs can generate significant delays in adjustments. This condition holds when returns are *i.i.d.* However, return predictability implies large benefit for market timing, which makes agents plan more frequently. The inattentive investors rebalance their portfolios infrequently, and follow a unique strategy mixing the buy-and-hold and hedging demands. Major events draw immediate attention to the market ex post, but they actually make investors plan less frequently ex ante. Fixed information costs lead to big variations for stock holdings across households, but they have limited effects on stock market participation. This is because infrequent planning helps investors "dilute" the fixed costs.

*I thank Snehal Banerjee, Frederico Belo, George Constantinides, Lars Hansen, John Heaton, Nick Roussanov, Martin Schneider, Pietro Veronesi, participants of the GSB Asset Pricing Brownbag and the 5th Transatlantic Doctoral Conference at London Business School, and especially John Cochrane and Monika Piazzesi for helpful comments. All errors are my own. Correspondence to hchen5@ChicagoGSB.edu.

1 Introduction

Rational delay models argue that since it is costly to acquire and process information to make consumption and portfolio decisions, agents will respond to new information with delays. Such delays could potentially explain the low covariation between stock returns and consumption growth, which is at the heart of the equity premium puzzle (Mehra and Prescott 1985). Information costs also appear to be an attractive explanation for the stock market participation puzzle. If managing a stock portfolio incurs nontrivial fixed costs, that could keep those households with low net worth out of the stock market.

Making optimal consumption and investment decisions in a dynamic environment requires information as well as expertise. Even though a huge amount of financial information is freely available today, the task of processing such information can still be challenging. The flourishing business of individual financial advisory is proof that such information costs do matter. Unfortunately, there is not much evidence on how big these costs are. Agents can invest in an index fund and follow certain rule-of-thumb consumption plans, which would require very little information or financial knowledge. These "simple" consumption and investment plans will put an upper bound on the information costs.

Thus, it is important that small (perhaps undetectable) costs are sufficient to generate significant adjustment delays and non-participation. This condition may not hold in realistic environments. On the one hand, major market events will draw immediate attention from all investors, while time-varying investment opportunities can bring large opportunity costs for investors who plan infrequently. On the other hand, a buy-and-hold strategy seems like an easy way to avoid any information costs for investing.

This paper addresses these issues by studying a model of consumption and portfolio decisions when information is costly to process. Agents can invest in two assets, a stock and a riskfree bond, or put money in a checking account that pays no interest. They consume on the income from the financial assets. With lump-sum type information processing costs, agents plan for consumption and investment at endogenous frequencies. When expected returns are constant, the optimal planning frequency is constant, and small information costs can indeed generate significant delays in planning. The possibility of big price jumps actually makes agents plan less frequently, not more. When expected returns are time-varying, the benefit of market timing makes agents choose much higher planning frequencies. Finally, fixed information costs have limited effects on stock market participation, because inattentive investors can effectively "dilute" the fixed costs through infrequent planning.

The model assumes that whenever an agent acquires and processes new information to make consumption and investment decisions, she pays lump-sum information costs. To avoid excessive information costs, the agent makes "long-term" plans, and remains inattentive to new information within each planning period. Consequently, she will respond to changes in the financial market with a lag. In the aggregate, such delays lead to low contemporaneous covariations between equity returns and consumption growth. The effectiveness of this mechanism crucially depends on the planning frequencies agents choose. A main goal of this paper is to investigate what determines the planning frequency in a realistic environment.

Longer planning horizons require agents to hold more money in the checking account (as opposed to investing in stocks and bonds), and prevent them from rebalancing their portfolios for extended periods. I refer to the first type of costs as checking account costs, the second as unbalanced portfolio costs. The trade-off between the information costs and these two types of opportunity costs determines the planning horizon.

When returns are *i.i.d.*, the main determinant of planning frequency is checking account costs, the size of which depends on the level of expected returns on stocks and the portfolio weight. In this case, planning horizon rises quickly for low levels of information processing costs. With calibrated parameters, information costs as low as 0.01% of the net worth can generate delays of two quarters or more. This result is essentially a restatement of Cochrane (1989): the utility costs of "near-rational" behaviors are quite small under standard settings.

When stock prices take small and infrequent jumps, it has little effect on the planning horizon. But a potential big jump can make agents plan much less frequently. At first look, this result is contrary to our intuition for major events: a common criticism of models of inattention is that major events will draw immediate attention from inattentive investors, thus forcing them to plan more frequently. While that is true ex post, price jumps make investors hold less stocks ex ante, which reduces the costs of unbalanced portfolios and leads to longer planning horizons.

Time-varying investment opportunities can significantly reduce delays. Using parameters calibrated with the procedure of Campbell and Viceira (1999), and information costs of 0.01% of net worth, the model produces planning horizons that are less than one quarter most of the time. Moreover, planning horizon will be long when the conditional equity premium is low, short when the premium is high, and particularly short in a region where portfolio weights are highly sensitive to changes in the conditional equity premium. The reason is that market timing raises the utility costs of "near-rational" behaviors, especially at times when a small change in investment opportunities requires significant adjustments in portfolio choices.

The portfolio choices for inattentive investors are unique. Since inattentive investors do not rebalance their portfolios within each planning period, they will never take leveraged or short positions. The portfolio rule in the case of constant expected returns coincides with that of a buy-and-hold investor with finite horizon, although the investment horizon will be endogenous. When there are no price jumps, the investment horizon is just the planning horizon; if there are price jumps, the horizon will also depend on the arrival time of jumps, which is random. When returns are predictable, the portfolio strategy consists of two parts: a buy-and-hold demand, and an extra hedging demand.

With fixed information costs, the planning horizons and portfolio weights become different across households with different net worth. However, long planning horizons make the per-period fixed costs of participation smaller, so only those households with extremely low wealth will stay out of the stock market. This could explain why there is such a lack of rebalancing activities in individual 401(K) plans (see, e.g., Ameriks and Zeldes 2001).

Several papers have explored the effects of infrequent decision making on consumption and investment. Lynch (1996) studies an OLG model where agents make decisions at fixed frequencies and without synchronizing with each other. He shows that such a setup can deliver the low correlation between equity returns and consumption growth in the data. In a continuous-time setup, Gabaix and Laibson (2001) find that estimations of the coefficient of relative risk aversion from the standard consumption Euler equation could generate a multiplicative bias of $6D$, where D is the delay in quarters. Both results are clearly sensitive to agents' planning frequencies, which are exogenous in these studies. An understanding of what determines planning frequencies helps us assess the success of these theories.

There is a large literature on consumption and portfolio allocation with adjustment delays. Examples include models of portfolio transaction costs (e.g., Constantinides (1986), Davis and Norman (1990), Shreve and Soner (1994)), and consumption adjustment costs (e.g., Grossman and Laroque (1990) and Marshall and Parekh (1999)). These models assume that agents process newly arrived information immediately. Typically, a "no-action zone" characterizes the optimal strategy: as long as one is not too far away from the optimal consumption level or portfolio position, no adjustment is necessary. While these models also generate delays, the "no-action zone" type decision rules make it difficult to quantify the exact magnitude of delays, especially when we move to richer settings.

From a theoretical point of view, lump-sum information processing costs lead to a dynamic structure for the portfolio problem that is completely different from that of standard portfolio problems with adjustment costs. The nature of the costs determines that agents will choose to receive and process information in "bundles" that arrive in discrete time, even though information is flowing continuously[1]. As a result, the consumption and port-

[1] This feature also distinguishes this model from models of information-processing capacity, including Sims (2003), Moscarini (2004), Peng (2005), Nieuwerburgh and Veldkamp (2006), where agents still constantly process information, but can choose the quality of information they receive.

folio choice problem has the "*discrete-continuous-time*" structure as in Duffie and Sun (1990). Duffie and Sun also interpret the transaction costs in their model as "the cost of adjusting the portfolio and the cost of processing information". They prove that the optimal planning frequency is constant in a setting with *i.i.d.* returns, CRRA preference, and costs proportional to the portfolio value. Using the same structure, Reis (2004) studies the endogenous planning frequency in a consumption problem with idiosyncratic labor income. He does not consider risky assets.

This paper is also related to the literature on dynamic asset allocations, led by the seminal work of Merton (1971). Recent studies have incorporated the empirical findings of return predictability and price jumps into the problem.[2] This paper extends this literature by looking at the portfolio implications of major market movement and time-varying expected returns in a model with information processing costs.

2 The Model

Consider an economy with infinitely lived agents. There are three securities available: a checking account that pays no interest, a riskfree asset (bond) with constant riskfree rate r_f, and a risky asset (stock). The return on the risky asset is given by

$$\frac{dP_t}{P_t} = \mu_t dt + \sigma dz_t + \xi dM_t, \tag{1}$$

where μ_t is the instantaneous expected return which could vary over time (capturing the time-varying investment opportunities), σ is the instantaneous standard deviation of diffusive returns, and z is a standard Brownian motion. M is a compensated Poisson process with constant arrival intensity λ. Let N be the corresponding standard poisson process, then

$$dM_t = dN_t - \lambda dt.$$

I model the jump size ξ as a binomial random variable taking value $\xi^u > 0$ or $\xi^d < 0$ with probability $1 - q$ and q respectively, conditional on the jump.

Jumps in returns capture the sudden and dramatic market movement following major events. It addresses an important criticism towards models of inattention: people do

[2] Studies on portfolio choice with time-varying expected return include Kim and Omberg (1996), Campbell and Viceira (1999), Liu (2005), Wachter (2002) and Campbell, Chan, and Viceira (2003), who provide analytical solutions to the portfolio choice (and in some cases, consumption) problem under various settings; Brennan, Schwartz, and Lagnado (1997), Lynch and Balduzzi (2000) and Lynch and Tan (2003), who study more realistic environments using numerical methods; Kandel and Stambaugh (1996), Brennan (1998), Barberis (2000), Brennan and Xia (2001), and Xia (2001), who study the interaction between return predictability and parameter uncertainty for Bayesian investors. Models that study the effect of jumps include Merton (1971), Liu, Longstaff, and Pan (2003) and Das and Uppal (2004). These models do not consider information processing costs.

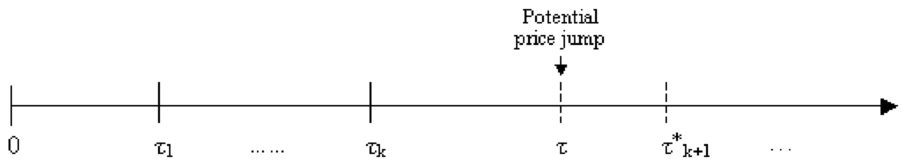

Figure 1: Decision Timeline of an Inattentive Investor.

not remain inattentive to major events. Among the most noticeable examples are the October 1987 market crash and the September 11, 2001 terrorist attacks, which captured everybody's attention immediately. To study the impact of such events, I model agents to be inattentive to small market movements, but respond immediately when a jump arrives.

Investors have CRRA preferences with a power coefficient $\gamma > 1$. Their objective is to maximize the expected utility over an infinite consumption stream. Investors are endowed with initial wealth W_0, and their only source of income is the assets they invest in. I rule out the possibility of borrowing from the checking account, but allow for short positions in stocks and bonds.

I assume that when an agent acquires and processes new information, she pays a lump-sum cost. With such costs (no matter how small), it is no longer optimal for an investor to pay attention to new information all the time. Instead, the optimal strategy is to make "long-term" plans: agents choose their planning horizons, and decide how to invest and consume during each planning period. These plans take into account the time variation in investment opportunities. They can be interrupted by big price jumps, at which times agents immediately make new plans.

The timeline in Figure 1 illustrates the idea. The ex post planning dates are denoted by $\tau_1, \cdots, \tau_k$. Suppose an agent is deliberating on a new plan at time τ_k. Based on all the information available at τ_k, she first chooses the planning horizon (time between τ_k and the next planning date, τ^*_{k+1}). Then she will decide on how to consume and invest between τ_k and τ^*_{k+1}. In order to pay for consumption without making intermediate transactions (which incurs additional information costs), she will need to make a sufficient withdrawal from the portfolio at τ_k, and put it in her checking account. If stock prices do not jump between τ_k and τ^*_{k+1}, the agent will follow through the original plan, and τ^*_{k+1} will indeed be the next planning date ex post, i.e., $\tau_{k+1} = \tau^*_{k+1}$. However, if a jump arrives at τ (before τ^*_{k+1}), she will immediately stop and replan, making τ the ex post planning date: $\tau_{k+1} = \tau$. The key feature is that an agent remains inattentive to new information between any two consecutive ex post planning dates.

Formally, an agent solves the following problem:

$$U(c) = \max_{\{c_t\},\left\{\tau^*_{k+1}, Y_{\tau_k}, \alpha_{\tau_k}\right\}} E\left[\int_0^\infty e^{-\rho t} u(c_t)\, dt\right],$$

where $\{c_t\}$ is the consumption stream; $\{\tau^*_{k+1}\}$ are the ex-ante planning dates; $\{Y_{\tau_k}\}$ are the withdrawals at the beginning of each planning period; $\{\alpha_{\tau_k}\}$ are the fraction of wealth invested into stocks after the withdrawals.

Since the investor cannot borrow from her checking account, she faces the following budget constraint within each planning period:

$$Y_{\tau_k} \geq \int_{\tau_k}^{\tau^*_{k+1}} c_s\, ds. \tag{2}$$

The problem is a recursive one. Knowing how an agent behaves in one planning period is sufficient to characterize the solution of the entire problem. Let's denote the planning horizon by $d = \tau^*_{k+1} - \tau_k$ ($d \geq 0$), and reset the starting point of time to τ_k, the beginning of the planning period under consideration. τ is now the arrival time of the first price jump since τ_k. I also drop the time subscripts for withdrawals (Y) and portfolio weights (α).

The possibility of early plan termination makes it necessary to keep track of the investor's wealth. However, the subtle point about an inattentive investor is that she does not keep track of her financial wealth all the time. Instead, I define a new wealth process, $\{\tilde{W}_t\}$, which is the wealth from the beginning of the current planning period minus the consumption expenditure since then:

$$\tilde{W}_t \triangleq W_0 - \int_0^t c_s\, ds, \qquad t \in [0, d \wedge \tau], \tag{3}$$

where $d \wedge \tau$ is the next ex post planning date. I will call $\tilde{W}_t$ the "book value of wealth", which does not reflect changes in wealth due to changes in the market value of assets. At any time t before the investor plans again, one can compute the total wealth, or the "market value of wealth", by adding ΔW_t, the net change in the value of the portfolio, back to the book value of wealth:

$$W_t = \tilde{W}_t + \Delta W_t.$$

By the definition of τ, there are no jumps between 0 and t^- for $t \leq d \wedge \tau$. So,

$$\Delta W_{t^-} = (W_0 - Y)\left[e^{r_f t} + \alpha(R_{t^-} - e^{r_f t}) - 1\right], \tag{4}$$

where R_{t^-} is the gross return of one share of stock between 0 and t, excluding the potential price jump at t. An application of Ito's Formula to (1) gives

$$R_{t^-} \triangleq \frac{P_{t^-}}{P_0} = e^{\int_0^t \mu_s ds - \frac{1}{2}\sigma^2 t + \sigma(z_t - z_0)}. \tag{5}$$

The value of ΔW_t will depend on whether a jump occurs at t. Define an indicator function $1_{\{\tau=t\}}$, then

$$\Delta W_t = (W_0 - Y)\left[e^{r_f t} + \alpha(R_t(1 + 1_{\{\tau=t\}}\xi) - e^{r_f t}) - 1\right]. \tag{6}$$

I first model information costs as proportional to total wealth, occurring in lump-sum at the beginning of each planning period. Thus, the investor's total wealth (net of information processing costs) for the next planning period is:

$$W_{d\wedge\tau} = (\tilde{W}_{d\wedge\tau} + \Delta W_{d\wedge\tau})(1 - \kappa). \tag{7}$$

Keeping information costs proportional to wealth is technically convenient as it makes the optimal planning horizon state-independent. Besides, this assumption captures the higher value of time for wealthy agents and the increasing complexity of the financial decisions they face. In Section 4, I relax this assumption and consider the effects of fixed information costs.

Suppose an investor has already chosen the planning horizon d and withdrawal amount Y. She chooses the consumption plan that maximizes the utility over the planning period:

$$H(d, Y) \triangleq \max_{\{c_t\}} E_0\left[\int_0^{d\wedge\tau} e^{-\rho s} u(c_s)\, ds\right],$$

subject to the budget constraint

$$Y \geq \int_0^d c_t\, dt.$$

τ is an exponentially distributed random variable, denoting the arrival of price jumps.

Using the first order condition and the budget constraint at equality (it is not optimal to withdraw more than one needs to consume), we can solve for the optimal consumption c_t:

$$c_t = \frac{\frac{\rho+\lambda}{\gamma} Y}{1 - e^{-\frac{\rho+\lambda}{\gamma} d}} e^{-\frac{\rho+\lambda}{\gamma} t}, \qquad t \in [0, d\wedge\tau). \tag{8}$$

This equation says that, given the planning horizon and initial checking account balance, the consumption path will be fully deterministic until the next planning date. Consumption level is declining over time at rate $(\rho + \lambda)/\gamma$. The more impatient the agent, the more she will sacrifice her future consumption to boost today's consumption. On the other hand, lower elasticity of intertemporal substitution makes her prefer a smoother consumption path, hence a lower rate of decline. Moreover, the effect of jumps on consumption is as if it raises the time preference parameter.

The deterministic consumption path within each planning period is an attractive feature. It suggests that, with sufficiently long planning horizons, the model can generate low covariation between stock returns and consumption growth. Next, I study how the planning horizon is determined.

3 Optimal Planning Horizon

Obviously, the horizon depends on the level of information processing costs. It is also affected by the potential arrival of major events, and the time-varying investment opportunities. In order to pin down the effects of these different sources, I first consider an environment with constant expected return and price jumps, then a second one with time-varying expected return and no jumps.

3.1 The Impact of Major Events

Let the expected return on the stock be fixed at μ. There are two state variables for an inattentive investor: the "book value of wealth" and time. Let $V\left(\tilde{W}_t, t\right)$ be the indirect utility function for the investor. The law of motion for $\tilde{W}_t$ is:

$$d\tilde{W}_t = -c_t dt, \qquad \tilde{W}_0 = W_0. \tag{9}$$

The Hamilton-Jacobi-Bellman (HJB) equation for the indirect utility function V is:

$$\begin{aligned} \rho V\left(\tilde{W}_t, t\right) &= \max_{c_t, Y, \alpha, d} \Big\{ u\left(c_t\right) + \lambda \Big[E_0 \Big[V\Big(\left(\tilde{W}_t + \Delta W_t\right)(1-\kappa), 0\Big)\Big] - V\left(\tilde{W}_t, t\right)\Big] \\ &\quad + V_W\left(\tilde{W}_t, t\right)(-c_t) + V_t\left(\tilde{W}_t, t\right)\Big\}. \end{aligned} \tag{10}$$

Furthermore, V satisfies the value matching condition on the optimal planning date:

$$V\left(\tilde{W}_{d^*}, d^*\right) = E_0 V\Big(\left(\tilde{W}_{d^*-} + \Delta W_{d^*-}\right)(1-\kappa), 0\Big). \tag{11}$$

It says that, if no jump has occurred before d^*, then based on the information the agent has from the beginning of the planning period, she should be indifferent between the continuation value and the value of stopping at the optimally chosen planning date.

The information structure is worth emphasizing. The expectations in the HJB equation and the boundary condition are both at time 0, which implies that the investor is making her decisions only using information available at the beginning of the planning period. Even though $\{c_t\}$ varies over time, the consumption path is fully determined by Y and d (see equation (8)), both chosen at time 0.

With proportional information costs and isoelastic utility function, the value function should be homogeneous in wealth[3]. Thus, I guess that V has the separable form:

$$V\left(\tilde{W}_t, t\right) = A(t) \frac{\tilde{W}_t^{1-\gamma}}{1-\gamma}. \tag{12}$$

Since the planning horizon d^* is endogenous, one side of the boundary is a priori unknown, making the system above a free-boundary problem. However, the homogeneity of the problem implies that the optimal planning horizon will be constant. Therefore, we can transform the problem into a standard boundary-value problem, first solving the system for $A(t)$ by treating α and d as parameters (rewriting $A(t)$ as $A(t;\alpha,d)$), and then looking for α^* and d^* that maximize the value function. The following proposition characterizes the solution.

Proposition 1 *With CRRA utility and stock prices following the jump-diffusion process in (1), the portfolio rule (α^*) and planning horizon (d^*) are minimizers of the function $A(0;\alpha,d)$:*

$$(\alpha^*, d^*) = \arg\min_{(\alpha,d)} A(0;\alpha,d), \tag{13}$$

and the optimal cash holdings (Y) is:

$$Y = \frac{1 - e^{-\frac{\rho+\lambda}{\gamma}d^*}}{\frac{\rho+\lambda}{\gamma}} A(0;\alpha^*,d^*)^{-\frac{1}{\gamma}} W_0. \tag{14}$$

Function $A(t;\alpha,d)$ solves the following ordinary differential equation:

$$(\rho+\lambda)A(t;\alpha,d) = Q(t)A(0;\alpha,d) + \gamma A(t;\alpha,d)^{\frac{\gamma-1}{\gamma}} + A'(t;\alpha,d) \tag{15}$$

with boundary condition

$$A(d;\alpha,d) = A(0;\alpha,d)(1-\kappa)^{1-\gamma} E_0\left\{\left[e^{rd} + \alpha\left(R_d - e^{rd}\right)\right]^{1-\gamma}\right\}. \tag{16}$$

The coefficient $Q(t)$ is given in Appendix A.

Proof. See Appendix A. ■

In standard frictionless models, agents usually keep their consumption-wealth ratio constant. This requires continuous adjustment of consumption to any wealth shocks,

[3] Take any trajectory $(\tilde{W}_t, c_t)$ satisfying all the constraints. Then, for any $v > 0$, $(v\tilde{W}_t, vc_t)$ also satisfies the constraints. Since the utility is homogenous of degree $(1-\gamma)$ in c_t, V must be homogenous of degree $(1-\gamma)$ in $\tilde{W}_t$.

big or small. With lump-sum information costs, agents follow deterministic consumption paths within each planning period (see equation (8)), which deviate from the first-best values. Consumption is only adjusted on planning dates to reflect the accumulative changes in wealth since the previous planning period. With a constant planning horizon, $A(0;\alpha^*,d^*)$ will be constant. Hence, according to equation (14), the ratio between the total consumption expenditure for the entire period and the total wealth will also be constant at the beginning of each planning period.

From the HJB equation, one can derive the first order condition for the planning horizon d. At the optimum, agents equalize the net benefit of changing the plan at time d and that of sticking to the old plan a little longer. This trade-off is solely based on information available at time 0. Thus, as long as there are no major events, agents will not change their plan in the middle of a planning period. However, if there is a major market movement, agents will immediately know that their wealth has changed significantly (although the exact amount of the change is still unknown). By setting the jump size sufficiently large, we can safely assume that it will be optimal for them to pay the costs and make a new plan.

Benchmark case: $\lambda = 0$

If λ is 0, stock prices become continuous. This special case is studied by Duffie and Sun (1990). The results in this case serve as a benchmark for the general cases in the following sections.

For an arbitrary planning horizon d, the optimal portfolio weight is:

$$\alpha(d) = \arg\sup_{\alpha} E\left[\frac{[e^{rd} + \alpha(R_d - e^{rd})]^{1-\gamma}}{1-\gamma}\right]. \tag{17}$$

Thus, the portfolio weight in the case of constant expected returns with no price jumps is the one that maximizes the expected utility over wealth at the end of the planning period. Because no intermediate rebalancing takes place, agents behave like buy-and-hold investors with finite horizon, except that their investment horizons will be endogenous in this model. The Inada condition implies that α must satisfy the following condition:

$$\Pr\left(e^{rd} + \alpha\left(R_d - e^{rd}\right) \geq 0\right) = 1,$$

which guarantees that the net wealth remaining on the next planning date will not fall below zero. This condition restricts the portfolio weight to be between 0 and 1.

Hence, even though there are no explicit short-selling constraints for stocks and bonds, agents will never take on leveraged or short positions. The reason lies in the "inability" of agents to adjust their portfolios within a planning period. When stock prices are continuous, the variance of returns goes to 0 as the time step becomes infinitely small, making the evolution of stock price deterministic. Thus, agents have complete control over their wealth by rebalancing continuously. With information processing costs, agents cannot afford continuous rebalancing anymore, which means their wealth can drop significantly before they trade again. If the portfolio weight is between 0 and 1, the worst-case scenario is to lose all the wealth invested in stocks. This feature is also in the model of Longstaff (2005), where agents face exogenous trading "blackout" periods due to market illiquidity. Liu, Longstaff, and Pan (2003) find similar mechanism at work when stock prices take jumps.

A plot of the portfolio weight as a function of planning horizons (not included in the paper) reveals that the effect of infrequent planning on the initial portfolio allocation is tiny. This is because doubling the planning horizon is equivalent to holding the horizon fixed while doubling the riskfree rate and the mean and variance of stock returns. So, stocks look almost as attractive as before, except for the differences in higher order moments. In fact, if we were in a mean-variance world, the portfolio weight should not change at all. The picture will be quite different when we look at asset allocations with respect to total wealth (including the checking account balance). As d increases, agents will hold more of their wealth as cash, leaving less money available for investment. Hence, the weight of stocks relative to total wealth will decrease.

Panel A of Figure 2 plots the planning horizon as a function of the information costs κ in the benchmark case (denoted "DS" for "Duffie and Sun" setting). I consider three different levels of risk aversion, $\gamma = 4, 6, 10$. The information costs parameter κ ranges from 0 to 4 basis points. To put these numbers into perspective, consider a household with net worth of \$1 million. $\kappa = 1\,\mathrm{bp}$ means that it costs the household \$100 each time to plan for consumption and investment. In the plots, the planning horizon always starts at 0 (implying continuous planning), then increases with κ in a concave fashion. As γ increases, the planning horizon increases. This is because more risk-averse agents hold less of their wealth in stocks, which makes the opportunity cost of keeping money in the checking account lower.

The planning horizon increases rapidly for low information costs. Costs as low as half a basis point of total wealth are able to induce adjustment delays of a quarter or more. This is good news for the delay theory. It shows that undetectably low lump-sum information processing costs could potentially generate sufficient delays to explain the low contemporaneous correlation between consumption growth and equity returns. It also suggests that modelling such costs could be an important consideration for dynamic consumption and portfolio choice problems.

Since the interest rate differential between bonds and checking account makes holding

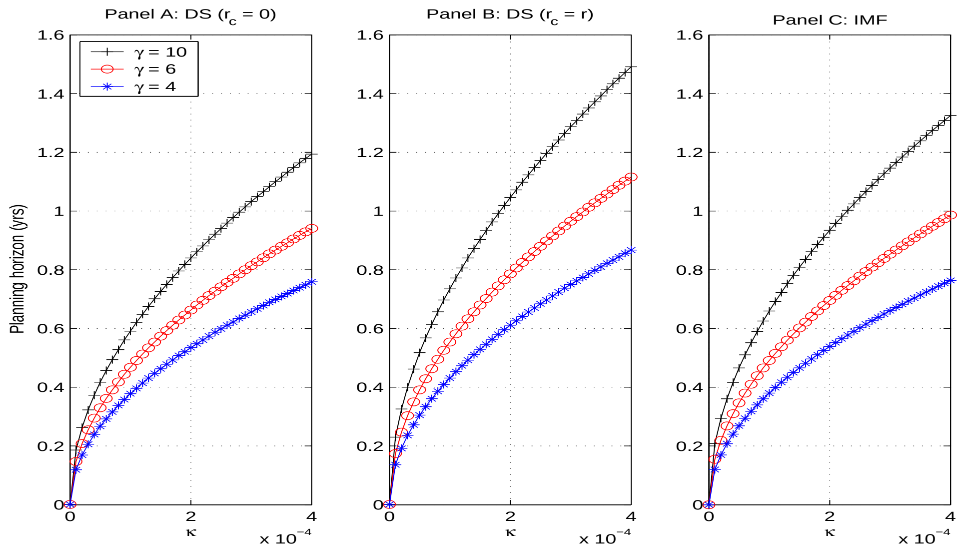

Figure 2: Planning Horizons for the Case of Constant Expected Returns without Price Jumps: Effects of Risk Aversion. The planning horizon is plotted as a function of the proportional information costs (κ) for three different values of the risk aversion coefficient (γ). Parameters are annualized values when applicable: $\mu = 0.08$, $\sigma = 0.15$, $r_f = 0.01$, and $\rho = 0.02$.

cash more costly, eliminating this difference should further increase the planning horizon. Reducing the costs of holding unbalanced portfolios should have a similar effect. To illustrate the magnitude of these effects, I conduct the following two thought experiments. First, I add the same constant riskfree rate to the checking account. Second, I follow Gabaix and Laibson (2001) and introduce an "individualized mutual fund" (IMF), which helps investors rebalance their portfolios continuously. One can show that the optimal portfolio weight for the IMF is always equal to the Merton's myopic demand: $\alpha = (\mu - r) / \gamma\sigma^2$.

Panel B and C of Figure 2 plot the results from these two experiments. The planning horizons are longer in both cases, but the differences are less prominent for the IMF case, especially when γ is small. This feature indicates that the costs of unbalanced portfolios are small in the benchmark case. A highly risk-averse agent holds more wealth in her checking account and invests less in stocks, both in absolute terms and relative to the myopic demand. As a result, they will benefit the most from an interest-paying checking account and the rebalancing services of the individualized mutual fund, which explains why the differences in planning horizons across panels become more significant when γ gets large.

Finally, I study how the level of expected return of the risky asset affects the planning horizon. Figure 3 plots the planning horizon as a function of κ for three different levels of expected returns, $\mu = 0.06$, 0.08, 0.10. The planning horizon drops significantly as μ rises,

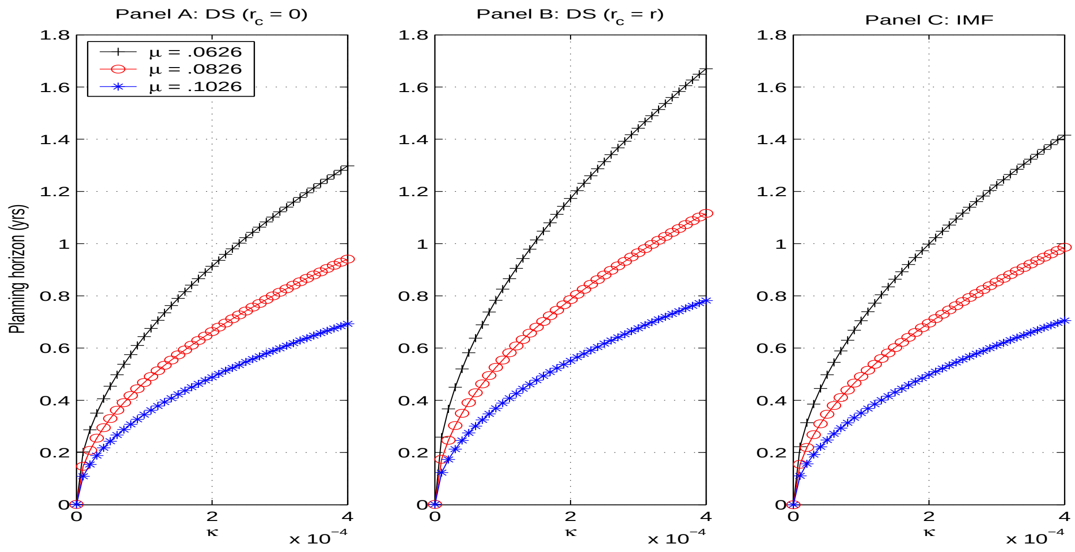

Figure 3: Planning Horizons for the Case of Constant Expected Returns without Price Jumps: Effects of Expected Excess Return. The planning horizon is plotted as a function of the proportional information costs (κ) using three different values for the expected return of the stock (μ). Parameters are annualized values when applicable: $\gamma = 6$, $\sigma = 0.15$, $r_f = 0.01$, and $\rho = 0.02$.

which is mainly because higher expected returns increase the opportunity costs of holding money in the checking account. Planning horizons are again longer when I eliminate the interest rate differential or introduce the IMF. But the differences are small when μ is large, implying that the effect of high expected return on stocks is dominating the other two. The significant impact of expected returns on the planning horizon provides motivation to study the effects of information processing costs in an environment with time-varying investment opportunities. Since portfolio weights will change with the expected returns, we expect the costs of holding unbalanced portfolios to play a much more important role in determining the planning horizon.

General case: $\lambda \neq 0$

For an inattentive investor, the large instantaneous returns due to price jumps is almost the same as the accumulation of small returns from a purely continuous price process. What is different is that the possibility of price jumps makes her investment horizon uncertain. As a result, the optimal strategy will still be a buy-and-hold strategy, but with a random investment horizon, which is jointly determined by the planning horizon the agent chooses, and the arrival time of price jumps. The planning horizon is again determined by the trade-off between information costs and the two types of opportunity costs of infrequent planning: checking account costs and unbalanced portfolio costs.

Table 1: DJIA 30 Single Day Percentage Changes

Upward Jumps	$> 5\%$	$> 10\%$	$> 15\%$	$> 20\%$	$> 30\%$
10/01/1928 to 08/20/2004	52	6	1	0	0
	(0.68)	(0.08)	(0.01)	(0.00)	(0.00)
01/01/1941 to 08/20/2004	5	1	0	0	0
	(0.08)	(0.02)	(0.00)	(0.00)	(0.00)

Downward Jumps	$< -5\%$	$< -10\%$	$< -15\%$	$< -20\%$	$< -30\%$
10/01/1928 to 08/20/2004	61	4	1	1	0
	(0.80)	(0.05)	(0.01)	(0.01)	(0.00)
01/01/1941 to 08/20/2004	11	1	1	1	0
	(0.17)	(0.02)	(0.02)	(0.02)	(0.00)

Note: Numbers without brackets are counts of single-day DJIA 30 index movement exceeding specified sizes. Numbers in brackets are the corresponding frequency estimates (number of times per year). Data source: http://www.djindexes.com.

Choosing the relevant range for the parameters of the jump component is tricky. Daily returns for the CRSP value-weighted index are not available until 1962. It is difficult to estimate the frequency of those rare major market movements using such a short sample. In Table 1, I summarize the frequencies of major single-day percentage changes of the Dow Jones Industrial Average 30 Index since October, 1928 (when the DJIA index first expanded to 30 stocks). The DJIA index is more volatile than the aggregate stock market, yet big single-day changes are still quite rare, especially when the volatile period of the 1930s is excluded. Based on these statistics, I set the relevant range of the jump size to $[-20\%, 20\%]$, and consider the arrival rate of the Poisson process to be between 0 and 0.2 (once every 5 years).

For an investor with one unit of wealth, Figure 4 plots her indirect utility for different choices of portfolio weights (α) and planning horizons (d). The most conspicuous feature of the plot is the flatness of the indirect utility function in the direction of planning horizons (see also the contour plot), which is in sharp contrast with the concavity of the graph in the direction of portfolio weights. An investor can choose a significantly longer planning horizon without suffering much loss in utility. Thus, when an inattentive investor is choosing among "near-rational" plans, it is much more important to get the portfolio weight right than to pick the right planning horizon.

Figure 5 plots the portfolio weight as a function of the jump size. To see the difference between downward and upward jumps, I consider one-sided jumps. I use two different levels of jump frequency: $\lambda = 0.1$ and $\lambda = 0.2$, which correspond to one jump every 10 and 5 years respectively. With downward jumps, the portfolio weight drops as the jump size increases. When jumps are in the upward direction, the portfolio weight first increases

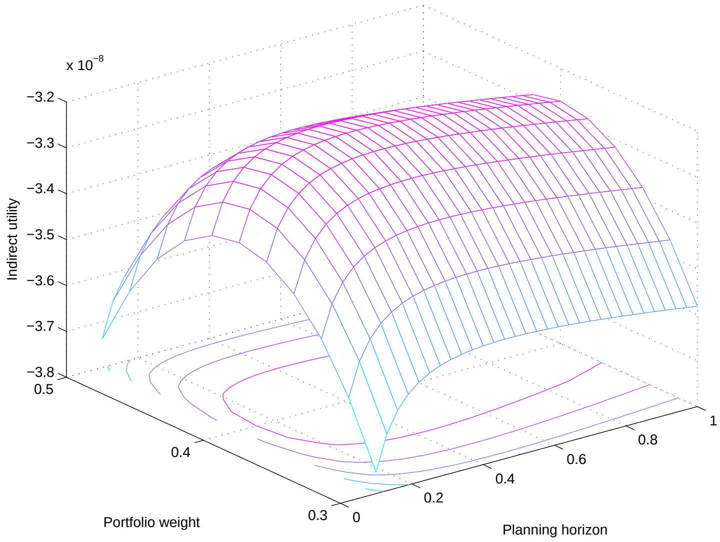

Figure 4: Indirect Utility Function for the Case of Constant Expected Returns with Price Jumps. This graph plots the indirect utility as a function of the planning horizon d (in years) and portfolio weight α. Parameters are annualized values when applicable: $r_f = 0.01$, $\mu = 0.08$, $\sigma = 0.15$, $\xi^u = 10\%$, $\xi^d = -10\%$, $q = 0.5$, $\lambda = 0.05$, $\kappa = 0.0001$, $\gamma = 6$ and $\rho = 0.02$.

slightly, reaching its maximum when the jump size is about 3%, then drops. Moreover, the portfolio weight drop faster in the downward direction.

The eventual decrease of the portfolio weight in both directions is due to the "variance effect". Equation (1) corrects the drift term through the compensated Poisson process so that the expected return will be unaffected by changes in the jump size. But an increase of the jump size in either direction makes return more volatile, which makes investors hold less wealth in stocks. In addition, Liu, Longstaff, and Pan (2003) point out the "skewness effect": upward (downward) jumps make the return distribution more positively (negatively) skewed, a feature that agents with CRRA utility prefer (dislike). This skewness effect causes the decline of the weight to be slower in the upward direction. The maximum portfolio weight of an inattentive investor occurs when the jump size is positive, as opposed to 0 in the model of Liu, Longstaff, and Pan (2003) (Figure 1)). Since inattentive agents do not trade continuously, the gross return over a period of time is already positively skewed even when the jump size is zero. The skewness rises rapidly when the jump size starts to increase. Thus, the initial increase in the portfolio weight is due to the skewness effect dominating the variance effect. As the jump size keeps rising, the variance effect starts to dominate again, leading to the eventual decline in the portfolio weight.

Using one-sided jumps to study the impact of major market movements on the planning horizon is unsatisfactory, because inattentive investors should respond to both good

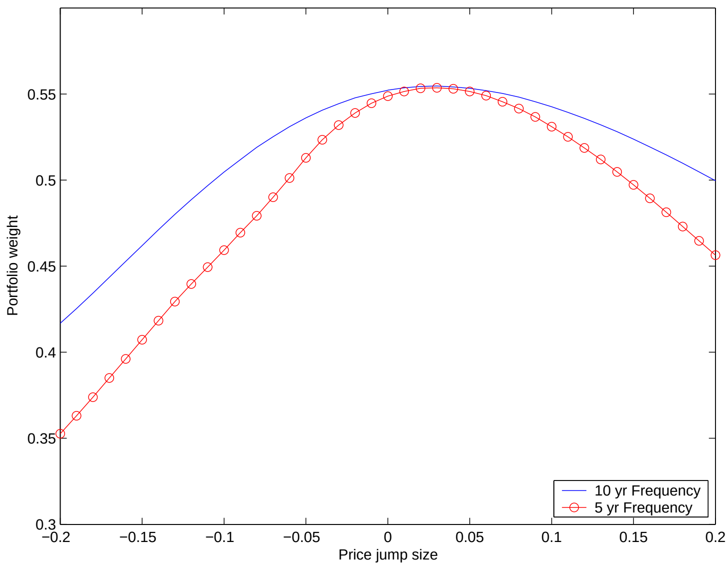

Figure 5: Portfolio Weights for the Case of Constant Expected Returns with Price Jumps. The graph plots the portfolio weight as a function of jump size ξ for the case of one-sided jumps ($q = 0$ or 1). Two different levels of jump frequency are considered: $\lambda = 0.1$ and $\lambda = 0.2$. Other parameter values: $\mu = 0.08$, $\sigma = 0.15$, $r_f = 0.01$, $\kappa = 0.0001$, $\gamma = 6$ and $\rho = 0.02$.

and bad major news. In Figure 6, I consider the case where upward and downward jumps are equally likely to happen ($q = 0.5$). The three panels plot the portfolio weight, planning horizon and cash holdings as functions of jump frequency (λ) for different jump sizes, from a mild 5% to an extreme value of 65%. In Panel A, the portfolio weight starts at the optimal level of the no-jump case, then decreases with λ. The bigger the jump size, the faster the decline. The drop of the portfolio weight is again due to the variance effect. More frequent jumps and bigger jump sizes both drive up the variance of returns.

Panel B shows that the planning horizon increases with jump frequency, although at a slow pace when jump frequency is low. The bigger the jump size, the faster the rise in the planning horizon. Panel C shows that the cash holdings increase with jump frequency as well. The increase in the planning horizon is closely related to the decline in portfolio weight and increase in cash holdings. Because agents hold less wealth in stocks, the costs of unbalanced portfolios are smaller, which makes investors choose longer planning horizons and keep more wealth in their checking account. This reduces the amount of wealth available for investment, hence lowering the costs of unbalanced portfolios further and causing the planning horizon to rise even more.

When the jump size required to trigger investors' immediate responses is small, or when the frequency of such jumps is low, the impact of price jumps on the portfolio weight and planning horizon is almost negligible. However, big jumps matter a lot. To illustrate this point, I consider an extreme case where $\xi = 65\%$. Jumps of such magnitude

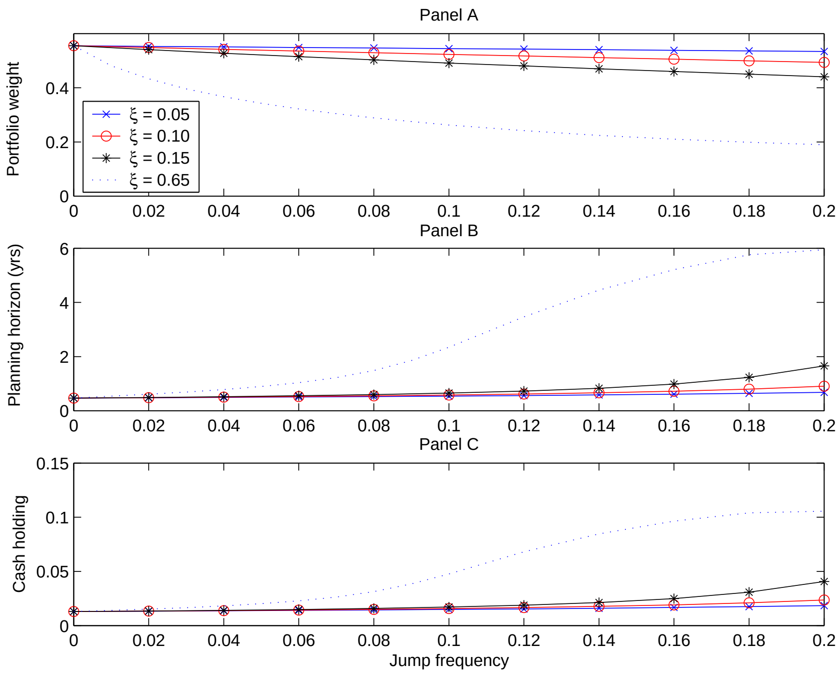

Figure 6: Portfolio Weights (α), Cash Holdings Relative to Total Wealth (Y/W), and Planning Horizons (d) for the Case of Constant Expected Returns with Price Jumps. The three panels graph the controls as functions of the arrival rate λ for the case of two-sided symmetric jumps ($q = 0.5$). Four different jump sizes are considered: $\xi = 0.05$, 0.10, 0.15.and 0.65. Other parameter values: $\mu = 0.08$, $\sigma = 0.15$, $r_f = 0.01$, $\kappa = 0.0001$, $\gamma = 6$ and $\rho = 0.02$.

can cause agents invest significantly less in stocks and remain inattentive for a much longer period (more than 2 years if such jumps happen once every 10 years). A 65% jump that occurs every 10 years will be too extreme for the US stock market, but it is much less so for emerging markets and individual stocks. Many studies have documented poor diversification of individual investors' portfolios. For these investors, the possibility of extreme stock price movements could be an important consideration.

3.2 The Impact of Time-varying Expected Returns

A large body of research[4] has documented time variations in the expected returns of stocks. This feature has profound implications for the dynamic asset allocation problem, making market timing (or strategic asset allocation) a first-order consideration for investors. The comparative statics in the case of constant expected returns demonstrates that changes in the level of expected returns can have big impact on the planning horizon. This prompts

[4]See, for example, Keim and Stambaugh (1986), Campbell (1987), Campbell and Shiller (1988), Fama and French (1988), and more recently Lettau and Ludvigson (2001), Menzly, Santos, and Veronesi (2004), Campbell and Yogo (2005) and Cochrane (2006).

the following question: are the sizable delays generated by undetectably small information costs robust to return predictability?

To focus on the effect of return predictability, I drop the jump component of returns in equation (1) by setting $\lambda = 0$. Define $x_t = \mu_t - r_f$ as the conditional expected excess return (equity premium) of the risky asset. Following several studies of portfolio choice with time-varying expected returns (such as Kim and Omberg (1996), Campbell and Viceira (1999) and Wachter (2002)), I assume that x_t is directly observable to investors, and it follows a mean-reverting process,

$$dx_t = \phi \left(\bar{x} - x_t\right) dt + \sigma_x dw_t. \tag{18}$$

where w is a standard Brownian motion. The instantaneous correlation between the asset-return and equity-premium process is

$$E_t \left[dz_t \, dw_t\right] = \rho_{xp} dt. \tag{19}$$

When $\rho_{xp} \in (-1, 1)$, the market is incomplete.

Within each planning period, the consumption plan is a special case of equation (8), with $\lambda = 0$. The accumulative utility from consumption over the planning period is:

$$H\left(d, Y\right) = \frac{Y^{1-\gamma}}{1-\gamma} \left(1 - e^{-\frac{\rho d}{\gamma}}\right)^{\gamma}. \tag{20}$$

$H\left(d, Y\right)$ is increasing in Y and decreasing in d. In the limit, $H\left(d, Y\right) \to \frac{Y^{1-\gamma}}{1-\gamma}$ as $d \to \infty$, which applies to an agent who plans only once in life.

Next, an agent chooses the optimal amount of withdrawal, portfolio weight and planning horizon in a discrete-time dynamic programming problem. There are two state variables for an agent's consumption and portfolio choice decision: wealth W and the conditional equity premium x. Let $V(W, x)$ be the value function of an agent at the beginning of a planning period. Then, the agent solves the following problem:

$$V\left(W, x\right) = \max_{Y, \alpha, d} \left\{ \frac{Y^{1-\gamma}}{1-\gamma} \left(1 - e^{-\frac{\rho d}{\gamma}}\right)^{\gamma} + e^{-\rho d} E\left[V\left(W', x'\right)\right] \right\}. \tag{21}$$

The first term inside the Bellman equation is the accumulative utility from consumption over the planning period. The second term is the continuation value. W' and x' are the state variables at the beginning of the next planning period.

From equation (6) and (7), we get

$$W' = \left(1 - \kappa\right)\left(W - Y\right)\left[e^{r_f d} + \alpha\left(R_d - e^{r_f d}\right)\right]. \tag{22}$$

The uncertainty in next period's wealth comes entirely from the realized asset returns.

The equity premium $\langle x_t \rangle$ is a Gaussian process with well-known properties. Conditional on information at time 0,

$$x_t = e^{-\phi t} x_0 + \bar{x}\left(1 - e^{-\phi t}\right) + \int_0^t e^{-\phi(t-s)} \sigma_x \, dw_s. \tag{23}$$

This equation shows that given x_0, the conditional equity premium x_t will be normally distributed for any time t. I pick parameters such that x_t is stationary, with the unconditional distribution:

$$x_t \xrightarrow{d} N\left(\bar{x}, \frac{\sigma_x^2}{2\phi}\right). \tag{24}$$

Next, the gross return on the risky asset from 0 to t is a special case of equation (5). Define r_t as the log (continuously compounded) return for one share of stock from 0 to t. Then,

$$r_t = \int_0^t \left(\mu_s - \frac{1}{2}\sigma^2\right) ds + \int_0^t \sigma \, dz_s = \left(r_f - \frac{1}{2}\sigma^2\right) t + \int_0^t x_s ds + \int_0^t \sigma \, dz_s. \tag{25}$$

Because agents do not pay attention to the market all the time, they need to act upon two discretely sampled processes of the asset return and equity premium, with the sampling frequency equal to their planning frequencies. The advantage of setting up the model in continuous-time is that we get a convenient characterization of the joint distribution of the two processes for *arbitrary* planning horizons. The following proposition expresses this joint distribution as a simple VAR system.

Proposition 2 *When stock prices do not take jumps, and the conditional equity premium follows the Ornstein-Uhlenbeck process in (18), agents who choose to remain inattentive between* 0 *and* d *effectively face the following processes of the asset return and equity premium:*

$$r_d = A(d)x_0 + B(d) + C(d)\,\varepsilon_{r,d}, \tag{26}$$

$$x_d = e^{-\phi d} x_0 + \bar{x}\left(1 - e^{-\phi d}\right) + \sigma_x \sqrt{\frac{1 - e^{-2\phi d}}{2\phi}} \varepsilon_{x,d}, \tag{27}$$

where

$$\begin{aligned}
A(d) &= \frac{1-e^{-\phi d}}{\phi}, \\
B(d) &= \left(r_f - \frac{1}{2}\sigma^2 + \bar{x}\right) d - \frac{1-e^{-\phi d}}{\phi}\bar{x}, \\
C(d) &= \sqrt{\frac{\sigma_x^2}{\phi^2}\left[t - \frac{2\left(1-e^{-\phi d}\right)}{\phi} + \frac{1-e^{-2\phi d}}{2\phi}\right] + \sigma^2 d + \frac{2\rho_{xp}\sigma\sigma_x}{\phi}\left[d - \frac{1-e^{-\phi d}}{\phi}\right]}.
\end{aligned}$$

The errors $\varepsilon_{r,d}$ and $\varepsilon_{x,d}$ are jointly normal with mean 0*, standard deviation* 1*, and correlation coefficient*

$$\rho_{xr,d} = \frac{\frac{\sigma_x^2}{2}\left(\frac{1-e^{-\phi d}}{\phi}\right)^2 + \rho_{xp}\sigma\sigma_x\frac{1-e^{-\phi d}}{\phi}}{\sqrt{\frac{\sigma_x^2}{2\phi}\left(1-e^{-2\phi d}\right)}C(d)}. \tag{28}$$

Proof. See Appendix B. ■

Equation (26) shows that both the conditional mean and variance of the log return depend on the length of time the agent chooses to remain inattentive, and so does the correlation between the error terms. When $\rho_{xp} = 0$, the instantaneous shocks dz and dw are independent, but the correlation between the shocks is not zero. This is because when the time step is not infinitesimal, shocks that affect the equity premium also indirectly affect the log return through the equity premium.

Equations (21), (22) and (26), (27) provide a complete description of the problem. Combining the first order condition with respect to Y and the Envelope condition with respect to W, we get

$$V_W(W,x) = \left(1 - e^{-\frac{\rho}{\gamma}d^*}\right)^\gamma Y^{-\gamma}. \tag{29}$$

With proportional information costs and isoelastic preference, the value function remains homogeneous in wealth, but the planning horizon will now depend on the conditional equity premium x. Because of this dependence, we can no longer use the old trick of first treating d as a parameter and then looking for d^* that maximizes the value function. Instead, we need to impose the optimality condition for d together with other constraints. Therefore, I guess the value function as

$$V(W,x) = K(x)\frac{W^{1-\gamma}}{1-\gamma}. \tag{30}$$

Substituting this guessed value function into equation (29), we can compute the withdrawal-wealth ratio:

$$y \triangleq \frac{Y}{W} = K(x)^{-\frac{1}{\gamma}}\left(1 - e^{-\frac{\rho}{\gamma}d^*}\right). \tag{31}$$

Since the planning horizon only depends on x and not W, the withdrawal-to-wealth ratio will also depend on x only. After substituting W' with the budget equation (22), the Bellman equation (21) becomes:

$$K(x) = \min_{\alpha,d} \left\{ y^{1-\gamma} \left(1 - e^{-\frac{\rho}{\gamma}d}\right)^{\gamma} + e^{-\rho d} (1-y)^{1-\gamma} (1-\kappa)^{1-\gamma} E\left[K(x') \left[e^{r_f d} + \alpha \left(e^{r_d} - e^{r_f d}\right)\right]^{1-\gamma}\right] \right\} \tag{32}$$

For a given value of the conditional equity premium, one can further characterize the portfolio weight. Let $\alpha(d)$ denote the optimal portfolio weight corresponding to an arbitrary planning horizon d. Since α only shows up inside the expectation of the new Bellman equation, the weight must be:

$$\begin{aligned} \alpha(d) &= \arg\min_{\alpha} E\left\{K(x') \left[e^{r_f d} + \alpha \left(e^{r_d} - e^{r_f d}\right)\right]^{1-\gamma}\right\} \\ &= \arg\max_{\alpha} E\left\{K(x') \frac{\left[e^{r_f d} + \alpha \left(e^{r_d} - e^{r_f d}\right)\right]^{1-\gamma}}{1-\gamma}\right\}. \end{aligned}$$

Again, the Inada condition implies that the portfolio weight will be between 0 and 1. By scaling the above equation, we get the following more intuitive expression:

$$\begin{aligned} \alpha(d) &= \arg\max_{\alpha} E\left[K(x') u(W')\right] \\ &= \arg\max_{\alpha} \left\{E\left[K(x')\right] E\left[u(W')\right] + cov\left(K(x'), u(W')\right)\right\}. \end{aligned} \tag{33}$$

When we fix the planning horizon, the conditional distribution of x' is no longer affected by any other controls. Thus, the first term in the above equation is proportional to the expected utility over wealth available at the end of the planning period. This implies that part of the portfolio strategy will be a buy-and-hold strategy with investment horizon equal to the planning horizon. However, due to return predictability, this buy-and-hold strategy will not be the same as in the case when expected returns are constant. Intuitively, this is because the mean-reversion of expected returns and the highly negative correlation between shocks to stock prices and expected return make the stock less risky in the long run, thus inducing more demand for stocks. Naturally, this hedging demand will be less important when the planning horizon is short.

The covariance term in (33) adds a new dimension to the portfolio strategy of inattentive investors. Agents not only care about the wealth they will possess for the next planning period, but also the investment opportunities available at that time. Since inattentive investors will never short stocks, the investment opportunity is good only when expected returns are highly positive. In the data, the correlation between shocks for the asset return and the conditional equity premium is close to -1. Thus, even though a series of negative shocks on expected returns can deteriorate the perspective of future investment opportunities, the positive shocks on realized returns at the same time make the agent wealthier. Thus, the stock is a natural hedge for uncertainties in the investment opportunities. This generates an additional positive hedging demand [5].

Finally, if x is constant, the covariance term drops out, and the portfolio weight becomes the same as in the case of constant expected returns.

Calibration

To calibrate the processes of asset return (1) without jumps and the equity premium (18), I follow the approach of Campbell and Viceira (1999). First, I use the maximum likelihood method for restricted VARs (see Hamilton (1994)) to estimate the following VAR(1) model, which uses the dividend-price ratio as a predictor for stock returns.

$$\begin{pmatrix} d_{t+1} - p_{t+1} \\ r_{t+1} - r_f \end{pmatrix} = \begin{pmatrix} \beta_0 \\ \theta_0 \end{pmatrix} + \begin{pmatrix} \beta_1 \\ \theta_1 \end{pmatrix} (d_t - p_t) + \begin{pmatrix} \varepsilon_{1,t+1} \\ \varepsilon_{2,t+1} \end{pmatrix}, \tag{34}$$

where ε's have covariance matrix Ω. Next, I calibrate the parameters of the continuous time model by matching the conditional moments of the two models. For ease of comparison with the estimation results in the literature, I set the time unit to quarterly for both the VAR and the continuous time model.

I use quarterly U.S. financial data for the sample period 1947.Q1 to 2003.Q4. For stock returns, I use quarterly returns (cum-dividend) of the CRSP value-weighted market portfolio (including the NYSE, AMEX and NASDAQ). Riskfree rates are based on the 90 Day T-Bill return series from CRSP. Rates of inflation are constructed from the U.S. Department of Labor CPI series. Following Campbell, Chan, and Viceira (2003), I construct the dividend payout series using the sum of dividend payments over the past year [6]. From 1942 to 1951, the Fed pegged interest rates to help finance WWII. The pegging was dropped in 1951. To see whether this change matters for the calibration, I consider a sub-sample covering 1952.Q2 to 2003.Q4. Ang and Bekaert (2005) show that predictability by the dividend yield is weak in the 1990's, while Campbell and Yogo (2005) show that the test of stock return predictability is sensitive to whether the sample period includes data after 1994. Hence, I conduct a third calibration covering 1947.1 to 1994.4, which largely coincides with the sample period of Campbell and Viceira (1999).

Table 2 provides some summary statistics of the data for the calibrations. In the three

[5] Mathematically, when $\gamma > 1$, the multiplicative term $W^{1-\gamma}/(1-\gamma)$ in the value function $V(W, x)$ is negative. Thus, we should expect $K(x)$ to be decreasing in x: the higher the equity premium, the better off the agent will be. In addition, the negative correlation between the equity premium and realized asset returns implies that when x' drops, W' likely will rise, and more so when α is higher. This induces a positive covariance between $K(x')$ and $u(W')$. Hence, the covariance term is positive and increases with α.

[6] The monthly dividend payments are calculated using the monthly CRSP value-weighted return including dividends and the associated price index series excluding dividends.

Table 2: SUMMARY STATISTICS: CALIBRATION OF RETURN PREDICTABILITY

	Full Sample 1947.1–2003.4	Sub-sample 1 1952.2–2003.4	Sub-sample 2 1947.1–1994.4
$E(r^n_{f,t} - \pi_t) + \frac{1}{2}\sigma^2(r^n_{f,t} - \pi_t)$	0.014	0.018	0.012
$\sigma(r^n_{f,t} - \pi_t)$	0.019	0.016	0.020
$E(r^n_t - r^n_f) + \frac{1}{2}\sigma^2(r^n_t - r^n_f)$	0.072	0.065	0.069
$\sigma(r^n_t - r^n_f)$	0.163	0.167	0.157
Sharpe Ratio	0.44	0.39	0.44
$E(d_t - p_t)$	-3.41	-3.47	-3.27
$\sigma(d_t - p_t)$	0.41	0.38	0.26

NOTE: $r^n_{f,t}$ = log nominal 90-day T-bill return, π_t = log inflation rate, r^n_t = log nominal CRSP value-weighted return, $d_t - p_t$ = log dividend price ratio. All values are annualized when applicable.

samples, the unconditional equity premium ranges from 6.5% to 7.2% per year, and the Sharpe ratio from 0.39 to 0.44. Table 3 reports the calibration results of both the restricted $VAR(1)$ model and its continuous time counterpart. The small magnitude of ϕ reflects strong persistence of the dividend-price ratio. Return predictability as measured by R^2 is indeed the strongest in the 1947.1–1994.4 period. In this sub-sample, the estimated unconditional equity premium (continuously compounded) is 1.17% per quarter, or 4.68% pear year. This number drops nearly by half in the two samples that include the post 1994 period as a result of the weaker predictability of returns using price-dividend ratios in those periods.

Adding additional predicting variables, for example, the consumption-price ratio (Menzly, Santos, and Veronesi (2004)), nominal short rate (Ang and Bekaert (2005)), or *cay* (Lettau and Ludvigson (2001)), can help improve on predictability. However, to keep the calibration straightforward, I simply use the parameters calibrated from the 1947.Q1–1994.Q4 period where predictability is the strongest. While this choice undoubtedly has the danger of overstating the return predictability in the data, it helps to more clearly demonstrate the potential impact of time-varying investment opportunities on the consumption and portfolio decisions of inattentive investors. Another advantage of using this sub-sample is that it renders results comparable to those of Campbell and Viceira (1999).

Planning horizon, consumption and asset allocation

To solve the Bellman equation (32), I apply a multivariate quadrature approximation (see e.g., Judd (1998)) for the two serially correlated shocks of the asset return and equity premium. According to the calibration, the unconditional standard deviation of the equity premium is 1.94% per quarter, or 3.88% per year. Thus, I pick $[-7\%, 16\%]$ (roughly three

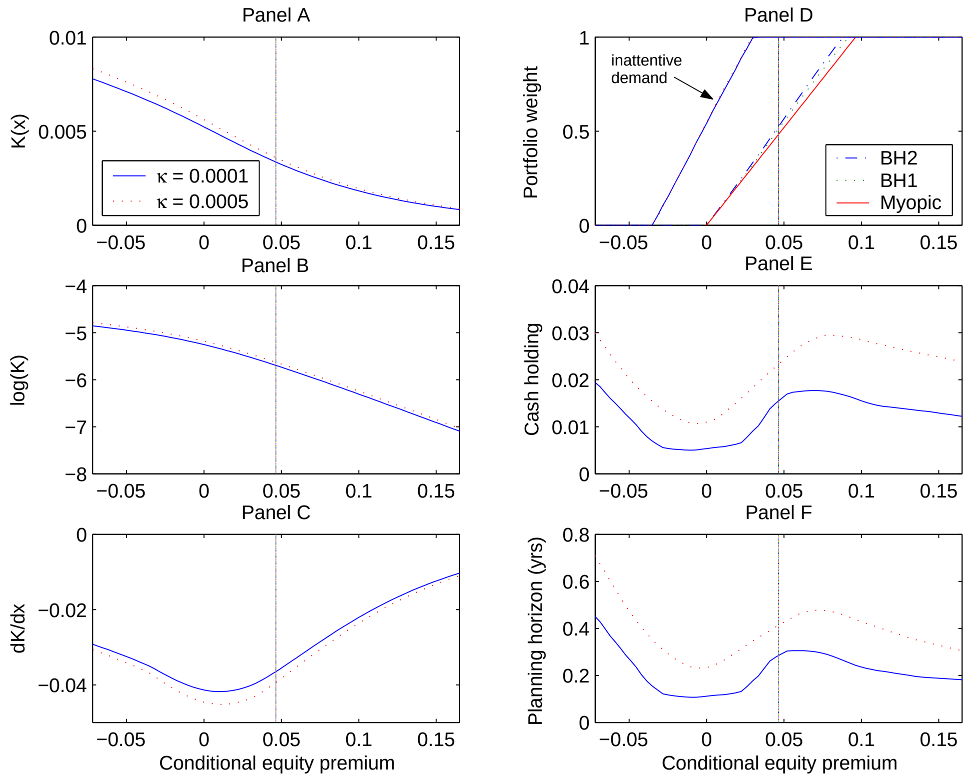

Figure 7: Value Function and Controls for the Case of Time-Varying Expected Returns. Panel A-F graph $K(x)$, $\log(K(x))$, dK/dx, portfolio weights, cash holding relative to total wealth and planning horizon as functions of the conditional equity premium. Two levels of information processing costs are considered: $\kappa = 0.0001$ and $\kappa = 0.0005$. BH1 and BH2 in Panel D denote the buy-and-hold demand for the two cost levels, respectively. The vertical lines mark the location of the unconditional equity premium. Preference parameters are $\gamma = 4$ and $\rho = 0.02$, while other parameters are calibrated using data from Sub-sample II (Table 3).

standard deviations away from the mean) as the relevant range for the grid for x, with grid points corresponding to the roots of order n Chebyshev polynomials. In each iteration, I use 30 quadrature nodes for each of the two shocks to compute the expectation inside the Bellman equation, and then use grid search to find the optimal controls. The iterations continue until the value function and the policy functions converge. For robustness, I compute the results with two procedures, one using standard grid search, the other using the collocation method.

Figure 7 shows the properties of $K(x)$ and the controls. The dotted line denotes the case where $\kappa = 5$ bp, the solid line for $\kappa = 1$ bp. Panel A shows that, as expected, $K(x)$ is decreasing in x. The more positive K gets, the worse-off the agent, which is why the dotted line stays above the solid one. $K(x)$ declines faster around the region where x is close to 0. In Panel C, the big dip of the first derivative dK/dx clearly demonstrates this feature. Hence, the neighborhood of $x = 0$ is where agents' utility is the most sensitive to the variations in expected returns.

In Panel D, I plot the portfolio weights for an investor facing information processing costs. As benchmarks, I also plot the myopic demand and the buy-and-hold demand for the two levels of costs. Since inattentive investors do not take leveraged or short positions, the weights are bounded between 0 and 1. The portfolio weights for the two levels of costs are practically on top of each other, implying the limited effect information costs have on the portfolio choice decision. The buy-and-hold demand is only slightly higher than the myopic demand. On the other hand, the difference between the total inattentive demand and the myopic demand is huge (when the weight is between 0 and 1). It must be that the extra hedging demand from the covariance term in equation (33) is much bigger than the hedging demand in the buy-and-hold part, making the portfolio strategy significantly different from the buy-and-hold strategy.

Another important feature in Panel D is the sharp increase of the portfolio weight in the region where the equity premium is around zero. Portfolio weights become the most sensitive to the conditional equity premium in this region, and mistiming the market could be the most costly here. This sensitivity plays a crucial role in determining the planning horizon, which is plotted in Panel F. The horizons are short compared to the case of constant expected returns. Even for costs of 5 basis points, the planning horizon is shorter than half a year in most cases, which is much smaller than the levels we see in Figure 2 and 3. This sends out a warning to the delay theory: with time-varying investment opportunities, low levels of information processing costs can no longer generate significant delays in consumption adjustment.

The planning horizon is relatively long when expected returns are low, and short when expected returns are high, confirming the comparative statics in the case of constant expected return. But more importantly, the planning horizon stays at a low level when the equity premium is around 0 , the same region where portfolio weights are highly sensitive to the equity premium. To understand this feature, recall the two main costs of infrequent planning: the opportunity costs of leaving money in the checking account, and the costs of holding unbalanced portfolios. The checking account costs explain why people plan more frequently when expected returns are high: agents are willing to pay the costs in order to fully exploit high equity returns.

The unbalanced-portfolio costs are the reason for the short planning horizon when the equity premium is near 0. Because inattentive investors do not rebalance within each planning period, the high sensitivity of portfolio weights means that small changes in the equity premium can quickly lead to large deviations of the portfolio weight from the optimal level. Thus, the costs of unbalanced portfolios are high, forcing agents to plan more frequently. The opposite is true outside of the region: the portfolio weights simply stay at 1 or 0 when expected returns are either high or low. For local changes in

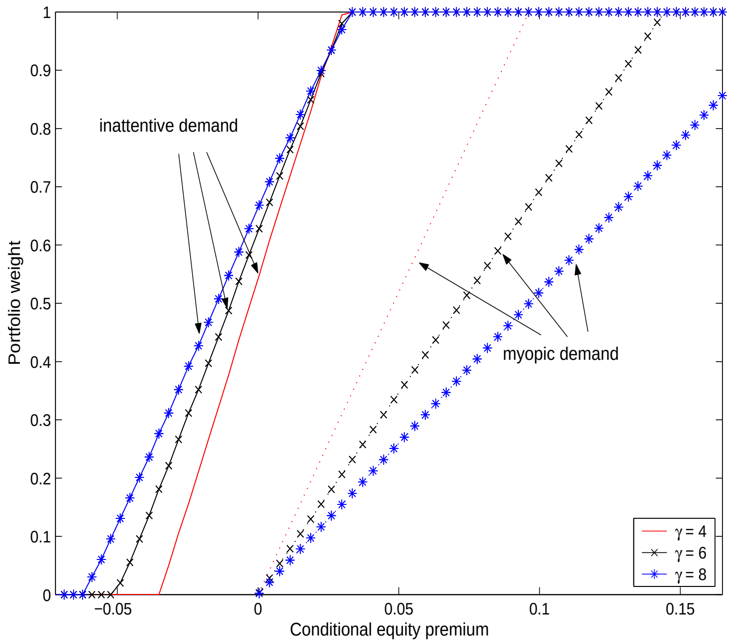

Figure 8: Portfolio Weights for the Case of Time-Varying Expected Returns. Both the myopic demand and total demand for inattentive agents are plotted for three risk aversion parameters: $\gamma = 4$, 6 and 8. $\rho = 0.02$, $k = 0.0001$. Other parameters are from the Subsample II calibration (Table 3).

x, rebalancing is unnecessary. Hence, the costs of holding unbalanced portfolios become relatively unimportant again.

Finally, Panel E plots the ratio of withdrawal (cash holding) to wealth. The pattern closely follows that of the planning horizon. It simply says that when agents plan less frequently, they need to hold a bigger checking account balance to pay for consumption.

Figure 8 provides additional information about the portfolio strategy. I plot the portfolio weights for inattentive investors with different levels of risk aversion. Interestingly, in the region where the portfolio weights are strictly between 0 and 1, more risk-averse agents actually hold more stocks. The reason is these agents have higher hedging demand in this region. If portfolio weights are not bounded from above, we will be able to see less risk-averse agents eventually holding more stocks as the equity premium keeps increasing (which is already happening when portfolio weights are close to 1).

Finally, to see what types of consumption and investment strategies this model will prescribe for the real data, I plot in Figure 9 the planning horizons, portfolio weights and consumption-wealth ratio corresponding to the series of conditional equity premium predicted by historical dividend-price ratio data from 1947.1 to 1994.4. In Panel B, the pattern of portfolio weights resembles the results in Campbell and Viceira (1999): the weights are high in the period before 1960 and from the mid 70s to the mid 80s, but low in the 60s and 90s. The obvious difference from Campbell and Viceira (1999) is that

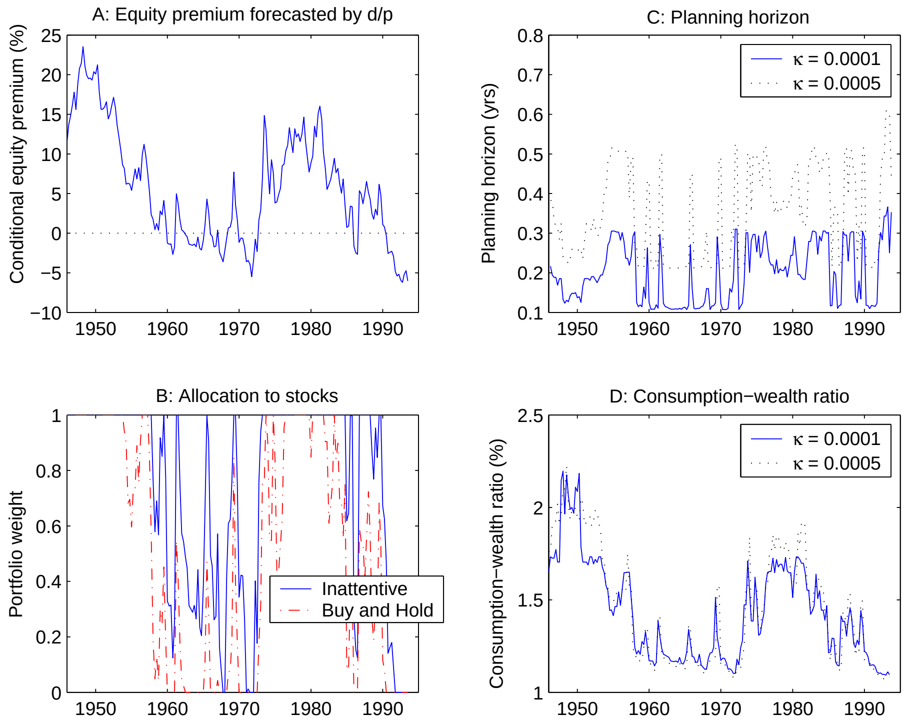

Figure 9: Planning Horizon, Portfolio Weight and Consumption-Wealth Ratio for the Series of Equity Premium Predicted by Historical Data. Panel A graphs the time series of conditional equity premium predicted by historical dividend-price ratio (based on the VAR). Panel B-D graph the corresponding portfolio weights, planning horizons and consumption-wealth ratios. $\gamma = 4$, $\rho = 0.02$, $k = 0.0001$, and other parameters are from the Sub-sample II calibration (Table 3).

portfolio weights are between 0 and 1 for inattentive investors. According to Panel A, the 1960s would be a period of low expected returns (x near 0). This is a region where the portfolio weight sensitivity is high, which results in the big fluctuations of the portfolio weights for the same period.

Panel C shows again that the planning horizons are short when expected returns are time varying. Even for a relatively high level of information processing costs ($\kappa = 5$ bp), the planning horizons are still less than two quarters most of the time. The graph shows substantial variations in the planning horizons across time, with particularly short horizons in the 60s and early 70s and relatively long ones from the mid 70s to the 80s. According to the delay theory, such variations provide a source of time variation for the covariance between stock returns and consumption growth. Interestingly, Duffee (2005) estimates the conditional correlation to be relatively high from the 60s to the mid 70s while low between the mid 70s and the mid 90s, which appears to be consistent with the prediction of the delay theory.

4 Fixed Information Costs and Market Participation

In previous sections I assume that the information processing costs are proportional to the wealth of an inattentive investor. However, under some circumstances, such costs do not increase with wealth. For example, it is quite common in practice for wealthy households to hire financial advisors. Such advisory fees are like fixed costs. When individual investors acquire knowledge about financial planning through books or training courses, the expenses are like fixed costs as well. In this section, I study the impact of fixed information costs on the planning horizons and portfolio allocations across households.

Vissing-Jorgensen (2002) uses per-period stock market participation costs to explain why a lot of households do not hold stocks despite the apparent high returns on stocks. Her interpretation of the participation costs is similar to the information processing costs I consider: it is also about the time spent following the market and determining the optimal trades. However, the per-period participation costs interpretation ignores an important alternative way of investing. Besides the choices of holding a constantly rebalanced portfolio and staying out of the market completely, inattentive investors have an additional choice: following "near-rational" investment strategies.

Fixed information costs make the planning horizons wealth dependent. Non-participation will arise, the trivial case being when households do not have enough wealth for a one-time fixed costs payment. Those households with more wealth than the fixed costs will have an option to participate. Whether they exercise this option or not depends on whether the benefit of participation exceeds the costs.

Let V be the value function of any agent. In addition, define V^P to be the value function for those who already chose to participate this period, and V^N the value function of non-participants. A cut-off wealth level $\underline{W}$ will characterize the optimal strategy: only those agents with wealth exceeding $\underline{W}$ will participate in the stock market. Because there are no other sources of income, the wealth of non-participants is always decreasing over time. Thus, they will never be able to participate in future periods. Participants can also leave the financial market if they get hit by a sufficiently negative wealth shock. The Bellman equation for a participant is:

$$V^P(W) = \max_{Y,\alpha,d} \left\{ \frac{Y^{1-\gamma}}{1-\gamma} \left(1 - e^{-\frac{\rho d}{\gamma}}\right)^{\gamma} + e^{-\rho d} E\left[V(W')\right] \right\} \tag{35a}$$

$$s.t. \quad W' = (1-\kappa)(W - L - Y)\left[e^{rd} + \alpha\left(R_d - e^{rd}\right)\right], \tag{35b}$$

where R_d is the gross return given by equation (5), L is the fixed information costs (in units of consumption goods). The costs occur on each planning date, *before* the agent makes withdrawal Y. In addition, I assume that agents do not have to pay any costs if they decide to keep all their wealth as cash and never plan again (a one-time "free planning").

A non-participant follows her consumption plan forever,

$$V^N(W) = \lim_{d \to \infty} H(d, W) = \frac{W^{1-\gamma}}{1-\gamma}. \tag{36}$$

The option to stay out of the financial market implies

$$V(W) = \max\left\{V^N(W), V^P(W)\right\}.$$

Clearly, agents with wealth below L can never participate. As wealth rises, there is a $\underline{W}$ such that

$$V^N(\underline{W}) = V^P(\underline{W}),$$

and

$$V(W) = \begin{cases} V^N(W) & \text{if } W < \underline{W} \\ V^P(W) & \text{if } W \geq \underline{W} \end{cases} \tag{37}$$

I solve the system numerically and present the results in Figure 10. In the figure, we see rich agents plan more frequently, invest more in stocks, and keep less of their wealth as cash. On the other hand, poor agents have long planning horizons (could be well over 10 years), hold significantly less wealth in stocks and keep a heavy chunk of wealth as cash. The big differences for stock holdings in the cross section confirm the finding of Peress (2004), who also uses information costs to explain why wealthier households hold more of their wealth in stocks in a three-period general equilibrium model.

For those agents with extremely low wealth ($W < \underline{W}$), it is optimal not to participate in the financial market. Instead they keep all their wealth as cash and never plan again. As W rises, agents start participating in the market, but still plan infrequently (see Panel A), and hold significantly less wealth in stocks than wealthy agents (see Panel B). They are the long-term buy-and-hold investors. As wealth keeps rising, the impact of fixed costs becomes smaller and smaller. In the limit, the planning horizon, portfolio weight and cash holding all converge to the levels when there are no fixed costs.

Finally, the lower bound of wealth for participation is close to the level of fixed costs, which we can see from Panel B: portfolio weights are already positive for agents with low wealth. Agents are willing to give up a big chunk of their wealth to cover the fixed information costs, so that they can earn returns from stocks and bonds for extended periods. This is in contrast with the finding of Vissing-Jorgensen (2002) that, a modest per period participation cost is sufficient to explain the participation decisions of a large part of households.

There are a few reasons for this difference. Although the interpretations for the per-period participation costs in Vissing-Jorgensen (2002) and the information processing

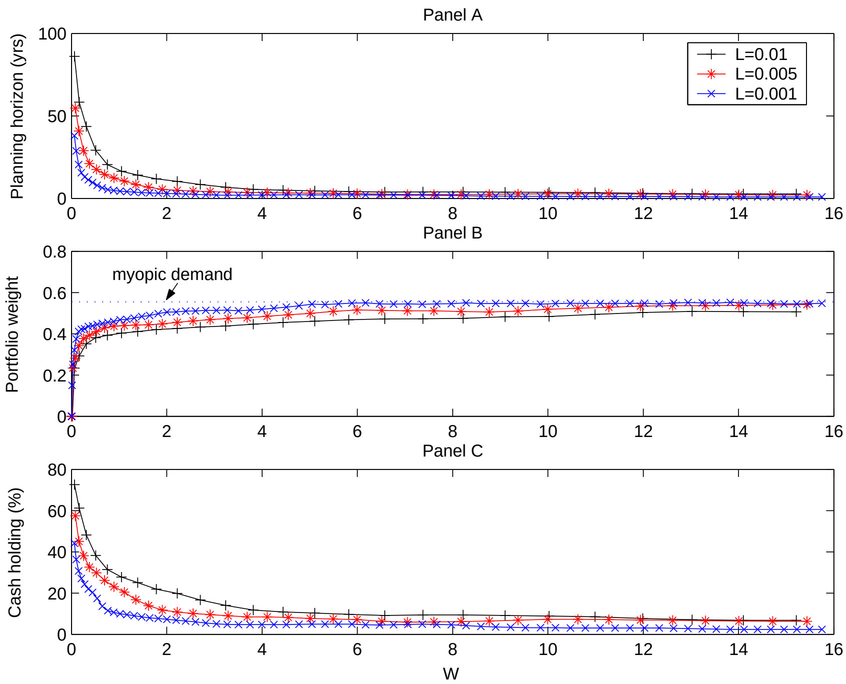

Figure 10: Planning Horizons, Portfolio Weights and Cash Holdings for the Case with Fixed Costs of Planning. Cash holdings are expressed as percentage of total wealth net of fixed costs. Three different levels of fixed cost L are considered. Parameters are annualized values when applicable: $\gamma = 6$, $\rho = 0.02$, $\mu = 0.08$, $\sigma = 0.15$, $r = 0.01$, and $\kappa = 0.0001$.

costs in this paper are similar, the two types of costs work quite differently. I assume that agents only have to pay the fixed costs once each *planning period*, as opposed to once *each period*. If following the market and reoptimizing are costly, why not form a long term investment plan (say 50 years) and then leave it alone? This way, the costs averaged to each period are much smaller. For example, if it costs \$30 per year to maintain an optimal portfolio, the costs should drop by half when investors only follow the market half the time. Several studies have documented the lack of rebalancing activities in individual 401(K) plans. For example, Ameriks and Zeldes (2001) find that 50% of the participants in the TIAA-CREF plans did not rebalance at all during the 10 year period from 1987 to 1996. Fixed information costs could be an explanation for this puzzling phenomenon.

A second reason for the small impact of fixed costs on participation is that I assume that non-participation excludes agents from holding either stocks or bonds, while in reality non-stockholders can still earn interest through their savings accounts. Thus, some of the attraction from participation might disappear if I eliminate the interest rate differential between the riskfree asset and the checking account. Another reason could be the assumption of infinite horizon. Holding a portfolio for 50 years without any adjustment is probably much less appealing to an agent with a life of 100 years, and they might find the value of participation much lower.

5 Concluding Remarks

This paper shows that, when returns are predictable, small information costs are unlikely to generate significant delays in consumption and investment adjustments. Fixed information costs also have limited effects on stock market participation. However, information costs do have interesting implications about portfolio choices. Lump-sum costs lead to infrequent portfolio rebalancing and a strategy mixing the buy-and-hold and hedging demands, while fixed costs generate significant variations for stock holdings across households.

The model is stylized and does not capture the complexity of the consumption and portfolio choice decisions in reality, which is the motivation for modelling information processing costs in the first place. In the case of constant expected returns, the only piece of information that an investor needs is the market value of a single risky asset. In addition, there are easy-to-implement consumption and portfolio rules in such a simple environment. Thus, neither the information nor the planning is likely to cost anything. To better justify the relevance of information processing costs, we will need to consider more sophisticated and realistic settings. Studying the settings of major market movements and time varying expected returns is a step towards that direction.

There are no additional sources of income besides capital gains in this model. Including dividend and labor income means that agents do not have to pay all their consumption from the cash withdrawals, which could make them plan less frequently. It will also be interesting to endogenize information costs and let agents decide how much time and effort they should spend to "produce" information. Huang and Liu (2006) consider the portfolio problem under such a setting.

In the case of return predictability, the planning horizons become significantly shorter. However, the conditional equity premium is not directly observable to investors. In addition, the degree of return predictability in the data is sensitive to the sample period one chooses. How would an investor who is skeptical about return predictability behave when facing information processing costs? Wachter and Warusawitharana (2006) find that even highly skeptical investors will still time the market, although the resulting portfolio weights are indeed less volatile. It would be useful to know how much the uncertainty about predictability will reduce the effects of time-varying expected returns on the planning horizons.

Finally, models of information costs can be useful tools to explain certain phenomena in individual investment behavior. These include how often different types of investors trade, and their preference for stocks that they might be more familiar with.

Appendix

A Proof of Proposition 1

Proof. When α and d are fixed, the value function becomes

$$V(W,t;\alpha,d) = A(t;\alpha,d)\frac{W^{1-\gamma}}{1-\gamma}.$$

From the continuous-time Bellman equation, we can derive the first order condition for c_t:

$$u'(c_t) = V_W\left(\tilde{W}_t, t; \alpha, d\right).$$

Combining this with the expression for the optimal consumption in terms of Y and t as given by equation (8), we get

$$Y = \frac{1-e^{-\frac{\rho+\lambda}{\gamma}d}}{\frac{\rho+\lambda}{\gamma}} A(0;\alpha,d)^{-\frac{1}{\gamma}} \tilde{W}_0. \tag{A.1}$$

With our guess of the value function, the Envelope condition with respect to $\tilde{W}_t$ becomes

$$\begin{aligned}
(\rho+\lambda)A(t;\alpha,d)\tilde{W}_t^{-\gamma} \quad = \quad & (1-\kappa)\lambda A(0;\alpha,d)\left\{E_0\left(1-q\right)\left(\left(\tilde{W}_t+\Delta W_t^u\right)(1-\kappa)\right)^{-\gamma}\right. \\
& \left. +q\left(\left(\tilde{W}_t+\Delta W_t^d\right)(1-\kappa)\right)^{-\gamma}\right\} \\
& +\gamma A(t;\alpha,d)^{\frac{\gamma-1}{\gamma}}\tilde{W}_t^{-\gamma} + A'(t;\alpha,d)\tilde{W}_t^{-\gamma}.
\end{aligned} \tag{A.2}$$

In order for the guess of the separable form of the value function to be correct, we need to be able to cancel $\tilde{W}_t^{-\gamma}$ terms so that $A(t;\alpha,d)$ does not depend on $\tilde{W}_t$.

$$\tilde{W}_t = \tilde{W}_0 - \int_0^t c_s\,ds = \tilde{W}_0 - Y\frac{1-e^{-\frac{\rho+\lambda}{\gamma}t}}{1-e^{-\frac{\rho+\lambda}{\gamma}d}}. \tag{A.3}$$

Define

$$F(t) \triangleq 1 - \frac{1-e^{-\frac{\rho+\lambda}{\gamma}t}}{\frac{\rho+\lambda}{\gamma}} A(0;\alpha,d)^{-\frac{1}{\gamma}}, \tag{A.4}$$

then from (A.1) and (A.3), we get

$$\tilde{W}_0 - Y = \frac{F(d)}{F(t)}\tilde{W}_t. \tag{A.5}$$

Next, using equation (6) and (A.5), we get

$$\tilde{W}_t + \Delta W_t^i = \tilde{W}_t\left[1+\frac{F(d)}{F(t)}\left[e^{rt}+\alpha\left(R_{t^-}(1+\xi^i)-e^{rt}\right)-1\right]\right].$$

Plug this back into equation (A.2) and cancel the $\tilde{W}_t$ terms, and we get the ordinary differential equation in Proposition 1, with

$$\begin{aligned} Q(t) \triangleq\ & \lambda (1-\kappa)^{1-\gamma} E_0 \left\{ (1-q) \left[1 + \frac{F(d)}{F(t)} \left[e^{rt} + \alpha \left(R_t \left(1 + \xi^u \right) - e^{rt} \right) - 1 \right] \right]^{-\gamma} \right. \\ & \left. + q \left[1 + \frac{F(d)}{F(t)} \left[e^{rt} + \alpha \left(R_t \left(1 + \xi^d \right) - e^{rt} \right) - 1 \right] \right]^{-\gamma} \right\}. \end{aligned} \tag{A.6}$$

As for the boundary condition, since we know that it holds for the optimal planning horizon, we can impose it for all d and then choose d^* from the subset of controls that satisfy the boundary condition. Replacing V by the guessed form gives

$$A(d;\alpha,d)\frac{\tilde{W}_d^{1-\gamma}}{1-\gamma} = E_0\left[A(0;\alpha,d)\frac{\left(\left(\tilde{W}_d+\Delta W_d^m\right)(1-\kappa)\right)^{1-\gamma}}{1-\gamma}\right],$$

where

$$\tilde{W}_d + \Delta W_d^m = \tilde{W}_d \left[1 + \frac{F(d)}{F(d)} \left[e^{rd} + \alpha \left(R_d - e^{rd} \right) - 1 \right] \right] = \tilde{W}_d \left[e^{rd} + \alpha \left(R_d - e^{rd} \right) \right].$$

Hence

$$A(d;\alpha,d) = A(0;\alpha,d)(1-\kappa)^{1-\gamma} E_0 \left\{ \left[e^{rd} + \alpha \left(R_d - e^{rd} \right) \right]^{1-\gamma} \right\}.$$

■

Numerical scheme To solve the ODE of (15) with boundary condition (16), I use the following procedure. For a given pair of (α, d), I first guess a value for $A(0;\alpha,d)$, then iterate on the ODE (using the fifth order Runge-Kutta method) until I reach $A(d;\alpha,d)$. Next, I check whether equation (16) holds. If yes, the right $A(t;\alpha,d)$ is found. Otherwise, start again with a different guess for $A(0;\alpha,d)$. Using this procedure, we can find the optimal portfolio rule and planning horizon that minimize $A(0;\alpha,d)$. Since the ordinary differential equation that determines $A(t;\alpha,d)$ does not depend on wealth, d^* will indeed be constant, thus validating our solution approach. This numerical procedure behaves well and locates the solution quickly.

B Proof of Proposition 2

Proof. From equation (23), one can express the conditional equity premium x_t as an AR(1) process,

$$x_t = e^{-\phi t} x_0 + \bar{x}\left(1 - e^{-\phi t}\right) + \sigma_x \sqrt{\frac{1 - e^{-2\phi t}}{2\phi}} \varepsilon_{x,t},$$

where $\varepsilon_{x,t}$ is a Gaussian white noise with unit variance.

The following lemma is useful in characterizing the conditional distribution of r_t.

Lemma 1 *Let $W(t)$ be a Brownian motion, and let $\delta(t)$ and $h(t)$ be two bounded and non-random functions.*

$$X(t) = \int_0^t \delta(u)\, dW(u), \quad Y(t) = \int_0^t h(u) X(u)\, du.$$

Then Y is a Gaussian process with mean function $m_Y(t) = 0$, and covariance function

$$\rho_Y(s,t) = \int_0^{s\wedge t} \delta^2(v) \left(\int_v^s h(y)\,dy\right)\left(\int_v^t h(y)\,dy\right) dv. \tag{B.1}$$

The proof of Lemma 1 is a direct application of stochastic Fubini Theorem (see, e.g., Protter (2003)). Using it, I characterize the conditional distribution of the log asset return r_t as follows.

Lemma 2 *Let $r_t = \left(r_f - \frac{1}{2}\sigma^2\right)t + \int_0^t x_s\,ds + \int_0^t \sigma\,dz_s$, where x follows the Ornstein-Uhlenbeck process in (18), and $E_t[dz\,dw] = \rho_{xp}dt$. Given x_0, r_t will be normally distributed, with mean and variance given by*

$$\begin{aligned}
m_r(t) &\triangleq E_0 r_t = \left(r_f - \frac{1}{2}\sigma^2\right)t + (x_0 - \bar{x})\frac{\left(1-e^{-\phi t}\right)}{\phi} + \bar{x}t, \\
V_r(t) &\triangleq E_0\left[(r_t - E_0 r_t)^2\right] \\
&= \frac{\sigma_x^2}{\phi^2}\left[t - \frac{2\left(1-e^{-\phi t}\right)}{\phi} + \frac{1-e^{-2\phi t}}{2\phi}\right] + \sigma^2 t + \frac{2\rho_{xp}\sigma\sigma_x}{\phi}\left[t - \frac{1-e^{-\phi t}}{\phi}\right].
\end{aligned}$$

Moreover, r_t can be rewritten as

$$r_t = A(t)\,x_0 + B(t) + C(t)\,\varepsilon_{r,t}, \tag{B.2}$$

where $\varepsilon_{r,t}$ is also a Gaussian white noise with unit variance, and

$$\begin{aligned}
A(t) &= \frac{1-e^{-\phi t}}{\phi}, \\
B(t) &= \left(r_f - \frac{1}{2}\sigma^2 + \bar{x}\right)t - \frac{1-e^{-\phi t}}{\phi}\bar{x}, \\
C(t) &= \sqrt{\frac{\sigma_x^2}{\phi^2}\left[t - \frac{2\left(1-e^{-\phi t}\right)}{\phi} + \frac{1-e^{-2\phi t}}{2\phi}\right] + \sigma^2 t + \frac{2\rho_{xp}\sigma\sigma_x}{\phi}\left[t - \frac{1-e^{-\phi t}}{\phi}\right]}.
\end{aligned}$$

Proof. Using equation (23), we can rewrite x_t as

$$x_t = \left[e^{-\phi t}x_0 + \left(1-e^{-\phi t}\right)\bar{x}\right] + e^{-\phi t}X_t,$$

where

$$X_t \triangleq \int_0^t \sigma_x e^{\phi u}\,dw_u = \int_0^t \delta(u)\,dw_u, \qquad \delta(t) \triangleq \sigma_x e^{\phi t}.$$

Thus,

$$\int_0^t x_s\, ds = Y_t + \left[(x_0 - \bar{x}) \frac{1}{\phi} \left(1 - e^{-\phi t}\right) + \bar{x} t \right],$$

with

$$Y_t \triangleq \int_0^t h(u)\, X_u\, du, \qquad h(t) \triangleq e^{-\phi t}.$$

According to Lemma 1, $m_Y(t) = 0$, and

$$\rho_Y(s,t) = \int_0^{s \wedge t} \left(\sigma_x e^{\phi v}\right)^2 \left(\int_v^s e^{-\phi y} dy\right) \left(\int_v^t e^{-\phi y} dy\right) dv.$$

Hence $\int_0^t x_s\, ds$ is conditionally normally distributed. Its mean is

$$E_0 \left(\int_0^t x_s\, ds \right) = (x_0 - \bar{x}) \frac{\left(1 - e^{-\phi t}\right)}{\phi} + \bar{x} t,$$

and its variance is

$$var_0 \left(\int_0^t x_s\, ds \right) = \rho_Y(t,t) = \frac{\sigma_x^2}{\phi^2} \left[t + \frac{1}{2\phi} \left(1 - e^{-2\phi t}\right) - \frac{2}{\phi} \left(1 - e^{-\phi t}\right) \right].$$

The term $\int_0^t \sigma\, dz_s$ is also Gaussian, with mean 0 and variance $\sigma^2 t$. The joint normality of z and w determines that $\int_0^t x_s\, ds$ and $\int_0^t \sigma\, dz_s$ will be jointly normal as well. Thus, we just need to compute the covariance between $\int_0^t x_s\, ds$ and $\int_0^t \sigma\, dz_s$ in order to get the conditional variance of r_t.

$$\begin{aligned} cov_0 \left(\int_0^t \sigma\, dz_s, \int_0^t x_s\, ds \right) &= E \left[\left(\int_0^t \sigma\, dz_s \right) Y_t \right] \\ &= E \left[\int_0^t \sigma\, dz_s \int_0^t \sigma_x e^{\phi s} \left(\int_s^t e^{-\phi u} du \right) dw_s \right] \\ &= \frac{\rho_{xp} \sigma \sigma_x}{\phi} \left[t - \frac{1}{\phi} \left(1 - e^{-\phi t}\right) \right]. \end{aligned}$$

The Stochastic Fubini Theorem is applied again at the second equality. Adding up the terms above leads to the resulting lemma. ■

Using equations (26-27), we can compute the conditional covariance between x_t and r_t:

$$\begin{aligned} cov_0(x_t, r_t) &= E_0 \left[\left(\int_0^t \sigma_x e^{-\phi(t-s)}\, dw_s \right) \left(Y_t + \int_0^t \sigma\, dz_s \right) \right] \\ &= E_0 \left[\left(\int_0^t \sigma_x e^{-\phi(t-s)}\, dw_s \right) \left(\int_0^t \sigma_x e^{\phi s} \left(\int_s^t e^{-\phi u} du \right) dw_s + \int_0^t \sigma\, dz_s \right) \right] \\ &= \frac{\sigma_x^2}{2} \left(\frac{1 - e^{-\phi t}}{\phi} \right)^2 + \rho_{xp} \sigma \sigma_x \frac{1 - e^{-\phi t}}{\phi}. \end{aligned}$$

The correlation between $\varepsilon_{r,t}$ and $\varepsilon_{x,t}$ is the covariance above divided by the standard deviations of r_t and x_t , hence the result in equation (28). This concludes the proof of Proposition 2. ■

References

Ameriks, J., and S. P. Zeldes, 2001, "How Do Household Portfolio Shares Vary With Age?," *Working paper, Columbia University.*

Ang, A., and G. Bekaert, 2005, "Stock Return Predictability: Is It There?," *Working paper, Columbia Business School.*

Barberis, N., 2000, "Investing for the Long Run when Returns are Predictable," *Journal of Finance*, 55, 225 – 264.

Brennan, M. J., 1998, "The Role of Learning in Dynamic Portfolio Decisions," *European Finance Review*, 1, 295–306.

Brennan, M. J., E. S. Schwartz, and R. Lagnado, 1997, "Strategic Asset Allocation," *Journal of Economic Dynamics and Control*, 21, 1377–1403.

Brennan, M. J., and Y. Xia, 2001, "Stock Price Volatility and the Equity Premium," *Journal of Monetary Economics*, 47, 249–283.

Campbell, J. Y., 1987, "Stock Returns and the Term Structure," *Journal of Financial Economics*, 18, 373–399.

Campbell, J. Y., Y. L. Chan, and L. M. Viceira, 2003, "A Multivariate Model of Strategic Asset Allocation," *Journal of Financial Economics*, 67, 41–80.

Campbell, J. Y., and R. J. Shiller, 1988, "Stock Prices, Earnings, and Expected Dividends," *Journal of Finance*, 43, 661–676.

Campbell, J. Y., and L. M. Viceira, 1999, "Consumption and Portfolio Decisions when Expected Returns are Time Varying," *Quarterly Journal of Economics*, 114, 433–95.

Campbell, J. Y., and M. Yogo, 2005, "Efficient Tests of Stock Return Predictability," *Forthcoming, Journal of Financial Economics.*

Cochrane, J. H., 1989, "The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives," *American Economic Review*, 79, 319–337.

Cochrane, J. H., 2006, "The Dog That Did Not Bark: A Defense of Return Predictability," *Working paper, University of Chicago.*

Constantinides, G. M., 1986, "Capital Market Equilibrium with Transaction Costs," *Journal of Political Economy*, 94, 842–862.

Das, S. R., and R. Uppal, 2004, "Systemic Risk and International Portfolio Choice," *Journal of Finance*, LIX, 2809–2834.

Davis, M. H. A., and A. R. Norman, 1990, "Portfolio Selection with Transaction Costs," *Mathematics of Operations Research*, 15, 676–713.

Duffee, G. R., 2005, "Time-Variation in the Covariance Between Stock Returns and Consumption Growth," *Journal of Finance*, LX, 1673–1712.

Duffie, D., and T.-s. Sun, 1990, "Transactions Costs and Portfolio Choice in a Discrete-Continuous-Time Setting," *Journal of Economic Dynamics and Control*, 14, 35–51.

Fama, E. F., and K. R. French, 1988, "Dividend Yields and Expected Stock Returns," *Journal of Financial Economics*, 22, 3–25.

Gabaix, X., and D. Laibson, 2001, "The 6D Bias and the Equity Premium Puzzle," *NBER Macroeconomic Annual*, pp. 257–311.

Grossman, S. J., and G. Laroque, 1990, "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods," *Econometrica*, 58, 25–51.

Hamilton, J. D., 1994, *Time Series Analysis*, Princeton University Press, New Jersey.

Huang, L., and H. Liu, 2006, "Rational Inattention and Portfolio Selection," *forthcoming, Journal of Finance.*

Judd, K. L., 1998, *Numerical Methods in Economics*, MIT Press, Cambridge, MA.

Kandel, S., and R. F. Stambaugh, 1996, "On the Predictability of Stock Returns: An Asset-Allocation Perspective," *Journal of Finance*, 53, 385–424.

Keim, D. B., and R. F. Stambaugh, 1986, "Predicting Returns in the Stock and Bond Markets," *Journal of Financial Economics*, 17, 357–390.

Kim, T. S., and E. Omberg, 1996, "Dynamic Nonmyopic Portfolio Behavior," *The Review of Financial Studies*, 9, 141–161.

Lettau, M., and S. Ludvigson, 2001, "Consumption, aggregate wealth and expected stock returns," *Journal of Finance*, 56, 815 – 849.

Liu, J., 2005, "Portfolio Selection in Stochastic Environments," *Working Paper.*

Liu, J., F. A. Longstaff, and J. Pan, 2003, "Dynamic Asset Allocation with Event Risk," *Journal of Finance*, LVIII, 231–259.

Longstaff, F. A., 2005, "Asset Pricing in Markets with Lliquid Markets," *Working Paper.*

Lynch, A. W., 1996, "Decision Frequency and Synchronization Across Agents: Implications for Aggregate Consumption and the Equity Premium," *Journal of Finance*, 51, 1479–1498.

Lynch, A. W., and P. Balduzzi, 2000, "Predictability and Transaction Costs: The Impact on Rebalancing Rules and Behavior," *Journal of Finance*, LV, 2285–2309.

Lynch, A. W., and S. Tan, 2003, "Explainning the Magnitude of Liquidity Premia: The Roles of Return Predictability, Wealth Shocks and State-Dependent Transaction Costs," *Working Paper.*

Marshall, D. A., and N. G. Parekh, 1999, "Can Costs of Consumption Adjustment Explain Asset Pricing Puzzles?," *Journal of Finance*, LIV, 623–654.

Mehra, R., and E. Prescott, 1985, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15, 145–161.

Menzly, L., J. Santos, and P. Veronesi, 2004, "Understanding Predictability," *Journal of Political Economy*, 112, 1–47.

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

Moscarini, G., 2004, "Limited Information Capacity as a Source of Inertia," *Journal of Economic Dynamics and Control*, 28, 2003–2035.

Nieuwerburgh, S. V., and L. Veldkamp, 2006, "Information Acquisition and Portfolio Under-Diversification," *Working paper, New York University.*

Peng, L., 2005, "Learning with Information Capacity Constraints," *Journal of Financial and Quantitative Analysis*, 40, 307–329.

Peress, J., 2004, "Wealth, Information Acquisition, and Portfolio Choice," *Review of Financial Studies*, 17, 879–914.

Protter, P. E., 2003, *Stochastic Integration and Differential Equations*, Springer-Verlag, New York, NY, 2nd edn.

Reis, R., 2004, "Inattentive Consumers," *Working Paper.*

Shreve, S. E., and H. M. Soner, 1994, "Optimal Investment and Consumption with Transaction Costs," *Annals of Applied Probability*, 4, 609–692.

Sims, C., 2003, "Implications of Rational Inattention," *Journal of Monetary Economics*, 50, 665–690.

Vissing-Jorgensen, A., 2002, "Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures," *Working Paper.*

Wachter, J., and M. Warusawitharana, 2006, "Predictable Returns and Asset Allocation: Should a Skeptical Investor Time the Market?," *Working paper, University of Pennsylvania.*

Wachter, J. A., 2002, "Portfolio and Consumption Decisions under Mean-Reverting Returns: An Exact Solution for Complete Markets," *Journal of Financial and Quantitative Analysis*, 37, 63–91.

Xia, Y., 2001, "Learning About Predictability: The Effects of Parameter Uncertainty on Dynamic Asset Allocation," *Journal of Finance*, 56, 205–246.

Table 3: Estimates of the Restricted VAR and the Diffusion Model

	VAR						
	β_0	β_1	θ_0	θ_1	Ω_{11}	Ω_{12}	Ω_{22}
Full Sample	-0.046	0.988	0.120	0.031	6.00E-03	-6.00E-03	6.52E-03
1947.1–2003.4	(0.043)	(0.013)	(0.045)	(0.013)	(5.63E-04)	(5.75E-04)	(6.12E-04)
	$R^2 = \begin{bmatrix} 0.965 \\ 0.024 \end{bmatrix}$						
	Continuous-time Model						
	σ	ϕ	$\bar{x}$	σ_x	ρ_{xp}	r_f	
	0.082	0.012	0.0055	0.0024	-0.96	0.0034	
	VAR						
	β_0	β_1	θ_0	θ_1	Ω_{11}	Ω_{12}	Ω_{22}
Sub-sample I	-0.072	0.981	0.114	0.029	6.18E-03	-6.37E-03	6.92E-03
1952.2–2003.4	(0.050)	(0.014)	(0.053)	(0.015)	(6.09E-04)	(6.36E-04)	(6.82E-04)
	$R^2 = \begin{bmatrix} 0.958 \\ 0.018 \end{bmatrix}$						
	Continuous-time Model						
	σ	ϕ	$\bar{x}$	σ_x	ρ_{xp}	r_f	
	0.084	0.019	0.0069	0.0023	-0.97	0.0044	
	VAR						
	β_0	β_1	θ_0	θ_1	Ω_{11}	Ω_{12}	Ω_{22}
Sub-sample II	-0.128	0.962	0.240	0.069	5.63E-03	-5.49E-03	5.85E-03
1947.1–1994.4	(0.069)	(0.021)	(0.070)	(0.021)	(5.76E-04)	(5.74E-04)	(5.98E-04)
	$R^2 = \begin{bmatrix} 0.918 \\ 0.052 \end{bmatrix}$						
	Continuous-time Model						
	σ	ϕ	$\bar{x}$	σ_x	ρ_{xp}	r_f	
	0.079	0.039	0.012	0.0054	-0.96	0.0030	

Note: Numbers in brackets are MLE standard errors. The estimates of the restricted VAR(1) model are based on quarterly log excess returns and price-dividend ratios of the CRSP value-weighted market portfolio. All parameters for the continuous-time model are quarterly. The VAR(1) model is:

$$\begin{pmatrix} d_{t+1} - p_{t+1} \\ r_{t+1} - r_f \end{pmatrix} = \begin{pmatrix} \beta_0 \\ \theta_0 \end{pmatrix} + \begin{pmatrix} \beta_1 \\ \theta_1 \end{pmatrix} (d_t - p_t) + \begin{pmatrix} \varepsilon_{1,t+1} \\ \varepsilon_{2,t+1} \end{pmatrix}.$$

The continuous-time model is:

$$\begin{aligned} \frac{dP_t}{P_t} &= \mu_t dt + \sigma dz_t^1 \\ dx_t &= \phi\left(\bar{x} - x_t\right) dt + \sigma_x dz_t^2 \end{aligned}$$

where $E[dz^1 dz^2] = \rho_{xp} dt$.