Optimal Consumption and Portfolio Allocation under Mean-Reverting Returns: An Exact Solution for Complete Markets

Jessica A. Wachter[*]

September 26, 2000

Abstract

This paper solves, in closed form, the optimal portfolio choice problem for an investor with utility over consumption under mean-reverting returns. Previous solutions either require approximations, numerical methods, or the assumption that the investor does not consume over his lifetime. This paper breaks the impasse by assuming that markets are complete. The solution leads to a new understanding of hedging demand and of the behavior of the approximate log-linear solution. The portfolio allocation takes the form of a weighted average and is shown to be analogous to duration for coupon bonds. Through this analogy, the notion of investment horizon is extended to that of an investor who consumes at multiple points in time.

I   Introduction

Consider the portfolio problem of an investor who trades continuously and maximizes expected utility of wealth at some future time T. How does this compare to the problem of an otherwise identical investor who, instead of maximizing utility over wealth at time T, maximizes expected utility of consumption between now and T?

As demonstrated by Merton (1971), when the distribution of asset returns varies in a deterministic way, or when the investor has logarithmic utility, both problems reduce to the single-period portfolio choice problem. The optimal portfolio is mean-variance efficient, as in single-period portfolio choice. The optimal ratio of consumption to wealth declines exponentially with the horizon.

Recent portfolio choice literature focuses on the case where the distribution of asset returns is *not* deterministic. Much of this literature examines the implication of mean reversion in asset returns for portfolio choice, e.g. Barberis (2000), Brennan, Schwartz and Lagnado (1997), Campbell and Viceira (1999), Kim and Omberg (1996). Except for Campbell and Viceira, these papers assume that the investor cares only about wealth at the end of some finite horizon, ignoring the possibility of consumption between now and the end of the horizon. Campbell and Viceira assume an infinitely-lived investor with utility over consumption, and, in order to circumvent the analytical and computational difficulties caused by consumption, derive an approximate solution. Balduzzi and Lynch (1999) assume an investor with utility over consumption and solve using numerical methods, but focus on utility cost implications rather than on portfolio strategies.

Surprisingly, nowhere in this literature is the relationship between portfolio strategies for the investor with and without utility over consumption addressed. The two popular models of investor preferences in the opening paragraph must be related. How does this relation map into portfolio allocations? This question is not only of theoretical interest. A key property emerging from the literature on predictability and portfolio choice is that, for levels of risk aversion exceeding those implied by logarithmic utility, allocation to stocks increases in the investor's horizon. This result has received much attention because it partially redeems the popular, but much-criticized advice of investment professionals. Previous literature pertains only to an investor with a single, fixed horizon. For example, if the investor is not only saving for retirement, but for education, and for a house, the results do not apply.

Intuitively, one would think that the more the value of consumption is weighted towards the present, the more the investor behaves as if he has a short-horizon. One might also think of the investor as saving for, say, three different future events, holding a separate portfolio for each one. The analytical results in this paper make this intuition precise, and at the same time, demonstrate its limits.

The first contribution of this paper is an exact, closed form solution for portfolio weights when the investor has utility over consumption, returns are predictable, and when markets are complete. Despite the intense interest in

this problem documented in the paragraphs above, this solution has never been derived. It is not because the case is unimportant. Generalizing to intermediate consumption allows the portfolio choice problem to be connected to decisions of actual investors in a way that assuming terminal wealth does not. And though assuming complete markets is restrictive, it still allows for a full analysis of the intertemporal problem under mean reversion. This is in contrast to previous exact solutions, such as those requiring unit elasticity of intertemporal substitution.

The solution is different from that for terminal wealth both in form and in magnitude. The allocation under terminal wealth, as well as that under the approximation of Campbell and Viceira (1999), is linear. But the actual allocation derived in this paper involves exponentials of nonlinear functions of the state variables. Moreover, in a calibration exercise, the allocation for terminal wealth is shown to be misleadingly high. The allocation for utility over consumption is in some cases less than half of the allocation for terminal wealth at the same horizon.

The second contribution of this paper is the economic interpretation of the solution and the insight it lends into the problem of multiperiod portfolio choice. The allocation is shown to take the form of a weighted average, where the averaged terms are stock allocations for investors with utility over terminal wealth and the weights depend on the present discounted value of consumption. This formula is shown to be exactly analogous to the duration formula for coupon bonds, which consists of a weighted average of the duration of the underlying zeros.

The analogy to coupon bonds is used to generalize the notion of horizon to the investor with utility over consumption. Besides the central question of investment horizon, the formula also has implications for the sign and the magnitude of hedging demand, the convergence of the solution at long horizons, and discrepancies between the log-linear solution and the actual solution.

The organization of the paper is as follows. Section II defines the optimization problem, both for utility over terminal wealth and over consumption, and lays out the assumptions on asset returns. Section III solves the portfolio choice problem using the martingale method of Cox and Huang (1989), Karatzas, Lehoczky, and Shreve (1987), and Pliska (1986). Section IV discusses the solution and clarifies why the assumption of complete markets is essential.

II   The consumption and portfolio choice problem under mean-reverting returns

This section lays out the assumptions on asset returns and states the two optimization problems of interest, namely that for the investor with utility over consumption, and that for the investor with utility over terminal wealth.

Let w_t denote the standard, one-dimensional Brownian motion. Assume

that the price S of the risky security follows the process

$$\frac{dS_t}{S_t} = \mu_t \, dt + \sigma \, dw_t, \tag{1}$$

and let

$$X_t = \frac{\mu_t - r_t}{\sigma},$$

where r_t is the riskless interest rate. The process X_t determines the price of risk in the economy, or the reward, in terms of expected return, of taking on a unit of risk. Assume that X_t follows an Ornstein-Uhlenbeck process:

$$dX_t = -\lambda_X (X_t - \bar{X}) \, dt - \sigma_X \, dw_t. \tag{2}$$

The volatilities σ and σ_X are assumed to be constant and strictly positive, and λ_X is assumed to be between 0 and 1. Note that the stock price and the state variable (X_t) are perfectly negatively correlated. To isolate the effects of time-variation in expected returns, the risk-free rate is assumed to be constant and equal to $r > 0$, but this assumption can be relaxed. Kim and Omberg (1996) allow imperfect correlation, and thus incomplete markets. Otherwise, however, their set-up is identical. Campbell and Viceira (1999) and Barberis (2000) work in discrete time, but their assumptions on the price process are essentially the same as those of Kim and Omberg. In the empirical applications of Campbell and Viceira and Barberis, X_t is taken to be the dividend-price ratio. The model is meant to capture the fact that the dividend-price ratio is strongly negatively correlated with contemporaneous returns (Barberis finds a correlation of -0.93), but is positively correlated with future returns. These assumptions on stock returns imply that returns are mean-reverting. That is, measured at longer horizons, returns have a lower variance than at shorter horizons. This is proved formally in Appendix E. Intuitively, price increases tend to be followed by lower future returns.

Merton (1971) considers a model that is analytically similar to the above, but he assumes that X_t is perfectly *positively* correlated with S_t, so the interpretation is quite different. Merton also solves for consumption and portfolio choice, but under the assumption that utility is exponential and the time horizon infinite. In line with the more recent literature described in the introduction, in what follows, the agent is assumed to have power utility and a finite horizon. As shown in Section 5, the results can be extended to the infinite horizon case by taking limits.

Two optimization problems are considered. In the first, the investor is assumed to care only about wealth at some finite horizon T. At each time, the investor allocates wealth between the risky asset (a stock) and a riskless bond. It is assumed that there are no transaction costs, and that continuous trading is possible. Let α_t denote the allocation to the risky asset. The investor solves:

$$\sup E \left[\frac{W_T^{1-\gamma}}{1-\gamma} \right] \tag{3}$$

4

$$\text{s.t.} \quad \frac{dW_t}{W_t} = (\alpha_t(\mu_t - r) + r)\,dt + \alpha_t\sigma\,dw_t. \tag{4}$$

In the second optimization problem, the investor cares about consumption between now and time T. At each time, besides allocating wealth between assets, the investor also decides what proportion of wealth to consume. The investor solves:

$$\sup E\left[\int_0^T e^{-\rho t}\frac{c_t^{1-\gamma}}{1-\gamma}\,dt\right] \tag{5}$$

$$\text{s.t.} \quad dW_t = (\alpha_t(\mu_t - r) + r)W_t\,dt + \alpha_t\sigma W_t\,dw_t - c_t\,dt \tag{6}$$
$$W_T \geq 0.$$

The first problem is that considered by Kim and Omberg (1996) and Brennan, Schwartz and Lagnado (1997). The second problem is a continuous-time, finite horizon version of Campbell and Viceira (1999). In both cases, γ is assumed to be greater than 1, i.e. investors have greater risk aversion than that implied by log utility. This assumption insures that a solution to these problems exist, and, as shown by the literature on the equity premium puzzle (see, e.g. Mehra and Prescott, 1985), it is the empirically relevant case.

The precise form of (5) is assumed for notational simplicity. As will be clear in the derivation that follows, introducing a bequest function involves no additional complications.

III An exact solution

The link between (3) and (5) is by no means apparent. In the first problem, the investor makes an allocation decision, subject to a linear budget constraint. This does not imply that the first problem is easy to solve. But it is less complicated than the second, in which the investor has two decisions to make at each time, and the budget constraint is nonlinear.

Previous literature only reinforces the differences between these problems. The papers that assume terminal wealth (e.g. Brennan, Schwartz, and Lagnado, 1997, Kim and Omberg, 1996) do not even mention intermediate consumption, much less hint at how their results might be generalized. Campbell and Viceira (1999) derive an approximate solution for the investor with utility over consumption. They discuss in detail a special case for which their solution is exact, namely the case where the investor has unit elasticity of intertemporal substitution. But this special case is much less rich than the one considered here because the ratio of consumption to wealth is nonstochastic. That is, the investor behaves myopically as far as the consumption decision is concerned.[1].

The assumption of complete markets turns out to be exactly what is needed to make the consumption problem tractable and, at the same time, relate it to

---

[1]Campbell and Viceira assume the utility function of Epstein and Zin (1989), a generalization of power utility that allows the elasticity of intertemporal substitution to be separated from risk aversion.

the terminal wealth problem. This result is most easily seen using the martingale method of Cox and Huang (1989), Karatzas, Lehoczky, and Shreve (1987), and Pliska (1986), and is laid out in Section III.A. Section III.B solves for the investor's wealth, as well as the consumption-wealth ratio; Section III.C uses the results of the previous two to derive the formula for the optimal allocation

A   The martingale method

The martingale method relies on the existence of a state-price density ϕ_t with the convenient property that:

$$E_t[\phi_s S_s] = \phi_t S_t \quad s > t.$$

The process ϕ_t can be interpreted as a system of Arrow-Debreu prices. That is, the value of ϕ_t in each state gives the price per unit probability of a dollar in that state. The price of the asset is given by the sum of its payoffs in each state, multiplied by the price of a dollar in that state, times the probability of the state occurring.

No arbitrage and market completeness imply that ϕ_t exists and is unique. In addition, under technical assumptions on the parameters, ϕ_t can be derived from the price processes (see Harrison and Kreps, 1979; and the textbook treatment of Duffie, 1996). Novikov's condition suffices:

$$E\left(\exp\left\{\frac{1}{2}\int_0^T X_t^2\,dt\right\}\right) < \infty. \tag{7}$$

When (7) applies, ϕ_t is given by:

$$\frac{d\phi_t}{\phi_t} = -r\,dt - X_t\,dw_t.$$

Using the state-price density ϕ, the dynamic optimization problem of Merton (1971) can be recast as a static optimization problem. In particular, budget constraints (4) and (6) are equivalent to the static budget constraints

$$E\left[W_T\phi_T\right] = W_0 \tag{8}$$

and

$$E\left[\int_0^T c_t\phi_t\,dt + W_T\phi_T\right] = W_0 \tag{9}$$

respectively. Equations (8) and (9) express the idea that consumption in different states can be regarded as separate goods. These equations state that the amount the investor allocates to consumption in each state multiplied by the price of consumption in that state must equal his total wealth. Proofs of this well-known result can be found in Cox and Huang (1998) and Karatzas and Shreve (1991).

6

The investor's optimal policies follow from setting the marginal utility of consumption (or terminal wealth) equal to marginal cost, as determined by the static budget constraints above. Strictly speaking, this is true as long as the solutions can first be shown to exist. This is done in Appendix A. For (3),

$$W_T^* = (k\phi_T)^{-\frac{1}{\gamma}}.\tag{10}$$

For (5),

$$c_t^* = (l\phi_t)^{-\frac{1}{\gamma}} e^{-\frac{1}{\gamma}\rho t}\tag{11}$$

and terminal wealth is zero. The constants k and l are Lagrange multipliers determined by substituting the optimal policy into the appropriate static budget constraint.

The aim is now to derive the portfolio policy for the agent with utility over consumption (5), and relate them to the portfolio policy for the agent with utility over terminal wealth, (3).

The portfolio policy is derived from the need for wealth to finance the consumption plan (11). Because markets are complete, any contingent payoff (satisfying certain regularity conditions) can be financed by dynamically trading in the existing assets, in this case a stock and a riskless bond. The first step to derive portfolio policies is to derive the process for the investor's wealth implied by (11). The value of wealth at time t is as follows:

$$W_t = \phi_t^{-1} E_t \left[\int_t^T \phi_s c_s^* \, ds \right].\tag{12}$$

Equation (12) can also be interpreted as the present discounted value of future consumption, where the discounting is accomplished by the state-price density.

Define a new variable

$$Z_t = (l\phi_t)^{-1}.\tag{13}$$

In order to solve the expectation (12), it is convenient to express W_t as a function of the variables X_t and Z_t. Because X_t and Z_t together form a strong Markov process, X_t and Z_t are all the investor needs to know to evaluate moments of Z_s at time t. Therefore define

$$\begin{align}
G(Z_t, X_t, t) &\equiv W_t \tag{14}\\
&= Z_t E \left[\int_t^T Z_s^{\frac{1}{\gamma}-1} e^{-\frac{\rho}{\gamma}s} \, ds \mid X_t, Z_t \right].\tag{15}
\end{align}$$

The second equation follows from substituting (13) and (11) into (12).

Now consider the investor with utility over terminal wealth. From (10), up to a constant, wealth is equal to:

$$F(Z_t, X_t, t; T) = Z_t E_t \left[Z_T^{\frac{1}{\gamma}-1} \mid X_t, Z_t \right].\tag{16}$$

Clearly there is nothing special about T; one could equally imagine an investor with utility over wealth at time $s < T$. For this investor, the wealth process equals

$$F(Z_t, X_t, t; s) = Z_t E_t \left[Z_s^{\frac{1}{\gamma} - 1} \mid X_t, Z_t \right].$$

When the expectation is brought inside the integral sign in (12), W_t becomes a sum of expressions of the form (16), adjusted to take into account the rate of time preference:

$$G(Z_t, X_T, t) = \int_t^T F(Z_t, X_t, t; s) e^{-\frac{\rho}{\gamma} s} \, ds. \tag{17}$$

The terms inside the integral in (17) equal the value of consumption at each point in time. Wealth is like a bond that pays consumption as its coupon; the total value of wealth is simply the sum over all future consumption values.

B Optimal wealth and the consumption-wealth ratio

Determining the precise functional form of F and G means solving the expectation (15). In the case of mean-reverting returns, the expectation can be solved in closed form.

As the investor's wealth is a tradeable asset, it must obey a no-arbitrage condition. Namely, the instantaneous expected return in excess of the riskfree rate must equal the market price of risk times the instantaneous variance. An differential equation analogous to that used to price bonds (and derived in the same manner, see e.g. Cox, Ingersoll, Ross, 1985) appears here. From Ito's lemma, the instantaneous expected return on the investor's wealth equals

$$\mathcal{L}G + G_t + Z_t^{\frac{1}{\gamma}} e^{-\frac{\rho}{\gamma} t}, \tag{18}$$

where $\mathcal{L}G = \frac{1}{2} G_{XX} \sigma_X^2 + \frac{1}{2} G_{ZZ} Z^2 X^2 - G_{XZ} Z X \sigma_X + G_X(-\lambda_X(X - \bar{X})) + G_Z Z(r + X^2)$. The last term in (18) comes from the consumption coupon that is payed each period. No arbitrage requires that

$$\mathcal{L}G + G_t + Z_t^{\frac{1}{\gamma}} e^{-\frac{\rho}{\gamma} t} - rG = (G_Z Z_t X_t - G_X \sigma_X) X_t. \tag{19}$$

G also obeys the boundary condition

$$G(Z_T, X_T, T) = 0.$$

This partial differential equation is solved by first "guessing" a general form for the solution. Equation (17) suggests that G can be written as an integral of functions F. Because $F(Z_t, X_t, t; t) = Z_t^{\frac{1}{\gamma}}$, it is reasonable to guess that G equals $Z_t^{\frac{1}{\gamma}}$ multiplied by a function of X_t. Finally, bond prices under an affine term structure can be expressed as exponentials of the underlying state

8

variables. Following the bond pricing literature, therefore, a reasonable guess for the form of a solution for each "coupon" F is the exponential of a polynomial:

$$G(Z_t, X_t, t) = Z_t^{\frac{1}{\gamma}} e^{-\frac{\rho}{\gamma} t} \int_0^{T-t} \Psi(X_t, \tau) \, d\tau, \qquad (20)$$

where

$$\Psi(X_t, \tau) \equiv \exp \left\{ \frac{1}{\gamma} \left(C(\tau) X_t^2 / 2 + B(\tau) X_t + A(\tau) \right) \right\}.$$

From the relationship between G and F described in (17), it is clear that:

$$F(Z_t, X_t, t; T) = Z_t^{\frac{1}{\gamma}} e^{\frac{\rho}{\gamma}(T-t)} \Psi(X_t, T - t).$$

Substituting (20) back into (19) and matching coefficients on X^2, X, and the constant term leads to a system of three differential equations in C, B, and A. The method for solving these equations is standard, and is discussed in Kim and Omberg (1996), so only a few words will be said here. The equation for C is known as a *Riccatti equation* and can be rewritten as:

$$\int_0^{\tau} \frac{dC}{cC^2 + bC + a} = \tau.$$

where

$$a = \frac{1 - \gamma}{\gamma}$$

$$b = 2 \left(\frac{\gamma - 1}{\gamma} \sigma_X - \lambda_X \right)$$

$$c = \frac{1}{\gamma} \sigma_X^2.$$

The solution for the integral can be found in integration tables. In Appendix B, it is shown that when $\gamma > 1$, $b^2 - 4ac > 0$. Defining

$$d = \sqrt{b^2 - 4ac}.$$

The solution is given by

$$C(\tau) = \frac{1 - \gamma}{\gamma} \frac{2(1 - e^{-d\tau})}{2d - (b + d)(1 - e^{-d\tau})}, \qquad (21)$$

$$B(\tau) = \frac{1 - \gamma}{\gamma} \frac{4\lambda_X \bar{X}(1 - e^{-d\tau/2})^2}{d[2d - (b + d)(1 - e^{-d\tau})]} \qquad (22)$$

The explicit solution for A is more complicated and can be found by integrating a polynomial in B and C:

$$A(\tau) = \int_0^{\tau} \left[\frac{1}{2} a B(\tau')^2 + \frac{1}{2} \sigma_X^2 C(\tau') + \lambda_X \bar{X} B(\tau') + (1 - \gamma) r - \rho \right] d\tau'. \qquad (23)$$

9

As mentioned above, very similar equations arise when studying term structure models (e.g. Cox, Ingersoll, Ross, 1985 and Duffie and Kan, 1996) and in the portfolio choice model of Kim and Omberg (1996)[2]. Appendix B demonstrates that G has the required derivatives for the equation for wealth (14) and the portfolio rule (26) to be valid.

While the investor's wealth depends on the variable Z_t, the ratio of wealth to consumption does not:

$$\frac{W_t}{C_t} = \int_0^{T-t} \exp\left\{\frac{1}{\gamma}\left(C(\tau)X_t^2/2 + B(\tau)X_t + A(\tau)\right)\right\} d\tau. \tag{24}$$

The analysis above allows the consumption wealth ratio to be expressed in closed form, using parameters that are all external to the model. This ratio provides a mechanism to examine how investors trade off between consumption and savings as a function of the state variable. The consumption vs. savings decision is at least as important an aspect of the multiperiod problem as the portfolio choice decision. Assuming utility over terminal wealth captures only the second aspect of the problem. Assuming utility over consumption captures both.

C   Optimal portfolio allocation

In order for the portfolio rule to finance the consumption plan, changes in the portfolio value must correspond one-to-one with changes in the value of future consumption. That is, the diffusion terms must be equal. This consideration determines α, the allocation to the risky asset:

$$\alpha_t G\sigma = G_Z Z_t X_t - G_X \sigma_X. \tag{25}$$

The right-hand side follows from applying Ito's lemma to the function G. The left hand side is the dollar amount invested in the risky asset multiplied by its variance.

Rearranging,

$$\alpha_t = \frac{G_Z Z_t}{G}\frac{X_t}{\sigma} - \frac{G_X}{G}\frac{\sigma_X}{\sigma}. \tag{26}$$

It follows immediately from (20) that $G_Z Z_t/G = 1/\gamma$. Moreover, X_t by definition equals the Sharpe ratio $(\mu_t - r)/\sigma$. Therefore, the first term is the myopic allocation as defined in Merton (1973), namely the allocation that investor would choose if he ignored changes in the investment opportunity set. This can also be seen directly by setting σ_X to zero: when σ_X is zero, the investment opportunity set is constant, and the second term disappears.

The second term, hedging demand is more complicated and interesting. Substituting in for G from (20) leads to the equation

$$\alpha_t = \frac{1}{\gamma}\left(\frac{\mu_t - r}{\sigma^2}\right) - \frac{\sigma_X}{\gamma\sigma}\frac{\int_0^{T-t}\Psi(X_t,\tau)(C(\tau)X_t + B(\tau))\,d\tau}{\int_0^{T-t}\Psi(X_t,\tau)\,d\tau}. \tag{27}$$

---

[2]More recently, Liu (1999) and Schroder and Skiadas (1999) examine conditions under which the portfolio choice problem reduces to solving a system of ordinary differential equations.

The derivation of the investor's allocation under terminal wealth is similar, and easier. The result can also be derived using dynamic programming (Kim and Omberg, 1996). The allocation under terminal wealth equals:

$$\tilde{\alpha}_t = \frac{1}{\gamma}\left(\frac{\mu_t - r}{\sigma^2}\right) - \frac{\sigma_X}{\gamma\sigma}\left(C(T-t)X_t + B(T-t)\right). \tag{28}$$

Along with the equation for the wealth-consumption ratio (24), (27) solves the consumption and portfolio choice problem. These formulas are a novel result and fill an important gap in the existing literature. What is equally if not more significant are the insights that can be derived from them: these are analyzed in the section that follows.

IV   Discussion

This section derives consequences of (27) and provides economic insight into the solution. Section IV.A discusses how (27) can be expressed as a weighted average, and what the implications of this form for the decision-making of investors. Section IV.B analyzes the sign of the hedging demand term and links it to the behavior of the consumption-wealth ratio. Section IV.C compares the solutions for terminal wealth and for consumption, and answers the question posed in the introduction, how is horizon to be interpreted in the case of utility over consumption? In addition, the convergence of the solution at long horizons is established. Section IV.D uses the formula for the portfolio choice rule to understand where and why the log-linear solution goes wrong. Finally, Section IV.E demonstrates why complete markets play such an important role.

A   A weighted average formula

At first glance, the difference between the allocation under terminal wealth and under consumption (27) appears large indeed. While myopic demand is the same in both cases, hedging demand is a linear function of X_t under terminal wealth, but a much more complicated, nonlinear function under consumption. However, a closer look reveals an intriguing relation. Hedging demand in (27) take the form of a weighted average. The functions that are averaged equal hedging demand in (28) for different values of the horizon. The weights depend on the functions Ψ.

To better understand this result, it is helpful to rewrite the portfolio allocation for the investor with utility over consumption as follows:

$$\alpha_t = \int_0^{T-t} \frac{\Psi(X_t, \tau)}{\int_0^{T-t} \Psi(X_t, \tau')\,d\tau'}\left[\frac{1}{\gamma}\left(\frac{\mu_t - r}{\sigma^2}\right) - \frac{\sigma_X}{\gamma\sigma}(C(\tau)X_t + B(\tau))\right]\,d\tau. \tag{29}$$

Note that the myopic term is also a weighted average, except in this case all the averaged terms are equal.

What does the function Ψ represent? From (24), the ratio of wealth to consumption equals

$$\frac{W_t}{C_t} = \int_0^{T-t} \Psi(X_t, \tau) \, d\tau. \tag{30}$$

Thus Ψ is the value, scaled by today's consumption, of consumption in τ periods. The weights in (29) correspond to the value of future consumption in each period.

More will be said about the economics behind (29) in Section IV.C. However, an immediate economic implication of (29) and (30) is that the investor with utility over consumption allocates wealth as if saving for each consumption event separately. To each future consumption event, the investor applies the terminal wealth analysis. Thus, it is correct to think of the investor as holding separate accounts for, say, retirement and a house. The allocation in the overall portfolio equals an average of the allocation in the "retirement" portfolio, the "house" portfolio, etc. The average is weighted by the amount the investor has saved in each of the portfolios.

B   Hedging demand and the consumption-wealth ratio

Before discussing horizon effects, it is necessary to establish whether mean reversion increases or decreases the demand for stocks, relative to the case of constant investment opportunities. As discussed in Section III.C, the first term in the optimal allocation (27) gives the myopic demand, or the percent the investor would allocate to stocks if investment opportunities were constant. The key term in analyzing this question is therefore the second, hedging demand. Campbell and Viceira (1999) and Balduzzi and Lynch (1999) demonstrate, for particular parameter values, that hedging demand under utility for consumption is positive and is quite substantial. However, there is no general result available, even for the approximate analytical solution of Campbell and Viceira (1999).

While (27) may first appear complicated, it can be used to prove a result on hedging demand that holds for all parameter values:

Property 1 *For $\gamma > 1$, mean reversion increases the demand for stocks whenever the risk premium, $\mu_t - r$ is greater than zero. Equivalently, when $\gamma > 1$ and the risk premium is positive, hedging demand is positive.*

In Appendix C, it is shown that $B(\tau)$ and $C(\tau)$ are positive when $\gamma > 1$. Property 1 follows from this result, from the equation for the optimal allocation (27), and from the fact that X_t has the same sign as the risk premium.

Kim and Omberg (1996) demonstrate a similar result for terminal wealth. But, as discussed in the introduction, it is by no means obvious what the equivalent result is when consumption is included. Moreover, ignoring consumption makes it impossible to embed the problem in a general equilibrium framework. And yet the general equilibrium implications are interesting. The result suggests that mean reversion should increase individuals' demand for stocks, and

therefore, in a general equilibrium model, make the equity premium puzzle of Mehra and Prescott (1985) harder to solve.

While a similar result holds in the terminal wealth case, it is more difficult to interpret in economic terms. The intuition behind hedging demand, first given by Merton (1973) and frequently repeated, is that the additional demand for stocks is used to hedge changes in the investment opportunity set. More precisely, an increase in X_t can affect current consumption relative to wealth in two ways. By increasing investment opportunities, an increase in X_t means, in effect, that the consumer has more wealth. This is known as the income effect, and it causes wealth to rise relative to consumption. But there is also a substitution effect: putting money aside is more powerful, the greater the investment opportunities. When $\gamma > 1$, the income effect dominates. Namely, the consumption-wealth ratio rises when investment opportunities are high and falls when investment opportunities are low. To keep consumption stable, the investor must choose his portfolio to have more wealth in states with poorer investment opportunities. Finally, because stocks pay off when investment opportunities are poor, the investor with $\gamma > 1$ will hold more of them relative to the myopic case.

The intuition for hedging demand is inseparable from the assumption of an investor who consumes. A consequence is that it ties hedging demand to the consumption-wealth ratio. When hedging demand is positive, the consumption-wealth ratio must be increasing; when it is negative, the ratio must be decreasing. This is what the next property shows.

Property 2 *The consumption-wealth ratio is increasing in X_t when the risk premium is positive and $\gamma > 1$.*

The derivative of the ratio of wealth to consumption with respect to X_t equals the negative of hedging demand (see Equation 26). Thus, whenever hedging demand is positive, the wealth-consumption ratio is falling in X_t and the consumption-wealth ratio is thus rising in X_t.

Figures 1 and 2 plot hedging demand and the consumption-wealth ratio for $\gamma = 10$, $\gamma = 4$, $\gamma = 1$, and horizons T equal to 30, 10, and 5 years. The parameters for these plots come from Barberis (2000) and Campbell and Viceira (1999). In these figures, and in all the ones that follow, the vertical lines indicate +/- two standard deviations from the mean. [3] As Properties 1 and 2 state, the consumption-wealth ratio is increasing when $X_t > 0$, and the hedging demand is also positive in this case. For most values of $X_t < 0$, the consumption-wealth ratio is decreasing, and hedging demand is negative. This makes sense: for $X_t < 0$, decreases in X_t represent an increase in the investment opportunity set, because the investor can short stocks. However, there is a region below zero for which hedging demand is positive and the consumption-wealth ratio is decreasing. Kim and Omberg (1996) report a similar result for utility over

---

[3] In monthly terms, they are: $r = 0.0036$, $\bar{X} = 0.0965$, $\sigma = 0.0436$, $\sigma_X = 0.0404$, $\lambda_X = 0.0423$, $\rho = 0.0043$.

terminal wealth and offer an explanation based on the asymmetry in the distribution for X_t. When X_t is below zero, it must pass through zero to return to its long-run mean. In other words, for X_t negative but close to zero, increases may actually represent improvements in the investment opportunity set.

Figure 2 shows that the consumption-wealth ratio is non-monotonic in γ. The ratio for $\gamma = 10$ lies between that for $\gamma = 4$ and $\gamma = 1$. As explained by Campbell and Viceira (1999), this effect arises from the fact that γ acts as both the coefficient of relative risk aversion and the inverse of the elasticity of intertemporal substitution. It is also interesting to observe, in Figure 1, the hedging demand is still quite high, even at $\gamma = 10$. While myopic demand declines at the rate of $1/\gamma$, hedging demand remains high, even when risk aversion is large.

C   What is the meaning of the investor's horizon?

This section addresses one of the central question raised in the introduction. How should the investor's horizon be interpreted when the investor has utility over consumption?

The relation between the investment horizon and allocation is addressed by Brennan, Schwartz and Lagnado (1997), Barberis (2000), and Kim and Omberg (1996). Kim and Omberg, whose approach is analytical, state the following result:

Property 3 *For the investor with utility over terminal wealth and $\gamma > 1$, the optimal allocation increases with the investment horizon as long as the risk premium is positive. (Kim and Omberg, 1996)*

A short proof is contained in the Appendix.

Property 3 has a nice ring to it; it states that investors with longer horizons should invest more in stocks than investors with shorter horizons. This appears to fit with the advice of investment professionals that allocation to stocks should increase with the investor's horizon.

At closer inspection, however, Property 3 appears to be more of a mathematical curiosity than a useful tool for investors. Actual investors do not consume all their wealth at a single date. Even assuming that the account in question is a retirement account, the date of retirement is not an appropriate measure of horizon. The horizon may in fact be much longer, and is clearly should be determined by the timing of consumption after the investor enters retirement. If the investor plans to dip into savings for major expenditures before retirement, the answer may be still more inaccurate.

Figure 3 shows how misleading it can be to assume terminal wealth. The allocation for the investor with utility over consumption and $T = 30$ lies below the allocation for the investor with utility over wealth and $T = 10$, for reasonable values of X_t. The investor with utility over consumption and $T = 30$ has, in effect, a horizon of less than 10. The allocation for utility over consumption is typically less than half of that for utility over terminal wealth.

14

While the discrepancy between the solutions are greater, the greater the value of T, it is still large for all but the very smallest values of T. Figure 4 plots the allocation against the horizon for utility over consumption and terminal wealth, and for X equal to its mean, and one unconditional standard deviation above and below its long-term mean. Only for the very shortest horizons are the allocations close at all.

Clearly a more general notion of horizon is needed. Fortunately the analogy to fixed income developed in Sections III.A and III.B provides just such a notion. From (24), it follows that

$$\frac{W_t}{C_t} = \int_0^{T-t} \Psi(X_t, \tau) \, d\tau.$$

The investor's wealth is analogous to a coupon bond that pays in units of consumption. The value of wealth is simply the sum of the underlying "zeros", namely bonds that pay optimal consumption at each date. The natural measure of horizon for bonds is duration, which equals the negative of the sensitivity of the bond to changes in the interest rate. Here, the appropriate state variable is not the interest rate, but X_t. The duration for wealth with respect to X_t is given by

$$-\frac{\partial F/\partial X}{F} = -\left(C(\tau)X_t + B(\tau)\right)$$

in the case of terminal wealth and

$$-\frac{\partial G/\partial X}{G} = -\frac{\int_0^{T-t} \Psi(X_t, \tau)(C(\tau)X_t + B(\tau)) \, d\tau}{\int_0^{T-t} \Psi(X_t, \tau) \, d\tau}$$

in the case of consumption. Comparing with the portfolio allocation (27) demonstrates the following:

Property 4 *Hedging demand equals the duration of the investor's consumption stream with respect to X_t.*

Property 3 and Figure 4 show that duration has the same properties one would expect. Namely, the duration for the investor with terminal wealth increases monotonically with the horizon. Moreover, the duration for the investor with utility over consumption equals the weighted average of the duration of the underlying zeros, namely the values of consumption at each future date.

Thus, the more consumption is weighted towards the present, the more the investor's allocation is shrunk towards the myopic allocation, just like intermediate coupons shrink the duration of a coupon-paying bond.[4] The further out consumption goes, the higher the duration of the consumption stream and the greater the allocation. Given this analogy, the following property becomes completely natural.

---

[4]Because the myopic allocation is the same at all horizons, it does not affect the horizon analysis. Therefore, all the statements in the paragraph above can be said to apply equally to the allocation itself.

Property 5 *For the investor with utility over consumption and $\gamma > 1$, the optimal allocation increases with the investment horizon as long as the risk premium is positive. Moreover, the solution always lies below that for terminal wealth.*

It is useful to prove the second statement first. Every element in the average is less than the allocation to terminal wealth at T, and thus the whole average is less. That is,

$$-C(T-t) > -\frac{\int \Psi(X_t, \tau) C(\tau) \, d\tau}{\int \Psi(X_t, \tau) \, d\tau}, \tag{31}$$

and the same for $B(\tau)$. To prove the first statement, note that the effect of "adding" more consumption at the end of the horizon pulls up the overall average. Formally, the derivative of $-\int \Psi(X_t, \tau) C(\tau) / \int \Psi(X_t, \tau)$ equals

$$-\frac{\Psi(X_t, T-t)}{\int_0^{T-t} \Psi(X_t, \tau) \, d\tau} \left(C(T-t) - \frac{\int \Psi(X_t, \tau) C(\tau) \, d\tau}{\int \Psi(X_t, \tau) \, d\tau} \right).$$

From (31), the derivative is always positive.

Besides generalizing the notion of horizon to account for intermediate consumption, Property 4 also helps to understand the horizon result both in the case of consumption and terminal wealth. While the reasoning behind the income effect and substitution effect can explain why hedging demand is positive or negative, it does not explain why the effect is more powerful at longer horizons. When X_t rises (assuming it is already positive), future consumption becomes less valuable relative to current consumption because the discount rate is higher. This is the income effect - because there are better investment opportunities, the investor needs to put aside less today for future consumption. The income effect is larger at longer horizons because the value of consumption further out is more sensitive to changes in the discount rate. On the other hand, when X_t rises, the investor can afford to consume more in the future because there is time for wealth set aside now to increase. This is the substitution effect. When $\gamma > 1$, the income effect wins out. Namely, the further out the consumption is, the more ratio of wealth to consumption falls when X_t rises, and the greater is hedging demand.

Finally, the question of whether the results extend to infinite horizons is addressed. The key question is whether the portfolio rule converges as the horizon approaches infinity. Barberis (2000) and Brandt (2000) note this property in their numerical solutions. From Figures 1 and 2, it is evident that convergence also occurs in the model considered here. The plots for $T = 30$ appear to be closer to those for $T = 10$ than the plots for $T = 10$ are to $T = 5$. Moreover, in Figure 4, convergence is noticeable even at $T = 10$. While the numerical results in the papers of Barberis and Brandt are strongly suggestive, they cannot demonstrate that convergence is guaranteed at all relevant parameter values. Based on the closed-form solutions for the portfolio choice rule (27), it is possible to demonstrate just such a result.

Property 6 *As the investor's lifetime approaches infinity, the allocation to stocks converges to a finite limit.*

The proof is contained in Appendix D. Besides its inherent interest, an additional benefit of this result is that it allows the exact solution (27) to be compared to the approximate infinite-horizon solution of Campbell and Viceira. This is done in the following section.

D   Non-linearities in the solution

Campbell and Viceira (1999) solve an infinite-horizon version of the intertemporal consumption and portfolio choice problem by taking a log-linearized approximation of the budget constraint. Like the allocation for terminal wealth, (28), the allocation that Campbell and Viceira find is linear. In the model of Campbell and Viceira, the linearity occurs as a direct result of the approximation; it is "hard-wired" into the model.

The framework in this paper can be used to address the performance of the log-linear approximation, and why it works well in some cases and not in others. Figure 1 demonstrates that hedging demand (and therefore the overall allocation) is close to a linear function near $\bar{X}$. For large values of X hedging demand appears to flatten out. For these values, the allocation to stocks is actually less sensitive to changes in the state variable than the analysis of Campbell and Viceira (1999) would imply.

This finding is consistent with that of Campbell, Cocco, Gomes, Maenhout, and Viceira (1998). Campbell et al. solve the infinite-horizon consumption problem numerically and find that the exact, numerical solution flattens out for large values of X_t. However, because their analysis is purely numerical, Campbell et al. cannot shed light on why the discrepancy occurs.

In contrast, the exact, closed-form solution (27) can help to understand the 0 between the log-linear and the actual solution. From (27), it follows that there are two ways changes in X_t can affect the portfolio rule. The first is directly, through $C(\tau)X_t + B(\tau)$, just as in the linear case. The second is indirectly, through changes in the weights $\Psi(X_t, \tau)$. This first effect is what the change would be if the solution were actually linear, namely, if $\Psi(X_t, \tau)$ were a constant:

$$\frac{d\alpha_t^{\text{lin}}}{dX_t} = \frac{1}{\gamma\sigma}\left(1 - \sigma_X \frac{\int \Psi(X_t, \tau)C(\tau)}{\gamma \int \Psi(X_t, \tau)}\right). \tag{32}$$

This term is always positive because $C(\tau) < 0$ when $\gamma > 1$ (Appendix C).

However, $\Psi(X_t, \tau)$ is not a constant in X_t. The difference between the true derivative and (32) equals

$$\frac{d\alpha_t}{dX_t} - \frac{d\alpha_t^{\text{lin}}}{dX_t} = \frac{\sigma_X}{\gamma\sigma}\left[\left(\frac{\int \Psi(X_t, \tau)(C(\tau)X_t + B(\tau))}{\gamma \int \Psi(X_t, \tau)}\right)^2 - \frac{\int \Psi(X_t, \tau)(C(\tau)X_t + B(\tau))^2}{\gamma \int \Psi(X_t, \tau)}\right]. \tag{33}$$

This term corresponds to the effect of changes in X_t on the weights and is always negative because the square is a convex function.

Figure 5 plots the "linear term", (32), and the derivative itself (the sum of (32) and (33)) for $\gamma = 10$. The figure shows that (32) is nearly a constant, demonstrating that it indeed represents a linear effect on changes on α_t. However, the derivative itself slopes down dramatically. The effect is more dramatic, the higher the value of T.

The fixed income analogy of the previous section is also useful in understanding this dramatic downward slope. The duration for coupon bonds decreases as the interest rate increases (see, e.g. Campbell, Lo and MacKinlay, 1997, chap. 10). This is because increases in the interest rate decrease the value of long-term bonds more than the value of short-term bonds. The bonds with the higher duration (long-term bonds) therefore receive less weight when interest rates fall. As in the previous section, it is useful to think of $\Psi(X_t, \tau)$ as the value of a bond paying in units of consumption at horizon τ. Increases in X_t cause the value of $\Psi(X_t, \tau)$ to decrease more, the greater the value of τ. The portfolio weight α is given by a weighted average of $C(\tau)X_t + B(\tau)$, where the weights are like the values of discount bonds. Thus (33) arises because increases in X_t decrease the weights on the terms with higher values of $C(\tau)X_t + B(\tau)$.

E   The role of complete markets

A theme of the previous sections is the difference between the solutions when the investor has utility over consumption, and when the investor has utility over terminal wealth. It is not correct to solve for one, and assume that it will "look like" the solution to the other. There is another important difference between the two solutions. It is not possible, under incomplete markets, to obtain a closed-form solution for the investor with utility over consumption. It is only possible under terminal wealth. The purpose of this section is to make it clear exactly why complete markets are required.

In Section III the problem was solved by first solving for consumption and then deriving the portfolio rules from the need to finance consumption. Such portfolio rules exist because wealth satisfies (19), or equivalently, that wealth follows a martingale under the equivalent martingale measure. As stated by Pliska (1986), when markets are complete, any process that is a martingale under the equivalent martingale measure and satisfies certain regularity conditions can be financed by trading in the underlying securities. When markets are incomplete, the set of consumption rules that can be financed is more difficult to describe. It is no longer possible to solve for consumption first.

To illustrate this point, the martingale method as outlined by Cox and Huang (1989) will be used to "solve" the problem defined by Kim and Omberg (1996), namely the incomplete markets version of the problem defined in this paper.

Assume a two-dimensional Brownian motion $w_t = (w_1(t) \ w_2(t))'$. The equations defining the processes S_t and X_t are identical to those in the first section, but for one important difference. The terms σ and σ_X are now 1-by-2 vectors. The covariance equals $-S\sigma\sigma_X'$, and the variances of S and X are $S^2\sigma\sigma'$ and $\sigma_X\sigma_X'$ respectively. Let ρ_{mX} denote the correlation. The analogous relationship

between X and μ is:

$$X_t = \frac{(\mu_t - r)}{(\sigma\sigma')^{\frac{1}{2}}}.$$

Unlike the case of complete markets, this case does not have a well-defined price-of-risk process. Any process η_t satisfying

$$\sigma\eta_t = -(\mu_t - r).$$

is a price-of-risk process and, because this is a single equation in two unknowns, there are infinitely many solutions. Each of these solutions corresponds to a different equivalent martingale measure, and hence a different process ϕ. The following is a natural choice for η:

$$\eta_t = \sigma'(\sigma\sigma')^{-1}(\mu_t - r).$$

X_t and η_t are related as follows:

$$\eta_t = \sigma'(\sigma\sigma')^{-1/2}X_t. \tag{35}$$

By the same reasoning used to derive (19), the investor's wealth satisfies:[5]

$$\mathcal{L}F + F_t - rF = F_Z Z\eta'\eta - F_X \sigma_X \eta \tag{36}$$

with boundary condition

$$F(Z_T, X_T, T) = Z_T^{\frac{1}{\gamma}}.$$

The trial solution is

$$F(Z_t, X_t, t) = Z_t^{\frac{1}{\gamma}} \exp\{C(T-t)X_t^2/2 + B(T-t)X_t + A(T-t)\} \tag{37}$$

When is substituted into the equation above, a system of ordinary differential equations similar to those found in Section III results. The difference is that,

$$\begin{aligned}
a &= \frac{1-\gamma}{\gamma} \\
b &= 2\left(\frac{\gamma-1}{\gamma}\sigma_X - \lambda_X\right) \\
c &= \frac{1}{\gamma}\sigma_X^2.
\end{aligned}$$

The problem is, this implies a *different* portfolio allocation that that found by Kim and Omberg (1996). The difference lies in the parameter a. Kim and Omberg find the value of this parameter to be

$$\sigma_X\sigma_X'\left(1 + \frac{1-\gamma}{\gamma}\rho_{mX}^2\right).$$

---

[5]The same notation is used as in Section III.A, except that the last argument of F clearly equals T and is therefore suppressed.

the solutions agree if and only if markets are complete.

Where does the reasoning of Section III break down? Since the solution above does satisfy differential equation (36), it must be the solution to the static optimization problem. The only way that it does not correspond to a solution to (5) is if there are no portfolio rules to finance it. A portfolio rule finances wealth F if and only if

$$-F_X(t)\sigma_X + F_Z(t)Z_t\eta_t = F(t)\alpha_t\sigma. \qquad (38)$$

This is a system of two equations in one unknown. A solution exist only if one of the equations equals a constant times the other, that is, if markets are perfectly correlated. When there is a single asset, state-variable risk and market risk cannot be perfectly hedged at the same time.

V   Conclusion

This paper demonstrates that, under mean reversion and complete markets, the multiperiod consumption and portfolio allocation problem can be solved in closed form. This question has interested the literature for some time: there are numerous papers that solve the problem for terminal wealth, and give numerical solutions for consumption, as well as approximate analytical solutions. But a closed-form solution as eluded the literature.

As has been shown, the solution is more than a complicated formula. It can be expressed as a weighted average that is analogous to the duration formula for coupon bonds. It can be used to resolve questions that are posed by but not solved in the current literature, such as the sign of hedging demand, the reason for inaccuracies in the log-linear solution, and whether the solution converges at long horizons. It raises and resolves issues that have been ignored unjustly, such as how horizon results are to be interpreted in the most realistic case, i.e. when the investor has utility over consumption.

This paper has chosen to focus on the case where investment returns are mean-reverting. However, the methods in this paper are more general and can be applied to other portfolio choice problems. In particular, results from this paper have already been used in recent work by Brennan and Xia (2000) and Chacko and Viceira (2000).

The solution does require that markets be complete, unlike the solution for terminal wealth. Though it generalizes the preferences, it requires more specific assumptions on the data-generating process. It does give researchers seeking an analytical solution a choice: either require terminal wealth, use log-linearized approximate solutions, or assume complete markets. In the case of mean-reverting returns, assuming complete markets is realistic, while requiring terminal wealth is highly misleading. Ultimately the choice of the least evil is up to the researcher, but it is important to realize that there is a choice to be made.

Appendix

A Proof that solutions to (3) and (5) exist, under assumption (7) and $\gamma > 1$.

The proof has two steps. First, it is necessary to show that the Lagrange multipliers k and l in (10) and (11) exist and are finite. Substituting the optimal policies into the budget constraints yields:

$$k = W_0^{-\gamma} \left(E \left(\phi_T^{1-\frac{1}{\gamma}} \right) \right)^{\gamma}$$

and

$$l = W_0^{-\gamma} \left(E \int_0^T \phi_t^{1-\frac{1}{\gamma}} e^{-\frac{1}{\gamma}\rho t} \, dt \right)^{\gamma}.$$

By Jensen's inequality, it suffices to show that $E\phi_t$ is finite and continuous for all t.

By definition,

$$\phi_t \leq \exp \left\{ - \int_0^t X_t \, dw_t - rdt \right\}.$$

Because $\int_0^t X_t \, dw_t$ is Gaussian (see, e.g. Duffie, 1996, Appendix E), the right hand side is lognormal. Therefore, the expectation exists. Moreover, the mean and variance of $\int_0^t X_t \, dw_t$ are continuous functions of time, implying that the expectation is continuous.

The second step is that the optimal policies satisfy regularity conditions. In particular, it is enough to show that

$$E[\phi_t^{-p\frac{1}{\gamma}}] < \infty$$

for some $p > 1$ (see Cox and Huang, 1989). Choose $p < \gamma$. Then

$$E[\phi_t^{-p\frac{1}{\gamma}}] \leq$$

$$\left(E \exp \left\{ q\frac{p}{\gamma} \left(\int_0^t X_s \, dw_s - rt \right) \right\} \right)^{\frac{1}{q}} \left(E \exp \left\{ q'\frac{p}{\gamma}\frac{1}{2} \int_0^t X_s^2 \, ds \right\} \right)^{\frac{1}{q'}},$$

where $\frac{1}{q} + \frac{1}{q'} = 1$ and $q, q' > 1$, by the Cauchy-Schwartz inequality. Choose q' so that $q'\frac{p}{\gamma} < 1$. The first term on the right hand side is finite because of the lognormality described above. The second term is finite because of Jensen's inequality, and (7).

B Existence and regularity conditions for G

The expression $b^2 - 4ac$ is greater than zero for $\gamma > 1$. First note that

$$\eta^2/4 = \lambda_X^2 - \left(\frac{1-\gamma}{\gamma} \right) (\sigma_X^2 - 2\sigma_X\lambda_X) \tag{39}$$

21

is greater than

$$\frac{\gamma - 1}{\gamma}\lambda_X^2 + \frac{\gamma - 1}{\gamma}\sigma_X^2 - \left(\frac{\gamma - 1}{\gamma}\right)2\sigma_X\lambda_X.$$

But this equals $\frac{\gamma-1}{\gamma}(\lambda - \sigma_X)^2$, hence it is always greater than zero.

The results require G_{XZ}, G_{ZZ}, G_{XX} and G_t to to be continuous. From (20) and (23), it follows that the continuity of B' and C' is sufficient. To establish the continuity of B' and C', it suffices to show

$$2\eta - (b + \eta)(1 - e^{-\eta\tau}) > 0$$

for $0 \leq \tau \leq T$. When $\gamma > 1$, $\eta < \beta$ so the expression is strictly positive.

C   Proof that the functions $B(\tau)$, $C(\tau)$ and their derivatives have sign $1 - \gamma$.

As shown above,

$$2\eta - (b + \eta)(1 - e^{-\eta\tau}) > 0$$

for the parameter values of interest. The statement for B and C follows directly. To prove the statement for the derivatives, note that the derivative of C with respect to τ equals

$$(2\eta\left(\frac{1-\gamma}{\gamma}\right)e^{-\eta\tau}(2\eta - (b+\eta)(1-e^{-\eta\tau})) - (-2\eta(b+\eta)e^{-\eta\tau})\left(\frac{1-\gamma}{\gamma}\right)(1-e^{-\eta\tau}))$$

divided by $(2\eta - (b + \eta)(1 - e^{-\eta\tau}))^2$. This reduces to

$$2\eta^2\left(\frac{1-\gamma}{\gamma}\right)e^{-\eta\tau}/(2\eta - (b + \eta)(1 - e^{-\eta\tau}))^2.$$

Similarly, the derivative of B is equal to

$$4\left(\frac{1-\gamma}{\gamma}\right)\lambda_X\bar{X}(1 - e^{-\eta\tau/s})\eta e^{-\eta\tau/2}\eta(2\eta - (b + \eta)(1 - e^{-\eta\tau}))$$

$$+ \quad 4\left(\frac{1-\gamma}{\gamma}\right)\lambda_X\bar{X}(1 - e^{-\eta\tau/2})^2(\eta(b + \eta)\eta e^{-\eta\tau})^2$$

divided by $2\eta - (b + \eta)(1 - e^{-\eta\tau})^2$. Because $b + \eta$ is greater than zero, both terms have the correct sign.

D   Proof that the allocation converges at long horizons

It suffices to show that

$$\frac{\int^{T-t}\Psi(X_t, \tau)C(\tau)\,d\tau}{\int^{T-t}\Psi(X_t, \tau')\,d\tau'}$$

and the same equation for $B(\tau)$ converge to finite limits. First, from Appendix C, $C(\tau)$ and $B(\tau)$ are monotonic in τ. Therefore, their averages over τ are monotonic in T. To prove convergence, it suffices to show that these sequences are bounded above. From Appendix C, it follows that

$$0 \leq \frac{\int_0^{T-t} \Psi(X_t, \tau) |C(\tau)| \, d\tau}{\int_0^{T-t} \Psi(X_t, \tau)} \leq |C(T)|.$$

The corresponding inequality holds for $|B(T)|$. It follows immediately from (22) and (21), that $B(T)$ and $C(T)$ are bounded above.

E   The univariate process for the stock price

Recall that the process for X_t is given by:

$$dX_t = -\lambda_X (X_t - \bar{X}) dt - \sigma_X dw_t.$$

It is well known that the solution to the above equation has the following form:

$$
\begin{aligned}
X_t &= e^{-\lambda_X t} \left[X_0 + \lambda_X \bar{X} \int_0^t e^{\lambda_X s} \, ds - \sigma_X \int_0^t e^{\lambda_X s} dw_s \right] \\
&= e^{-\lambda_X t} X_0 + \bar{X}(1 - e^{-\lambda t}) - \sigma_X \int_0^t e^{-\lambda_X (t-s)} \, dw_s.
\end{aligned}
$$

From Ito's lemma it follows that:

$$\ln S_t = \ln S_0 + \int_0^t \left(\sigma X_t + \sigma r - \frac{\sigma^2}{2} \right) dt + \sigma w(t) \tag{40}$$

Substituting the expression for X_t from above, we have:

$$
\begin{aligned}
\ln S_t &= \ln S_0 + \sigma \int_0^t \left(e^{-\lambda_X s} X_0 + \bar{X}(1 - e^{-\lambda_X s}) \right) \, ds + t\sigma r \\
&\quad - t\frac{\sigma^2}{2} - \sigma \sigma_X \int_0^t e^{-\lambda_X s} \int_0^s e^{\lambda_X u} \, dw_u \, ds + \sigma w(t)
\end{aligned}
$$

The only tricky part is calculating the Ito integral. Let $I_{[a,b]}(x)$ denote the function that equals 1 when $x \in [a, b]$ and zero otherwise. For all values of s and u in $[0, t]$,

$$I_{[0,s]}(u) = I_{[u,t]}(s).$$

23

The Ito integral is equal to:

$$\int_0^t e^{-\lambda_X s} \int_0^t I_{[0,s]}(u) e^{\lambda_X u} \, dw_u \, ds$$

$$= \int_0^t e^{\lambda_X u} \int_0^t I_{[u,t]}(s) e^{-\lambda_x s} \, ds \, dw_u$$

$$= \int_0^t e^{\lambda_X u} \frac{1}{\lambda_X} \left(e^{-\lambda_X t} - e^{-\lambda_X u} \right) \, dw_u$$

$$= \frac{1}{\lambda_X} \int_0^t \left(1 - e^{-\lambda_X (t-u)} \right) \, dw_u$$

Note that the drift term inside the integral is equal to:

$$-\frac{1}{\lambda_X} X_0 \left(e^{-\lambda_X t} - 1 \right) + \sigma \bar{X} t + \frac{1}{\lambda_X} \bar{X} \left(e^{-\lambda_X t - 1} \right)$$

Putting it all together yields:

$$\ln S_t = \ln S_0 + \sigma r t + \frac{\sigma}{\lambda_X} X_0 \left(1 - e^{-\lambda_X t} \right) + \sigma \bar{X} \left(t - \frac{1}{\lambda_X} \left(e^{-\lambda_X t} - 1 \right) \right)$$

$$- \frac{t \sigma^2}{2} + \sigma w(t) - \frac{\sigma \sigma_X}{\lambda_X} \int_0^t (1 - e^{-\lambda_X (t-s)}) \, dw_S$$

Unsurprisingly, the univariate process for $\ln S_t$ is mean-reverting. We can see this by comparing the instantaneous variance σ to the variance over a finite length of time t. For Brownian motion, the variance would increase at a linear rate. For t sufficiently large, the variance of the process above is smaller than the variance of a Brownian motion.

References

Balduzzi, Pierluigi, and Anthony W. Lynch, 1999, Transaction costs and predictability: some utility cost calculations, *Journal of Financial Economics* 52, 47–78.

Barberis, Nicholas, 2000, Investing for the long run when returns are predictable, *Journal of Finance* 55, 225–264.

Brandt, Michael W., 1999, Estimating portfolio and consumption choice:a conditional Euler equations approach, *Journal of Finance* 54, 1609–1645.

Brennan, Michael J., Eduardo S. Schwartz, and Ronald Lagnado, 1997, Strategic asset allocation, *Journal of Economic Dynamics and Control* 21, 1377–1403.

Brennan, Michael J., and Yihong Xia, 2000, Dynamic asset allocation under inflation, University of Pennsylvania.

Campbell, John Y., Joao Cocco, Francisco Gomes, Pascal J. Maenhout, and Luis M. Viceira, 1998, Stock market mean reversion and the optimal equity allocation of a long-lived investor, Harvard University.

Campbell, John Y., A. Craig MacKinlay, and Andrew W. Lo, 1997, *The Econometrics of Financial Markets*. (Princeton University Press Princeton, NJ).

Campbell, John Y., and Luis M. Viceira, 1999, Consumption and portfolio decisions when expected returns are time-varying, *Quarterly Journal of Economics* 114, 433–495.

Chacko, George, and Luis Viceira, 2000, Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets, Harvard University.

Cox, John C., and Chi-Fu Huang, 1989, Optimal consumption and portfolio policies when asset prices follow a diffusion process, *Journal of Economic Theory* 39, 33–83.

Cox, John C., and Chi-Fu Huang, 1991, A variational problem arising in financial economics, *Journal of Mathematical Economics* 20, 465–487.

Cox, John C., Jonathan E. Ingersoll, Jr., and Stephen A. Ross, 1985, A theory of the term structure of interest rates, *Econometrica* 53, 385–407.

Duffie, Darrell, 1996, *Dynamic Asset Pricing Theory*. (Princeton University Press Princeton, NJ).

Duffie, Darrell, and Rui Kan, 1996, A yield-factor model of interest rates, *Mathematical Finance* 6, 379–406.

Epstein, Lawrence, and Stanley Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption and asset returns, *Econometrica* 57, 937–968.

Harrison, Michael, and David Kreps, 1979, Martingales and multiperiod securities markets, *Journal of Economic Theory* 20, 381–408.

Karatzas, Ioannis, John P. Lehoczky, and Steven E. Shreve, 1987, Optimal portfolio and consumption decisions for a small investor on a finite horizon, *SIAM Journal of Control and Optimization* 25, 1557–1586.

Karatzas, Ioannis, and Steven E. Shreve, 1998, *Methods of Mathematical Finance.* (Springer-Verlag New York, NY).

Kim, Tong Suk, and Edward Omberg, 1996, Dynamic nonmyopic portfolio behavior, *Review of Financial Studies* 9, 141–161.

Liu, Jun, 1999, Portfolio selection in stochastic environments, Stanford University.

Mehra, Rajnish, and Edward C. Prescott, 1985, The equity premium: a puzzle, *Journal of Monetary Economics* 15, 145–161.

Merton, Robert C., 1971, Optimum consumption and portfolio rules in a continuous-time model, *Journal of Economic Theory* 3, 373–413.

Merton, Robert C., 1973, An intertemporal capital asset pricing model, *Econometrica* 41, 867–887.

Pliska, Stanley R., 1986, A stochastic calculus model of continuous trading: optimal portfolios, *Mathematics of Operations Research* 11, 239–246.

Schroder, Mark, and Costis Skiadas, 1999, Optimal consumption and portfolio selection with stochastic differential utility, *Journal of Economic Theory* 89, 68–126.

Figure 1: Hedging demand for the investor with utility over consumption as a function of X_t for $\gamma = 10$ (solid), $\gamma = 4$ (dash), and $\gamma = 1$ (dots).
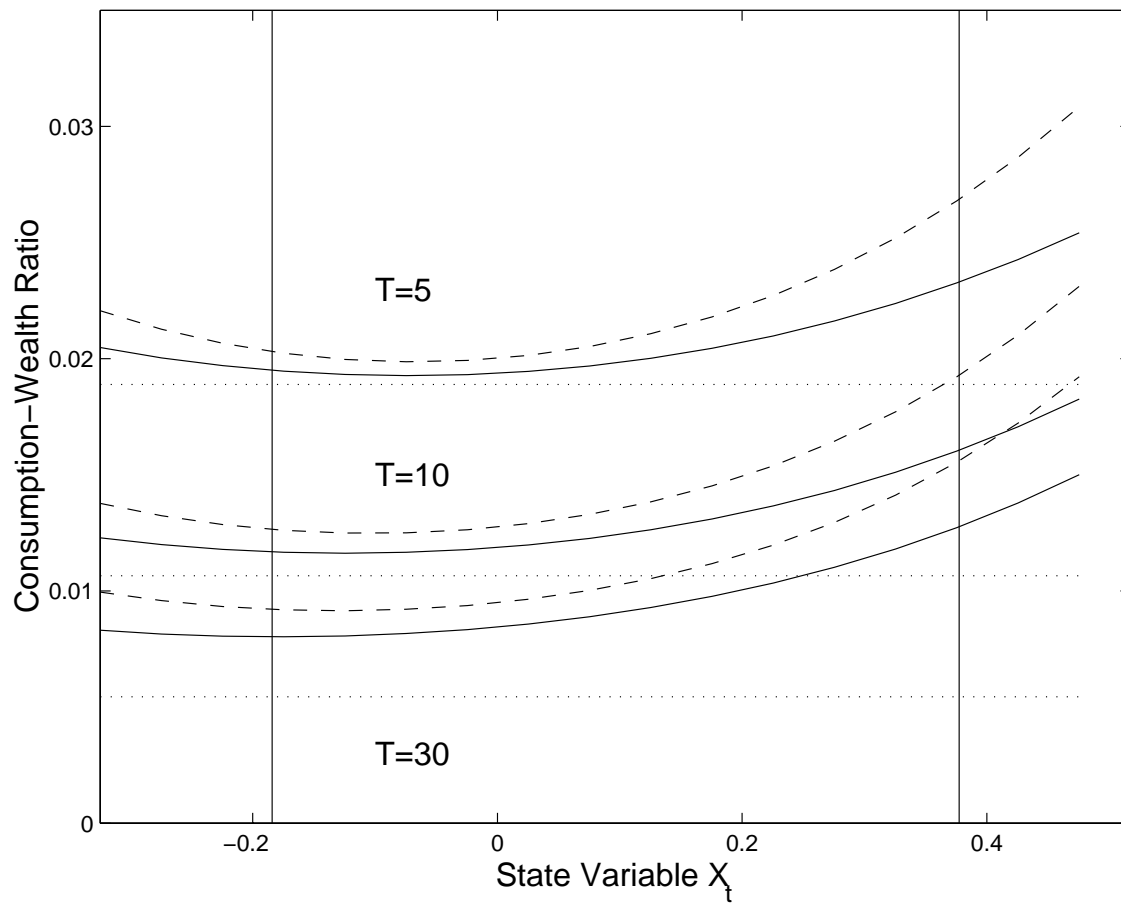
Figure 2: The consumption-wealth ratio as a function of X_t for $\gamma = 10$ (solid), $\gamma = 4$ (dash), and $\gamma = 1$ (dots).

Figure 3: Optimal allocation as a function of X_t for utility over consumption (circles) and over terminal wealth for $\gamma = 10$.

Figure 4: Optimal allocation as a function of horizon for $\gamma = 10$ and utility over consumption (circles) and utility over terminal wealth. For the center graph, $X = \bar{X}$, for the top and bottom graphs $X = \bar{X} + \sigma_X \left(1 - (1 - \lambda_X)^2\right)^{-\frac{1}{2}}$ and $X = \bar{X} - \sigma_X \left(1 - (1 - \lambda_X)^2\right)^{-\frac{1}{2}}$.
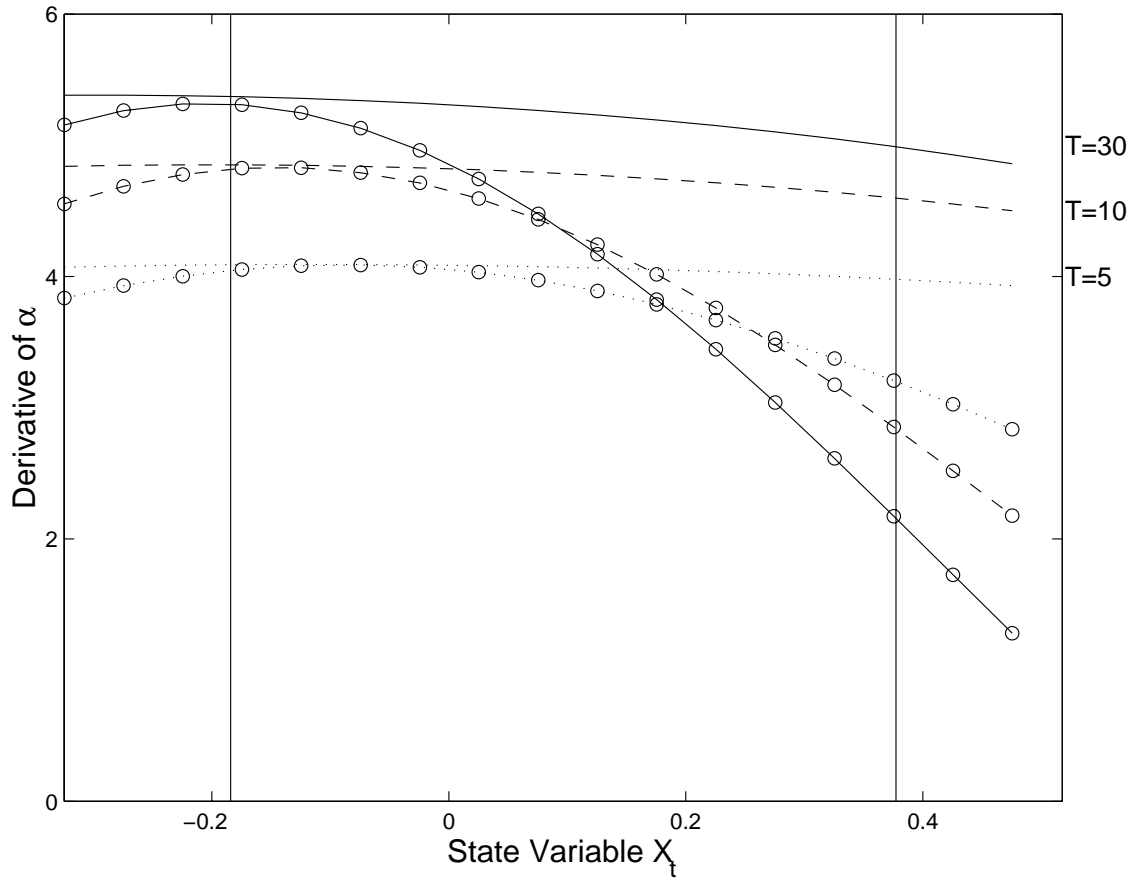
Figure 5: The derivative of the optimal allocation with respect to X_t (circles) and the linear term in the derivative (see Eq. 32) for $\gamma = 10$.