

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

Jessica A. Wachter*

November 19, 2001

Forthcoming, *Journal of Financial and Quantitative Analysis*

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.

*NYU, Stern School of Business, Department of Finance. Address: 44 West 4th Street, Suite 9-190, New York, NY 10012. Phone: (212)998-0799. Email: jwachter@stern.nyu.edu. I thank Nicholas Barberis, Stephen Brown (the editor), John Campbell, John Cochrane, George Chacko, Gary Chamberlain, Benjamin Friedman, Andre Perold, Stephen Ross, Jeremy Stein, Luis Viceira, an anonymous referee, and seminar participants at Boston College, Columbia University, Cornell University, Dartmouth College, Duke University, Harvard Business School, Massachusetts Institute of Technology, New York University, University of California at Berkeley, University of California at Los Angeles, University of Chicago, University of Southern California, and Yale University for helpful comments. I thank the National Science Foundation and Lehman Brothers for financial support.

I Introduction

Recent empirical work demonstrates that stock returns are predictable by scaled-price measures. For example, Campbell and Shiller (1988) and Fama and French (1989) demonstrate that the dividend-price ratio predicts excess returns on the market. There is related evidence that stock returns exhibit mean-reversion. Poterba and Summers (1988) demonstrate that the variance of stock returns is reduced at longer horizons.

What are the consequences of predictability in returns for portfolio choice? A recent literature, beginning with Brennan, Schwartz and Lagnado (1997), Campbell and Viceira (1999), and Kim and Omberg (1996), analyzes the consequences of this empirical fact for financial decisions. Barberis (2000) generalizes the set-up to account for parameter uncertainty, while Balduzzi and Lynch (1999) account for transaction costs. Brandt (1999) derives a portfolio rule that is robust to distributional assumptions. Xia (2001) documents the effect of learning about predictability.

These papers, with the exception of Campbell and Viceira (1999) and Kim and Omberg (1996) solve the investor's problem numerically. Campbell and Viceira assume an infinitely-lived investor, and derive an approximate analytical solution. Kim and Omberg (1996) derive an exact solution by specializing the problem to the case where the investor cannot consume over his lifetime, and has utility over terminal wealth only.

This paper departs from the literature above by deriving an exact, closed-form solution for portfolio weights when the investor has utility over consumption. In order to solve the problem in closed form, it is necessary to assume that the predictor variable and the stock return are perfectly negatively correlated. In this setting, the perfect negative correlation implies that markets are complete. In the case of return predictability, perfect negative correlation is realistic: Empirical studies have found this correlation to be close to -1. Moreover, even perfectly correlated continuous-time processes are imperfectly correlated when measured in discrete time. Thus assuming perfect negative correlation lends insight into the problem investors actually face.

Assuming utility over consumption allows the portfolio choice problem to be connected to the decisions of investors in a way that assuming terminal wealth does not. 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, the allocation to stocks increases with the investment horizon. This result has received much attention because it redeems the popular, but much criticized advice of investment professionals. While this result is well-understood for the case of terminal wealth, it is not clear how to interpret it for the case of utility over consumption. For example, when an investor is saving not only for retirement, but for education and a house, previous results cannot be applied.

Intuitively, one would think that the more consumption is weighted toward 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 solution has a novel economic interpretation. The portfolio 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 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 approximate solution of Campbell and Viceira (1999) and the actual solution.

This paper is related to recent work by Liu (1999). Independently and concurrently, Liu shows that the investor's problem can be characterized by a series of ordinary differential equations under several sets of assumptions. This paper considers a case to which Liu's assumptions do not apply. Liu also provides an economic analysis, but focuses on the case of an investor with utility over terminal wealth rather than consumption.

The organization of this 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 gives an economic interpretation of the solution, while Section V discusses how the solution might be extended to the case of incomplete markets.

II The Consumption and Portfolio Choice Problem under Mean-Reverting Returns

This section lays out the assumptions on asset returns and states the optimization problem of interest.

Let w_t denote the standard, one-dimensional Brownian motion. Assume that the price S of the risky security follows the process

$$(1) \quad \frac{dS_t}{S_t} = \mu_t dt + \sigma dw_t,$$

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:

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

The volatilities σ and σ_X are assumed to be constant and strictly positive, and λ_X is assumed to be greater than or equal to 0. 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 $r \geq 0$, but this assumption can be relaxed. These assumptions are like those in Kim and Omberg (1996), except that Kim and Omberg allow for imperfect correlation, and thus incomplete markets. In the empirical applications of Campbell and Viceira (1999) and Barberis (2000), 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. As shown in Appendix C, these assumptions imply that stock returns are mean-reverting for realistic parameter values.

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:

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

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

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:

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

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

$$W_T \geq 0.$$

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 More recently, Xia (2001) numerically compares the allocation for utility over terminal wealth to that for utility over consumption. The analytical results below confirm these results.

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 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.D uses the results of the previous two to derive the formula for the optimal allocation

¹Campbell 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.

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:

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

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

$$(8) \quad E [W_T \phi_T] = W_0$$

and

$$(9) \quad E \left[\int_0^T c_t \phi_t dt \right] = W_0$$

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 (1991) and Karatzas and Shreve (1998).

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),

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

For (5),

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

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 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:

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

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

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

By Ito's Lemma

$$(14) \quad \frac{dZ_t}{Z_t} = (r + X_t^2) dt + X_t dw_t.$$

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

$$(15) \quad G(Z_t, X_t, t) \equiv W_t$$

$$(16) \quad = Z_t E \left[\int_t^T Z_s^{\frac{1}{\gamma}-1} e^{-\frac{\rho}{\gamma}s} ds \mid X_t, Z_t \right].$$

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:²

$$(17) \quad F(Z_t, X_t, t; T) = Z_t E_t \left[Z_T^{\frac{1}{\gamma}-1} \mid X_t, Z_t \right].$$

² Assuming both investors start with the same initial wealth W_0 , the constant is given by

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

Alternatively, initial wealth could be chosen for each investor so that the constant equals one.

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 (17), adjusted to take into account the rate of time preference:

$$(18) \quad G(Z_t, X_T, t) = \int_t^T F(Z_t, X_t, t; s) e^{-\frac{\rho}{\gamma}s} ds.$$

The terms inside the integral in (18) 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 (16). 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. A 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

$$(19) \quad \mathcal{L}G + \frac{\partial G}{\partial t} + Z_t^{\frac{1}{\gamma}} e^{-\frac{\rho}{\gamma}t},$$

where $\mathcal{L}G = \frac{1}{2} \frac{\partial^2 G}{\partial X^2} \sigma_X^2 + \frac{1}{2} \frac{\partial^2 G}{\partial Z^2} Z^2 X^2 - \frac{\partial^2 G}{\partial Z \partial X} Z X \sigma_X + \frac{\partial G}{\partial X} (-\lambda_X (X - \bar{X})) + \frac{\partial G}{\partial Z} Z (r + X^2)$. The last term in (19) comes from the consumption coupon that is paid each period. No arbitrage requires that

$$(20) \quad \mathcal{L}G + \frac{\partial G}{\partial t} + Z_t^{\frac{1}{\gamma}} e^{-\frac{\rho}{\gamma}t} - rG = \left(\frac{\partial G}{\partial Z} Z_t X_t - \frac{\partial G}{\partial X} \sigma_X \right) X_t.$$

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 (18) 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 F 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 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:

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

where

$$H(X_t, \tau) \equiv \exp \left\{ \frac{1}{\gamma} (A_1(\tau) X_t^2 / 2 + A_2(\tau) X_t + A_3(\tau)) \right\}.$$

From the relationship between G and F described in (18), it is clear that

$$(21) \quad G(Z_t, X_t, t) = Z_t^{\frac{1}{\gamma}} e^{-\frac{\rho}{\gamma}t} \int_0^{T-t} H(X_t, \tau) d\tau.$$

Substituting (21) back into (20) and matching coefficients on X^2 , X , and the constant term leads to a system of three differential equations in A_1 , A_2 , and A_3 . The method for solving these equations is standard, and is discussed in Kim and Omberg (1996). The equation for A_1 is known as a *Riccati equation* and can be rewritten as

$$\int_0^\tau \frac{dA_1}{b_1 A_1^2 + b_2 A_1 + b_3} = \tau,$$

where

$$(22) \quad b_1 = \frac{1-\gamma}{\gamma}$$

$$(23) \quad b_2 = 2 \left(\frac{\gamma-1}{\gamma} \sigma_X - \lambda_X \right)$$

$$(24) \quad b_3 = \frac{1}{\gamma} \sigma_X^2.$$

The solution for the integral can be found in integration tables. In Appendix B, it is shown that when $\gamma > 1$, $b_2^2 - 4b_1b_3 > 0$. Defining

$$\theta = \sqrt{b_2^2 - 4b_1b_3},$$

the solution is given by

$$(25) \quad A_1(\tau) = \frac{1-\gamma}{\gamma} \frac{2(1-e^{-\theta\tau})}{2\theta - (b_2 + \theta)(1-e^{-\theta\tau})},$$

$$(26) \quad A_2(\tau) = \frac{1-\gamma}{\gamma} \frac{4\lambda_X \bar{X} (1-e^{-\theta\tau/2})^2}{\theta[2\theta - (b_2 + \theta)(1-e^{-\theta\tau})]}.$$

The explicit solution for A_3 is more complicated and can be found by integrating a polynomial in A_1 and A_2 :

$$(27) \quad A_3(\tau) = \int_0^\tau \frac{1-\gamma}{2} \frac{1-\gamma}{\gamma} A_2^2 + \frac{1}{2} \sigma_X^2 A_1 + \lambda_X \bar{X} A_2 + (1-\gamma)r - \rho.$$

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). 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. Appendix B demonstrates that G has the required derivatives for the equation for wealth (15) and the portfolio rule (34) to be valid.

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

$$(28) \quad \frac{W_t}{c_t} = \int_0^{T-t} \exp \left\{ \frac{1}{\gamma} (A_1(\tau)X_t^2/2 + A_2(\tau)X_t + A_3(\tau)) \right\} d\tau.$$

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 Indirect Utility

While the indirect utility function is not required as an intermediate step in the martingale method, it is nonetheless an interesting quantity to compute. The indirect utility function can be computed directly using dynamic programming, as discussed in Section V. Alternatively, it can be derived from the results of the previous section, as shown below.

Cox and Huang (1989) show that the indirect utility function $J(W_t, X_t, t)$ and the function for wealth $G(Z_t, X_t, t)$ are related via the following formula:

$$(29) \quad \frac{\partial J}{\partial W} = \frac{1}{G^{-1}(W_t, X_t, t)} = \frac{1}{Z_t},$$

where G^{-1} is defined implicitly by $G^{-1}(G(Z_t, X_t, t), X_t, t) = Z_t$. From (21) and $W_t = G(Z_t, X_t, t)$, it follows that

$$\frac{\partial J}{\partial W} = W_t^{-\gamma} e^{-\rho t} \left(\int_0^{T-t} H(X_t, \tau) \right)^\gamma.$$

Finally, the boundary condition $J(W_T, X_T, T) = W_T^{1-\gamma}/(1-\gamma)$ implies

$$(30) \quad J(W_t, X_t, t) = \frac{W_t^{1-\gamma}}{1-\gamma} e^{-\rho t} \left(\int_0^{T-t} H(X_t, \tau) \right)^\gamma.$$

The indirect utility function is an alternative way of summarizing future investment opportunities. As (30) shows, the indirect utility function and the wealth-consumption ratio

are closely related: In fact, the term in the indirect utility function involving X_t is simply the wealth-consumption ratio raised to the power γ .

How is indirect utility assuming consumption related to indirect utility assuming terminal wealth? The indirect utility function (30) can be rewritten as:

$$(31) \quad J(W_t, X_t, t) = \frac{W_t^{1-\gamma}}{1-\gamma} e^{-\rho t} \left(\int_0^{T-t} (\exp \{A_1(\tau)X_t^2/2 + A_2(\tau)X_t + A_3(\tau)\})^{\frac{1}{\gamma}} d\tau \right)^\gamma.$$

By analogous reasoning, the indirect utility assuming terminal wealth is given by

$$(32) \quad \frac{W_t^{1-\gamma}}{1-\gamma} \exp \{A_1(\tau)X_t^2/2 + A_2(\tau)X_t + A_3(\tau)\}.$$

In contrast to the equation for wealth, indirect utility for consumption is not a simple sum of indirect utility for terminal wealth. Rather, for $\gamma > 1$, a concave transformation is applied to terms of the form (32). This concave transformation penalizes extreme values and rewards consistency.

D Optimal Portfolio Allocation

Following the martingale method, the optimal portfolio rule is derived from (21), the equation for optimal wealth. Alternatively, dynamic programming could be used with the investor's indirect utility (30). While the methods produce identical results, the former is more instructive.

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:

$$(33) \quad \alpha_t G \sigma = \frac{\partial G}{\partial Z} Z_t X_t - \frac{\partial G}{\partial X} \sigma_X.$$

The right-hand side follows from applying Ito's lemma to the function G , and from (2) and (14). The left hand side is the dollar amount invested in the risky asset multiplied the volatility.

Rearranging,

$$(34) \quad \alpha_t = \left(\frac{\partial G}{\partial Z} \frac{Z}{G} \right) \frac{X_t}{\sigma} - \left(\frac{\partial G}{\partial X} \frac{1}{G} \right) \frac{\sigma_X}{\sigma}.$$

It follows immediately from (21) that $\frac{\partial G}{\partial Z} \frac{Z}{G} = 1/\gamma$. 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 (21) leads to the equation

$$(35) \quad \alpha_t = \frac{1}{\gamma} \left(\frac{\mu_t - r}{\sigma^2} \right) - \frac{\sigma_X \int_0^{T-t} H(X_t, \tau) (A_1(\tau)X_t + A_2(\tau)) d\tau}{\int_0^{T-t} H(X_t, \tau) d\tau}.$$

The derivation of the investor's allocation under terminal wealth is similar, and easier. The allocation under terminal wealth equals:

$$(36) \quad \frac{1}{\gamma} \left(\frac{\mu_t - r}{\sigma^2} \right) - \frac{\sigma_X}{\gamma\sigma} (A_1(T-t)X_t + A_2(T-t)).$$

Together, equations (28) and (35) solve the investor's optimal consumption and portfolio choice problem. The economic consequences of these equations are explored in Section IV.

IV Discussion

This section derives consequences of (35) and provides economic insight into the solution. Section IV.A discusses how (35) 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.

A A Weighted Average Formula

At first glance, the difference between the allocation under terminal wealth and under consumption (35) 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 (35) take the form of a weighted average. The functions that are averaged equal hedging demand in (36) for different values of the horizon. The weights depend on the functions H .

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

$$(37) \quad \alpha_t = \int_0^{T-t} \frac{H(X_t, \tau)}{\int_0^{T-t} H(X_t, \tau') d\tau'} \left[\frac{1}{\gamma} \left(\frac{\mu_t - r}{\sigma^2} \right) - \frac{\sigma_X}{\gamma\sigma} (A_1(\tau)X_t + A_2(\tau)) \right] d\tau.$$

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

What does the function H represent? From (28), the ratio of wealth to consumption equals

$$(38) \quad \frac{W_t}{c_t} = \int_0^{T-t} H(X_t, \tau) d\tau.$$

Thus H is the value, scaled by today's consumption, of consumption in τ periods. The weights in (37) correspond to the value of future consumption in each period.

More will be said about the economics behind (37) in Section IV.C. However, an immediate economic implication of (37) and (38) 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.D, the first term in the optimal allocation (35) 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 (35) may first appear complicated, it can be used to prove a result on hedging demand that holds for any set of parameter values, such that X_t is perfectly negatively correlated with stock returns.

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 B, it is shown that $A_1(\tau)$ and $A_2(\tau)$ are negative when $\gamma > 1$. Property 1 follows from this result, from the equation for the optimal allocation (35), and from the fact that X_t has the same sign as the risk premium. \square

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 affects current consumption relative to wealth in two directions. By increasing investment opportunities, an increase in X_t allows the investor to afford more consumption. This is known as the income effect, and it causes consumption to rise relative to current wealth. 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.

This analysis ties hedging demand to the consumption-wealth ratio. When hedging demand is positive, the consumption-wealth ratio must be increasing in X_t ; 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 34). Thus, whenever hedging demand is positive, the wealth-consumption ratio is falling in X_t and the consumption-wealth ratio is rising in X_t . \square

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 are given in Table 1. Appendix D describes the details of this calibration. The integrals in (28) and (35) are computed via 40-point Gauss-Legendre quadrature. 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 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.

Hedging demand is also related to indirect utility. Merton (1971) shows that α_t is also characterized by

$$\alpha_t = \frac{1}{\gamma} \left(\frac{\mu_t - r}{\sigma^2} \right) + \frac{\partial^2 J}{\partial W \partial X} \left(\frac{\partial^2 J}{\partial W^2} \right)^{-1} \frac{\sigma_X}{\sigma},$$

where J is the indirect utility function. It is straightforward to show that this equation, combined with (30) produces the same function for α as given by (35). Thus, hedging demand can be interpreted as the investor's attempts to hedge marginal (indirect) utility of wealth. That is, the investor wants more wealth in states when the marginal utility is higher. This interpretation is closely related to the one given above: the states with high marginal utility are none other than those where the consumption-wealth ratio is low.

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), Brandt (1999), 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)*

Appendix B shows that $\frac{\partial A_1}{\partial \tau}$ and $\frac{\partial A_2}{\partial \tau}$ have negative sign when $\gamma > 1$. Property 3 follows from this and the equation for the optimal allocation (35). \square

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 quantitatively compares the allocation assuming utility over consumption with the allocation assuming utility over terminal wealth. The allocation for consumption and $T = 30$

lies below the allocation for terminal wealth and $T = 10$ as long as the risk premium is positive. The investor with utility over consumption and $T = 30$ has, in effect, a horizon of less than 10. The allocation for consumption is typically less than half of that for terminal wealth.

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 (28), it follows that

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

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

$$-\frac{\partial F}{\partial X} \frac{1}{F} = -(A_1(\tau)X_t + A_2(\tau))$$

in the case of terminal wealth and

$$-\frac{\partial G}{\partial X} \frac{1}{G} = -\frac{\int_0^{T-t} H(X_t, \tau)(A_1(\tau)X_t + A_2(\tau)) d\tau}{\int_0^{T-t} H(X_t, \tau) d\tau}$$

in the case of consumption. Comparing with the portfolio allocation (35) 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.³ The further out consumption goes, the higher the duration of the

³Because 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.

consumption stream and the greater the allocation. Property 5 is a natural consequence of this analogy.

Property 5 *For the investor with utility over consumption and $\gamma > 1$, the optimal allocation increases with the investment horizon as long as $\mu_t - r > 0$ and X_t is perfectly negatively correlated with S_t . 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,

$$(39) \quad -A_1(T-t) > -\frac{\int H(X_t, \tau) A_1(\tau) d\tau}{\int H(X_t, \tau) d\tau},$$

and the same for $A_2(\tau)$. 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 $-\int H(X_t, \tau) A_1(\tau) / \int H(X_t, \tau)$ equals

$$-\frac{H(X_t, T-t)}{\int_0^{T-t} H(X_t, \tau) d\tau} \left(A_1(T-t) - \frac{\int H(X_t, \tau) A_1(\tau) d\tau}{\int H(X_t, \tau) d\tau} \right).$$

From (39), the derivative is always positive. \square

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) notes this property in his numerical solution. 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 result of Barberis is strongly suggestive, it cannot demonstrate that convergence is guaranteed at all relevant parameter values. Based on the closed-form solutions for the portfolio choice rule (35), 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 B. Besides its inherent interest, an additional benefit of this result is that it allows the exact solution (35) 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, (36), 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 (2001). 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 (35) can help to understand the discrepancies between the log-linear and the actual solution. From (35), it follows that there are two ways changes in X_t can affect the portfolio rule. The first is directly, through $A_1(\tau)X_t + A_2(\tau)$, just as in the linear case. The second is indirectly, through changes in the weights $H(X_t, \tau)$. This first effect is what the change would be if the solution were actually linear, namely, if $H(X_t, \tau)$ were a constant:

$$(40) \quad \frac{d\alpha_t^{\text{lin}}}{dX_t} = \frac{1}{\gamma\sigma} \left(1 - \sigma_X \frac{\int H(X_t, \tau) A_1(\tau)}{\gamma \int H(X_t, \tau)} \right).$$

This term is always positive because $A_1(\tau) < 0$ when $\gamma > 1$ (Appendix B).

However, $H(X_t, \tau)$ is not a constant in X_t . The difference between the true derivative and (40) equals

$$(41) \quad \frac{d\alpha_t}{dX_t} - \frac{d\alpha_t^{\text{lin}}}{dX_t} = \frac{\sigma_X}{\gamma^2\sigma} \left[\left(\frac{\int H(X_t, \tau)(A_1(\tau)X_t + A_2(\tau))}{\int H(X_t, \tau)} \right)^2 - \frac{\int H(X_t, \tau)(A_1(\tau)X_t + A_2(\tau))^2}{\int H(X_t, \tau)} \right].$$

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”, (40), and the derivative itself (the sum of (40) and (41)) for $\gamma = 10$. The figure shows that (40) 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 $H(X_t, \tau)$ as the value of a bond paying in units of consumption at horizon τ . Increases in X_t cause the value of $H(X_t, \tau)$ to decrease more, the greater the value of τ . The portfolio weight α is given by a weighted average of $A_1(\tau)X_t + A_2(\tau)$, where the weights are like the values of discount bonds. Thus (41) arises because increases in X_t decrease the weights on the terms with higher values of $A_1(\tau)X_t + A_2(\tau)$.

V The Case of Incomplete Markets

It is easy to lose sight of where, in the above arguments, the assumption of complete markets is 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 (20), i.e. wealth is a martingale under the equivalent martingale measure. As Pliska (1986) points out, the defining feature of complete markets is that any (well-behaved) process that is a martingale under the equivalent martingale measure can be financed by trading in the underlying securities. When markets are incomplete, this equivalence no longer holds. Thus the set of consumption rules that can be financed is more difficult to describe, and the technique of solving first for consumption is no longer valid.

To illustrate where the method fails, let w_t is a two-dimensional Brownian motion. Following convention, w_t is a 2-by-1 vector, and σ and σ_X are 1-by-2 vectors. The instantaneous volatility of S_t is therefore $\sqrt{\sigma\sigma'}$. By analogy, $X_t = (\mu_t - r)/\sqrt{\sigma\sigma'}$. The instantaneous correlation between the process for S_t and that of X_t equals $\sigma\sigma'_X/(\sqrt{\sigma\sigma'}\sqrt{\sigma_X\sigma'_X})$. Any process η_t satisfying

$$(42) \quad \sigma\eta_t = (\mu_t - r).$$

is a valid price-of-risk process and, because this is a single equation in two unknowns, there are infinitely many solutions. Each value for η corresponds to a different state-price density ϕ , where

$$\frac{d\phi_t}{\phi_t} = -r dt - \eta_t dt.$$

By the same reasoning used to derive (20), the investor's wealth satisfies:

$$(43) \quad \mathcal{L}G + \frac{\partial G}{\partial t} + Z_t^{\frac{1}{\gamma}} e^{-\frac{\rho}{\gamma}t} - rG = \left(\frac{\partial G}{\partial Z} Z\eta' - \frac{\partial G}{\partial X} \sigma_X \right) \eta$$

with boundary condition $G(Z_T, X_T, T) = 0$.

It is tempting to choose η such that (42) is satisfied, and such that this equation can be solved in closed form. For example,

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

is a natural choice. Under this choice of η , the solution takes the same form as (21), where the functions A_i take slightly different, but nonetheless tractable forms.

Because this solution solves (43), it is a solution to the investor's static optimization problem. However, it does not correspond to the solution to the dynamic problem because there is no portfolio rule to finance it. The portfolio rule α_t must satisfy

$$(44) \quad \alpha_t G \sigma = \frac{\partial G}{\partial Z} Z_t \eta_t - \frac{\partial G}{\partial X} \sigma_X.$$

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 the variables are perfectly correlated. When there is a single asset, state-variable risk and market risk cannot be perfectly hedged at the same time.

There is a choice for η such that the two-dimensional system (44) can be reduced to one dimension. Choose

$$(45) \quad \eta = \sigma'(\sigma\sigma)^{-1}(\mu_t - r) + \nu,$$

where

$$\nu' = \frac{\partial G}{\partial X} \left(\frac{\partial G}{\partial Z} Z \right)^{-1} (\sigma_X - \sigma_X \sigma'(\sigma\sigma)^{-1} \sigma).$$

(45) represents the unique choice of η that solves (42), such that (44) can be reduced to one dimension. The equation for α becomes

$$\alpha_t = \frac{\partial G}{\partial Z} \frac{Z}{G} \frac{\mu_t - r}{\sigma\sigma'} - \frac{\partial G}{\partial Z} \frac{\sigma_X \sigma'}{\sigma\sigma'}$$

He and Pearson (1991) demonstrate rigorously that the static problem is equivalent to the dynamic problem under this choice for η .

This fix is less helpful than it first appears. When (45) is substituted into the differential equation for G , The solution can no longer take the form of (21). In fact, the problem under incomplete markets is much harder to solve.⁴

Can further progress be made in the incomplete markets case using the dynamic programming approach? The dynamic programming approach (like the martingale approach extended using the equations above) delivers a closed-form solution for the case of terminal wealth. But, as above, it is not possible to extend the solution to the case of consumption. The solution to the intertemporal problem by dynamic programming requires solving a nonlinear partial differential equation for the value function J . The differential equation is derived by Merton

⁴The case of utility over terminal wealth encounters no such difficulties, because the solution is the exponential of a quadratic function of X_t , rather than the integral of such functions.

(1971):

$$\begin{aligned}
(46) \quad & \frac{1}{1-\gamma} \left(\frac{\partial J}{\partial W} e^{\rho t} \right)^{\frac{\gamma-1}{\gamma}} e^{-\rho t} + \frac{\partial J}{\partial t} \\
& - \frac{1}{2} \left(-\frac{\partial J}{\partial W} \left(\frac{\partial^2 J}{\partial W^2} \right)^{-1} \frac{X_t}{(\sigma\sigma')^{\frac{1}{2}}} - \frac{\partial^2 J}{\partial W \partial X} \left(\frac{\partial^2 J}{\partial W^2} \right)^{-1} \left(\frac{\sigma\sigma'_X}{\sigma\sigma'} \right)^{\frac{1}{2}} \right)^2 \frac{\partial^2 J}{\partial W^2} \sigma\sigma' \\
& + \frac{\partial J}{\partial W} rW - \frac{\partial J}{\partial X} \lambda_X (\bar{X} - X) + \frac{1}{2} \frac{\partial^2 J}{\partial X^2} \sigma_X \sigma'_X = 0.
\end{aligned}$$

This equation has a solution of the form (30) if and only if $(\sigma\sigma'_X)^2 = (\sigma_X\sigma'_X)(\sigma\sigma')$, namely, if markets are complete.⁵

This result is not surprising. In this setting, the dynamic problem (solveable by dynamic programming) is equivalent to the static problem (solveable by the martingale method). Thus the portfolio and consumption policies must be the same, nomatter which method is applied. Moreover, Section C demonstrates that the indirect utility function can be derived from the wealth function and vica-versa. Thus a closed-form solution for indirect utility can be easily transformed into a closed-form solution for investor's wealth.

Despite the lack of a closed-form solution, equation (43) does show that the conclusions above are robust to small amounts of market incompleteness. While a thorough examination of this issue is beyond the scope of this paper, some informal conclusions can be drawn. Because $\nu'\nu$ and $\sigma_X\nu$ are smooth functions of the correlation, and because, and $\eta'\nu = \sigma\nu = 0$, (43) can be rewritten so that the coefficients are smooth functions of the correlation. Thus, for values of the correlation sufficiently close to -1, the solutions will be close to the solution for complete markets, and the intuition developed in the preceding sections will still apply.

VI 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 papers that solve the problem for terminal wealth, give numerical solutions for consumption, or derive 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

⁵By contrast, in the case of terminal wealth, the first term in the differential equation drops out. A solution of the form (32) is therefore possible, even when markets are incomplete. Liu (1999) gives a clear account of this result.

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 incorrectly ignored, 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. Recent examples include 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 because the correlation in the data is nearly perfect. 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

Throughout the Appendix γ is assumed to be greater than 1.

A Proof that solutions to (3) and (5) exist.

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,

$$\phi_t \leq \exp \left\{ - \int_0^t X_s dw_s - r t \right\}.$$

Because $\int_0^t X_s dw_s$ is Gaussian (see, e.g. Duffie, 1996, Appendix E), the right hand side is lognormal, implying that the expectation exists and is a continuous function of time.

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

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

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 Properties of the solution

Existence and differentiability.

For $G(Z_t, X_t, t)$, it is sufficient that $b_2^2 - 4b_1b_3 > 0$. It follows from (22) - (24) that

$$\begin{aligned} b_2^2 - 4b_1b_3 &= 4 \left(\lambda_X^2 - \left(\frac{1-\gamma}{\gamma} \right) (\sigma_X^2 - 2\sigma_X\lambda_X) \right) \\ &> 4 \left(\frac{\gamma-1}{\gamma} \lambda_X^2 + \frac{\gamma-1}{\gamma} \sigma_X^2 - \left(\frac{\gamma-1}{\gamma} \right) 2\sigma_X\lambda_X \right) \\ &= 4 \frac{\gamma-1}{\gamma} (\lambda_X - \sigma_X)^2 \geq 0 \end{aligned}$$

For the differential equation (20) to describe the solution, $\frac{\partial^2 G}{\partial X \partial Z}$, $\frac{\partial^2 G}{\partial Z^2}$, $\frac{\partial^2 G}{\partial X^2}$ and $\frac{\partial G}{\partial t}$ must exist and be continuous. From (21), it suffices that A_1 , A_2 and A_3 be continuously differentiable. From (27) A_3 is continuously differentiable as long as A_1 and A_2 are. Therefore, from (25) and (26), it suffices to show that

$$(47) \quad 2\theta - (b_2 + \theta)(1 - e^{-\theta\tau}) > 0$$

for $0 \leq \tau \leq T$, where $\theta = \sqrt{b_2^2 - 4b_1b_3}$. From (22) - (24), it follows that

$$\begin{aligned} \theta &\geq 2\sqrt{\frac{\gamma-1}{\gamma}} |\lambda_X - \sigma_X| \\ &> 2\frac{\gamma-1}{\gamma} (\lambda_X - \sigma_X) \\ &> 2 \left(\frac{\gamma-1}{\gamma} \lambda_X - \sigma_X \right) = b_2. \end{aligned}$$

(47) follows immediately.

Proof that the functions $A_1(\tau)$, $A_2(\tau)$ and their derivatives have sign $1 - \gamma$.

Define $a(\tau) = 2\theta - (b_2 + \theta)(1 - e^{-\theta\tau})$. As shown above, $a(\tau) > 0$. The statement for A_1 and A_2 follows immediately from (25) and (26). Moreover,

$$\begin{aligned} \frac{dA_1}{d\tau} &= \frac{1-\gamma}{\gamma} \left(\frac{2\theta e^{-\theta\tau} a(\tau) + 2(1 - e^{-\theta\tau})(b_2 + \theta)\theta e^{-\theta\tau}}{a(\tau)^2} \right) \\ &= \frac{1-\gamma}{\gamma} \left(\frac{2\theta^2 e^{-\theta\tau}}{a(\tau)^2} \right), \end{aligned}$$

and

$$\begin{aligned} \frac{dA_2}{d\tau} &= \frac{1-\gamma}{\gamma} \frac{4\lambda_X \bar{X}}{\theta} \left(\frac{(1 - e^{-\frac{\theta\tau}{2}})(e^{-\frac{\theta\tau}{2}})\theta a(\tau) + (1 - e^{-\theta\tau}2)^2(b_2 + \theta)e^{-\theta\tau}\theta}{a(\tau)^2} \right) \\ &= \frac{1-\gamma}{\gamma} \frac{4\lambda_X \bar{X}}{\theta} \left(\frac{e^{-\frac{\theta\tau}{2}}\theta(1 - e^{-\frac{\theta\tau}{2}}) \left(2\theta - (b_2 + \theta)(1 - e^{-\frac{\theta\tau}{2}}) \right)}{a(\tau)^2} \right). \end{aligned}$$

The same reasoning used to show (47), implies that $2\theta - (b_2 + \theta)(1 - e^{-\frac{\theta\tau}{2}}) > 0$.

Proof that the allocation converges at long horizons

It suffices to show that

$$\frac{\int_0^{T-t} H(X_t, \tau) A_1(\tau) d\tau}{\int_0^{T-t} H(X_t, \tau') d\tau'}$$

and the same equation for $A_2(\tau)$ converge to finite limits. First, $A_1(\tau)$ and $A_2(\tau)$ 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. Because $A(\tau)$ is negative and increasing,

$$0 \leq \frac{\int_0^{T-t} H(X_t, \tau) |A_1(\tau)| d\tau}{\int_0^{T-t} H(X_t, \tau) d\tau} \leq |A_1(T)|.$$

The corresponding inequality holds for $|A_2(T)|$. It follows immediately from (26) and (25), that $A_1(T)$ and $A_2(T)$ are bounded above.

C The discrete-time process

The joint process for S_t and X_t can be rewritten in vector form as

$$\begin{pmatrix} d \log S_t \\ dX_t \end{pmatrix} = \left(\begin{pmatrix} 0 & \sigma \\ 0 & -\lambda_X \end{pmatrix} \begin{pmatrix} \log S_t \\ X_t \end{pmatrix} + \begin{pmatrix} r - \frac{\sigma^2}{2} \\ \lambda_X \bar{X} \end{pmatrix} \right) dt + \begin{pmatrix} \sigma \\ -\sigma_X \end{pmatrix} dw_t$$

As shown in Duffie (1999, Appendix E), this linear equation differential equation implies that

$$\begin{pmatrix} \log S_t \\ X_t \end{pmatrix} \sim N \left(\begin{pmatrix} m_S(t) \\ m_X(t) \end{pmatrix}, \begin{pmatrix} V_{SS}(t) & V_{SX}(t) \\ V_{SX}(t) & V_{XX}(t) \end{pmatrix} \right),$$

where the means and the variances satisfy ordinary differential equations. The solutions for the mean is given by:

$$\begin{aligned} m_X(t) &= \bar{X} + e^{-\lambda_X t} (X_0 - \bar{X}) \\ m_S(t) &= S_0 + \left(\sigma \bar{X} + r - \frac{\sigma^2}{2} \right) t + \frac{\sigma}{\lambda_X} (1 - e^{-\lambda_X t}) (X_0 - \bar{X}). \end{aligned}$$

The solution for the variance-covariance matrix is given by:

$$\begin{aligned} V_{XX}(t) &= \frac{\sigma_X^2}{2\lambda_X} (1 - e^{-2\lambda_X t}) \\ V_{SX}(t) &= \frac{\sigma\sigma_X^2}{2\lambda_X^2} (1 - e^{-\lambda_X t})^2 - \frac{\sigma\sigma_X}{\lambda_X} (1 - e^{-\lambda_X t}) \\ V_{SS}(t) &= 2\sigma \int_0^t V_{12}(s) ds + \sigma^2 t \end{aligned}$$

When

$$(48) \quad \frac{\sigma_X}{2\lambda_X} < 1,$$

$V_{SX}(t) < 0$ for all values of t . Therefore, under (48), $V_{SS}(t) < \sigma^2 t$. That is, the variance of stock returns grows less than linearly with the horizon. This implies that stock returns are mean-reverting.

When $\lambda_X \approx 0$,⁶

$$(49) \quad m_X(t) \approx \bar{X} + (1 - \lambda_X t)(X_0 - \bar{X})$$

$$(50) \quad m_S(t) \approx S_0 + (\sigma X_0 + r)t.$$

$$(51) \quad V_{XX}(t) \approx \sigma_X^2 t$$

$$(52) \quad V_{SX}(t) \approx \frac{\sigma\sigma_X^2}{2}t^2 - \sigma\sigma_X t$$

$$(53) \quad V_{SS}(t) \approx \frac{\sigma^2\sigma_X^2}{3}t^3 - \sigma\sigma_X t^2 + \sigma^2 t.$$

D Parameter values

The parameters are taken from Barberis (2000, Table 2) for the 1952-1995 sample. This section describes how these parameters are translated into model parameters.

The discrete-time process for S_t and X_t is approximated by (49)-(53) above. The time interval $t = 1$ is taken to be a month. For $t = 1$, the first term in (52), and the first two terms in (53) are of second order (estimates of the volatilities σ and σ_X are on the order of .01). Therefore the approximate discrete-time process is given by:

$$(54) \quad \Delta(\log S_{t+1}) = \sigma X_t + r - \frac{1}{2}\sigma^2 + \sigma\epsilon_{t+1}$$

$$(55) \quad X_{t+1} = \lambda_X \bar{X} + (1 - \lambda_X)X_t - \sigma_X \epsilon_{t+1},$$

where $\epsilon_{t+1} \sim N(0, 1)$.

Barberis (2000) computes a VAR for stock returns and the dividend-price ratio. He finds that the correlation between these two series is -0.935. In what follows, this value is set to -1. With this modification, the VAR can be rewritten as:

$$(56) \quad \Delta(\log S_{t+1}) = a_{11} + a_{12}Y_t + \sigma w_{t+1}$$

$$(57) \quad Y_{t+1} = a_{21} + a_{22}Y_t - \sigma_Y w_{t+1}.$$

Comparing (54) with (56), it follows that X_t and Y_t are related by

$$Y_t = \frac{\sigma X_t + r - \frac{\sigma^2}{2} - a_{11}}{a_{12}}.$$

⁶For $\lambda_X \approx 0$, $\frac{1-e^{-\lambda_X t}}{\lambda_X} \approx t$.

Substituting into (57) produces

$$X_{t+1} = \frac{a_{12}a_{21}}{\sigma} + \frac{a_{22} - 1}{\sigma} \left(r - \frac{\sigma^2}{2} - a_{11} \right) + a_{22}X_t - \frac{a_{12}\sigma Y}{\sigma} w_{t+1}.$$

Comparing with (55) allows the underlying parameters to be written in terms of the regression parameters:

$$\begin{aligned} \lambda_X &= 1 - a_{22} \\ \sigma_X &= \frac{a_{12}\sigma Y}{\sigma} \\ \bar{X} &= \frac{a_{12}a_{21}}{\sigma(1 - a_{22})} + \frac{1}{\sigma} \left(a_{11} + \frac{\sigma^2}{2} - r \right). \end{aligned}$$

The riskfree rate is the average real return on the 3-month Treasury Bill over the sample period 1952-1995, in monthly terms.

Barberis (2000) assumes utility over terminal wealth, so by definition, there is no rate of time preference. Campbell and Viceira assume an annualized discount rate of 0.062, translating into a monthly rate of 0.0052. Table 1 describes the resulting parameter values.


References

- Balduzzi, P., and A. W. Lynch. “Transaction Costs and Predictability: Some Utility Cost Calculations.” *Journal of Financial Economics*, 52 (1999), 47–78.
- Barberis, N. “Investing for the Long Run when Returns are Predictable.” *Journal of Finance*, 55 (2000), 225–264.
- Brandt, M. W. “Estimating portfolio and consumption choice: A conditional Euler equations approach.” *Journal of Finance*, 54 (1999), 1609–1645.
- Brennan, M. J., E. S. Schwartz, and R. Lagnado. “Strategic Asset Allocation.” *Journal of Economic Dynamics and Control*, 21 (1997), 1377–1403.
- Brennan, M. J., and Y. Xia. “Dynamic Asset Allocation under Inflation.” (2001). Forthcoming, *Journal of Finance*.
- Campbell, J. Y., J. Cocco, F. Gomes, P. J. Maenhout, and L. M. Viceira. “Stock Market Mean Reversion and the Optimal Equity Allocation of a Long-lived Investor.” (2001). Harvard University.
- Campbell, J. Y., A. C. MacKinlay, and A. W. Lo. *The Econometrics of Financial Markets*. Princeton, NJ: Princeton University Press (1997).
- Campbell, J. Y., and R. Shiller. “The Dividend-price Ratio and Expectations of Future Dividends and Discount Factors.” *Review of Financial Studies*, 1 (1988), 195–227.
- Campbell, J. Y., and L. M. Viceira. “Consumption and Portfolio Decisions when Expected Returns are Time-varying.” *Quarterly Journal of Economics*, 114 (1999), 433–495.
- Chacko, G., and L. Viceira. “Dynamic Consumption and Portfolio Choice with Stochastic Volatility in Incomplete Markets.” (2000). Harvard University.
- Cox, J. C., and C.-F. Huang. “Optimal Consumption and Portfolio Policies when Asset Prices Follow a Diffusion Process.” *Journal of Economic Theory*, 39 (1989), 33–83.
- . “A Variational Problem Arising in Financial Economics.” *Journal of Mathematical Economics*, 20 (1991), 465–487.
- Cox, J. C., Jonathan E. Ingersoll, Jr., and S. A. Ross. “A Theory of the Term Structure of Interest Rates.” *Econometrica*, 53 (1985), 385–407.
- Duffie, D. *Dynamic Asset Pricing Theory*. Princeton, NJ: Princeton University Press (1996).

- Duffie, D., and R. Kan. “A Yield-factor Model of Interest Rates.” *Mathematical Finance*, 6 (1996), 379–406.
- Epstein, L., and S. Zin. “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns.” *Econometrica*, 57 (1989), 937–968.
- Fama, E. F., and K. R. French. “Business Conditions and Expected Returns on Stocks and Bonds.” *Journal of Financial Economics*, 29 (1989), 23–49.
- Harrison, M., and D. Kreps. “Martingales and Multiperiod Securities Markets.” *Journal of Economic Theory*, 20 (1979), 381–408.
- He, H., and N. D. Pearson. “Consumption and Portfolio Policies with Incomplete Markets and Short-sale Constraints: The Infinite Dimensional Case.” *Journal of Economic Theory*, 54 (1991), 259–304.
- Karatzas, I., J. P. Lehoczky, and S. E. Shreve. “Optimal Portfolio and Consumption Decisions for a Small Investor on a Finite Horizon.” *SIAM Journal of Control and Optimization*, 25 (1987), 1557–1586.
- Karatzas, I., and S. E. Shreve. *Methods of Mathematical Finance*. New York, NY: Springer-Verlag (1998).
- Kim, T. S., and E. Omberg. “Dynamic Nonmyopic Portfolio Behavior.” *Review of Financial Studies*, 9 (1996), 141–161.
- Liu, J. “Portfolio Selection in Stochastic Environments.” (1999). Stanford University.
- Mehra, R., and E. C. Prescott. “The Equity Premium: A Puzzle.” *Journal of Monetary Economics*, 15 (1985), 145–161.
- Merton, R. C. “Optimum Consumption and Portfolio Rules in a Continuous-time Model.” *Journal of Economic Theory*, 3 (1971), 373–413.
- . “An Intertemporal Capital Asset Pricing Model.” *Econometrica*, 41 (1973), 867–887.
- Pliska, S. R. “A Stochastic Calculus Model of Continuous Trading: Optimal Portfolios.” *Mathematics of Operations Research*, 11 (1986), 239–246.
- Poterba, J. M., and L. H. Summers. “Mean Reversion in Stock Prices: Evidence and Implications.” *Journal of Financial Economics*, 22 (1988), 27–59.
- Schroder, M., and C. Skiadas. “Optimal Consumption and Portfolio Selection with Stochastic Differential Utility.” *Journal of Economic Theory*, 89 (1999), 68–126.


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

FIGURE 1
Hedging Demand for the Investor with Utility over Consumption




Hedging demand as a function of the Sharpe ratio X_t for risk aversion $\gamma = 10$ (solid), $\gamma = 4$ (dash), and $\gamma = 1$ (dots), and for horizons $T = 30$, $T = 10$, and $T = 0$ years. The vertical lines indicate $\bar{X} \pm \frac{2\sigma_X}{\sqrt{1-(1-\lambda_X)^2}}$. Parameter values are derived from discrete-time counterparts in Appendix D and given in Table 1.

FIGURE 2
The Consumption-Wealth Ratio




The consumption-wealth ratio as a function of the Sharpe ratio X_t for risk aversion $\gamma = 10$ (solid), $\gamma = 4$ (dash), and $\gamma = 1$ (dots), and for horizons $T = 30, 10$, and 0 years. The vertical lines indicate $\bar{X} \pm \frac{2\sigma_X}{\sqrt{1-(1-\lambda_X)^2}}$. Parameter values are derived from discrete-time counterparts in Appendix D and given in Table 1.

FIGURE 3
Allocation as a Function of the Sharpe Ratio




Optimal allocation as a function of the Sharpe ratio X_t for utility over consumption (circles) and over terminal wealth, for horizons $T = 30, 10,$ and 0 years. Risk aversion $\gamma = 10$. The vertical lines indicate $\bar{X} \pm \frac{2\sigma_X}{\sqrt{1-(1-\lambda_X)^2}}$. Parameter values are derived from discrete-time counterparts in Appendix D and given in Table 1.

FIGURE 4
Allocation as a Function of Horizon



Optimal allocation as a function of horizon for utility over consumption (circles) and utility over terminal wealth. Risk aversion $\gamma = 10$. For the center plot, the Sharpe ratio $X = \bar{X}$. For the top and bottom plots, $X = \bar{X} + \frac{\sigma_X}{\sqrt{1-(1-\lambda_X)^2}}$ and $X = \bar{X} - \frac{\sigma_X}{\sqrt{1-(1-\lambda_X)^2}}$. Parameter values are derived from discrete-time counterparts in Appendix D and given in Table 1.

FIGURE 5
Non-Linearities in the Optimal Allocation



The derivative of the optimal allocation with respect to X_t (circles) compared to the derivative ignoring the effects of X_t on the value of future consumption (see Eq. 40), for horizons $T = 30$, 10 and 0 years. Risk aversion $\gamma = 10$. The vertical lines indicate $\bar{X} \pm \frac{2\sigma_X}{\sqrt{1-(1-\lambda_X)^2}}$. Parameter values are derived from discrete-time counterparts in Appendix D and given in Table 1.

TABLE 1
Parameter Values

Rate of time preference ρ	0.0052
Riskfree rate r	0.0014
Volatility of stock prices σ	0.0436
Volatility of X_t σ_X	0.0189
Mean reversion parameter λ_X	0.0226
Unconditional mean of X_t \bar{X}	0.0788

Parameters are calculated based on those in Barberis (2000) and Campbell and Viceira (1999). Appendix D describes the details of this calculation. All parameters are in monthly units.