

NBER WORKING PAPERS SERIES

ASSET RETURNS AND INTERTEMPORAL PREFERENCES

Shmuel Kandel

Robert F. Stambaugh

Working Paper No. 3633

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 1991

The authors are grateful to Andrew Abel, John Campbell, Ravi Jagannathan, Robert Lucas, Rajnish Mehra, Charles Plosser, two anonymous referees, and participants in workshops at UCLA, MIT, the University of Minnesota, the University of Pennsylvania, Stanford University, Tel Aviv University, and Washington University. The first author gratefully acknowledges support from the Batterymarch Fellowship. Portions of this work were completed while the first author was a visiting associate professor at the Wharton School of the University of Pennsylvania. This paper is part of NBER's research program in Financial Markets and Monetary Economics. Any opinions expressed are those authors and not those of the National Bureau of Economic Research.

NBER Working Paper #3633
February 1991

ASSET RETURNS AND INTERTEMPORAL PREFERENCES

ABSTRACT

A representative-agent model with time-varying moments of consumption growth is used to analyze implications about means and volatilities of asset returns as well as the predictability of asset returns for various investment horizons. A comparative-statics analysis using non-expected-utility preferences indicates that, although risk aversion is important in determining the means of both equity returns and interest rates, implications about the volatility and the predictability of equity returns are affected primarily by intertemporal substitution. Lower elasticities of intertemporal substitution are associated with greater variance in the temporary component of equity prices.

Shmuel Kandel
Recanati Graduate School
of Business Administration
Tel Aviv University
Ramat-Aviv 69978
Tel Aviv, ISRAEL

Robert F. Stambaugh
The Wharton School
University of
Pennsylvania
Philadelphia, PA 19104
and
National Bureau
of Economic Research
Cambridge, MA 02138

1. Introduction

During the past decade, much empirical scrutiny has focused on models that represent an asset's price as a representative agent's rational expectation of future payoffs multiplied by marginal rates of substitution between present and future consumption. In the standard model, the representative agent is assumed to maximize time-additive expected utility, $E_t(\sum_{\tau=0}^{\infty} \beta^{\tau} U(c_{t+\tau}))$, where c_t is consumption, $U(\cdot)$ is a current period utility function, and E_t denotes the expectation conditioned on information at time t [e.g., Lucas (1978)]. These models have been investigated using a variety of econometric approaches, but even the basic sample moments of asset returns seem to raise formidable challenges to the standard model. Mehra and Prescott (1985) conclude that the sample means of interest rates and equity returns for the period from 1889-1978 cannot be reconciled with the model. Grossman and Shiller (1981) conclude that, especially for the latter part of that period, the sample variance of equity returns is higher than what the model seems able to accommodate.

Restrictions on investor preferences are important in interpreting the evidence. For example, Mehra and Prescott (1985) assume that $U(\cdot)$ exhibits constant relative risk aversion equal to α , and they restrict their investigation to values of β less than unity and to values of α less than ten. Using the same specification of preferences, Grossman and Shiller (1981) refer to "implausibly high" estimates of α required to produce perfect foresight discounted dividends with volatility as high as that observed in stock prices.

When preference parameters are allowed to exceed traditionally held bounds, the standard time-additive model is capable of producing implied values for the basic moments of returns that match sample estimates. Kandel and Stambaugh (1990) show that holding β less than unity but relaxing the upper bound on α allows the time-additive model to approximate sample means

and variances of equity returns and interest rates. Benninga and Protopapadakis (1990) obtain a similar result holding α around ten but allowing β to exceed unity. Cecchetti, Lam, and Mark (1990c) also relax the restrictions on α and β and conclude that the time-additive model produces implications about basic moments of interest rates and equity returns that are consistent with the data. These studies all specify a time-varying distribution of consumption growth, and they treat equity as a levered claim on the endowment consumption stream.

As emphasized by Hall (1988), the preference framework of time-additive expected utility implies that the coefficient of relative risk aversion, α , is the reciprocal of the elasticity of intertemporal substitution, η . Thus, an upper bound on α entails a lower bound on η , but economists who view a given value of α as too high to be a plausible degree of risk aversion may not necessarily view $1/\alpha$ as an unreasonably low elasticity of intertemporal substitution. For example, Hall concludes from his empirical analysis that the elasticity of intertemporal substitution "is unlikely to be much above 0.1 and may well be zero," but he does not endorse a correspondingly high coefficient of relative risk aversion.

In this study, we analyze implications about various moments of asset returns using a class of non-expected-utility preferences modeled previously by Epstein and Zin (1989a, 1989b) and Weil (1989, 1990). These preferences allow us to investigate the separate roles of risk aversion (α) and intertemporal substitution (η) in producing implications about a given moment of asset returns. We find for example, that while both α and η are important in determining the means of equity returns and interest rates as well as the volatility of interest rates, the volatility of equity returns is determined primarily by η : lower values of η are associated with higher volatilities of

equity returns.

The separate roles played by α and η in determining moments of asset returns offer new insights into the extent to which various empirical challenges hinge on beliefs about reasonable values of preference parameters. For example, while discussions about reasonable upper bounds on risk aversion may be relevant for the "equity premium puzzle," such discussions appear to have little bearing on the "excess volatility" issue. Given the results of our investigation, the relevant discussion about preferences in the latter case would presumably center instead on whether certain values of η were too low to be economically plausible elasticities of intertemporal substitution.

Sample estimates of the predictability of equity returns, especially returns for multi-year investment horizons, also present challenges to asset pricing theory. For example, long-horizon returns exhibit seemingly large negative sample autocorrelations [e.g., Fama and French (1988a)], and some authors have suggested that "mean reversion" of these estimated magnitudes is not accommodated easily by standard models of rational asset pricing [e.g., Poterba and Summers (1988)]. Cecchetti, Lam, and Mark (1990a) find that, although the sample autocorrelations of long-horizon returns are below the true values implied by a version of the model with $\alpha = 1.7$, the estimates nevertheless lie close to the median of the simulated small-sample distribution. Kandel and Stambaugh (1989, 1990) show that the time-additive model implies autocorrelations as low as the sample estimates when α exceeds the traditional upper bounds.

We find that the specification of preferences affects implications about the predictability of asset returns primarily through η . That is, obtaining a given degree of predictability of returns depends primarily on specifying a sufficiently low value of η (simultaneously requiring high risk aversion in

the time-additive case). We show that, when the distribution of consumption growth rates varies through time, equity prices contain a temporary (stationary) component, and the persistence of the temporary component depends on the persistence in the moments of the consumption process. For a given consumption process, the volatility of the temporary price component is decreasing in η but fairly insensitive to α .

The pricing model is developed in section 2, and there we provide analytic expressions for the implied moments of asset returns to be examined. Section 3 constructs an initial numerical example in the time-additive case in which the implied moments of returns approximate empirical counterparts. This example specifies the stochastic process for consumption growth rates as well as the preference parameters α and β . Our objective here is not to perform statistical inference but simply to provide a reference point for the subsequent comparative-statics analysis. Section 4 conducts that analysis by recomputing the implied return moments for alternative values of α and η using the more general non-expected-utility preferences. Section 5 analyzes the behavior of the "temporary" component of equity prices and examines the role of persistence in the moments of consumption growth.

It appears that a fairly high value of relative risk aversion (α) is required to match the first moments of asset returns, especially when β is held below unity. In our numerical analysis, for example, we find that with β slightly below unity the first moments are matched for $\alpha = 29$. This value of α exceeds those traditionally viewed as economically plausible. In section 6 we reconsider several of the arguments often made against such high values and suggest that these arguments are less than persuasive. Section 7 briefly reviews our conclusions.

2. The Pricing Model

2.1 Intertemporal Preferences

The representative consumer's utility for future consumption is specified in the recursive form analyzed previously by Epstein and Zin (1989a, 1989b), Epstein (1988), Weil (1989, 1990), and Kocherlakota (1990), who build on the earlier work of Kreps and Porteus (1978) and Selden (1978). These preferences relax the standard assumptions that consumers maximize the expected value of a time-additive utility function. The infinitely lived consumer maximizes lifetime utility V_t , the recursive structure of which is given by

$$V_t = \left[c_t^{\frac{\eta-1}{\eta}} + \beta [E_t(V_{t+1}^{1-\alpha})]^{\frac{\eta-1}{\eta(1-\alpha)}} \right]^{\frac{\eta}{\eta-1}} \quad (1)$$

where $0 < \beta < 1$, $0 < \alpha \neq 1$, and $0 < \eta \neq 1$.¹

The parameter α can be interpreted as the coefficient of relative risk aversion in atemporal gambles: the certainty equivalent of V_{t+1} is computed as the expected value of a utility function with constant relative risk aversion α . The parameters η and β reflect intertemporal substitution and time preference. When future utility is deterministic, then β is the rate of time preference and η is the elasticity of intertemporal substitution. Throughout the paper we will refer to α as "risk aversion" and η as "intertemporal substitution."

We investigate properties of asset returns as determined by the Euler equation derived by Epstein and Zin (1989a) and Weil (1989),

¹Although (1) is not defined for $\alpha = 1$ and $\eta = 1$, these special cases can be included within the same general framework [e.g., Epstein and Zin (1989b) and Weil (1989)].

$$\beta^{\frac{\eta(\alpha-1)}{1-\eta}} E_t \left\{ (c_{t+1}/c_t)^{\frac{1-\alpha}{1-\eta}} (1 + R_{A,t+1,1})^{\frac{\alpha\eta-1}{1-\eta}} (1 + R_{k,t+1,1}) \right\} = 1 \quad (2)$$

where $R_{A,t+1,1}$ is the one-period rate of return from time t to time $t+1$ on the consumer's optimal portfolio (aggregate wealth) and $R_{k,t+1,1}$ is the rate of return on any asset k . In the case of time-additive expected utility, where $\alpha = 1/\eta$, equation (2) simplifies to the more familiar expression,

$$\beta E_t \{ (c_{t+1}/c_t)^{-\alpha} (1 + R_{k,t+1,1}) \} = 1 \quad (3)$$

Our numerical investigations in later sections will first analyze the special case in (3) and then consider the general form in (2).

2.2 The Consumption Process

We employ the endowment framework of Lucas (1978), wherein aggregate consumption equals aggregate output in each period. Let λ_{t+1} denote unity plus the one-period growth rate in the representative agent's consumption, i.e., $\lambda_{t+1} = c_{t+1}/c_t$. Hamilton's (1989) approach to modeling changes in regime is used to characterize the process for λ_{t+1} . Assume

(i) Conditional on information at time t , the logarithmic consumption growth rate, $\ln(\lambda_{t+1})$, is distributed normally with mean μ_t and standard deviation σ_t .

(ii) The pair (μ_t, σ_t) follows a joint stationary Markov process with a finite number of states, S . Let s_t denote the state for (μ_t, σ_t) , and let Φ denote the $S \times S$ transition matrix with (i, j) element

$$\phi_{ij} = \text{Prob}(s_{t+1} = j \mid s_t = i) \quad (4)$$

Let π denote the S-vector of steady-state probabilities.

(iii) Given s_t , the distribution of s_{t+1} is independent of λ_{t+1} (which is drawn from the distribution determined by s_t).

These assumptions imply that the economy follows a Markov process with an infinite number of states, each represented as (c, i) , where, at time t , $c = c_t$ and $i = s_t$. There is a continuum of values for consumption and its growth rate but only a finite number of values for i , which represents the state for the conditional moments of the consumption growth rate. We follow Mehra and Prescott (1985) in assuming stationarity of the consumption growth rate, not the level. Abel (1988) presents an asset pricing model with time-varying conditional moments in which the level of consumption is stationary. Versions of the above Markov regime-switching model for the growth rate have been used in expected-utility frameworks for asset pricing by Cecchetti, Lam, and Mark (1990a, b, c) and Kandel and Stambaugh (1989, 1990).

2.3 Asset Prices

We use this model to derive prices of various types of financial claims. As is shown in propositions 1 through 3 below, the risky assets considered have the property that, in state (c, i) , their prices are homogeneous of degree one in c . The prices of riskless bonds in state (c, i) are shown to depend only on i and not on c .

Proposition 1.² The price of aggregate wealth (the claim on total future consumption) when the economy is in state (c, i) at time t is given by

$$P_A(c, i) = c \cdot w(i) \quad (5)$$

² Proofs of the propositions are available from the authors upon request.

where $w(i)$ is the i -th element of the vector w satisfying

$$w(i)^{\frac{\eta(\alpha-1)}{1-\eta}} = \sum_{j=1}^S \phi_{ij} \cdot E(\lambda_{t+1}^{1-\alpha} | i) \cdot (\beta \cdot [1+w(j)])^{\frac{\eta(\alpha-1)}{1-\eta}} \quad (6)$$

and $E(\cdot | i)$ denotes the conditional expectation given that $s_t = i$.

It is important to note that the equilibrium valuation in (5) requires that the representative consumer's utility be finite, which in turn restricts choices of the model's parameters, α , η , β , Φ , and (μ_i, σ_i) , $i = 1, \dots, S$.³

Equity is defined as a claim on aggregate wealth net of a one-period risky bond; the risky bond pays at time $t+1$ either (i) a fraction θ of aggregate wealth at time t or (ii) total aggregate wealth plus consumption at time $t+1$, whichever is less. The fraction θ can be viewed as the degree of leverage in the economy.

Proposition 2. The price of levered equity when the economy is in state (c, i) at time t is given by

$$P_L(c, i) = q(i) \cdot c \quad (7)$$

where

$$q(i) = w(i) \cdot \left[\sum_{j=1}^S \beta^{\frac{\eta(\alpha-1)}{1-\eta}} \cdot \phi_{ij} \cdot [(1+w(j))/w(i)]^{\frac{1-\eta\alpha}{\eta-1}} \cdot E(\min[\lambda_{t+1}^{1-\alpha} \cdot (1+w(j)), \lambda_{t+1}^{-\alpha} \cdot \theta \cdot w(i)] | i, j) \right] \quad (8)$$

and $E(\cdot | i, j)$ denotes the conditional expectation given $s_t = i$ and $s_{t+1} = j$.

Since the expectation in (8) is conditioned on both i and j , the conditional

³Kocherlakota (1990) analyzes the existence and uniqueness of the valuation operator for these preferences in Lucas-type endowment economies.

expectation is simply that of the random variable λ_{t+1} . Thus, it is straightforward to compute this expectation by integrating the lognormal distribution, where $\ln \lambda_{t+1}$ has conditional moments (μ_i, σ_i) .

An N-period riskless bond is defined as a certain claim on one unit of consumption to be received N periods in the future. The prices of such bonds are given in the following proposition.

Proposition 3. The price of a riskless claim on one unit of consumption received at time t+N when the economy is in state (c, i) at time t is given by

$$P_F(c, i; N) = (\Psi^N \iota)_i \quad (9)$$

where ι denotes an S-vector of ones and Ψ is the $S \times S$ matrix with (i,j) element

$$\psi_{ij} = \beta^{\frac{\eta(\alpha-1)}{1-\eta}} \cdot \phi_{ij} \cdot E(\lambda_{t+1}^{-\alpha} | i) \cdot [(w(j)+1)/w(i)]^{\frac{\eta\alpha-1}{1-\eta}} \quad (10)$$

Since proposition 3 reveals that the prices of riskless bonds do not depend on c, we suppress the "c" in (9) and denote the price of an N-period riskless bond when the economy is in state (c, i) by $P_F(i; N)$.

2.4 Rates of Return

We now consider rates of return on riskless bonds and levered equity for investment horizons of various lengths. Let the S-vector $R_F(N)$ contain the yields on an N-period riskless bond in each of the S states for (μ_t, σ_t) .

The i-th element of $R_F(N)$ is given by

$$R_F(i; N) = P_F(i; N)^{-1/N} - 1 \quad (11)$$

The unconditional mean and variance of the N-period yield are given by

$$\bar{R}_F(N) = \pi' R_F(N) , \quad (12)$$

and

$$\bar{V}_F(N) = \pi' \left[[R_F(N) - \bar{R}_F(N) \iota] * [R_F(N) - \bar{R}_F(N) \iota] \right] , \quad (13)$$

where "*" denotes a Hadamard matrix product.⁴

From the definition of equity and proposition 2, the one-period rate of return on equity when the economy moves from state (c, i) at time t to state (cλ_{t+1}, j) at time t+1 can be written as

$$R_{L,t+1,1} = \frac{\max(0, \lambda_{t+1} [1 + w(j)] - \theta w(i))}{q(i)} - 1 \quad (14)$$

A special case of (14) occurs for $\theta = 0$, in which case the return simplifies to the one-period return on total aggregate wealth:

$$R_{A,t+1,1} = \frac{\lambda_{t+1} [1 + w(j)]}{w(i)} - 1 \quad (15)$$

The N-period return on levered equity, covering the investment horizon from time t to t+N, is defined as

$$R_{L,t+N,N} = \prod_{n=1}^N (1 + R_{L,t+n,1}) - 1 \quad (16)$$

It follows from (14) and the assumed consumption process that the conditional distribution of equity returns when the economy is in state (c, depends only on i, the state for (μ_t, σ_t) . The following propositions give analytical expressions for the conditional means and variances of returns fo

⁴If A and B are m×n matrices and A*B = C, then $c_{ij} = a_{ij} \cdot b_{ij}$.

various investment horizons.

Proposition 4. The conditional means of the N-period equity returns in the S states for (μ_t, σ_t) are given by the S-vector

$$E_L(N) = \Gamma^N \mu_t \quad (17)$$

where Γ^N denotes the N-th matrix power of Γ , Γ is an SxS matrix with (i, j) element

$$\gamma_{ij} = \phi_{ij} \frac{E(\max[0, \lambda_{t+1} \cdot (1 + w(j)) - \theta \cdot w(i)] | i, j)}{q(i)} \quad (18)$$

and $w(i)$ and $q(i)$ are defined in (6) and (8).

Proposition 5. The conditional variances of the N-period equity returns in each of the S states for (μ_t, σ_t) are given by the S-vector

$$V_L(N) = \Xi^N \mu_t - [(\Gamma^N \mu_t) * (\Gamma^N \mu_t)] \quad (19)$$

where Ξ is an SxS matrix with (i, j) element

$$\xi_{ij} = \phi_{ij} \frac{E((\max[0, \lambda_{t+1} \cdot (1 + w(j)) - \theta \cdot w(i) \cdot (1 + w(j))])^2 | i, j)}{[q(i)]^2} \quad (20)$$

and $w(i)$ and $q(i)$ are defined in (6) and (8).

As explained following proposition 2, it is straightforward to compute the conditional expectations in (18) and (20) by integrating the lognormal distribution for λ_{t+1} , where $\ln \lambda_{t+1}$ has conditional moments (μ_i, σ_i) .

Unconditional moments of returns are obtained by combining the

conditional moments of returns, $E_L(N)$ and $V_L(N)$, with the steady-state probabilities. The unconditional mean and variance of $R_{L,t+1,N}$, denoted as $\bar{E}_L(N)$ and $\bar{V}_L(N)$, are given by

$$\bar{E}_L(N) = \pi' E_L(N) \quad (21)$$

and

$$\bar{V}_L(N) = \pi' V_L(N) + \pi' \left[[E_L(N) - \bar{E}_L(N)] * [E_L(N) - \bar{E}_L(N)] \right] \quad (22)$$

The first-order autocorrelation of equity returns for various investment horizons can also be calculated from the above unconditional moments:

$$\begin{aligned} \text{corr}(R_{L,t+N,N}, R_{L,t,N}) &= \frac{E(1 + R_{L,t+N,2N}) - [E(1 + R_{L,t+N,N})]^2}{\text{var}(R_{L,t+N,N})} \\ &= \frac{\bar{E}_L(2N) - 2\bar{E}_L(N) - [\bar{E}_L(N)]^2}{\bar{V}_L(N)} \end{aligned} \quad (23)$$

We define the predictability of the N-period equity return as the ratio of the variance of the conditional expected return to the unconditional variance of the return. This ratio, denoted Ω_N , is calculated from the conditional and unconditional moments defined above:

$$\Omega_N = \frac{\pi' \left[[E_L(N) - \bar{E}_L(N)] * [E_L(N) - \bar{E}_L(N)] \right]}{\bar{V}_L(N)} \quad (24)$$

Note that Ω_N is the implied upper bound on the goodness of fit, R_N^2 , in any linear projection of the N-period equity return on predetermined variables. Specifically, $R_N^2 = \bar{\rho}_N \cdot \Omega_N$, where $\bar{\rho}_N$ is the multiple correlation coefficient between the conditional expected N-period equity return and the predetermined variables.

3. An Initial Example with Time-Additive Expected Utility

Our objective is to use the pricing model to investigate the properties of asset returns implied by alternative specifications of preferences and consumption processes. We begin by constructing an example in the more familiar framework of time-additive expected utility. This initial example serves as a point of departure for our analysis in section 4, where the separate effects of risk aversion and intertemporal substitution are examined using the generalized form of non-expected-utility preferences.

3.1 The Consumption Process

The stochastic process for consumption in the model is defined by the discrete-state Markov process for the conditional moments of consumption growth, (μ_t, σ_t) . We construct our example with a four-state process in which each moment can take two distinct values: μ_t equals μ^+ or μ^- and σ_t equals σ^+ or σ^- . For simplicity, we assume equal unconditional probabilities of the four states, and we assume that μ_t and σ_t evolve independently of each other. The transition matrix Φ is then specified by two values, ρ_μ and ρ_σ , the first-order autocorrelations of μ_t and σ_t . If the states, as numbered in ascending order, are (μ^+, σ^-) , (μ^-, σ^-) , (μ^+, σ^+) , and (μ^-, σ^+) , then

$$\Phi = .25 \left[\begin{array}{c} \left[\begin{array}{cc} 1 + \rho_\sigma & 1 - \rho_\sigma \\ 1 - \rho_\sigma & 1 + \rho_\sigma \end{array} \right] \otimes \left[\begin{array}{cc} 1 + \rho_\mu & 1 - \rho_\mu \\ 1 - \rho_\mu & 1 + \rho_\mu \end{array} \right] \end{array} \right] \quad (25)$$

The example is constructed so that a single period corresponds to one month. Panel A of table 1 reports the values for the monthly Markov process for (μ_t, σ_t) , and panel B reports the implied moments for simple annual growth rates. We specify the values $\mu^{+/-}$ and $\sigma^{+/-}$, the conditional moments for the monthly logarithmic growth rates, so that the implied mean and the

implied standard deviation for the annual simple growth rates equal those reported by Mehra and Prescott (1985) based on data for the period from 1889 through 1978.

The values for the consumption process are chosen not only to match the two unconditional annual moments of consumption growth but also to allow the model to produce desired implications about moments of asset returns. For example, many choices of μ^+ and μ^- are consistent with a given mean of the simple growth rate (which also depends on the σ 's). For a given choice of preferences, however, changing μ^+ and μ^- changes the implied distribution of asset returns, and some choices of μ^+ and μ^- may not even allow the equilibrium valuation in proposition 1.

In the example constructed, the conditional moments of consumption growth do not vary greatly in relation to realized growth rates. For example, annual growth rates have an implied R-squared of .0016 (the variance of the mean rate divided by the variance of the realized rate). Nevertheless, the implications about moments of asset returns will differ significantly from the case in which consumption growth rates are identically and independently distributed (i.i.d.).

3.2 The Equity Premium

To complete the specification of the numerical example in the case of time-additive utility, we must give values for the preference parameters α and β . In addition, since the "equity" in our model is levered, we must also specify the degree of leverage, θ . The example is constructed with $\alpha = 29$, $\beta = .9978$, and $\theta = 0.44$.

We select these parameters in order that, when coupled with the assumed consumption process, the implied average one-year interest rate (0.80%) and one-year equity return (6.18%) match the estimates used by Mehra and Prescott

as empirical benchmarks. (Since our model is specified on a monthly basis, a one-year return is a 12-period return.) Mehra and Prescott also report estimates of the standard deviations of the annual real interest rate and the equity return. When we match the equity-return standard deviation of 16.54%, we obtain an interest-rate standard deviation of 4.28%, somewhat less than the estimate of 5.67% reported by Mehra and Prescott. We return to this point below in our discussion of more general preferences.

Specifying relative risk aversion (α) equal to 29 violates the upper bound of 10 that Mehra and Prescott set for this parameter a priori. The issue of "high" relative risk aversion will be discussed in more detail in section 6. Matching the first moments of interest rates and equity returns with a high value of α (and $\beta < 1$) is not unique to the model of consumption growth considered here. In fact, Mehra and Prescott state that values of α greater than 10 allow one to match the two return averages by making small changes in their consumption process. We can demonstrate this claim with a simple example. If, in their model, the probability of remaining in the current state is set to 0.47 instead of 0.43, it is easily verified that the two return averages are matched for $\beta = 0.999$ and $\alpha = 34.23$.

3.3 Autocorrelations and Predictability of Equity Returns

Stambaugh (1986) and Fama and French (1988a) report autocorrelations of five-year equity returns ranging from -0.2 to -0.5. As Fama and French observe, these estimates are larger in absolute magnitude than those for shorter horizons, thus producing a U-shaped pattern of autocorrelations with respect to investment horizon. Similar U-shaped patterns appear in the two sets of estimated first-order autocorrelations plotted in figure 1. The first set of estimates uses monthly returns on the value-weighted portfolio of stocks on the New York Stock Exchange for the period from December 1926

through December 1985.⁵ The estimated autocorrelation is the slope coefficient in a regression of the current return on the lagged return, and the regression uses overlapping observations in the same manner as Fama and French (1988a).⁶ The second set of sample autocorrelations, shown for investment horizons of 12 months, 24 months, etc., uses annual returns on Standard & Poor's Composite Index for the period from 1891 through 1985 (again using regressions with overlapping observations).⁷

Figure 1 also displays the first-order autocorrelations of equity returns in the example of the pricing model [equation (23)]. The implied autocorrelations are negative and exhibit a U-shaped pattern, beginning at $-.08$ for a one-month horizon, declining to $-.21$ at a 30-month horizon, and then increasing toward zero for longer horizons. This pattern is obtained with persistence parameters $\rho_{\mu} = 0.94$ and $\rho_{\sigma} = 0.20$. As we demonstrate in section 5, alternative specifications of ρ_{μ} and ρ_{σ} can produce different patterns.

As indicated by the estimates shown in figure 1, sample autocorrelations of equity returns for short investment horizons, such as one month, are positive and in the range of 0.1 to 0.2 [e.g., Fama and Schwert (1977) and Lo and Mackinlay (1988)]. We were unable to find parameter specifications for the model that result in positive autocorrelations at short horizons but negative autocorrelations at longer horizons. The Markov regime-

⁵These data are obtained from the Center for Research in Security Prices (CRSP) at the University of Chicago.

⁶The estimates are not bias adjusted. Fama and French (1988a) report simulation evidence suggesting that the bias in the estimated autocorrelations is, in general, not severe when the true autocorrelations are similar to those displayed in the figure.

⁷These data are obtained from Wilson and Jones (1987) for the period prior to 1926 and from CRSP for the period thereafter.

switching process used here assumes that unexpected consumption growth does not impact the change in the conditional moments of consumption growth, and this could be important in restricting the model's ability to match the positive autocorrelations at short horizons. The independence between μ_t and σ_t used in this example does not appear to play a critical role. Kandel and Stambaugh (1990) obtain a similar pattern of implied autocorrelations from the time-additive expected-utility version of the model using a process in which μ_t and σ_t are correlated.

Figure 2 displays the implied predictability measure Ω_N for various investment horizons (N). We see that Ω_N is humped with respect to N , beginning at .038 for one-month returns, rising to .091 for 29-month returns, and then declining gradually toward zero for longer return horizons. Recall that Ω_N provides an upper bound on R_N^2 , the R-squared in a regression of the N -month return on predetermined variables.

We also show in figure 2 the sample values of R_N^2 obtained by regressing equity returns on three predetermined financial variables. The regressions use monthly observations, so the returns overlap for investment horizons greater than one month. Returns are computed using the value-weighted portfolio of NYSE stocks, and the predetermined variables have been used in previous studies to predict equity returns: (i) the dividend-price ratio for the NYSE value-weighted portfolio, (ii) the difference between Moody's Aaa yield and the yield on a U.S. Treasury Bill with maturity closest to one month, and (iii) the difference between Moody's Baa and Aaa yields.⁸ As Keim and Stambaugh (1986) observe using regressions with similar predetermined

⁸These three variables are the same as those used by Fama and French (1989), except that we were unable to obtain their corporate bond yields and instead used Moody's yield series. The dividend-price ratio is the sum of dividends for the previous twelve months divided by the price at the end of the most recent month.

variables, the R^2 is small for one-month equity returns--about .02 in this case--but, as Fama and French (1989) observe, the R^2 increases to more than .25 at a 48-month horizon.⁹

As discussed in section 2, R_N^2 is the product of Ω_N and $\bar{\rho}_N$, the multiple correlation coefficient between the conditional expected equity return and the predetermined variables. To gain some insight into the magnitude of $\bar{\rho}_N$ that might obtain in the equilibrium model analyzed here, we computed the multiple correlation coefficient between the expected equity return $E_L(i; N)$ and two predetermined variables, $dp(i) = 1/q(i)$ and $term(i) = [R_F(i; 240) - R_F(i; 1)]$. The first variable, $dp(i)$, is similar to a dividend-price ratio, except that aggregate consumption appears in place of dividends on (levered) equity. The variable $term(i)$ is the difference in yields between twenty-year and one-month bonds, but these are yields on real (index) bonds, not the nominal bonds generally used in empirical work. Since these two variables depend only on i when the economy is in state (c, i) , it is straightforward to compute their multiple correlation with $E_L(i; N)$. This value exceeds 0.99 for all investment horizons in the example constructed here and in the alternative cases constructed in section 4. Thus, it appears that the implied predictive ability from a projection of equity returns on a set of predetermined financial variables can be close to the upper bound Ω_N . Given that $dp(i)$ and $term(i)$ have only crude counterparts in the empirical work, however, we do not analyze implications about the individual partial correlation coefficients.

4. Separating Risk Aversion and Intertemporal Substitution

In the time additive case, since risk aversion (α) is linked to

⁹The sample R-squared values provide upward biased estimates of the true R-squared, due primarily to the autocorrelation in the residuals caused by the use of overlapping observations.

intertemporal substitution (η) by the relation $\alpha = 1/\eta$, the high risk aversion specified in the previous example implies a correspondingly low value of intertemporal substitution. In order to discover the role that each of these parameters plays in obtaining the desired moments of asset returns, we analyze in this section the general non-expected-utility versions of the preferences in (1). As discussed earlier, these preferences break the link between risk aversion and intertemporal substitution found in the time-additive case.

In this section we consider four values of risk aversion ($\alpha = 1/2, 2, 10, 29$) and four values for intertemporal substitution ($\eta = 2, 1/2, 1/10, 1/29$). Four of the sixteen combinations represent time-additive utility where $\alpha = 1/\eta$, with the pairing ($\alpha = 29, \eta = 1/29$) being that used in the example in the previous section. The consumption process, the rate of time preference ($\beta = 0.9978$), and the degree of leverage ($\theta = 0.44$) are the same as those used in that initial example.

4.1 First Moments of Equity Returns and Interest Rates

The implied unconditional moments of annual returns for each combination of α and η are presented in table 2. The third and fourth columns contain the implied unconditional means of the interest rate and the excess return on equity. These columns correspond to tables 1 and 2 of Weil (1989), who uses the Mehra-Prescott (1985) two-state Markov process for the realized consumption growth rate. For the pairs of α and η considered here and by Weil, the mean interest rate is decreasing in both α and η , whereas the mean equity premium is increasing in α and decreasing in η .

The last row of table 2 contains the example introduced in section 3, where the first moments are matched to the empirical benchmarks using a high value of α ($= 29$) in the framework of time-additive expected utility. Weil

also observes that a high value of α is required to match the first moments. He concludes, however, that although such a fit can be achieved when α and η are separated, specifying a very high α in the time-additive model necessarily specifies a very low η and thus leads to a "counterfactual prediction of an extremely high risk-free rate" (p. 409).¹⁰ While this analysis of the effects of α on the interest rate in the time-additive model is valid for low values of α , high values of α produce different effects. For sufficiently high values of α , the interest rate decreases in α .

The effects of α and η on the average interest rate can perhaps be seen most easily for the one-period rate in the special case where consumption growth rates are independently and identically distributed. Combining propositions 1 and 3 when there is only one state gives, after simplifying,

$$r_F = -\ln P_F(1; 1) = -\ln \beta + \frac{1}{\eta} \mu - \frac{1}{2} \left[\frac{\alpha(\eta + 1) - 1}{\eta} \right] \sigma^2. \quad (26)$$

In the special case of time-additive expected utility, the interest rate simplifies to

$$r_F = -\ln \beta + \alpha \mu - \frac{1}{2} \alpha^2 \sigma^2, \quad (27)$$

which is obtained by substituting $1/\alpha$ for η in equation (26).¹¹

¹⁰ A similar point is made by Kocherlakota (1990), who uses the same preferences and the Mehra-Prescott consumption process and concludes that a high equity premium also requires a high interest rate for α and $1/\eta$ between 0.5 and 11.5, whether or not $\alpha = 1/\eta$.

¹¹ One could alternatively begin with the familiar expression in equation (27) and then obtain equation (26) by applying theorem 2 from Kocherlakota (1990), who shows that with i.i.d. uncertainty the valuation operator for the recursive-utility economy can be obtained from that for the expected-utility economy by appropriately redefining the rate of time preference. The comparative statics results for changing α in these two economies will differ, however.

Equation (27) is a special case of the "mean-variance" representation of interest rates derived by Breeden (1986), while (26) presents a similar mean-variance characterization for non-expected-utility preferences. In either case, the mean (μ) term reflects the positive relation between expected consumption growth and the interest rate, and the slope of that relation is $1/\eta$. This role of intertemporal substitution in the relation between interest rates and expected consumption growth is well known [e.g., Hall (1988)] and, in the absence of the variance term, would indeed result in a positive relation between α and r_F in the time-additive case.

The negative relations between the interest rate and consumption uncertainty indicated by the variance (σ^2) terms in equations (26) and (27) essentially reflect the "precautionary savings" motive discussed by Leland (1968) and others, wherein an agent with convex marginal utility faced with consumption uncertainty saves as a precaution against future shortfalls.¹² Since consumption is exogenous in our endowment economy, a higher shadow price for savings is reflected as a lower interest rate. Equation (26) indicates that this negative variance effect is increasing in α and decreasing in η for $\alpha > 1$, and the coefficient is proportional to α^2 in the time-additive case in (27). Thus, as α increases, the sign of the relation between r_F and α in the time-additive case switches from positive to negative as the precautionary-savings motive overtakes the effect of positive expected consumption growth. It is clear that the low interest rates obtained with $\alpha = 29$ in table 2 arise from the negative effect of variance.

4.2 Volatility of Equity Returns

The implied unconditional standard deviations of the annual interest rate

¹²See Caballero (1990) for a recent analysis.

and the annual equity return are reported in the fifth and sixth columns of table 2. The implied volatility of equity returns decreases in η for the four values considered. For a given η , however, the implied equity volatility is virtually unaffected by α . Observe, for example, that a value of $\eta = 1/29$ implies an equity volatility approximately equal to the empirical benchmark of 16.5% for both $\alpha = 1/2$ and $\alpha = 29$. Thus, low intertemporal substitution appears to be the essential requirement for obtaining a sufficiently high volatility of equity returns for a given consumption process.

The high implied equity volatilities obtained with the low values of η are consistent with a strong preference for temporally smooth consumption. In those cases, changes in the conditional moments of consumption growth induce investors either to invest or to disinvest in order to smooth consumption. Prices of equity in each state must rise or fall sufficiently to offset these effects and thereby induce the representative investor to hold the endowed aggregate supply of equity. For example, an investor desiring smooth consumption will attempt to invest more when expected consumption growth is low, and thus equity prices are driven higher in those states. Such effects result in greater variation in equity prices across states and thus greater volatility of equity returns.

That η rather than α is important in determining equity volatility is only partially evident from previous studies. The underlying economic behavior has been described, but this behavior is often labeled as risk aversion. For example, LeRoy and LaCivita (1981) describe how a representative agent with a greater desire for smooth consumption will attempt to invest more in the high-endowment state and to disinvest in the low-endowment state, thereby driving a larger wedge between equity prices in the

two states.¹³ LeRoy and LaCivita identify an investor who desires smooth consumption as being risk averse. Michener (1982) points out that the key effect of risk aversion in determining the volatility of asset prices is the desire for smooth consumption.

Since previous studies have analyzed volatility in a framework of time-additive expected utility, it would seem to make little difference whether those discussions were formulated in terms of α or η . There is obviously no mathematical distinction, but previous discussions of volatility have often led researchers to ponder high risk aversion rather than low intertemporal substitution. For example, Grossman and Shiller (1981) conclude that an "implausibly high" value of risk aversion is required to obtain a series of discounted perfect-foresight dividends with volatility as high as that observed in actual stock prices since the early 1950's. LeRoy and LaCivita (1981, p. 546) conclude that "there remains the empirical question of whether the existing degree of risk aversion is sufficient to account for the observed dispersion in asset prices." If economic intuition about the reasonableness of high risk aversion does not necessarily provide intuition about the reasonableness of correspondingly low intertemporal substitution, then it would seem useful to focus on the characteristic of preferences that is more relevant to the empirical issue at hand.

4.3 Volatility of Interest Rates

In the framework of time-additive expected utility, LeRoy and LaCivita (1981) conclude that the volatility of interest rates is increasing in risk aversion, and their conclusion is easily examined in our model. It is straightforward to show, using proposition 3, that equation (27) extends to

¹³This occurs as long as risk aversion exceeds unity.

more than one state:

$$r_F(i) = -\ln P(i; 1) = -\ln \beta + \alpha \mu_i - \frac{1}{2} \alpha^2 \sigma_i^2 \quad (28)$$

Thus, in the time-additive case,

$$\text{var}(r_F) = \alpha^2 \left[\text{var}(\mu_i) + \frac{\alpha^2}{4} \text{var}(\sigma_i^2) - \alpha \text{cov}(\mu_i, \sigma_i^2) \right], \quad (29)$$

which is monotonically increasing in α if $\text{cov}(\mu_i, \sigma_i^2) \leq 0$. (Recall that, in the numerical example analyzed here, $\text{cov}(\mu_i, \sigma_i^2)$ is zero.) In the four cases of time-additive utility in table 2 (where $\alpha = 1/\eta$), the volatility of the interest rate increases in α . When σ_i^2 is constant, then $\text{var}(r_F) = \alpha^2 \text{var}(\mu_i)$, and the role of α (or $1/\eta$) corresponds to commonly held intuition. That is, an agent with a desire for smooth consumption will attempt to save when expected consumption growth is low (the "low" state) and dissave when expected growth is high (the "high state"). To induce the representative agent to consume the endowment each period, the interest rate must be lower in the low state and higher in the high state than would be required with less desire for smoothness. LeRoy and LaCivita (1981) present a similar argument, representing the desire for smoothness as risk aversion, while Hall (1988) characterizes a similar analysis explicitly in terms of intertemporal substitution.

In the more general non-expected-utility framework, the implied volatility of the interest rate appears to depend importantly on both α and η , and neither parameter has a monotonic effect across the ranges considered. For the lower values of α , the interest-rate volatility is decreasing in η , and this would seem to correspond to the common intuition outlined above in the simplest case. The highest interest-rate volatilities occur for the

highest values of risk aversion, and in those cases the volatility is increasing in η . For the level of risk aversion that appears necessary to match the desired first moments, that is $\alpha = 29$, we see that low intertemporal substitution is required to obtain both a sufficiently low interest-rate volatility and a sufficiently high equity-return volatility. Recall that, in the time-additive example in section 3, the implied volatility of the interest rate (4.3%) is slightly lower than the empirical benchmark (5.7%). By making η higher than $1/\alpha$ and making slight changes to the other parameters of the model, we can match the first and second moments of both the interest rate and the equity return.

4.4 Autocorrelations and Predictability of Equity Returns

The separate effects of α and η on the autocorrelations of equity returns are displayed in figure 3. For larger values of η , the autocorrelations are negative but close to zero for all horizons. As η declines, the pattern of autocorrelations becomes more U-shaped with respect to investment horizon. In all cases, the effects of α on the implied autocorrelations appear to be negligible. For example, the autocorrelations giving the U-shaped pattern in the time-additive example of section 3 (cf. figure 1) obtain whether $\alpha = 1/2$ or $\alpha = 29$. Similarly, the autocorrelations are essentially flat at zero for $\eta = 1/2$, whether $\alpha = 1/2$ or $\alpha = 29$. Thus, it appears that low intertemporal substitution is the key to generating the patterns of negative autocorrelations found in sample estimates.

Figure 4 displays the effects of α and η on the implied predictability measure Ω_N . Both α and η exert significant effects on Ω_N , although η appears to be more important. For $\eta = 1/2$, the implied Ω_N is virtually zero for all investment horizons, whether $\alpha = 1/2$ or $\alpha = 29$. Lower values of η produce larger values of Ω_N , particularly at the longer horizons, and α has a greater

effect in these cases. For a given η , Ω_N is decreasing in α .

A comparison of figures 3 and 4 reveals the extent to which lagged returns capture variation in expected returns. Consider, for example, the case in which $\alpha = 1/2$ and $\eta = 1/29$. As shown in figure 4, Ω_N begins at 0.02 for a one-month horizon and reaches a maximum of 0.23 for a thirty-month horizon. In contrast, the squared autocorrelation of the thirty-month return--the R^2 when regressing that return on its own lag--is only about .05. (The autocorrelation shown in figure 3 reaches a value of -0.22 for a thirty-month horizon.) Performing the same type of comparison for other combinations of α and η reveals that, in general, lagged returns reveal only a portion of the information about expected returns.

Intertemporal substitution appears to play a key role in determining the extent to which future equity returns can be predicted either by past returns or by complete knowledge of the current state. In essence, a lower value of η implies a greater degree of predictability. These patterns seem intuitively reasonable, in that an investor with a stronger preference for smooth consumption would be less inclined to alter his intertemporal pattern of consumption in order to exploit (and thereby reduce) the variation in expected returns.

Our analysis also demonstrates that the predictability in equity returns implied by low values of η can be considerably greater than the predictability of consumption growth rates. Recall from the discussion in section 3 that the predictability (maximum R-squared) for annual consumption growth rates in the process used here is less than 0.002. In contrast, figure 4 shows that the implied predictability of annual returns is about 0.04 for $\eta = 1/10$ and ranges from 0.07 to 0.17 for $\eta = 1/29$.

5. Permanent and Temporary Components of Equity Prices

Some additional insight into the volatility and the predictability of equity returns can be gained by analyzing the "permanent" and "temporary" components of equity prices. Taking the natural log of the equity price in (7) gives

$$p_t = \ln P_{L,t} = \ln c_t + \ln q_t \quad , \quad (30)$$

where time subscripts replace the previous notation in section 2. From the analysis in that section, it is clear that $\ln c_t$ is nonstationary and thus contains a "permanent" component, whereas $\ln q_t$ is stationary, or "temporary."

Fama and French (1988a) and Poterba and Summers (1988) consider similar decompositions of prices into nonstationary and stationary components. The intuition developed in their framework assigns distinct roles to the variance and the persistence of the stationary component. The greater is the variance of the stationary component, the greater are the variance, the predictability, and the negative autocorrelations of returns at some investment horizons. The horizons at which these effects are most evident depend on the persistence of the stationary component: high persistence in the stationary component makes its effects most evident at longer horizons.

In the standard framework described above, the nonstationary component is modeled as a random walk with increments that are uncorrelated at all leads and lags with changes in the stationary component. The nonstationary component here, $\ln c_t$, is not a random walk, and in general its increments are correlated with lagged changes in the stationary component, $\ln q_t$. Moreover, Fama and French (1988a) and Poterba and Summers (1988) consider a first-order autoregressive process for the stationary component, but $\ln q_t$ generally will not admit this type of linear representation. Although the properties of

the nonstationary and stationary components in (30) do not conform precisely to those of the standard framework, we wish to explore the extent to which the above intuition developed in that framework applies to our model.

We first consider the effects of changing the preference parameters α and η . Table 3 displays the implied standard deviations and autocorrelations of $\ln q_t$ for the various combinations of α and η considered earlier, where the consumption process is the same as that used in the previous sections. The standard deviation of $\ln q_t$ is increasing in α and decreasing in η , but the changes induced by decreasing η are much larger. The autocorrelations of $\ln q_t$ are virtually unaffected by changes in α for low values of η (1/10 and 1/29) or by changes in η for low values of α (1/2 and 2). For other values of the parameters the correlations are only slightly affected by changes in either α or η .

The results in table 3 indicate that the preference parameters affect the stationary component $\ln q_t$ primarily through an inverse relation between η and the variance of $\ln q_t$. This variance effect, when combined with the general intuition described above for stationary components of prices, is consistent with the previously discussed role of η in determining the volatility of equity returns (table 2) as well as the autocorrelations and the predictability measures of returns (figures 3 and 4). The inverse relation between η and the volatility of equity returns follows from the observation that lowering η increases the variance of the stationary component but does not affect its persistence. Recall from figures 3 and 4 that the autocorrelations are U-shaped and the predictability measures are humped for all of the combinations of α and η . In essence, lowering η raises the maximum magnitudes but does not significantly change the horizons at which they occur. This result also follows from the observation that lowering η raises the

variance of the stationary component but does not change its persistence.

Table 4 displays the standard deviations and the autocorrelations of the stationary component $\ln q_t$ implied by different amounts of persistence in the consumption moments μ_t and σ_t . Each of the parameters ρ_μ and ρ_σ is given three alternative values, 0.9, 0.7, and 0.0, while the preference parameters α and η are held constant at the values for the initial time-additive example in section 3. The standard deviation as well as the autocorrelations of $\ln q_t$ are increasing in both ρ_μ and ρ_σ , although the greater of ρ_μ and ρ_σ essentially appears to determine the autocorrelation of $\ln q_t$.

Figures 5 and 6 display the first-order autocorrelations and predictability measures of equity returns implied for five of the above combinations of ρ_μ and ρ_σ . The combinations with high persistence of the stationary component, where either $\rho_\mu = 0.9$ or $\rho_\sigma = 0.9$, produce patterns similar to those obtained in the initial example in section 3. For example, with $\rho_\mu = 0.9$ and $\rho_\sigma = 0.9$, the return autocorrelations begin at -0.05 for a one-month horizon, reach a minimum of -0.24 at a 19-month horizon, and then increase toward zero for longer horizons. The implied predictability measures in that case also reach a maximum at a 19-month horizon.

Lowering ρ_μ and ρ_σ tends to make the absolute magnitudes of the autocorrelations and the predictability measures larger at short horizons and smaller at long horizons. This effect conforms to the intuition discussed earlier: lowering ρ_μ and ρ_σ also lowers the persistence of the stationary component and thereby shortens the investment horizons at which the effects of the stationary component are most evident. For example, with $\rho_\mu = \rho_\sigma = 0.7$, the implied autocorrelations of returns reach their minimum of -0.23 at a horizon of only six-months. For $\rho_\mu = \rho_\sigma = 0$, the autocorrelations increase monotonically from -0.25 for a one-month horizon, and the predictability

measures decrease monotonically from a one-month value of 0.13.

6. On High Relative Risk Aversion

The value for relative risk aversion of $\alpha = 29$ used in the previous examples exceeds levels traditionally thought to be reasonable. As noted earlier, Mehra and Prescott (1985) restrict α to be ten or less in their investigation. A central argument for a lower value of α , cited by Mehra and Prescott and many others, relies on the estimated "price of risk," as computed by Friend and Blume (1975). Friend and Blume use the relation wherein ω , the fraction of wealth placed in risky assets, is given by

$$\omega = \frac{1}{\alpha} \cdot \left[\frac{E(R_S) - R_F}{\text{var}(R_S)} \right] \quad (31)$$

where R_S is the rate of return on an equity index. The price of risk, the bracketed quantity in (31), is defined as the ratio of the expected excess return on the risky asset (equity) to the variance of the risky asset's return. With an estimated price of risk of about 1.7, and the proportion of wealth in risky assets between 0.5 and 0.8, Friend and Blume arrive at an estimate of α between 2 and 3. While this calculation is appropriate in the case of time-additive expected utility where consumption growth rates (and thus returns) are i.i.d., equation (31) generally does not obtain in other settings.

In fact, the implied price of risk for equity in our original example in section 3 equals the sample estimate of 2.26 for the Mehra-Prescott period of 1889-1978, since the model is calibrated to match the average interest rate, the average equity return, and the variance of the equity return (cf. table 2). In other words, the Friend and Blume approach applied to interest rates and stock returns generated by our time-additive expected-utility example with

$\alpha = 29$ would lead one to infer that α is in the range of 3 to 4. Kocherlakota (1988) demonstrates a similar result using $\alpha = 13.7$ and a value of β greater than unity ($\beta = 1.14$ for annual periods).

A closely related argument often made against higher values of α is that investors with this degree of risk aversion would place fractions of wealth in the risky asset that would be lower than what we observe in practice. Based on equation (31), for example, with a price of risk equal to 2.26 and $\alpha = 29$, the fraction of wealth placed in risky assets would be only 0.08. Again, since (31) generally fails when returns are not i.i.d., this calculation is misleading. In an i.i.d. environment, an investor will (optimally) consume a fixed fraction of wealth, thereby equating the variances of the growth rates of consumption and wealth. A high α leads the investor to prefer a low variance of consumption (in the time-additive expected-utility case), thereby dictating a low variance of wealth and a low choice of ω . When returns are not i.i.d., however, optimal consumption may be smoothed relative to wealth and thereby be consistent with higher fractions of wealth in the risky asset than (31) would imply.¹⁴ A deeper analysis of this issue would require a model with endogenous investment in a risky production technology, rather than the exchange-type economy modeled here, but the pitfalls inherent in using (31) seem evident.

High α values might also appear to be inconsistent with patterns across countries in interest rates and average rates of consumption growth. A familiar relation from neoclassical growth theory expresses the interest rate r_F as

¹⁴ A related point is made by Black (1990), who demonstrates that relative risk aversion for the derived utility of wealth can be substantially less than relative risk aversion for the direct single-period utility of consumption.

$$r_F = \rho + \alpha E(\ln \lambda) \quad (32)$$

where ρ [$= -\ln \beta$] is the rate of time preference and λ is the growth rate of real consumption. Assume that α and ρ are identical across countries. Consider that, for the period from 1957 through 1987, average real consumption growth was 8.2% in Korea but only 3.2% in the U.S., and suppose that these sample means reflect expected growth rates.¹⁵ With $\alpha = 29$, equation (32) implies that this 5% difference in expected growth rates requires the Korean interest rate to be 145% higher than the U.S. interest rate.¹⁶ However, equation (32) ignores the effects of volatility of consumption growth discussed in section 4. For the same 1957-87 period, the standard deviations of annual growth rates are 4.4% in Korea and 2.3% in the U.S., a difference of 2.1%. If, for example, growth rates in a given country are independently lognormally distributed and agents maximize time-additive expected utility, then equation (27) gives the interest rate in each country (treating each country as an isolated endowment economy). With $\alpha = 29$, the difference in standard deviations would have to be about 4% to equate interest rates in the two countries (ignoring imprecision in the estimated means reported above). This simple exercise is certainly not intended to confirm a high value of α , but it does illustrate the pitfalls in arguments based on the more familiar expression in (32).

Inferences about α are perhaps most elusive when pursued in the introspective context of thought experiments. It seems possible in such experiments to choose the size of a gamble so that any value of α seems

¹⁵The data used in this example are obtained from various issues of International Financial Statistics published by the International Monetary Fund.

¹⁶We are grateful to Robert Lucas for suggesting this example.

unreasonable. To illustrate the difficulty, assume that an investor with constant relative risk aversion is faced with a gamble in which X dollars are won or lost with equal probability. Figures 7 and 8 illustrate, for small and large gambles, the relation between the fraction of X paid to avoid the gamble and the ratio of X to the investor's wealth. Consider an investor whose current wealth is \$75,000. On one hand, if $X = \$25,000$, representing 33% of the investor's wealth, then $\alpha = 30$ implies that the investor will pay about \$24,000 to avoid the gamble. This seemingly large payment suggests a lower value for risk aversion, such as $\alpha = 2$, which implies a more reasonable payment of \$8,333 to avoid this gamble.¹⁷ On the other hand, if $X = \$375$, representing 0.5% of the investor's wealth, then $\alpha = 2$ implies that the investor will pay only \$1.88 to avoid the gamble. This seemingly small payment suggests a higher value for risk aversion, such as $\alpha = 30$, which implies a more reasonable payment of \$28 to avoid the gamble.

We follow others in specifying preferences in which certainty equivalents for atemporal gambles are computed as expected utilities of functions exhibiting constant relative risk aversion. Although this framework may be problematic in the context of thought experiments involving both small and large gambles, the extent to which this framework is crucial to the implications of the asset pricing model is less clear. Some clues may lie in a recent study by Epstein and Zin (1989c). They obtain a higher implied equity premium using a specification of preferences in which calculations of certainty equivalents in atemporal gambles do not obey the von Neumann-Morgenstern axioms. In this alternative specification, small gambles command premiums similar to what high values of α imply in figure 7, but large gambles

¹⁷ Mankiw and Zeldes (1990) offer such an example as an argument against high values of α .

command lower premiums than what such high values of α imply in figure 8. Their results, when viewed together with ours, suggest that the key ingredient in obtaining a high equity premium is a high aversion to small gambles, such as in the graph for $\alpha = 30$ in figure 7, whether or not the aversion to large gambles is as high as the graph for $\alpha = 30$ in figure 8.

7. Conclusions

We derive implications about asset returns in a model where moments of consumption growth vary through time and a representative agent maximizes non-expected-utility intertemporal preferences. These preferences allow us to conduct a comparative-statics analysis that separates the roles of risk aversion, α , and intertemporal substitution, η , in determining asset prices and moments of returns.

The average interest rate and equity premium both decrease with η , so as Weil (1989) concludes, lowering η for a given α does not offer an advantage in matching the Mehra and Prescott (1985) empirical benchmarks. We also observe, however, that increasing α for a given η raises the equity premium while lowering the interest rate, and these effects allow the model to match both of the empirical benchmarks even in the special case of time-additive expected utility. High risk aversion depresses the price of risky assets, but it raises the price of safe ones through a precautionary savings motive. The value of risk aversion required to match both moments can be high by traditional standards: we use a value of $\alpha = 29$ in our example. Discarding such values as unreasonable, however, is more difficult than suggested by arguments often offered in this regard.

Separating α and η reveals distinctly different roles for these parameters in determining the volatility of equity returns. Equity volatility is decreasing in η , and we find that specifying the relatively low

value of $\eta = 1/29$ in our model gives an implied volatility equal to the empirical benchmark of Mehra and Prescott. In essence, a stronger preference for temporally smooth consumption increases the volatility of wealth that is consistent with a given volatility of consumption. Previous discussions of equity volatility often assign a key role to risk aversion, but we find that varying α over a wide range produces virtually no effect on the implied volatility of equity returns.

For a given consumption process, the extent to which returns can be predicted by knowledge of either past returns or conditional expected returns depends primarily on η . A lower value of η implies a greater degree of predictability, and we demonstrate that this effect can be viewed as arising from a temporary price component whose volatility decreases in η . We also find that greater persistence in either the conditional mean or the conditional volatility of consumption growth produces greater persistence in the temporary component, thereby lengthening the investment horizons at which the effects of the temporary component are most evident.

The stochastic process for consumption growth rates in our analysis includes time-varying first and second moments, but the variation in these moments is modest and would probably be difficult to detect in small samples. For example, variation in the expected annual growth rate accounts for less than one percent of the variation in the realized annual growth rate. Nevertheless, a low value of η produces equity returns that are substantially more volatile and predictable than would be consistent with an i.i.d. process for consumption growth rates.

REFERENCES

- Abel, Andrew B., 1988, "Stock Prices under Time-Varying Dividend Risk: An Exact Solution in an Infinite-Horizon General Equilibrium Model," Journal of Monetary Economics 22, 375-394.
- Benninga, Simon and Aris Protopapadakis, 1990, "Leverage, Time Preference, and the Equity Premium Puzzle," Journal of Monetary Economics 25, 49-58.
- Black, Fischer, 1990, "Mean Reversion and Consumption Smoothing," Review of Financial Studies 3, 107-114.
- Breeden, Douglas T., 1986, "Consumption, Production, Inflation, and Interest Rates: A Synthesis," Journal of Financial Economics 16, 3-39.
- Caballero, Ricardo J., 1990, "Consumption Puzzles and Precautionary Savings," Journal of Monetary Economics 25, 113-136.
- Cecchetti, Stephen G., Pok-sang Lam, and Nelson C. Mark, 1990a, "Mean Reversion in Equilibrium Asset Prices," American Economic Review 80, 398-418.
- Cecchetti, Stephen G., Pok-sang Lam, and Nelson C. Mark, 1990b, "Evaluating Empirical Tests of Asset Pricing Models: Alternative Interpretations," American Economic Review 80, 48-51.
- Cecchetti, Stephen G., Pok-sang Lam, and Nelson C. Mark, 1990c, "The Equity Premium and the Risk Free Rate: Matching the Moments," working paper, Ohio State University
- Epstein, Larry G., 1988, "Risk Aversion and Asset Prices," Journal of Monetary Economics 22, 177-192.
- Epstein, Larry G. and Stanley E. Zin, 1989a, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," Econometrica 57, 937-969.
- Epstein, Larry G. and Stanley E. Zin, 1989b, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis," forthcoming, Journal of Political Economy.
- Epstein, Larry G. and Stanley E. Zin, 1989c, " 'First-Order' Risk Aversion and the Equity Premium Puzzle," Working Paper, University of Toronto and Carnegie-Mellon University.
- Fama, Eugene F. and Kenneth R. French, 1988a, "Permanent and Temporary Components of Stock Prices," Journal of Political Economy 96, 246-273.
- Fama, Eugene F. and Kenneth R. French, 1988b, "Dividend Yields and Expected Stock Returns," Journal of Financial Economics 22, 3-25.
- Fama, Eugene F. and Kenneth R. French, 1989, "Business Conditions and Expected Returns on Stocks and Bonds," Journal of Financial Economics 25, 23-49.

- Fama, Eugene F. and G. William Schwert, 1977, "Asset Returns and Inflation," Journal of Financial Economics 5, 115-146.
- Friend, Irwin and Marshall E. Blume, 1975, "The Demand for Risky Assets," American Economic Review 65, 900-922.
- Grossman, Sanford and Robert J. Shiller, 1981, "The Determinants of the Variability of Stock Market Prices," American Economic Review 71, 222-227.
- Hall, Robert E., 1988, "Intertemporal Substitution in Consumption," Journal of Political Economy 96, 339-357.
- Hamilton, James D., 1989, "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," Econometrica 57, 357-384.
- Kandel, Shmuel, and Robert F. Stambaugh, 1989, "Modeling Expected Stock Returns for Long and Short Horizons," working paper, University of Chicago and University of Pennsylvania.
- Kandel, Shmuel, and Robert F. Stambaugh, 1990, "Expectations and Volatility of Consumption and Asset Returns," Review of Financial Studies 3, 207-232.
- Keim, Donald B. and Robert F. Stambaugh, 1986, "Predicting Returns in the Stock and Bond Markets," Journal of Financial Economics 17, 357-390.
- Kocherlakota, Narayana R., 1988, "In Defense of the Time and State Separable Utility-Based Asset Pricing Model," Working Paper, Northwestern University.
- Kocherlakota, Narayana R., 1990, "Disentangling the Coefficient of Relative Risk Aversion from the Elasticity of Intertemporal Substitution: An Irrelevance Result," Journal of Finance 45, 175-190.
- Kreps, David M. and Evan L. Porteus, 1978, "Temporal Resolution of Uncertainty and Dynamic Choice Theory," Econometrica 46, 185-200.
- Leland, Hayne E., 1968, "Savings and Uncertainty: The Precautionary Demand for Savings," Quarterly Journal of Economics 82, 465-473.
- Leroy, Stephen F. and C. J. LaCivita, 1981, "Risk Aversion and the Dispersion of Asset Prices," Journal of Business 54, 535-547.
- Lo, Andrew W. and A. Craig Mackinlay, 1988, "Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test," Review of Financial Studies 1, 137-158.
- Lucas, Robert E. Jr., 1978, "Asset Prices in an Exchange Economy," Econometrica 46, 1429-1445.
- Mankiw, N. Gregory and Stephen P. Zeldes, 1990, "The Consumption of Stockholders and Non-Stockholders," Working Paper, Harvard University and the University of Pennsylvania.

- Mehra, Rajnish and Edward C. Prescott, 1985, "The Equity Premium: A Puzzle," Journal of Monetary Economics 15, 145-162.
- Michener, Ronald W., 1982, "Variance Bounds in a Simple Model of Asset Pricing," Journal of Political Economy 90, 166-175.
- Poterba, James A. and Lawrence H. Summers, 1988, "Mean Reversion in Stock Prices: Evidence and Implications," Journal of Financial Economics 22, 27-59.
- Selden, Larry, 1978, "A New Representation of Preferences Over 'Certain x Uncertain' Consumption Pairs: The 'Ordinal Certainty Equivalent' Hypothesis," Econometrica 46, 1045-1060.
- Stambaugh, Robert F., 1986, "Discussion," Journal of Finance 41, 601-602.
- Weil, Philippe, 1989, "The Equity Premium Puzzle and the Riskfree Rate Puzzle," Journal of Monetary Economics 24, 401-421.
- Weil, Philippe, 1990, "Nonexpected Utility in Macroeconomics," Quarterly Journal of Economics 105, 29-42.
- Wilson, Jack W. and Charles P. Jones, 1987, "A Comparison of Annual Common Stock Returns: 1871-1925 with 1926-85," Journal of Business 60, 239-258.

Table 1

The Consumption Process Used in the Analysis

A. The Markov Process for the Conditional Moments
of the Monthly Growth Rate of Consumption

State	Logarithmic Growth Rate (%)		Probability of moving to state				Unconditional Probability
	Mean (μ)	Std. Dev. (σ)	1	2	3	4	
1	0.162	0.875	0.582	0.018	0.388	0.012	0.25
2	0.130	0.875	0.018	0.582	0.012	0.388	0.25
3	0.162	1.130	0.388	0.012	0.582	0.018	0.25
4	0.130	1.130	0.012	0.388	0.018	0.582	0.25

Autocorrelation of μ : 0.94Autocorrelation of σ : 0.20Correlation between μ and σ : 0

B. Implied Unconditional Moments for Annual Series

	Mean (%)	Standard Deviation (%)	First-Order Autocorrelation
Growth Rate (Simple)	1.830	3.570	0.001
Conditional Expected Growth Rate	1.830	0.142	0.476
Conditional Standard Deviation of the Growth Rate	3.567	0.047	0.005

Table 2

Implied Moments of Annual Returns

Preference Parameters		Unconditional Mean (%)		Unconditional Std. Deviation (%)	
Risk Aversion	Intertemporal Substitution	Interest Rate	Equity Premium	Interest Rate	Equity Return
1/2	2	2.45	0.11	0.07	6.42
1/2	1/2	5.23	0.12	0.13	6.59
1/2	1/10	21.34	0.21	0.90	8.62
1/2	1/29	70.08	0.77	1.11	16.63
2	2	2.31	0.45	0.49	6.43
2	1/2	4.94	0.46	0.29	6.59
2	1/10	20.12	0.59	0.51	8.58
2	1/29	65.44	1.27	0.83	16.47
10	2	1.63	2.19	2.74	6.49
10	1/2	3.46	2.22	2.52	6.61
10	1/10	13.84	2.47	1.59	8.43
10	1/29	42.80	3.52	0.65	15.91
29	2	0.40	5.96	8.27	6.65
29	1/2	0.43	6.74	8.06	6.67
29	1/10	0.56	5.82	6.94	8.24
29	1/29	0.80	6.17	4.28	16.54

Note: The consumer is assumed to maximize lifetime utility V_t , the recursive structure of which is given by

$$V_t = \left[c_t^{\frac{\eta-1}{\eta}} + \beta [E_t(V_{t+1}^{1-\alpha})]^{\frac{\eta-1}{\eta(1-\alpha)}} \right]^{\frac{\eta}{\eta-1}}$$

where risk aversion is denoted by α , and intertemporal substitution is denoted by η . The logarithmic consumption growth rate obeys a conditional normal distribution with moments following the Markov process given in panel A of table 1. All of the cases set the rate of time preference $\beta = .9978$ (monthly), and equity is defined with the leverage parameter $\theta = 0.44$. (That is, levered equity accounts for roughly sixty percent of aggregate wealth.)

Table 3

Behavior of the Stationary Component of Equity Prices
for Alternative Preference Parameters

$$\ln P_t = \ln c_t + \ln q_t$$

Preference Parameters		Behavior of Stationary Component ($\ln q_t$) ^a				
Risk Aversion (α)	Intertemporal Substitution (η)	$\sigma(\ln q_t)$ ^c	$\rho(\ln q_t, \ln q_{t-j})$ ^b			
			j=1	j=12	j=24	j=48
1/2	2	0.1	.940	.476	.226	.051
1/2	1/2	0.2	.940	.476	.226	.051
1/2	1/10	1.8	.940	.476	.226	.051
1/2	1/29	4.1	.940	.476	.226	.051
2	2	0.1	.940	.476	.226	.051
2	1/2	0.2	.940	.476	.226	.051
2	1/10	1.8	.940	.476	.226	.051
2	1/29	4.2	.940	.476	.226	.051
10	2	0.1	.922	.464	.221	.050
10	1/2	0.2	.937	.474	.226	.051
10	1/10	1.9	.940	.476	.226	.051
10	1/29	4.8	.940	.476	.226	.051
29	2	0.2	.801	.386	.184	.042
29	1/2	0.2	.918	.462	.220	.050
29	1/10	2.2	.939	.475	.226	.051
29	1/29	6.9	.937	.474	.226	.051

^a P_t denotes the price of equity. The log of consumption, $\ln c_t$, obeys the process described in table 1.

^bThe j-th order autocorrelation is for j months.

^cThe numbers in this column are multiplied by 100.

Table 4

Behavior of the Stationary Component of Equity Prices
for Alternative Degrees of Persistence in the
Conditional Moments of Consumption Growth

$$\ln P_t = \ln c_t + \ln q_t$$

Autocorrelations of the Consumption Moments		Behavior of Stationary Component $(\ln q_t)^a$				
		$\sigma(\ln q_t)^c$	$\rho(\ln q_t, \ln q_{t-j})^b$			
ρ_μ	ρ_σ			j-1	j-12	j-24
0.9	0.9	9.9	.900	.282	.080	.006
0.9	0.7	4.7	.845	.209	.058	.005
0.9	0.0	4.0	.899	.282	.080	.006
0.7	0.9	9.1	.897	.278	.079	.006
0.7	0.7	2.7	.700	.014	.000	.000
0.7	0.0	1.1	.686	.014	.000	.000
0.0	0.9	9.0	.900	.282	.080	.006
0.0	0.7	2.5	.699	.014	.000	.000
0.0	0.0	0.2	.000	.000	.000	.000

^a P_t denotes the price of equity. The preference parameters are held constant at $\alpha = 29$ and $\eta = 1/29$.

^bThe j-th order autocorrelation is for j months.

^cThe numbers in this column are multiplied by 100.

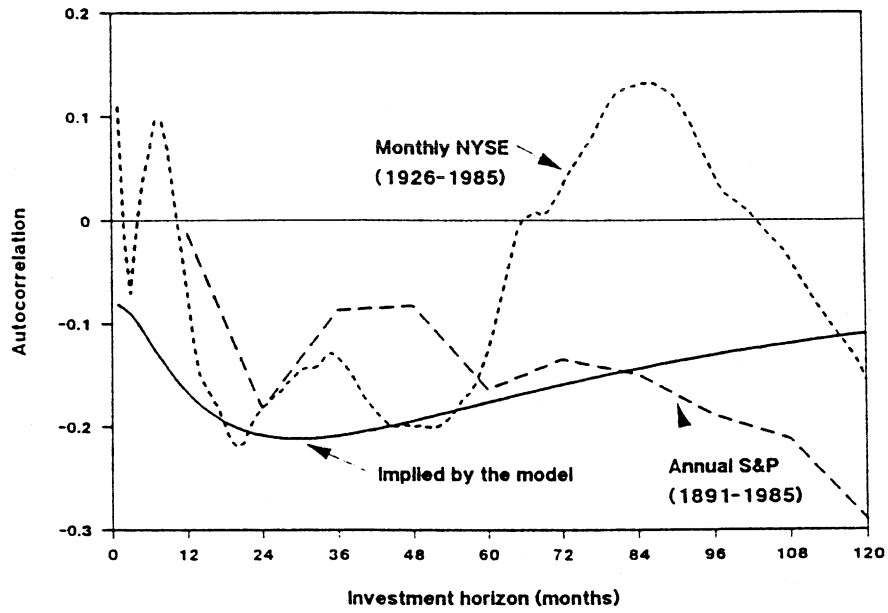


Figure 1. First-order autocorrelations of equity returns for various investment horizons.

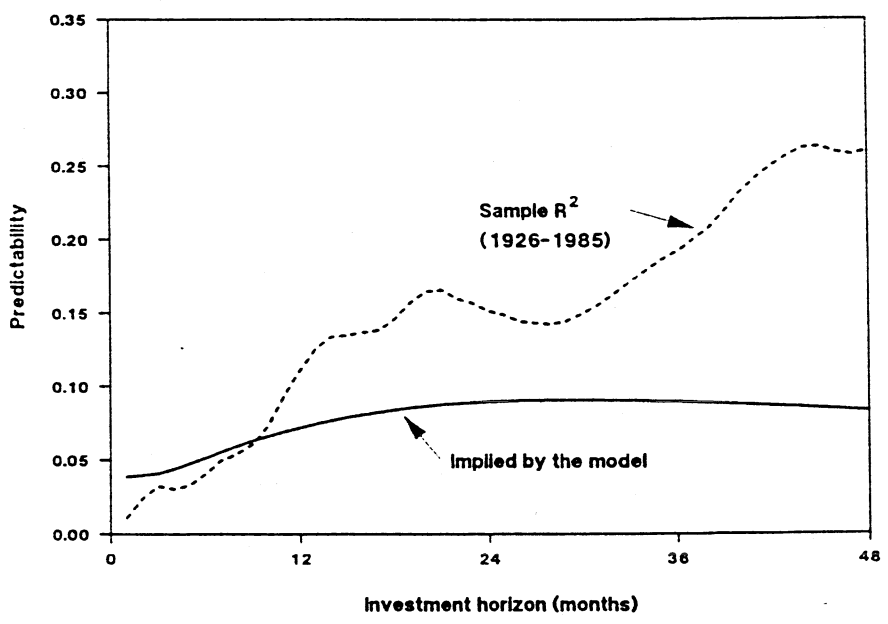


Figure 2. The predictability measure Ω_N implied by the model for various investment horizons, and the R^2 in regressions of the value-weighted NYSE return for various investment horizons on a dividend-price ratio, a term spread, and a default spread.

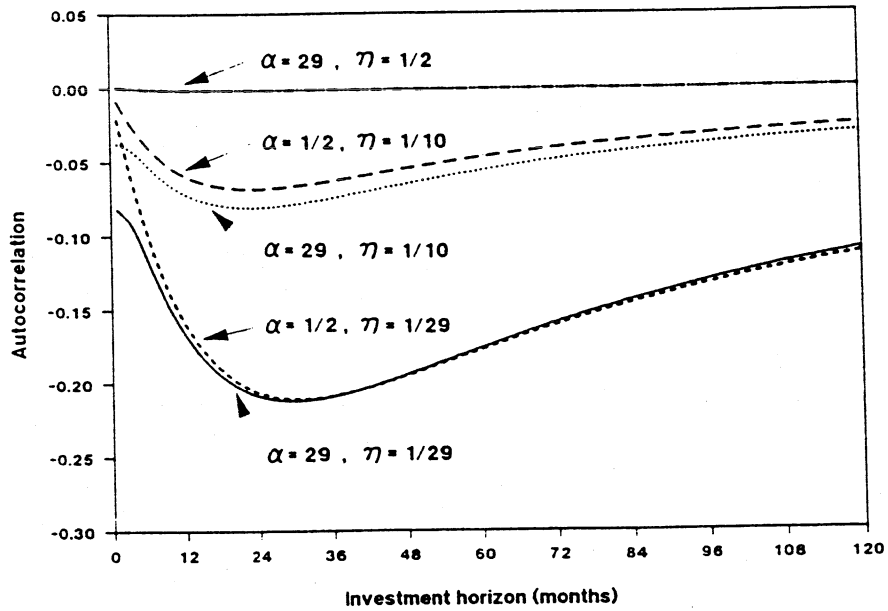


Figure 3. First order autocorrelations of equity returns for various investment horizons. The values plotted are those implied by the equilibrium model for alternative values of relative risk aversion (α) and the elasticity of intertemporal substitution (η).




Figure 4. The predictability measure Ω_N for various investment horizons. The values plotted are those implied by the equilibrium model for alternative values of relative risk aversion (α) and the elasticity of intertemporal substitution (η).




Figure 5. First-order autocorrelations of equity returns for various investment horizons. The values plotted are those implied by alternative values of ρ_μ and ρ_σ , the autocorrelations of the conditional mean and variance of the monthly consumption growth rate. The preference parameters in each case are $\alpha = 29$ and $\eta = 1/29$.

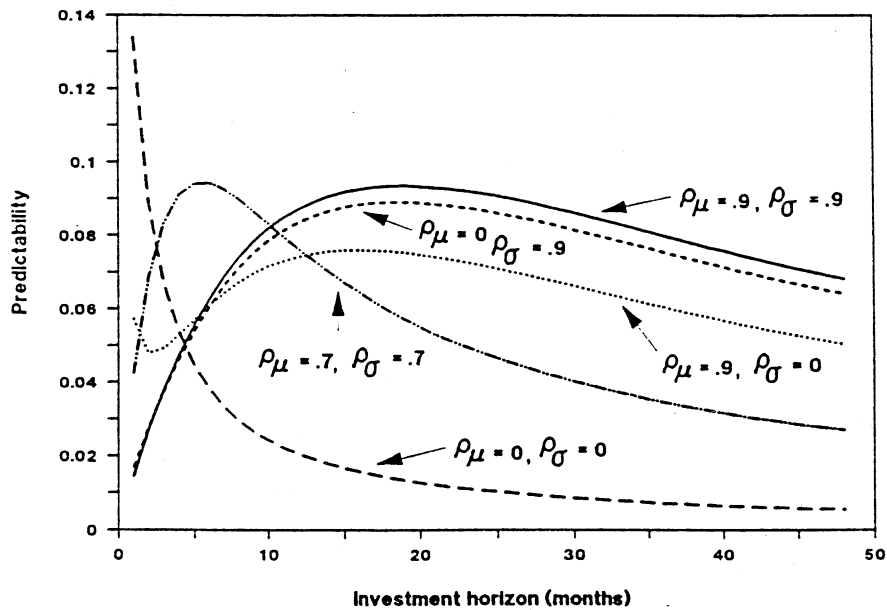


Figure 6. The predictability measure Ω_N for various investment horizons. The values plotted are those implied by alternative values of ρ_μ and ρ_σ , the autocorrelations of the conditional mean and variance of the monthly consumption growth rate. The preference parameters in each case are $\alpha = 29$ and $\eta = 1/29$.




Figure 7. Small gambles with constant relative risk aversion (RRA). The amount gambled is won or lost with equal probabilities.

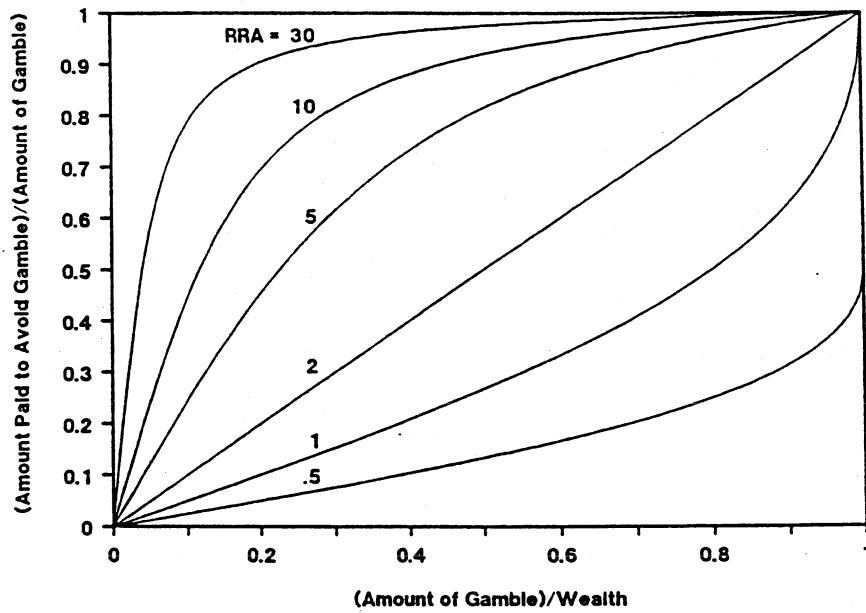


Figure 8. Large gambles with constant relative risk aversion (RRA). The amount gambled is won or lost with equal probabilities.