Persistence, Predictability, and Portfolio Planning*

Michael J. Brennan and Yihong Xia[†]

February 28, 2005

*We thank Robert Arnott of First Quadrant, Antti Ilmanen of Citigroup, Charles Stunkard of Wilshire Associates, and Richard Grinold of Barclays Global Investors for providing us with time series estimates of expected returns on common stocks. We thank Utpal Bhattacharya, Marshall Blume, John Campbell, Francisco Gomes, Robert Hodrick, Ravi Jagannathan, Shmuel Kandel, Craig MacKinlay, Amir Yaron, Moto Yogo, and seminar participants at the 2004 ESF/SCSS-PESC Exploratory Workshop, the 2004 SIFR conference, Carnegie Mellon University, Copenhagen Business School, Indiana University, Melbourne Business School, University of Pennsylvania, and University of Vienna for helpful comments. Brennan and Xia both acknowledge financial support from Inquire Europe and Xia acknowledges the Hirtle, Callaghan & Co. - Arthur D. Miltenberger Research Fellowship from the Rodney L. White Center for Financial Research at the Wharton School.

[†]Michael Brennan is Emeritus Professor at the Anderson School, UCLA, 110 Westwood Plaza, Los Angeles, CA 90095-1481. Yihong Xia is an assistant professor at the Wharton School of University of Pennsylvania, 2300

Abstract

We use a model of stock price behavior in which the expected rate of return on stocks follows an Ornstein-Uhlenbeck process to show that levels of return predictability that cause large variation in valuation ratios and offer significant benefits to dynamic portfolio strategies are hard to detect or measure by standard regression techniques, and that the R^2 from standard short run predictive regressions carry little information about either long run predictability or the value of dynamic portfolio strategies.

We propose a new approach to portfolio planning that uses forward-looking estimates of *long run* expected rates of return from dividend discount models. We show how such long run expected rates of return can be used to estimate the instantaneous expected rate of return under the assumption that the latter follows an Ornstein-Uhlenbeck process. Simulation results using four different estimates of long run rates of return on US common stocks suggest that this approach may be valuable for long horizon investors.

Steinberg Hall-Dietrich Hall, Philadelphia, PA 19104-6367. Phone: (215) 898-3004. Fax: (215) 898-6200. E-mail: yxia@wharton.upenn.edu.

1 Introduction

Hakansson (1970) and Merton (1971) extended the principles of dynamic portfolio optimization to situations of stochastic (time-varying) investment opportunities over thirty years ago, but interest in this topic languished until the 1990's under the soothing influence of the efficient markets paradigm and the belief that the investment opportunity set, if not exactly constant, was sufficiently close to constant that myopic portfolio strategies that ignored time variation in investment opportunities could be safely employed. However, the debate about excess stock price volatility,[1] the evidence of mean reversion in stock prices,[2] and of the predictive power of instruments such as the dividend yield, the book-to market ratio, the term spread and the short term interest rate,[3] have revived interest in dynamic portfolio theory in recent years. This interest has been further stimulated by the behavior of stock prices in the late 1990's, which drew attention to the implications of valuation ratios such as the market dividend yield and the book to market ratio for expected future returns of equity securities,[4] and therefore for portfolio strategies.

The dynamic portfolio models that have been developed and calibrated to US stock returns by Brennan, Schwartz, and Lagnado (1997), Barberis (2000), Campbell and Viceira (2001), and Xia (2001), among others, suggest that time variation in investment opportunities may have significant effects on optimal portfolio strategies for long lived investors, and that failure to take account of time variation in expected returns may carry significant costs. The analysis of Kandel and Stambaugh (1996) suggests that short horizon investors also can benefit from predictability, even if it is highly uncertain. These models all rely on statistical relations between (excess) returns on stocks and (usually one of) a set of predictor variables or instruments such as those mentioned above.

However, doubt has been cast on the usefulness of dynamic portfolio models by the weakness and instability of the estimated relations between stock returns and the instruments. It has also been suggested that *in sample* 'predictability' may be an illusion, the result of uncertainty and learning about the parameters of the underlying stochastic process for dividends.[5] Consistent with this suggestion, Goetzmann and Jorion (1993) argue that long run predictive regressions lack power in small samples,[6] while Bossaerts and Hillion (1999) find no evidence of *out of sample* excess

[1]See, for example, Shiller (1981), LeRoy and Porter (1981), Kleidon (1986), Marsh and Merton (1986).
[2]See, for example, Fama and French (1988b).
[3]See, for example, Keim and Stambaugh (1986), Fama and French (1988a) etc.
[4]See, for example, Lewellen (2003).
[5]See Brav and Heaton (2002) and Lewellen and Shanken (2002).
[6]Campbell (2001) shows that long-horizon regression tests have serious size distortions when asymptotic critical

return predictability at the monthly frequency in 14 countries using the standard instruments such as lagged returns, interest rates, and dividend yields. Goyal and Welch (2003) demonstrate that the dividend yield has no out of sample predictive power for annual stock returns during the period 1926 to 2002. They report that the dividend yield "is primarily forecasting future market returns over horizons greater than 10 years. It does not forecast market returns (or dividend growth) over horizons less than 5 years." In a more recent paper, Goyal and Welch (2004) show similar results for a wide range of variables[7] and conclude that "for all practical purposes, the equity premium has not been predictable, and any belief about whether the stock market is now too high or too low has to be based on a theoretical prior." These results of Goyal and Welch seem to be consistent with the finding of Philips *et al.* (1996) that investment managers who followed Tactical Asset Allocation strategies during the period 1988-1994 failed to outperform static strategies.

On the other hand, there is continuing evidence of predictability in equity returns, especially over *long horizons*. Fama and French (1988b), Poterba and Summers (1988), Lo and MacKinlay (1988), Cochrane (1999), and Campbell and Yogo (2003) all find evidence of temporary components in stock prices that becomes stronger at long horizons. Informal evidence of long run predictability is further apparent from the long swings and upward trend in valuation ratios during the 20th Century - swings that do not appear to have been mirrored in changing dividend growth rates.[8] The apparent paradox of strong evidence of long run predictability and only weak, if any, evidence of out of sample short run predictability is resolved by recognizing the importance of the time series behavior of the predictor variable which determines the time variation in the conditional mean return. This time series behavior has implications both for the estimation of the relation and for the economic importance of a given short run relation.

In this paper we demonstrate, first, that levels of return predictability that are of first order economic significance for long term investors are likely to be hard to detect from short run returns using the standard econometric methods that have been employed, even though they imply major deviations of valuation ratios from their long run 'equilibrium' levels. In showing that (in-sample) predictability may be hard to find when it exists, our paper complements the analyses of Goyal and Welch (2003) and Bossaerts and Hillion (1999) which show that predictability may appear to exist for certain instruments even when no true (out-of-sample) predictability exists. Secondly, we show how long run measures of expected return that are anchored in fundamentals and are derived by

---

values are used, but some versions of such tests have power advantages remaining after size is corrected.

[7]Campbell and Thompson (2004) show that the findings of Goyal and Welch (2004) are no longer true if sensible restrictions are imposed on the signs of predictive coefficients and return forecasts.

[8]Campbell and Shiller (1998, 2001) and Fama and French (2002), among others, find that dividend yields have little or no forecast power for subsequent dividend growth in the years since 1870.

comparing current stock prices with expected future dividends may be used to construct portfolio policies for long run investors. We provide time series estimates of instantaneous expected returns derived from these Dividend Discount Model measures of long run expected returns, and provide evidence that these measures would have been useful to long horizon investors.

2 Detecting and Exploiting Predictability

In this section we use the simplest model of stock return predictability, in which the proportional drift of the stock price follows a standard Ornstein-Uhlenbeck process, to analyze two issues. The first is the likelihood of being able to detect return predictability using standard approaches if the stock price drift is perfectly observable. The second issue that we address is the value of following an optimal strategy that takes account of the time variation in the drift parameter, relative to the value of an unconditional strategy that treats the drift as constant. We conclude that levels of time-variation in expected returns that are very important for investors with long horizons (20 years or more), may well be undetectable using the standard statistical approaches and, even if detectable, are likely to be estimated with considerable imprecision.

2.1 Implications of a Model of Stock Price Behavior

Our basic framework relies on the following continuous time model of the dynamics of the stock price with dividends reinvested which we denote by P:

$$\frac{dP}{P} = (\alpha + \beta\mu)\, dt + \sigma_P dz_P, \tag{1}$$

$$d\mu = \kappa\,(\bar{\mu} - \mu)\, dt + \sigma_\mu dz_\mu, \tag{2}$$

where dz_P and dz_μ are correlated standard Brownian motions with the correlation $\rho_{P\mu}$. In this formulation, μ_t is to be thought of as a perfect *signal* of the drift of the asset price process, $\alpha + \beta\mu_t$; however, for the most part we shall assume that $\alpha = 0$, and $\beta = 1$; then μ can be interpreted as the drift of the stock price.

The discrete time equivalent of the Ornstein-Uhlenbeck process (2) is the AR(1) process:

$$\mu_t - \bar{\mu} = e^{-\kappa\Delta}\,(\mu_{t-\Delta} - \bar{\mu}) + \eta_t. \tag{3}$$

where Δ is the time interval between observations of μ.

Let $R(t, t+\tau)$ denote the cumulative return on the stock over the interval $[t, t+\tau]$, then

$$
\begin{aligned}
R(t, t+\tau) &\equiv \ln P(t+\tau) - \ln P(t) = \int_t^{t+\tau} \left(\alpha + \beta\mu(s) - \frac{1}{2}\sigma_P^2\right) ds + \int_t^{t+\tau} \sigma_P dz_P, \\
&= \left(\alpha + \beta\bar{\mu} - \frac{1}{2}\sigma_P^2\right)\tau + \beta\left(\mu_t - \bar{\mu}\right)\frac{1 - e^{-\kappa\tau}}{\kappa} \\
&+ \frac{\beta\sigma_\mu}{\kappa}\int_t^{t+\tau}\left(1 - e^{-\kappa(t+\tau-s)}\right) dz_\mu(s) + \int_t^{t+\tau} \sigma_P dz_P(s). \qquad (4)
\end{aligned}
$$

The conditional variance of the τ-period return, $R(t, t+\tau)$, which we denote by $\mathrm{Var}(R(\tau))$, depends on the volatility of the return process σ_P, the volatility of the signal σ_μ, the horizon τ, and the mean reversion parameter κ, and can be written as:

$$
\mathrm{Var}(R(\tau)) = \frac{\beta^2\sigma_\mu^2}{\kappa^2}\left[\tau - 2\frac{1 - e^{-\kappa\tau}}{\kappa} + \frac{1 - e^{-2\kappa\tau}}{2\kappa}\right] + \sigma_P^2\tau + \frac{2\beta\sigma_P\sigma_\mu\rho_{P\mu}}{\kappa}\left[\tau - \frac{1 - e^{-\kappa\tau}}{\kappa}\right]. \qquad (5)
$$

Since the variance of μ_s conditional on μ_t ($s \geq t$), $\sigma^2(\mu_s|\mu_t) = \frac{\sigma_\mu^2}{2\kappa}\left(1 - e^{-2\kappa(s-t)}\right)$, is decreasing in κ, one might conjecture that the conditional variance of the cumulative return $\mathrm{Var}(R(\tau))$ would also be decreasing in the mean reversion intensity κ. The following expression for the derivative of the variance with respect to κ shows that this conjecture is true only if the correlation between innovations in P and innovations in μ is positive. When the correlation is negative (as we shall typically assume it is), the sign of the derivative is indeterminate:

$$
\begin{aligned}
\frac{\partial \mathrm{Var}}{\partial \kappa} &= -\frac{\sigma_\mu^2\beta^2}{\kappa^3}\left[2\tau - 4\frac{1 - e^{-\kappa\tau}}{\kappa} + \frac{1 - e^{-2\kappa\tau}}{\kappa} + \frac{2\kappa\tau e^{-\kappa\tau} - 2 + 2e^{-\kappa\tau}}{\kappa} - \frac{2\kappa\tau e^{-2\kappa\tau} - 1 + e^{-2\kappa\tau}}{2\kappa}\right] \\
&- \frac{2\beta\sigma_P\sigma_\mu\rho_{P\mu}}{\kappa^2}\left[\tau - \frac{1 - e^{-\kappa\tau}}{\kappa} + \frac{\kappa\tau e^{-\kappa\tau} - 1 + e^{-\kappa\tau}}{\kappa}\right]\begin{cases} < 0 & \text{if } \rho_{P\mu} \geq 0, \\ \text{indeterminate} & \text{if } \rho_{P\mu} < 0 \end{cases}.
\end{aligned}
$$

To analyze the role of μ_t as a predictor of the cumulative return, $R(t, t+\tau)$, re-write equation (4) as

$$
R(t, t+\tau) = a_\tau + b_\tau\mu_t + \epsilon, \qquad (6)
$$

where

$$
\begin{aligned}
a_\tau &\equiv \left(\alpha + \beta\bar{\mu} - \frac{1}{2}\sigma_P^2\right)\tau - \beta\bar{\mu}\frac{1-e^{-\kappa\tau}}{\kappa}, && (7)\\
b_\tau &\equiv \beta\frac{1-e^{-\kappa\tau}}{\kappa}, && (8)\\
\epsilon &\equiv \frac{\beta\sigma_\mu}{\kappa}\int_t^{t+\tau}\left(1-e^{-\kappa(t+\tau-s)}\right)dz_\mu(s) + \int_t^{t+\tau}\sigma_P dz_P(s),
\end{aligned}
$$

and ϵ is independent of μ_t.

Note first that b_τ increases with τ as long as $\beta > 0$ and that $b_\tau < \beta\tau\ \forall\tau > 0$. In addition, $b_\tau \to \frac{\beta}{\kappa}$ as the horizon goes to infinity, so that an innovation in μ has a permanent effect on the stock price. As $\tau \to 0$, $b_\tau \to 0$, and $b_\tau/\tau \to \beta$.

Secondly, b_τ is decreasing in the intensity of the mean reversion parameter κ. As $\kappa \to \infty$ the process for μ tends to a constant; deviations of μ from $\bar{\mu}$ are transient and, as a result, $b_\tau \to 0$. On the other hand, as $\kappa \to 0$, the process for μ approaches a random walk, and $b_\tau \to \beta\tau$, or $b_\tau/\tau \to \beta$ so that an increase in μ increases the expected rate of return in all future periods by the same amount. Thus, as has been pointed out by Cochrane (1999) and others, the degree of persistence in μ has major implications for the predictive role of μ in long horizon regressions, and two signals that differ in their persistence parameter κ may have the same predictive importance for returns at one horizon, but quite different importance at a different horizon. This suggests, as we shall see, that it may be dangerous to attempt to assess the relevance of predictive variables by examining their power at a single horizon.

A simple measure of the predictive power of a signal μ_t is the R^2 from the regression (6) at different return intervals τ. The unconditional distribution of the predictive variable, μ_t, is normal with mean $\bar{\mu}$ and variance $\frac{\sigma_\mu^2}{2\kappa}$, and the unconditional variance of $R(t, t+\tau)$ is simply $b_\tau^2 \mathrm{Var}(\mu_t) + \mathrm{Var}(\epsilon)$. The theoretical R^2 can be expressed explicitly in terms of the model parameters as:

$$
\begin{aligned}
R^2 &\equiv \frac{b_\tau^2 \mathrm{Var}(\mu_t)}{b_\tau^2\mathrm{Var}(\mu_t) + \mathrm{Var}(\epsilon)},\\
&= \frac{\left(\frac{1-e^{-\kappa\tau}}{\kappa}\right)^2}{\left(\frac{1-e^{-\kappa\tau}}{\kappa}\right)^2 + \frac{2}{\kappa}\left[\tau - 2\frac{1-e^{-\kappa\tau}}{\kappa} + \frac{1-e^{-2\kappa\tau}}{2\kappa}\right] + 2\kappa\tau\left(\frac{\sigma_P}{\beta\sigma_\mu}\right)^2 + 4\rho_{P\mu}\frac{\sigma_P}{\beta\sigma_\mu}\left[\tau - \frac{1-e^{-\kappa\tau}}{\kappa}\right]} \qquad (9)
\end{aligned}
$$

The expression for R^2 depends on four parameters: the cumulative return interval τ, the mean reversion intensity of the signal κ, the noise-to-signal ratio $\frac{\sigma_P}{\beta\sigma_\mu}$, and the correlation between the innovations to the return and to the signal $\rho_{P\mu}$. While all four parameters are important in determining the magnitude of the theoretical regression R^2, the dependence of R^2 on τ has received special attention in the literature. For example, Fama and Bliss (1987), Cochrane (1997, 1999), and Campbell (2001) have all explored the relation between R^2 and τ empirically finding that R^2 normally increases with the return interval τ. While Campbell, Lo and MacKinlay (1997) derive an approximate relation between a two-period and a one-period R^2 in a discrete model, our simple continuous time setup enables us to examine the relation between R^2 and all relevant parameters, including τ, in closed form.

To analyze formally the relation between R^2 and τ, differentiate R^2 with respect to τ:

$$\begin{aligned}\frac{\partial(R^2)}{\partial\tau} &= \frac{\frac{\sigma_\mu^2\beta^2}{2\kappa^2}\frac{1-e^{-\kappa\tau}}{\kappa}}{\left(b_\tau^2\mathrm{Var}\left(\mu_t\right)+\mathrm{Var}\left(\epsilon\right)\right)^2}\left[\frac{\sigma_\mu^2\beta^2}{\kappa^2}\left(2\kappa\tau e^{-\kappa\tau}+e^{-2\kappa\tau}-1\right)+\sigma_P^2\left(2\kappa\tau e^{-\kappa\tau}+e^{-\kappa\tau}-1\right)\right.\\ &+ \left.\frac{2\sigma_P\sigma_\mu\beta\rho_{P\mu}}{\kappa}\left(2\kappa\tau e^{-\kappa\tau}+e^{-2\kappa\tau}-1\right)\right],\end{aligned}$$

The expression outside the bracket is always positive but the term inside the bracket may be positive or negative depending on the parameter values. In general, R^2 first increases and then decreases with the return interval τ, and the sufficient condition for $\frac{\partial R^2}{\partial\tau} > 0$ is that $\tau < \tau^c$ where τ^c is the solution to the nonlinear equation $2\kappa\tau^c e^{-\kappa\tau^c} + e^{-2\kappa\tau^c} - 1 = 0$. The turning point at which R^2 starts to decrease with τ depends on all four parameter values. It may also be verified that $R^2 \to 0$ as $\tau \to 0$. Therefore, a variable that perfectly predicts the drift of the stock price will appear to have *no predictive power* if the predictive regressions are run using sufficiently short horizon returns. How short is "sufficiently" short depends on parameter values.

In addition to the hump-shaped relation between R^2 and τ, the signs of both $\frac{\partial R^2}{\partial\kappa}$ and $\frac{\partial R^2}{\partial\left(\frac{\sigma_P}{\beta\sigma_\mu}\right)}$ are also indeterminate. On the other hand, $\frac{\partial(R^2)}{\partial\rho_{P\mu}} < 0$, so that the theoretical R^2 attains its highest value at $\rho_{P\mu} = -1$. Since equation (9) is complicated, we shall explore the determinants of the R^2 numerically.

2.2 Detecting Predictive Power in Finite Samples

For our numerical analysis we choose the following 'global' parameters. We set the volatility of the annual cumulative return $\sqrt{\mathrm{Var}(R(1))} = 0.14$, which is approximately the annual volatility of real return on the S&P 500 index for the period 1950 to 2003. We set $\alpha = 0$ and $\beta = 1$ so that μ is the drift of the return process. With this in mind, we set the interest rate $r = 1\%$ and the long run mean $\bar{\mu} = 7\%$ which implies a mean equity premium of 6%. Finally, the volatility of the stationary distribution of μ is defined by $\nu_\mu \equiv \frac{\sigma_\mu}{\sqrt{2\kappa}}$. We set ν_μ at 4%:[9] this means that there is a 68% chance that the equity premium at any moment will lie between 2% and 6%, and a 6.7% probability that it will be negative.[10] This seems to us to represent a substantial amount of variation in the equity premium which we should have a reasonable chance of detecting. Within these global parameters we consider nine scenarios, the product of three values of the mean reversion parameter, κ, and three values of the correlation between innovations in μ and in returns, $\rho_{P\mu}$. The three values of κ are 0.02, 0.10, and 0.50 which correspond to half-lives for innovations in μ of 34.7, 6.9, and 1.4 years and therefore would seem to span the economically interesting range. The values of $\rho_{P\mu}$ are -0.90, -0.5, and 0.0. We restrict our attention to non-positive correlations because it is to be expected that, as innovations in expected returns are equivalent to innovations in discount rates, they should be negatively related to contemporaneous stock returns;[11] Xia (2001) reports a correlation of -0.93 between monthly innovations in the dividend yield and stock returns for the period January 1950 to December 1997 which is similar to Barberis (2000) who also uses the dividend yield as a predictor of stock returns.

Table 1 reports various statistics for our 9 different scenarios. Scenario (vi) is highlighed because it corresponds most closely with the empirical estimates presented below. The first four lines of the table report the parameters describing each scenario. The volatility of the drift, σ_μ, is set to be consistent with the exogenously chosen values of ν_μ and κ using the formula $\sigma_\mu = \nu_\mu\sqrt{2\kappa}$. The volatility of the instantaneous return σ_P is chosen to be consistent with $\sqrt{\mathrm{Var}(R(1))} = 0.14$ and the values of the other parameters using the quadratic equation (5).

Panel A reports the variance ratios implied by the nine different scenarios. The variance ratio for m months is defined by the expression:

$$VR_m = \frac{12}{m}\frac{\mathrm{Var}\left(R\left(t, t+\frac{m}{12}\right)\right)}{\mathrm{Var}(R(t, t+1))},$$

where $R\left(t, t+\frac{m}{12}\right)$ is the cumulative m-month return, and the variances for the returns are calculated using expression (5). Under a pure random walk the variance ratio will be equal to unity for all values of m. Mean reversion or negative autocorrelation in stock prices, such as would be introduced by changing discount rates, will cause the variance ratios to decline with the return calculation interval. The historical variance ratios are for real US equity returns. We report both the Poterba and Summers (1988) estimates for the period 1871-1985 and an updated calculation

[9] We shall find that this corresponds to the empirical estimates of this parameter reported below.

[10] Boudoukh *et al.* (1993) report evidence that the ex ante risk premium is negative in some states of the world.

[11] It is *possible* for innovations in discount rates to be positively correlated with current returns if the innovations in discount rates are highly correlated with innovations in cash flow news.

for the period 1871-2003. The standard errors of these historical estimates are of the order of 0.2. An interesting feature of the historical ratios is that they increase for the 2 year calculation period, which implies small positive autocorrelation at this horizon. Comparing the variance ratios under the nine scenarios with the historical ratios, we see that scenarios (i), (iv) and (vii), in which $\rho = 0.0$, produce the wrong pattern of variance ratios. Scenarios (v) and (viii) ($\rho = -0.5$ and $\kappa = 0.10, 0.50$) are the most consistent with the empirical data.

Panels B-E present, for each scenario, means and quantiles of the distribution of statistics for the regression (6) for one month and one year forecast periods. The theoretical (asymptotic) R^2 and predictive coefficient b_τ, calculated using definition (9) and (8), are reported as benchmarks. The means and quantiles of the distributions of the regression R^2 and estimated values of predictive coefficient, b, are calculated from Monte Carlo simulations of returns. Each simulation is for 70 years (840 months), roughly corresponding to the length of data available on the CRSP tape. Each scenario is simulated 2000 times, and the predictive regression (6) is run for one-month and one-year forecast horizons. The one year regressions use non-overlapping observations and the ordinary least squares estimates of the slope coefficient and its $t-$statistic are adjusted for bias using the Amihud-Hurvich (2004) approach. The Stambaugh (1999) bias-adjusted slope coefficients are indistinguishable from those obtained under the Amihud-Hurvich (2004) adjustment.

The theoretical R^2 for the one-month regressions are between 0.5% and 0.7%, and for all scenarios the average simulated R^2 lies between 0.3% and 0.9%. These values compare with those reported by Bossaerts and Hillion (1999) for a shorter sample period of between 2% and 9.8% using a combination of popular predictors. At the annual forecast horizon the theoretical R^2 rises to between 4.8% and 7.4%, while the median of the simulations yield estimates between 2.3% and 9.0%, which compares with the R^2 of 17% reported by Cochrane (1999) where overlapping observations of price-dividend ratio are used to predict annual excess stock returns. Note that even with non-overlapping observations, there is a small sample bias in R^2 which depends on κ and $\rho_{P\mu}$. The 25^{th} percentile of the R^2 distribution is around 3 percentage points below the mean, so the Goyal-Welch (2003) estimate of 5.83% for the period 1926-1990 is well within the range of estimates that are consistent with our scenarios.

Figure 1 plots the theoretical R^2 as a function of the return interval, τ, for the nine different scenarios.[12] Quite different patterns arise from the different sets of parameter values. Three features stand out: first, the generally humped shape of the graphs which is consistent with the

[12]Hodrick (1992) explores the implications of short run predictability for long horizon regressions in small samples using Monte Carlo simulation in a VAR framework.

theoretical analysis in the previous subsection; secondly, the small differences between the R^2 at the one year return interval for the different scenarios as compared with the large differences in R^2 for longer intervals; and thirdly, the importance of $\rho_{P\mu}$, as well as the persistence parameter κ, in determining the long horizon R^2. When $\kappa = 0.10$, for example, the R^2 at the 20-year interval ranges from 15.9% when $\rho_{P\mu} = 0$ to 46.1% when $\rho_{P\mu} = -0.9$.

Figures 2a and 2b plot the ten- and twenty-year theoretical R^2 against the one-year theoretical R^2. The one-year R^2 is independent of the correlation $\rho_{P\mu}$ by construction, since we set $\sqrt{\mathrm{Var}(R(1))}$ exogenously at 14%. The same R^2 at the one-year horizon can be associated with quite different R^2 at long horizons: when $\kappa = 0.1$, for example, the values of the one-year R^2 are all around 6.9%, but the R^2 at the twenty-year horizon for $\rho_{P\mu} = 0.0$ is 15.9% as compared to 46.1% for $\rho_{P\mu} = -0.9$. Thus, modest variation in the one-year R^2 between 6.5% and 7.5% is associated with enormous variation in values of R^2 at the ten- and twenty-year intervals - from under 20% to over 70% in the case of the 20-year R^2.

A high value of κ (0.5) greatly attenuates the long-horizon predictability of the return. On the other hand, the predictability of long run stock returns when expected returns are highly persistent depends strongly on the degree of negative correlation between the drift innovations and the stock return. For example, Scenario (i), which has a highly persistent expected return ($\kappa = 0.02$) but zero correlation ($\rho_{P\mu} = 0.0$), yields less return predictability at the ten- and twenty-year horizons than does Scenario (vi), which has less persistent expected returns ($\kappa = 0.10$) but a highly negative correlation ($\rho_{P\mu} = -0.9$). These examples show that, for a given unconditional volatility of the instantaneous expected return, the long run predictability of stock returns depends crucially on the correlation between innovations in expected returns and the realized stock return. This is not surprising since, in the absence of underlying cash flow uncertainty, the correlation would be perfect and long run returns would be perfectly predictable. Thus it is dangerous to infer anything about the long run predictability of stock returns from the R^2 of one-year regressions, particularly when the sampling variability of the R^2 is taken into account. For example, Table 1 Panel B shows that under Scenario (ii) there is a 50% chance of estimating an R^2 below 5% in an annual regression using 70 years data even though, as shown in Table 3 below, the theoretical R^2 for 20-year returns is almost 57%.

Panels D and E report statistics for the bias-adjusted coefficients of μ_t and their associated $t-$ratios from predictive regressions for one-month and one-year returns using the simulated data. For most of the scenarios, the median $t-$statistic is less than 2, and the predictive relation between the signal μ and the stock return would not be identified as significant in roughly two-thirds of the cases.[13] The shorter data samples that are typically used in empirical studies to identify potential signals of the stock price drift make it even less likely that a significant relation will be found, and if the drift is not a perfect linear function of the signal as we have assumed, then the likelihood of identifying a significant relation is further reduced. Similar results are obtained for the annual forecast period. Now, however, it is noticeable that both the theoretical and the estimated value of b decline for large κ which reduces the persistence of the signal.

The evidence in Table 1 shows that *for the stock price processes that we have considered* it is likely to be difficult to identify a significant relation between even a perfect signal of the expected return on stocks and the realized returns.[14] Even if a statistically significant relation is found, the standard error of the coefficient will be large, which makes the use of the estimate for portfolio planning purposes problematic.[15] Thus, the failure of Goyal and Welch (2004) to find any ex-post predictive power in the dividend yield for excess returns on common stocks is not surprising *even if the equity premium has a standard deviation as high as 4% and the dividend yield is a perfect signal of the expected return*. Throughout this section we have assumed that the standard deviation in the equity premium is 4%. While this seems to us to be quite large, we have found that it is broadly consistent with the levels of stock return predictability and negative autocorrelation in stock returns that previous researchers have found. We consider next whether this level of return predictability is of economic significance for variation in the level of stock prices.

3 Stock Price Variation and Variation in the Expected Returns

To assess the implications for the level of stock prices of time variation in expected returns of the magnitude discussed above, we employ a simple valuation model in which the instantaneous expected rate of return on the stock, μ, follows the O-U process described by equation (2). Since stock prices vary with dividend growth expectations as well as expected returns, we assume that the expected growth rate in dividends is constant in order to isolate the effects of expected returns on stock prices. Thus, the dividend on the stock is assumed to follow a geometric Brownian

[13] Of course, if the coefficient is not corrected for bias, the $t-$statistic are somewhat higher.

[14] Engstrom (2003) analyzes a related model in which μ is the dividend yield and β is a stochastic parameter. He also finds that predictive regressions of the form (6) have no power to reject the null of "no predictability."

[15] See Xia (2001) for an analysis of the effects of parameter uncertainty on optimal portfolio planning.

Motion with constant expected growth rate, g:

$$\frac{dD}{D} = gdt + \sigma_D dz_D. \tag{10}$$

The *ex-dividend* value of the stock, $V(D, \mu)$, is homogeneous of degree one in the level of the dividend, $V \equiv Dv(\mu)$, and its total return (capital gain plus dividend) may be written as:

$$\frac{Ddt + dV}{V} = \left[\frac{1}{2}\frac{v''}{v}\sigma_\mu^2 + \frac{v'}{v}\left[\kappa\left(\bar{\mu} - \mu\right) + \sigma_{D\mu}\right] + g + \frac{1}{v}\right] dt + \sigma_D dz_D + \frac{v'}{v}\sigma_\mu dz_\mu, \tag{11}$$

where $\sigma_{D\mu} \equiv \sigma_\mu \sigma_D \rho_{D\mu}$ is the covariance between the dividend growth and the expected return. On the other hand, the total stock return is also given by (1) in reduced form where P denotes the cum-dividend price and $\alpha = 0$ and $\beta = 1$ in the current setting.

Equating the drift term in (11) to the drift term in (1) yields the following ordinary differential equation (ODE) for the price-dividend ratio, v, as a function of the expected rate of return μ:

$$\frac{1}{2}v''{\sigma_\mu}^2 + v'\left[\kappa(\bar{\mu} - \mu) + \sigma_{D\mu}\right] + (g - \mu)v + 1 = 0. \tag{12}$$

The growth condition, $g + \frac{\sigma_\mu^2}{2\kappa^2} - \frac{\sigma_{D\mu}}{\kappa} < \bar{\mu}$, ensures that the above ODE has a solution. It can also be shown that $\nu(\mu) < 0$ so that a higher discount rate μ leads to a smaller price-dividend ratio. The boundary condition for $v(\mu)$ for high values of μ is $\lim_{\mu\to\infty} v(\mu) = 0$; and we impose a lower reflecting boundary on μ at μ^* such that $v'(\mu^*) = 0$.[16]

In order to explore the sensitivity of the price-dividend ratio, v, to the expected rate of return μ, when the expected rate of return follows the stochastic process (2), the partial differential equation was solved numerically for the values of κ and $\rho_{P\mu}$ that describe the scenarios used in Table 1. The real expected dividend growth rate, g, and the volatility of the real dividend growth rate, σ_D, are set to 0.86% and 8.52% respectively, which are the sample mean and volatility of the real dividend growth rate of the S&P 500 for the period from 1950 to 2002.[17] The parameter, $\rho_{D\mu}$, was then set so that σ_V was approximately equal to 14%. This is not possible for the scenarios in which $\rho_{P\mu}$ is 0 or -0.5. For the three remaining scenarios, there are two equal and opposite values of $\rho_{D\mu}$ which satisfy the condition $\sigma_V \approx 14\%$.

Table 2 shows, for the three scenarios in Table 1 for which the σ_V condition is satisfied, the effect of variation in μ on the price-dividend ratios and dividend yields. In scenario (vi) with

[16] See Feller (1951). In our simulations below, we set $\mu^* = -2.5\%$, which is about three standard deviations away from the long run $\bar{\mu} = 9\%$.

[17] The P/D ratio and dividend data are from Robert Shiller's web page http://www.econ.yale.edu/ shiller/data.htm.

$\kappa = 0.1$, which corresponds mostly closely to the empirical estimates which we shall present below, the price-dividend ratio increases by around 65% as μ moves from one standard deviation below its long run mean to one standard deviation above. The sensitivity of the stock price to μ is less in the high κ scenario (ix) and greater in the low κ scenario (iii).

The results in Tables 1 and 2 show that it is possible for changes in the expected rate of return, μ, to have very large effects on stock prices *without those changes being easily detected by standard statistical approaches*. For example, Table 1 shows that in Scenario (vi), in which $\kappa = 0.10$, the median bias-adjusted $t-$ratio on the predictor variable in regressions of one-year (one-month) returns on μ is only 1.74 (1.81) when 70 years of data are available. Yet Table 2 shows that in this same scenario a change in μ from one standard deviation below the mean to one standard deviation above the mean can increase the dividend yield by 60-68% without any change in dividend growth rate expectations.

4 Economic Significance of Predictability

Under the scenarios that we have described above, an investor faces a situation in which there is little short run predictability in stock returns, but there may be very significant long run predictability, depending on the persistence of the equity premium and the correlation of innovations in the premium with stock returns. Under these circumstances it is natural to ask whether there are likely to be significant costs for an investor who ignores time variation in expected returns. We answer this question by comparing the certainty equivalent wealth of an investor under an "unconditional" strategy with the certainty equivalent wealth under the "optimal" strategy.

The investor is assumed to have an iso-elastic utility function defined over end-of-period wealth at the investment horizon T:[18]

$$\max_x \mathrm{E}_0 \left[u \left(W_T \right) \right],$$

where

$$u \left(W_T \right) = \begin{cases} e^{-\rho T} \frac{W_T^{1-\gamma}}{1-\gamma} & \gamma > 1, \\ e^{-\rho T} \ln W_T & \gamma = 1. \end{cases}$$

[18]Brennan and Xia (2002) and Wachter (2002) show that the determination of an optimal consumption and investment program for an investor with time-additive utility can be represented as a time integral of optimal utility of terminal wealth problems.

The dynamic portfolio choice problem is subject to the dynamic budget constraint:

$$\frac{dW}{W} = [x(\alpha + \beta\mu - r) + r]\,dt + x\sigma_P dz_P,$$

with x defined as the proportion of wealth invested in the single risky asset whose stochastic process was given in equations (1-2). The risk free interest rate, r, and the stock return volatility, σ_P, are assumed to be constant for simplicity.[19] The "unconditional" strategy, which is based on the assumption that the equity premium is constant, allocates a constant fraction of wealth to the risky asset, but the allocation to the risky asset under the "optimal" strategy depends on μ_t as shown below.

Denoting the equity premium by $y_t \equiv \alpha + \beta\mu_t - r$, the investor's optimal dynamic policy[20] is to invest a fraction of his wealth, x_t^*, in the risky asset , where:

$$x_t^* = \frac{y_t}{\gamma\sigma_P^2} + \frac{\rho_{P\mu}\sigma_\mu}{\gamma\sigma_P}\left[B(\tau) + C(\tau)y_t\right], \tag{13}$$

and $\tau \equiv T - t$ is the time remaining to the investment horizon. The investor's indirect utility, $J(W,\mu,\tau) \equiv \max_{x_t} E_t\left[u\left(W_T\right)\right]$, is given by:

$$\begin{aligned} J(W,\mu,\tau) &= e^{-\rho t}\frac{W^{1-\gamma}}{1-\gamma}\exp\left\{A(\tau) + B(\tau)y_t + \frac{1}{2}C(\tau)y_t^2\right\} \\ &\equiv e^{-\rho t}\frac{W^{1-\gamma}}{1-\gamma}\phi^o(\mu,t), \end{aligned} \tag{14}$$

where $A(\tau)$, $B(\tau)$ and $C(\tau)$ are defined in Appendix A.

The optimal equity allocation, x_t^*, is a time-dependent function of the equity premium y. Xia (2001) shows that the allocation in general is a non-monotonic function of the investment horizon, and that if the return process is negatively correlated with the process for μ, i,e, $\rho_{P\mu} < 0$, which we have argued is likely to be case, then the optimal equity allocation x^* most likely increases with the investment horizon τ.

Let x^u denote the constant equity allocation under the unconditional strategy which treats the equity premium, $\bar{y} \equiv \alpha + \beta\bar{\mu} - r$, as a constant. Then

$$x^u \equiv \frac{\bar{y}}{\gamma\sigma_P^2}. \tag{15}$$

[19] For models of dynamic investment strategies with stochastic interest rates see Campbell and Viceira (2001) and Brennan and Xia (2002). For models of dynamic investment strategies with stochastic return volatility see Liu and Pan (2003).

[20] This was first derived by Kim and Omberg (1996); it is discussed in more detail in Xia (2001).

Although the optimal unconditional strategy is based on an equity premium that is assumed to be constant, $\bar{y}$, the evolution of the investor's wealth depends on the true dynamics of the equity premium $\mu - r$. As shown in Appendix B, the indirect utility function for the unconditional strategy under the true dynamics of the equity premium is given by:

$$\begin{aligned} J^u(W, \mu, \tau) &= e^{-\rho t} \frac{W^{1-\gamma}}{1-\gamma} \exp\{D(\tau) + E(\tau)\mu_t\} \\ &\equiv e^{-\rho t} \frac{W^{1-\gamma}}{1-\gamma} \phi^u(\mu, t), \end{aligned} \tag{16}$$

with $D(\tau)$ and $E(\tau)$ defined in Appendix B.

The Certainty Equivalent Wealth under the optimal strategy, $CEW^o(\mu, \tau)$, is defined as the amount of wealth to be received at the horizon with certainty that would give the investor the same expected utility as he receives under the optimal strategy with an initial \$1 of wealth to invest:

$$CEW^o(\mu, \tau) \equiv [\phi^o(\mu, \tau)]^{\frac{1}{1-\gamma}}. \tag{17}$$

The Certainty Equivalent Wealth for the unconditional strategy, $CEW^u(\mu, \tau)$, is defined analogously.

We shall analyze the ratio of the certainty equivalents under the two strategies, $CEWR^{ou}$, which is defined by

$$CEWR^{ou}(\mu, \tau) \equiv \frac{CEW^o(\mu, \tau)}{CEW^u(\mu, \tau)}.$$

The ratio, $CEWR$, is a measure of the value of taking account of variation in the equity premium.

Table 3a reports values of $CEWR^{ou}$ for different horizons across the nine different scenarios of the stochastic stock price processes, and for different initial values of μ_t when the coefficient of relative risk aversion, γ is equal to 5. We include scenarios (i), (iv), and (vii) in the table for completeness, although the pattern of declining variance ratios generated by these scenarios is inconsistent with the empirical evidence, and the zero value of $\rho_{\mu D}$ contrasts with the negative value we expect to be associated with discount rate effects. Therefore we restrict our discussion to the other six scenarios.

When $\mu_t = \bar{\mu}$, the optimal strategy offers no measurable advantage over a one-year horizon in any of the scenarios; a gain of up to 5% over a five-year horizon; a gain of between 2% and 15% over 10 years, depending on the scenario, and a gain of 7% to 57% over 20 years, again depending on the scenario. Under Scenario (v) the gain is 2% for a 5 year horizon, 6% for a 10-year horizon and 16% for a 20-year horizon. Focusing on the 20-year horizon, the gain from the optimal strategy when μ_t is one standard deviation above its mean may be substantial, particularly when $\rho_{P\mu} = -0.9$ (scenarios (iii) and (vi), and (ix)) and, for any given scenario, the increase in the $CEWR$ is greater for a one standard deviation increase in μ_t than is the reduction for a half standard deviation decrease in μ_t. Note that, while the R^2 of the predictive regressions was determined primarily by the value of κ, the influence of $\rho_{P\mu}$ on $CEWR$ is much greater; for example, comparing scenarios (iv) and (vi) at the 20-year horizon when $\mu_t = \bar{\mu}$, the advantage of the optimal strategy increases from 19% to 57% as $\rho_{P\mu}$ decreases from zero to -0.90. This is because a large negative correlation between stock returns and the expected rate of return reduces the volatility of long horizon returns; this effect is taken into account by the optimal strategy which invests more than the unconditional strategy in stocks for $\mu > \bar{\mu}$ when the horizon is long. The benefit of the optimal strategy is a non-monotone function of the persistence parameter κ, tending to be greatest for $\kappa = 0.10$ at the longer horizons.

The *short run* predictability of returns as measured by the one-year R^2 is not, however, necessarily associated with greater values of the optimal strategy relative to that of the unconditional one. This is illustrated in Figure 3, where the certainty equivalent ratio under the optimal and the unconditional strategies for a twenty-year horizon, $CEWR^{ou}$, is plotted against the one-year R^2 for the nine scenarios. Interestingly, there is also no clear relation between the advantage of the optimal strategy and the *long run* predictability of returns (twenty-year regression R^2) when $\mu = \bar{\mu}$: the correlation between the 20-year $CEWR$ across scenarios and the 20-year R^2 is -0.05. Less surprisingly, since μ carries more information about future investment opportunities when the R^2 is high, the 20-year $CEWR$ has a correlation of -0.48 with the 20-year R^2 when μ is one standard deviation below its mean and a correlation of 0.57 when μ is one standard deviation above its mean. Finally, it is interesting to note that there is also no relation between $CEWR^{ou}$ and the monthly or annual *out-of-sample* predictability as measured by the root mean squared error ($RMSE$).[21]

In summary, the gains to the optimal strategy can be very large for long horizon investors. For scenarios (v) and (vi) which correspond closely to the behavior of the expected rate of return that we shall extract from forecasts of long run rates of return, the gains run from 16% to 125% over a 20-year horizon, depending on the initial value of μ.

[21] The out of sample $RMSE$ is calculated from the simulated data used for Table 1; the predictive relation is estimated from the first 65 years of data, and the out of sample $RMSE$ is calculated from the differences between the predicted and the realized returns over the following 5 years. The results, which are not reported here, are available on request. Goyal and Welch (2003, 2004) use the out of sample $RMSE$ as a measure of the value of a predictive instrument.

The gains of the optimal strategy come from both market timing and hedging. Market timing is simply the variation in the equity allocation with the equity premium. Hedging is the additional allocation to equities that results from the negative correlation between innovations to rates of return and returns on the equity security, which implicitly recognizes that equities are not so risky in the long run.[22] The benefits of market timing, but not of hedging, are captured by the *myopic* rule, x^m:

$$x_t^m \equiv \frac{\alpha + \beta\mu_t - r}{\gamma\sigma_P^2} = \frac{y_t}{\gamma\sigma_P^2}. \tag{18}$$

The certainty equivalent wealth associated with the myopic strategy is calculated by evaluating the expected utility associated with the myopic strategy, and the details are given in Appendix C. Table 3b reports the certainty equivalent wealth ratios between the optimal and the myopic strategies $CEWR^{om}$. The results show that the hedging gains offered by the optimal strategy but not by the myopic strategy, are zero when $\rho_{P\mu} = 0$ and of the order of 6% at the 20-year horizon when $\rho_{P\mu} = -0.50$, but are as high as 38-66% when $\rho_{P\mu} = -0.90$.

For comparison, the certainty equivalent wealth associated with the optimal buy-and-hold strategy, x^b, which is described in Appendix D, was calculated numerically. The initial equity allocation of the buy-and-hold strategy depends on the value of μ_t, but subsequent changes in the allocation are determined entirely by the realized asset returns. The certainty equivalent wealth ratios between the optimal and the buy-and-hold strategies $CEWR^{ob}$, reported in Table 3c, show that the buy-and-hold strategy is extremely inefficient.[23] At the 20-year horizon the gains of the optimal strategy when $\kappa = 0.1$ are of the order of 38-81% when $\rho_{\mu D} = -0.50$ and 76-192% when $\rho_{\mu D} = -0.90$.

To this point, we have shown that time-variation in expected stock returns, which may be very difficult to detect by standard regression methods, may nevertheless imply both significant variation in stock prices that is unrelated to changes in cash flow expectations, and substantial potential gains to the use of dynamic portfolio strategies for long horizon investors. However, the gains that we have calculated assume that it is possible to observe the instantaneous expected return on the stock and, as we have seen, regression estimates of the relation between expected returns and even perfect instruments of it, are likely to be very imprecise. Therefore, in the next section, we explore the use of estimates of *long run* expected returns derived from a dividend discount model (DDM) as inputs to dynamic portfolio models.

[22] Stambaugh (1999) and Barberis (2000) compare myopic and buy-and-hold strategies.

[23] This is in contrast to the findings of Brennan and Torous (1999) who show that a buy-and-hold strategy performs well relative to a rebalancing strategy *when the investment opportunity set is treated as constant.*

5 Forecasts of Equity Returns

Regressions that attempt to predict stock returns from instruments such as the dividend yield or the interest rate lack strong theoretical restrictions on the regression coefficients and, as we have seen in Section 2, the data are likely to yield very imprecise estimates of the coefficients even when the instruments are perfect. It is not surprising therefore that such regressions have essentially no out of sample predictive power.

An alternative to the simple regression approach is to estimate the expected return on stocks by comparing the current level of the stock market with forecasts of future dividends on the market portfolio - the *Dividend Discount Model* (DDM) approach. The advantage of this approach, which has long been employed in "Tactical Asset Allocation" models, is that the expected return is estimated *directly* and that there is no need to estimate a regression coefficient relating the stock (excess) return to the predictor instruments. The offsetting disadvantage is that the rate of return estimated from the DDM is a *long run* internal rate of return and there is no reason to believe that this will be equal to the instantaneous expected rate of return even if the dividend forecasts are unbiased. Therefore, it is necessary to develop a model to derive the expected instantaneous rate of return from the long run expected rate of return.

5.1 Models and Estimation Procedure

We employ two models to convert the estimated DDM long run expected rate of return into an estimate of μ_t, the instantaneous expected rate of return. The first model assumes that the dividend growth rate g is a known constant. The second model assumes that the growth rate follows an Ornstein-Uhlenbeck process.[24] Both methods assume that the instantaneous expected rate of return follows an O-U process as in equation (2).

Our basic input data are direct estimates of the DDM *long run* 'expected rate of return,' k_t,[25] as defined by the discounted cash flow model:

$$P_t = \sum_{\tau=1}^{\infty} \frac{\mathrm{E}_t\left[D_{t+\tau}\right]}{\left(1+k_t\right)^{\tau}}, \tag{19}$$

where P_t is the level of the stock price index at time t, and $\mathrm{E}_t\left[D_{t+\tau}\right]$ is the dividend on the index

[24] Brennan and Xia (2001) assume a similar model for the dividend growth rate.

[25] There is no assurance that k_t, the solution to equation (19), will be equal to the expected rate of return except when all future dividends are known and discount rates are constant.

expected at time t to be paid at time $t+\tau$.

Model 1: constant g

When the dividend growth rate, g, is constant, there is a one-to-one correspondence between k_t and the price-dividend ratio, P/D, which is given by the (discrete time version of the) Gordon growth model:

$$v_t \equiv \frac{P_t}{D_t} = \frac{1+g}{k_t - g}. \tag{20}$$

Note that equation (20) rests on the assumption that the dividend expectations in (19) form a geometric series.

It was shown in Section 3 that if the expected rate of return, μ, follows an O-U process and the dividend growth rate is constant, then the price-dividend ratio $v(\mu)$ also satisfies the ordinary differential equation (ODE) (12). Since v is a monotonic function of μ, at each point in time there exists a μ_t whose implied v from the ODE is equal to the price-dividend ratio associated with k_t from equation (20) for a given value of g. This implies a (nonlinear) one-to-one mapping between k_t and μ_t for a given set of parameters for the stochastic process for μ, $\theta_\mu \equiv (\kappa, \nu_\mu, \bar{\mu}, \sigma_{D\mu})$. As a result, these parameters and the time series of μ_t can be estimated by an iterative process: starting with an initial value of $\theta = \theta^0$, μ_t is calculated from the time series of k_t; then θ_μ is estimated from the time series of μ_t, and a new μ_t series is re-calculated from k_t; the process continues until convergence is achieved. We denote the resulting Model 1 estimates by μ^1 and θ^1_μ.

This iterative procedure is essentially a non-linear Kalman filter in which the latent variable μ_t is a nonlinear function of the observable variable k_t. The transition equation for μ_t is the discrete-time equivalent of the O-U process:

$$\mu_{t+\Delta t} = (1 - e^{-\kappa_\mu \Delta t})\bar{\mu} + e^{-\kappa_\mu \Delta t}\mu_t + \epsilon \tag{21}$$

and the observation equation (which contains no observation error) is :

$$k_t = f(\mu_t) = g + \frac{D}{P} = g + \left(v(\mu_t, \theta)\right)^{-1}.$$

The mapping of v implied by the DDM to the v implied by the ODE is equivalent to solving the nonlinear function f numerically.

Model 2: stochastic g

To allow for the possibility that dividend growth rate expectations are stochastic, the instantaneous dividend growth rate is assumed to follow an Ornstein-Uhlenbeck process:

$$dg = \kappa_g \left(\bar{g} - g\right) dt + \sigma_g dz_g. \tag{22}$$

Two steps are then required to derive μ_t from k_t: first the instantaneous dividend growth rate, g_t, and the parameters of its stochastic process, $\theta_g \equiv (\kappa_g, \sigma_g, \bar{g})$, are estimated and then, given g_t and θ_g, the time series of μ_t and the parameters of its stochastic process, θ_μ are estimated.

With stochastic dividend growth, the DCF valuation formula that defines k_t can be written as:

$$\frac{P_t}{D_t} = \sum_{s=1}^{\infty} \frac{\Pi_{i=1}^{s} \mathrm{E}\left[(1 + g_{t+i})|g_t\right]}{(1 + k_t)^s}, \tag{23}$$

For a given set of parameters of the dividend growth process θ_g, $\mathrm{E}\left[(1 + g_{t+i})|g_t\right]$ is a known linear function of g_t. That is, equation (23) is used to solve for the time series g_t, given θ_g, the observed series of price-dividend ratios P/D, and the DDM expected rates of return k_t. The estimated time series g_t is then used to estimate θ_g, and the procedure iterates until there is a consistent set of g_t and θ_g, (g_t^*, θ_g^*).

This procedure can again be interpreted as a Kalman filter in which g is the latent variable with the transition equation

$$g_{t+\Delta t} = (1 - e^{-\kappa_g \Delta t})\bar{g} + e^{-\kappa_g \Delta t} g_t + \varepsilon,$$

while the observation equation relating the observable variables k and P/D to the unobservable variable g is given implicitly by (23).

Conditional on (g_t^*, θ_g^*) from the first step, μ_t is determined in the second step. The price-dividend ratio $v \equiv v(\mu, g; \theta_\mu, \theta_g^*)$ now satisfies the following two-state-variable partial differential equation:

$$\begin{aligned} 0 &= \frac{1}{2}\sigma_g^2 v_{gg} + \sigma_g \sigma_\mu \rho_{g\mu} v_{g\mu} + \frac{1}{2}\sigma_\mu^2 v_{\mu\mu} + \left[\kappa_g(\bar{g} - g) + \sigma_g \sigma_D \rho_{gD}\right] v_g \\ &+ \left[\kappa_\mu(\bar{\mu} - \mu) + \sigma_\mu \sigma_D \rho_{D\mu}\right] v_\mu + [g - \mu]v + 1. \end{aligned} \tag{24}$$

Equation (24) depends on the set of parameters θ_g^*, which was estimated in the first step, and the unknown θ_μ, which is to be estimated together with μ_t in the second step. For a given value of θ_μ, the PDE (24) is solved numerically for $v(\mu, g; \theta_\mu, \theta_g^*)$. Given g_t^* from the first step, a mapping between k_t and μ_t, $k_t = f(\mu_t, g_t)$, is defined by setting $v(\mu, g_t^*; \theta_\mu, \theta_g^*)$ equal to the observed price dividend ratio P/D. The time series of μ_t estimates are then used to estimate a new θ_μ and the process is iterated until convergence. This can be interpreted as another Kalman filter in which the transition equation of the latent variable μ_t is given by (21) while the nonlinear observation equation is:

$$P/D = v(\mu_t, g_t^*; \theta_\mu, \theta_g^*).$$

The resulting estimates from Model 2 are denoted by θ_μ^2 and μ^2.

5.2 Estimates

Quarterly data on real dividends and the price-dividend ratio for the S&P 500 index are obtained for the sample period 1950.1 to 2002.2. Four different DDM discount rate series, k_t, are used for illustrative purpose. The first two series are estimates of the real long run expected rate of return on equities that would have been assessed by investors at each date. The Arnott and Bernstein (2002) (A&B) series, k^1, is constructed by adding to the current dividend yield an estimate of the expected long run real growth rate in dividends which in turn is equal to a forecast of real GNP per capita growth less a 'dilution factor' - these are estimated using an average of the experience over the previous 40 years and the experience since the series began in year 1810. The Ilmanen (2003) (IL) series, k^2, is constructed by adding together estimates of the "dividend yield" and the long run growth rate of dividends. The former is calculated from a smoothed earnings yield multiplied by 59%. The latter is an average of 2% and the past 10, 20, 30, 40 and 50 years' geometric average real growth rate in corporate earnings.

Both the A&B and the IL series were constructed around 2002 for research purposes and use data or parameter values that may not have been available to market participants at the date of the forecast. However, the next two series, k^3 and k^4, were constructed in real time for use by investment professionals in asset allocation strategies. The series k^3 is provided by Wilshire Associates (WA), an investment management consulting firm, and the series k^4 is from Barclays Global Investors (BGI), a global investment management firm. They are both in *nominal* terms and are constructed by first aggregating I/B/E/S consensus estimates of growth rates in earnings per share out for five years for individual firms in the S&P 500 index and then letting the growth rates converge linearly to the economy-wide average growth rate over the next 10 years. The resulting growth rates are then used to project current dividends and to calculate an implied DDM rate of return from the current level of the S&P 500 index. These two nominal series are obtained from real-time forecasts of fundamentals that do not use any future information.

Table 4 reports the estimated parameters of the μ process for each series under the two assumptions about dividend growth. Several features stand out in the real series. First, the estimates of κ_μ, σ_μ, and $\bar{\mu}$ are largely unaffected by whether the expected dividend growth rate is assumed to be constant or not. Secondly, the estimates of both κ_μ and σ_μ are somewhat higher for the Ilmanen series than for the A&B series; the effect of these differences on ν_μ, the standard deviation of the stationary distribution of μ, are largely offsetting and the estimates are very close to the value of 4% that we have assumed in our simulations in Sections 2 and 3. Thirdly, the estimates of the correlation $\rho_{P\mu}$ are all between -0.884 and -0.981; this is what we should expect since μ is a proxy for the discount rate. Finally, most of the parameters are quite close to the values that were assumed or derived for Scenario (vi) in Table 2 (which assumes a constant expected dividend growth rate). The exception is $\rho_{\mu D}$ which is much higher in Scenario (vi). We suspect that this is because σ_D was set equal to 12.02% in Table 2 while its estimate is only 8.52% in Table 4.[26]

The estimates of the parameters of the nominal μ process are broadly similar to those estimated for the real models except that in both cases the Model 2 σ_D is over 20%, but this is offset by the much higher level of mean reversion of the dividend growth rate κ_g; in addition, $\rho_{\mu D}$ is positive in both nominal models, while it is negative in the two real models.

Figure 4a shows that when the A&B series k^1 is used as the input, the estimates of μ from the two models, $\mu^{1,1}$ and $\mu^{1,2}$, are virtually coincident, suggesting that there is little advantage in allowing for a stochastic dividend growth rate. Note that both estimates of the instantaneous expected rate of return, μ, are much more volatile than the underlying DDM k series from which they are derived: the maximum value of $\mu^{1,2}$ is over 16% in June 1982 and the minimum is around -3% in March 2000; in contrast, the maximum and minimum values of the DDM k, which fall exactly on the same two months, were respectively 7.9% and 2.1%.

Figure 4b plots the corresponding μ estimates from the Ilmanen series k^2. As in the previous case, the $\mu^{2,1}$ and $\mu^{2,2}$ series track each other closely except around the end of 1999 when the decline in the Model 2 estimate, $\mu^{2,2}$, is much more dramatic than that in $\mu^{2,1}$: $\mu^{2,2}$ reaches a minimum of -3.2% just before the end of the bull market in December 2000. While the two series of $\mu^{i,2}$ ($i = 1, 2$) generally track each other quite well, there are periods of significant differences. In the 1950's the Ilmanen estimate exceeds the A&B estimate by up to 3.8%. Significant differences of the opposite sign occur during the period 1983-1996; in September 1987 the Ilmanen estimate was less than 1% while the A&B estimate was over 6.5%.

[26] As noted above, 12.02% is the sample volatility of annual real dividend growth over the period 1872-2001, while 8.52% is the estimate of σ_D in equation (22) derived from quarterly data for the period 1950.1 to 2002.2 using the algorithm described under *Model 2* above.

Figure 5 plots the quarterly nominal estimates of k^i, $\mu^{i,1}$ and $\mu^{i,2}$ $(i = 3, 4)$ from BGI and WA for the shorter periods for which these series are available. These last two series differ from the previous two in three significant ways. First, they are truly *ex-ante*; secondly, they are *nominal* expected rates of return; thirdly, they are based on analysts' forecasts of earnings growth rates which are known to be upward biased. Not surprisingly, these nominal k series lie everywhere above the two real k series. Nevertheless, they show a similar pattern of increase during the 1970's, the BGI series k^3 (the WA series k^4) reaching a peak of 19.1% (18.8%) in 1981Q3 (1982Q2). After the peak, there is a prolonged decline in both series. Although for both series, $\mu^{i,1}$ and $\mu^{i,2}$ $(i = 3, 4)$ move largely in parallel, $\mu^{i,1}$ is everywhere above $\mu^{i,2}$ $(i = 3, 4)$ and the difference between them is much larger than that observed for the A&B and the IL real series. The two $\mu^{i,2}$ $(i = 3, 4)$ series track each other closely except for the period 1980-82 as well as 1999-2000, when the WA estimate $\mu^{4,2}$ exceeds the BGI estimate $\mu^{3,2}$ by 1-3%, and the period 2001-2002 when the WA estimate $\mu^{4,2}$ is below the BGI estimate $\mu^{3,2}$ by 2-5%.

5.3 Return Prediction

We have already shown that even when return predictability is of great economic importance, the statistical evidence of predictability may be weak and hard to detect using standard statistical methods. Therefore, even if the μ series contain valuable information for portfolio planning, the statistical evidence of their predictive power may be weak.

Panel A of Table 5 reports the results of regressions of quarterly real returns on the S&P 500 index on values of $\mu^{i,2}$ $(i = 1, 2)$ derived from the A&B and IL real DDM series, while Panel B reports the results of regressing quarterly nominal returns on the estimates of the nominal $\mu^{i,2}$ $(i = 3, 4)$ derived from the BGI and WA nominal DDM series.

We report results for the whole sample period from 1950.2 to 2002.2, and also the two approximate halves of the period, omitting the influential 1974.3 quarter when the real return on the S &P 500 was approximately *minus* 28%. For the two real μ series, the effect of omitting this quarter is to raise both the estimated coefficient towards its theoretical value of unity and to raise the regression R^2. While the regression coefficients are not significantly different from either zero or their theoretical value of unity, the point estimates are close to unity. The explanatory power of the predictive variable in both models is considerably greater in the first half of the sample period where the regression R^2 is around 4-5% and the predictive coefficient is statistically significant at the 5% level. In all other periods, however, regression R^2's are all below 2% and none of the predictive coefficients is significant at the 5% level.

Panel B reports the corresponding results for nominal returns using the predictors derived from the BGI and WA nominal DDM models. These models have greater predictive power than the two real models for the relevant sample periods, with R^2 at around 2-3%, despite the fact that they were made in real time and contain no 'look-ahead' bias. The estimates of the coefficients of the predicted return are little affected by the omission of 1974.3 and, as in the two previous cases, are close to the theoretical value of unity but are not significantly different from zero either.

The evidence from Table 5 is thus broadly consistent with the earlier observation that the predictive coefficient estimates, while close to their theoretical values, are associated with large estimation error and that the weak statistical evidence of predictive power of μ series at a quarterly horizon does not provide much information on their economic importance to investors with horizons of twenty years or longer.

5.4 Historical Simulations

In this section we report the results of simulating the optimal and unconditional policies using each of four μ series, A&B $\mu^{1,1}$, IL $\mu^{2,1}$, BGI $\mu^{3,2}$, and WA $\mu^{4,2}$, for a long horizon investor with a relative risk aversion coefficient, γ, of 5. Allocations to stocks were revised quarterly, and borrowing and short sale constraints were imposed. Under the unconditional strategies, the fraction of the portfolio that is allocated to the risky asset is determined from equation (15) subject to the constraint that $0 \leq x^u \leq 1$, and is constant over time. Under the *constrained* optimal strategy, the fraction of wealth allocated to the risky asset is determined by solving an optimal control problem whose value function depends on μ; it is described in Appendix A.

The return on the market portfolio is taken as the historical real (nominal) return on the S&P 500 index for each quarter and the riskless interest rate is taken as the realized real (nominal) return on a 30-day Treasury from CRSP for the real $\mu^{1,1}$ and $\mu^{2,1}$ (the nominal $\mu^{3,2}$ and $\mu^{4,2}$). The instantaneous return volatility, $\sigma_P = 15.7\%$ (17.1%), is determined from equation (5) by setting $\mathrm{Var}(R(0.25))$ equal to the sample volatility, 7.74% (8.40%), of quarterly S&P 500 index real (nominal) return, when either $\mu^{1,1}$ or $\mu^{2,1}$ ($\mu^{3,2}$ or $\mu^{4,2}$) is used. The same value of σ_P is used in both the optimal and the unconditional policies. The other parameters for the optimal policies are taken from the appropriate line of Table 4, and the constant unconditional equity premium $\bar{y}$ for the unconditional policy is calculated from the corresponding $\bar{\mu}$ from Table 4.[27] For both the

[27]The unconditional equity premium $\bar{y}$ for the unconditional strategy was calculated in three different ways: (1) using the same $\bar{\mu}$ as that used in calculating the optimal strategy; (2) setting $\bar{y}$ to the sample mean of the S&P 500 Index excess return during the whole sample period of 1929 to 2002; and (3) setting $\bar{y}$ to the gradually updated sample

optimal and the unconditional strategies under the real μ (nominal μ), the risk free rate r is set at a constant 1.1% (4.1%), which is the sample mean of the realized real (nominal) return of the 30-day Treasury Bill rate from 1950 to 2002.

Figures 6 and 7 summarize the results of the simulations under the real A&B $\mu^{1,1}$ and IL $\mu^{2,1}$. For each figure the investor is assumed to be concerned with maximizing the expected utility of wealth on the last date included in the figure. For example, Figure 6a describes the evolution of wealth of an investor who starts investing at the end of the first quarter (or equivalently the beginning of the second quarter) of 1950 with a horizon of the end of the first quarter of 1970; thus his initial horizon is 20 years and decreases each period. Investment decisions are assumed to be made at the beginning of each quarter based on the current value of μ. The figure shows that an investor, who would have followed the optimal strategy based on $\mu^{1,1}$ over this period, would have vastly outperformed his unconditional counterpart. Much the same pattern is visible in Figure 7a which is based on $\mu^{2,1}$. Figures 6b and 7b show that the optimal strategies continue to outperform the unconditional strategies but by a smaller margin over the subsequent 20-year period 1970-1990. Both the $\mu^{1,1}$ and $\mu^{2,1}$ based strategies lose more in the oil-price related bear market of 1974 on account of their more aggressive stock positions but more than make up for this by the end of the period. A 12-year horizon investor over the period 1990-2002 does better (about as well) under the optimal strategies than (as) under the unconditional strategy using $\mu^{1,1}$ ($\mu^{2,1}$) series. Finally, the simulations for an investor with a 52 year horizon starting in 1950 show that the 'optimal' investor ends up well ahead. Since the $\mu^{i,1}$ and $\mu^{i,2}$ ($i = 1, 2$) series are trivially different we do not report the results for the μ^2 series.

Both the A&B and Ilmanen DDM series are constructed by projecting historical growth rates of aggregate series such as profits or GDP. In contrast, the nominal BGI and WA DDM series constructed by investment professionals relies on bottom-up forecasts of individual firm growth rates. Figures 8 and 9 compare the (nominal) wealth outcomes from the unconditional vs. the optimal strategy derived from BGI $\mu^{3,2}$ and WA $\mu^{4,2}$ series for the 20-year period 1973-1993 and the 11-year period 1993-2003. In the first period, the wealth outcome under the optimal strategy is 53-67% higher than that under the unconditional strategy: it is noteworthy that these strategies do not suffer the large losses of the A&B and Ilmanen based strategies during 1974. For the

mean excess return with the initial value calculated from 1929 to 1949 (or 1972 for the BGI or WA series). The first approach ensures that differences between the wealth realized under the optimal and the unconditional strategies are not caused by different assumptions about the unconditional equity premium. Unconditional strategies based on (2) and (3) have similar realized wealth as that based on (1) except for the period 1950-1970, during which unconditional strategies based on (2) and (3) significantly outperform that based on (1) but still underperform the optimal strategy. Since the relative performance of the optimal strategy is consistent across the three unconditional strategies, we only report results of the unconditional strategy based on (1) and omit those based on (2) and (3) for brevity.

11-year investment period 1993-2003, the optimal strategies based on both $\mu^{3,2}$ and $\mu^{4,2}$ series underperform the unconditional ones by 10-20%. It is interesting to note that the unconditional strategy *out*performs the optimal strategies during the bubble period at the end of the 1990's, but that it substantially *under*performs as the bubble collapses during the period 2000-2003: it is precisely in such circumstances that we would expect that a dynamic strategy that takes account of the long run expected rate of return implicit in asset prices to do well. Figures 8c and 9c compare the outcome of the unconditional strategies with those of the optimal ones that start with a 30 year horizon in 1973. Over this longer horizon, the optimal strategies show a consistent advantage.

Figure 10 plots the portfolio allocations for the WA estimates of $\mu^{4,2}$ for the different sub-period strategies. The allocation to stocks under the unconditional strategy is about 48% while it varies between 0 and 100% under the optimal dynamic strategy. The myopic allocation rises from 0 in 1973 to 100% by 1980, and then declines irregularly to 0 by 1998. The optimal dynamic allocation is the sum of the myopic allocation and the hedge demand, which depends on both the investment horizon and μ. The hedge demand is always positive and is large at the beginning of each sub period when the investment horizon is long. For example, in September 1976, the myopic demand is 55% and the hedging demand is 45% for an investor with a twenty-five year horizon.

While we should be careful from inferring too much from these historical simulations which represent only a single sample path of stock prices for a single level of risk aversion, it is encouraging that the optimal dynamic strategies tend to outperform naive unconditional strategies, even when they are based on real time data.[28]

6 Conclusion

There is considerable disagreement about whether or not stock returns are time-varying and predictable. On the one hand, there is ample evidence of in-sample return predictability from regressions of stock returns on instruments such as the dividend yield and short term interest rate, and from variance ratio tests. On the other hand, it has been argued that in-sample predic-

[28] The instantaneous expected returns estimated from the DDM long run expected returns depend on parameters of the expected return μ and the dividend growth g processes, κ_μ, σ_μ, and $\bar{g}$ etc. These parameters were estimated using data from the whole sample period so that our estimates of instantaneous expected returns, even when they are based on real time DDM estimates, rely on future data. For the A&B and IL series, we also estimated μ for the period of 1950-2002 by first estimating the parameters using data only from 1900 to 1949, and the superior performance of the optimal strategy remains unchanged. We do not have long enough sample for the BGI and WA series, which are only available starting from 1973, to carry out this robustness check.

tive regressions perform poorly out of sample, so that the evidence of time variation should be discounted.

In this paper, we have shown that time variation in expected returns that implies both large variation in stock market valuation ratios and substantial gains to long term dynamic investment strategies is likely to be hard to detect by standard statistical methods. As a result, weak statistical evidence for return predictability does not in itself imply that return predictability is economically insignificant.

In view of the difficulty of estimating expected returns from regressions of realized returns on instruments such as the interest rate or dividend yield, we suggest that it is likely to be more productive to estimate the expected long run rate of return by comparing the current level of stock prices with forecasts of expected future dividends in the dividend discount model (DDM) paradigm. This forward-looking approach has the advantage that it does not rely on hard-to-estimate regression coefficients from past data. The disadvantage is that the rate of return that emerges from the dividend discount rate model is a *long run* expected rate of return. In order to use the DDM expected rate of return estimate in the dynamic portfolio planning, we show how the instantaneous expected rates of return can be estimated from the DDM long run expected rate of return under the assumption that the instantaneous expected rate of return follows an Ornstein-Uhlenbeck process; the technique is also extended to a setting in which the expected growth rate of dividends, instead of being constant, also follows an Ornstein-Uhlenbeck process.

Time series of expected rates of return are estimated for four time series of DDM expected rates of return. Two of them are historical 'back-casts,' and two are real time estimates provided by investment professionals. Simulations using realized S&P 500 index returns and the 30-day Treasury bill rates suggest that there may be significant benefits from the use of expected instantaneous rates of return derived from dividend discount model expected returns, even when the expected returns are estimated in real time. In this paper, however, we have examined the benefit of the optimal market timing strategy without taking account of misspecification of, or errors in estimating, the stochastic process for the expected instantaneous rate of return. Determining the sensitivity of DDM based dynamic portfolio strategies to errors in specifying and estimating the stochastic process for the instantaneous expected return is a task for future work.

Appendix

A. The Optimal Strategy

The investor is assumed to have an iso-elastic utility function defined over end-of-period wealth at the investment horizon T:

$$\max_x \mathrm{E}_0 \left[u\left(W_T\right) = \left\{ \begin{array}{ll} e^{-\rho T}\frac{W_T^{1-\gamma}}{1-\gamma} & \gamma > 1, \\ e^{-\rho T}\ln W_T & \gamma = 1. \end{array} \right. \right],$$

subject to the following dynamic budget constraint:

$$\frac{dW}{W} = \left[x\left(\alpha + \beta\mu - r\right) + r\right]dt + x\sigma_P dz_P,$$

where x is defined as the proportion of wealth invested in the single risky asset whose stochastic process was given in equations (1-2). The risk free interest rate, r, is assumed to be constant for simplicity.

Under the iso-elastic utility function, the indirect utility function, $J(W, \mu, t) \equiv \max_{x_t} E_t\left[u\left(W_T\right)\right]$, is homogeneous in W:

$$J(W, \mu, t) = e^{-\rho t}\frac{W^{1-\gamma}}{1-\gamma}\phi(\mu, t),$$

where ϕ satisfies the following Bellman equation:

$$\begin{aligned} 0 \quad &= \quad \max_x \left\{ \frac{1}{2}\sigma_\mu^2\phi_{\mu\mu} + \left[\kappa\left(\bar{\mu} - \mu\right) + (1-\gamma)x\sigma_P\sigma_\mu\rho_{P\mu}\right]\phi_\mu \right. \\ &+ \quad \left. \left[(1-\gamma)\left[x(\alpha + \beta\mu - r) + r\right] - \rho - \frac{1}{2}\gamma(1-\gamma)x^2\sigma_P^2\right]\phi + \phi_t \right\}, \end{aligned} \qquad \text{(A1)}$$

where $\phi_{\mu\mu} \equiv \frac{\partial^2\phi}{\partial\mu^2}$, $\phi_\mu \equiv \frac{\partial\phi}{\partial\mu}$, and $\phi_t \equiv \frac{\partial\phi}{\partial t}$ are the partial derivatives of ϕ with respect to μ or t.

Denote the equity premium by $y_t \equiv \alpha + \beta\mu_t - r$ and the remaining investment horizon by $\tau \equiv T - t$, then the investor's *unconstrained* optimal dynamic policy x_t^* is given by:

$$x_t^* = \frac{y_t}{\gamma\sigma_P^2} + \frac{\rho_{P\mu}\sigma_\mu}{\gamma\sigma_P}\left[B(\tau) + C(\tau)y_t\right],$$

and the function $\phi(\mu, t)$ under the unconstrained optimal strategy is reduced to:

$$\phi^o(\mu, \tau) = \exp\left\{A(\tau) + B(\tau)y_t + \frac{1}{2}C(\tau)y_t^2\right\}$$

where $A(\tau)$, $B(\tau)$ and $C(\tau)$ are solutions to a system of three ordinary differential equations with boundary conditions of

$$A(\tau) = 0, \quad B(\tau) = 0, \quad \text{and} \quad C(\tau) = 0 \quad \text{at} \quad \tau = 0.$$

The details of the equations are contained in Kim and Omberg (1996) for the general HARA utility and in Xia (2001) for the CRRA utility.

In particular, let

$$a_1 = (\beta\sigma_\mu)^2 \left(1 + \frac{1-\gamma}{\gamma}\rho_{P\mu}^2\right), \tag{A2}$$

$$a_2 = 2\left(\frac{(1-\gamma)\beta\sigma_\mu\rho_{P\mu}}{\gamma\sigma_P} - \kappa\right), \tag{A3}$$

$$a_3 = \frac{1-\gamma}{\gamma\sigma_P^2}, \tag{A4}$$

then $q \equiv a_2^2 - 4a_1a_3 > 0$ for all $\gamma \geq 1$, which is the condition for the well-behaved normal case.

In this paper, we focus on dynamic strategies under $\gamma > 1$, so $A(\tau)$, $B(\tau)$ and $C(\tau)$ are given by the normal solution:

$$C(\tau) = \frac{2a_3\left(1 - e^{-\eta\tau}\right)}{(\eta - a_2) + (\eta + a_2)e^{-\eta\tau}}, \tag{A5}$$

$$B(\tau) = \frac{4a_3\kappa(\alpha + \beta\bar{\mu} - r)\left(1 - e^{-\frac{\eta\tau}{2}}\right)^2}{\eta\left[(\eta - a_2) + (\eta + a_2)e^{-\eta\tau}\right]}, \tag{A6}$$

$$\begin{aligned} A(\tau) &= \left[a_3\left(\frac{2\kappa^2(\alpha + \beta\bar{\mu} - r)^2}{\eta^2} + \frac{(\beta\sigma_s)^2}{\eta - a_2}\right) + r(1-\gamma) - \rho\right]\tau \\ &+ \frac{4a_3\kappa^2(\alpha + \beta\bar{\mu} - r)^2\left[(2a_2 + \eta)e^{-\eta\tau} - 4a_2e^{-\frac{\eta\tau}{2}} + 2a_2 - \eta\right]}{\eta^3\left[(\eta - a_2) + (\eta + a_2)e^{-\eta\tau}\right]} \\ &+ \frac{2a_3(\beta\sigma_s)^2}{\eta^2 - a_2^2}\ln\left|\frac{(\eta - a_2) + (\eta + a_2)e^{-\eta\tau}}{2\eta}\right|, \end{aligned} \tag{A7}$$

where $\eta = \sqrt{q}$.

If the optimal allocation x^* is constrained to be between zero and one, i.e., no borrowing or short-sale is allowed, then no closed-form solution is available. Both the *constrained* optimal policy x^* and the associated function ϕ are solved numerically from the Bellman equation (A1) subject to the constraint $0 \leq x^* \leq 1$.

B. The Unconditional Strategy

If the investor believes that the equity premium is constant, then the investment opportunity set is constant from the investor's perspective. The unconditional strategy,

$$x^u = \frac{\alpha + \beta\bar{\mu} - r}{\gamma\sigma_P^2} = \frac{\bar{y}}{\gamma\sigma_P^2},$$

is based on the long run mean of the equity premium. However, the wealth process of the investor evolves according to the true dynamics of the equity premium given by equations (1-2):

$$\frac{dW}{W} = \left[x^u\left(\alpha + \beta\mu - r\right) + r\right]dt + x^u\sigma_P dz_P. \tag{B1}$$

The indirect utility function $J^u(W,\mu,\tau)$ is then given by:

$$\begin{aligned} J^u(W,\mu,\tau) &= \mathrm{E}_t\left[e^{-\rho T}\frac{W_T^{1-\gamma}}{1-\gamma}\right] = e^{-\rho T}\frac{W_t^{1-\gamma}}{1-\gamma}\mathrm{E}_t\left[\exp\left\{(1-\gamma)\left(\ln\frac{W_T}{W_t}\right)\right\}\right] \\ &= e^{-\rho T}\frac{W_t^{1-\gamma}}{1-\gamma}\exp\left\{\mathrm{E}_t\left[(1-\gamma)\left(\ln\frac{W_T}{W_t}\right)\right] + \frac{1}{2}\mathrm{Var}_t\left[(1-\gamma)\left(\ln\frac{W_T}{W_t}\right)\right]\right\}. \end{aligned}$$

Solving for $\ln\frac{W_T}{W_t}$ and its first two conditional moments from equation (B1) gives the following result:

$$J^u(W,\mu,\tau) = e^{-\rho t}\frac{W_t^{1-\gamma}}{1-\gamma}e^{D(\tau)+E(\tau)\mu_t}, \tag{B2}$$

where

$$\begin{aligned} D(\tau) &= (1-\gamma)\left[r + \frac{1}{2}\left(x^u\right)^2\left(\gamma\sigma_P^2 + (1-\gamma)\frac{\beta^2\sigma_\mu^2}{\kappa^2} + (1-\gamma)\frac{2\beta\sigma_{P\mu}}{\kappa}\right)\right]\tau \\ &- (1-\gamma)x^u\beta\left[\bar{\mu} + (1-\gamma)\frac{x^u\beta\sigma_\mu^2}{\kappa^2} + (1-\gamma)\frac{x^u\sigma_{P\mu}}{\kappa}\right]\frac{1-e^{-\kappa\tau}}{\kappa} \\ &+ \frac{1}{2}(1-\gamma)^2\frac{x_u^2\beta^2\sigma_\mu^2}{\kappa^2}\frac{1-e^{-2\kappa\tau}}{2\kappa} - \rho\tau, \qquad \text{(B3)} \\ E(\tau) &= (1-\gamma)x^u\beta\frac{1-e^{-\kappa\tau}}{\kappa}. \qquad \text{(B4)} \end{aligned}$$

C. The Myopic Strategy

If the investor has a short investment horizon, then it is optimal to adopt the myopic market-timing strategy,

$$x_{m,t} = \frac{\alpha + \beta\mu_t - r}{\gamma\sigma_P^2}.$$

The wealth process of the investor evolves according to the following dynamics:

$$\frac{dW}{W} = \left[x^m \left(\alpha + \beta\mu_t - r\right) + r\right] dt + x^m \sigma_P dz_P, \tag{C1}$$

and the investor's indirect expected utility function is given by

$$J^m(W, \mu, t) = \mathrm{E}_t \left[e^{-\rho T} \frac{W_T^{1-\gamma}}{1-\gamma}\right] = e^{-\rho t} \frac{W_t^{1-\gamma}}{1-\gamma} \phi^m(\mu, t), \tag{C2}$$

where W_T is determined via the wealth dynamics (C1).

Similar to the indirect utility function under the optimal strategy, we conjecture that the function ϕ^m is of the form

$$\phi^m = e^{A^m(\tau) + B^m(\tau) y + \frac{1}{2} C^m(\tau) y^2}. \tag{C3}$$

where $A^m(\tau)$, $B^m(\tau)$, and $C^m(\tau)$ are solutions to a system of ordinary differential equations (ODE) similar to the case of the optimal strategy:

$$\frac{dC^m}{d\tau} = a_1^m \left(C^m\right)^2 + a_2^m C^m + a_3^m, \tag{C4}$$

$$\frac{dB^m}{d\tau} = a_1^m B^m C^m + \frac{1}{2} a_2^m B^m + \kappa(\alpha + \beta\bar{\mu} - r) C^m, \tag{C5}$$

$$\frac{dA^m}{d\tau} = \frac{1}{2} a_1^m \left(C^m + \left(B^m\right)^2\right) + \kappa(\alpha + \beta\bar{\mu} - r) B^m + (1-\gamma) r - \rho, \tag{C6}$$

with the boundary conditions:

$$A^m(\tau) = 0, \quad B^m(\tau) = 0, \quad \text{and} \quad C^m(\tau) = 0 \quad \text{at} \quad \tau = 0.$$

This system of ODEs has exactly the same form of solutions as that under the optimal strategy, but with slightly different coefficients, (a_1^m, a_2^m, a_3^m). In particular,

$$a_1^m = \beta^2 \sigma_\mu^2 \neq a_1, \quad a_2^m = 2\left[\frac{(1-\gamma)\beta\sigma_\mu\rho_{P\mu}}{\gamma\sigma_P} - \kappa\right] = a_2, \quad \text{and} \quad a_3^m = \frac{1-\gamma}{\gamma\sigma_P^2} = a_3,$$

where a_2 and a_3 are given in (A3-A4) in Appendix A. Therefore, the solution to $A^m(\tau)$, $B^m(\tau)$, and $C^m(\tau)$ has exactly the same expression as that given by equations (A5-A7) in Appendix A except that a_1 is replaced by a_1^m and η is replaced by $\eta^m \equiv \sqrt{a_2^2 - 4a_1^m a_3}$.

D. The Optimal Buy-and-Hold Strategy

The dynamics of the stock price (with dividends reinvested) given in equations (1-2) implies that the stock price at the time T conditional on information at t is:

$$\begin{aligned} P_T &= P_t \exp\left\{ \left(\alpha + \beta\bar{\mu} - \frac{1}{2}\sigma_P^2\right)(T-t) + \beta\left(\mu_t - \bar{\mu}\right)\frac{1-e^{-\kappa(T-t)}}{\kappa} \right. \\ &+ \left. \beta\sigma_\mu \int_t^T \frac{1-e^{-\kappa(T-s)}}{\kappa} dz_\mu(s) + \sigma_P \int_t^T dz_P(s) \right\} \end{aligned} \tag{D1}$$

The optimal buy-and-hold strategy, x^b, which is the proportion of current wealth W_t invested in the stock, then solves the following optimization problem by normalizing W_t to one dollar:

$$\max_{x^b} \mathrm{E}_t \left[\frac{\left(x^b \frac{P_T}{P_t} + (1-x^b)e^{r(T-t)} \right)^{1-\gamma}}{1-\gamma} \right]. \tag{D2}$$

The strategy x^b, the indirect utility J^b, and the certainty equivalent wealth CEW^b are all solved numerically.

References

Amihud, Yakov, and Clifford M. Hurvich, 2003, Predictive Regressions: A Reduced-Bias Estimation Method, *Journal of Financial and Quantitative Analysis*, forthcoming.

Arnott, Robert D., and Peter L. Bernstein, 2002, What Risk Premium is "Normal"?, *Financial Analysts Journal* 58, 64-85.

Avramov, Doron, 2002. Stock return predictability and model uncertainty. *Journal of Financial Economics* 64, 423-458.

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

Black, Fisher, 1986. Noise. *Journal of Finance* 41, 529-543.

Bossaerts, Peter, and Pierre Hillion, 1999. Implementing statistical criteria to select return forecasting models: what do we learn? *Review of Financial Studies* 12, 405-428.

Boudoukh, Jacob, Matthew Richardson, and Tom Smith, 1993. Is the ex ante risk premium always positive? - A new approach to testing conditional asset pricing models. *Journal of Financial Economics* 34, 387-408.

Brav, Alon, and J. B. Heaton, 2002. Competing Theories of Financial Anomalies. *Review of Financial Studies* 15, 575-606.

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

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

Brennan, Michael J., and Walter N. Torous, 1999. Individual Decision Making and Investor Welfare. *Economic Notes* 28, 119-143.

Brennan, Michael J., and Yihong Xia, 2000. Stochastic interest rates and the bond-stock mix. *The European Finance Review* 4, 197-210.

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

Brennan, Michael J., and Yihong Xia, 2001. Assessing asset pricing anomalies. *Review of*

Financial Studies 14, 905-942.

Brennan, Michael J., and Yihong Xia, 2002. Dynamic Asset Allocation under Inflation. *The Journal of Finance* 57, 1201-1238.

Campbell, John Y., 2001, Why long horizons? A study of power against persistent alternative, *Journal of Empirical Finance* 8, 459-491.

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinaly, 1997, The Econometrics of Financial Markets, Princeton University Press, Princeton, New Jersey.

Campbell, John Y., and Robert J. Shiller, 1988. The dividend-price ratio and expectations of future dividends and discount factors. *The Review of Financial Studies* 1, 195-227.

Campbell, John Y., and Robert J. Shiller, 1998, Valuation ratios and the long run stock market outlook, *Journal of Portfolio Management*, 24, 11-26.

Campbell, John Y., and Robert J. Shiller, 2001, Valuation ratios and the long run stock market outlook: an update, forthcoming in *Advances in Behavioral Finance*, Volume II, Nicholas Barberis and Richard Thaler eds., Russell Sage Foundation, 2004.

Campbell, John Y., and Samuel B. Thompson, 2004, Predicting the Equity Premium Out of Sample: Can Anything Beat the Historical Average?, working paper, Harvard University.

Campbell, John Y., and Luis Viceira, 2001, Who should buy long term bonds?, *American Economic Review* 91, 99-127.

Campbell, John Y., and Motohiro Yogo, 2003, Efficient test of stock return predictability, working paper, Harvard University.

Cochrane, John. H., 1997, Where is the market going? Uncertain facts and novel theories, *Economic Perspectives* XXI (6) (Federal Reserve Bank of Chicago), 3-37.

Cochrane, John. H., 1999. New facts in finance. *Economic Perspectives* XXIII (3) (Federal Reserve Bank of Chicago), 36-58.

Engstrom, Eric, 2003, The Conditional Relationship Between the Equity Risk Premium and the Dividend Price Ratio, working paper, Columbia University.

Fama, Eugene F., and Robert R. Bliss, 1987, The information in long-maturity forward rates, *American Economic Review* 77, 680-692.

Fama, Eugene F., and Kenneth R. French, 1988a, Dividend yields and expected stock returns, *Journal of Financial Economics* 23, 3-25.

Fama, Eugene F., and Kenneth R. French, 1988b, Permanent and temporary components of stock prices, *Journal of Political Economy* 96, 246-273.

Fama, Eugene F., and Kenneth R. French, 2002, The Equity Premium, *Journal of Finance*, 57, 637-659.

Goetzmann, William N. and Philippe Jorion, 1993. Testing the predictive power of dividend yields. *The Journal of Finance* 48, 663-679.

Goyal, Amit, and Ivo Welch, 2003. Predicting the equity premium with dividend ratios. *Management Science* 49, 639-654.

Goyal, Amit, and Ivo Welch, 2004. A Comprehensive Look at The Empirical Performance of Equity Premium Prediction. Working paper, Yale University.

Hakansson, Nils H., 1970. Optimal Investment and Consumption Strategies Under Risk for a Class of Utility Functions. *Econometrica* 38, 587-607.

Hodrick, Robert J., 1992. Dividend yields and expected stock returns: Alternative procedure for inference and measurement. *Review of Financial Studies* 5, 357-386.

Ilmanen, Antti, 2003, Expected returns on stocks and bonds, *Journal of Portfolio Management* Winter 2003, 7-27.

Kandel, Shmuel, and Robert F. Stambaugh, 1996. On the Predictability of Stock Returns: An Asset-Allocation Perspective. *The Journal of Finance* 51, 385-424.

Keim, Donald B., and Robert F. Stambaugh, 1986. Predicting returns in the stock and bond markets. *Journal of Financial Economics* 17, 357-390.

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

Kleidon, Allan W., 1986. Variance Bounds Tests and Stock Price Valuation Models. *Journal of Political Economy* 94, 953-1001.

Lewellen, Jonathan W., 2002. Momentum and autocorrelation in stock returns. *Review of Financial Studies* 15, 533-563.

Lewellen, Jonathan W., 2003. Predicting returns using financial ratios. *Journal of Financial Economics*, forthcoming.

Lewellen, Jonathan W., and Jay A. Shanken, 2002. Learning, asset-pricing tests, and market efficiency. *The Journal of Finance* 57, 1113{1145.

Liu, Jun, and Jun Pan, 2003, Dynamic Derivative Strategies, *Journal of Financial Economics* 69,401-430.

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, 41-66.

Malkiel, Burton G., 1990. *A Random Walk Down Wall Street*, New York: Norton.

Malkiel, Burton G., 2003. The Efficient Market Hypothesis and Its Critics. *Journal of Economic Perspectives* 17, 59-82.

Marsh, Terry A., and Robert C. Merton, 1986. Dividend Variability and Variance Bounds Tests for the Rationality of Stock Market Prices. *American Economic Review* 76, 483-498.

Menzly, Lior, Tano Santos, and Pietro Veronesi, 2004, Understanding Predictability, *Journal of Political Economy* 112, 1-47.

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

Philips, Thomas K., Greg T. Rogers, and Robert E. Capaldi, 1996. Tactical asset allocation: 1977-94. *Journal of Portfolio Management* 23, 57-64.

Stambaugh, Robert F., 1999. Predictive regression. *Journal of Financial Economics* 54, 375-421.

Vuolteenaho, Tuomo, 2002, What Drives Firm-Level Stock Returns?, *Journal of Finance* 57, 233-264.

Wachter, Jessica, 2002. Portfolio and consumption decisions under mean-reverting returns: an exact solution for complete markets. *Journal of Financial and Quantitative Analysis* 37, 63-91.

Xia, Yihong, 2001. Learning about predictability: the effect of parameter uncertainty on dynamic asset allocation. *The Journal of Finance* 56, 205-246.

Table 1

Stock Return Model Parameters and Parameter Estimates

Panel A reports theoretical variance ratios calculated using equation (5) and the returns generated under each of the scenarios. The historical data for 1871-1985 are from Table 3 of Poterba and Summers (1988) and those for 1871-2002 are calculated using the real S&P 500 index returns from Robert Shiller's website. Panels B-C report the distributions of statistics from ordinary least squares predictive regressions with non-overlapping observations of the form

$$R(t, t+\tau) = a_\tau + b_\tau \mu_t + \epsilon$$

for $\tau = 1, 12$ months. The distributions are estimated from 2000 simulations of the return process given by equations (3) and (4) for 840 months using the parameter values for each of 9 scenarios. Panels D-E report the bias-corrected regression coefficients and their $t-$ratios using the Amihud-Hurvich (2004) approach where both the point estimate and the standard errors are adjusted. The Stambaugh (1997) correction yields almost the same point estimates and is omitted from the table. The $t-$ratios are reported in parentheses. The exogenous parameters are $\nu_\mu = 4\%$, $\bar{\mu} = 7\%$, $r = 1\%$, $\sqrt{\text{Var}(R(1))} = 14\%$, $\alpha = 0$, and $\beta = 1$.

Scenarios	(i)	(ii)	(iii)	(iv)	(v)	**(vi)**	(vii)	(viii)	(ix)		
κ	0.02	0.02	0.02	0.10	0.10	**0.10**	0.50	0.50	0.50		
σ_μ	0.008	0.008	0.008	0.018	0.018	**0.018**	0.040	0.040	0.040		
$\rho_{P\mu}$	0.00	-0.50	-0.90	0.00	-0.50	**-0.90**	0.00	-0.50	-0.90		
σ_P	0.140	0.142	0.144	0.140	0.144	**0.148**	0.139	0.147	0.155		
						Panel A: Variance Ratios				Historical	
										1871-1985	1871-2003
1 month(s)	1.00	1.03	1.05	1.00	1.05	**1.10**	0.98	1.10	1.20		
6	1.00	1.01	1.03	1.00	1.03	**1.05**	0.99	1.05	1.10		
1 year(s)	1.00	1.00	1.00	1.00	1.00	**1.00**	1.00	1.00	1.00	1.00	1.00
2	1.00	0.98	0.95	1.01	0.95	**0.90**	1.04	0.94	0.86	1.04	1.03
3	1.01	0.95	0.91	1.03	0.92	**0.82**	1.07	0.91	0.77	0.88	0.92
4	1.02	0.93	0.86	1.06	0.89	**0.75**	1.11	0.89	0.70	0.88	0.74
5	1.02	0.91	0.82	1.09	0.87	**0.69**	1.13	0.88	0.66	0.86	0.80
						Panel B: One-month R^2					
Asymptotic	0.7%	0.7%	0.6%	0.7%	0.6%	**0.6%**	0.7%	0.6%	0.5%		
25%	0.1%	0.2%	0.5%	0.2%	0.4%	**0.6%**	0.3%	0.3%	0.4%		
Median	0.2%	0.4%	0.7%	0.5%	0.6%	**0.8%**	0.6%	0.6%	0.7%		
Mean	0.3%	0.5%	0.7%	0.6%	0.7%	**0.9%**	0.7%	0.7%	0.7%		
75%	0.5%	0.7%	1.0%	0.9%	1.0%	**1.1%**	1.0%	1.1%	1.0%		
						Panel C: Twelve-month R^2					
Asymptotic	7.4%	7.4%	7.4%	6.9%	6.9%	**6.9%**	4.8%	4.8%	4.8%		
25%	0.6%	2.3%	5.4%	2.0%	3.7%	**6.4%**	1.7%	2.4%	3.3%		
Median	2.3%	5.0%	7.9%	5.1%	6.8%	**9.0%**	4.6%	5.2%	6.0%		
Mean	3.0%	6.9%	8.4%	6.6%	7.7%	**9.5%**	5.7%	6.4%	6.8%		
75%	5.8%	8.5%	11.0%	9.6%	10.8%	**12.1%**	8.4%	9.2%	9.5%		

Table 1 (continued)

	Panel D: One-month Bias-Corrected Regression Coefficient b								
Theoretical b_τ	0.083	0.083	0.083	0.083	0.083	**0.083**	0.082	0.082	0.082
	(2.39)	(2.36)	(2.33)	(2.38)	(2.32)	**(2.27)**	(2.36)	(2.23)	(2.14)
25%	0.038	0.041	0.037	0.058	0.053	**0.051**	0.057	0.056	0.054
	(0.58)	(0.65)	(0.65)	(1.32)	(1.28)	**(1.29)**	(1.56)	(1.49)	(1.44)
Median	0.082	0.082	0.081	0.085	0.081	**0.080**	0.081	0.080	0.078
	(1.27)	(1.26)	(1.22)	(2.05)	(1.88)	**(1.81)**	(2.29)	(2.10)	(1.94)
Mean	0.082	0.094	0.101	0.085	0.086	**0.089**	0.081	0.082	0.082
	(1.31)	(1.26)	(1.21)	(2.05)	(1.88)	**(1.81)**	(2.29)	(2.11)	(1.97)
75%	0.127	0.133	0.146	0.112	0.113	**0.117**	0.104	0.107	0.107
	(2.08)	(1.87)	(1.79)	(2.80)	(2.48)	**(2.30)**	(2.96)	(2.76)	(2.52)
	Panel E: Twelve-month Bias-Corrected Regression Coefficient b								
25%	0.453	0.491	0.419	0.654	0.595	**0.575**	0.495	0.481	0.464
	(0.56)	(0.65)	(0.62)	(1.22)	(1.20)	**(1.22)**	(1.11)	(1.09)	(1.08)
Median	0.987	0.980	0.933	0.974	0.924	**0.910**	0.778	0.761	0.741
	(1.28)	(1.22)	(1.18)	(1.93)	(1.78)	**(1.74)**	(1.83)	(1.74)	(1.66)
Mean	0.987	1.113	1.170	0.971	0.963	**1.005**	0.787	0.791	0.784
	(1.30)	(1.25)	(1.19)	(1.96)	(1.81)	**(1.77)**	(1.83)	(1.78)	(1.69)
75%	1.549	1.587	1.717	1.314	1.283	**1.337**	1.059	1.080	1.070
	(2.04)	(1.87)	(1.78)	(2.73)	(2.42)	**(2.29)**	(2.51)	(2.41)	(2.26)

Table 2

Stock Return Models and Valuation Ratios

This table reports valuation ratios for a security whose dividend follows:

$$\frac{dD}{D} = gdt + \sigma_D dz_D$$

and the expected rate of return on the security follows an Ornstein-Uhlenbeck process:

$$d\mu = \kappa(\bar{\mu} - \mu)dt + \sigma_\mu dz_\mu.$$

The parameters of the real dividend growth rate process, $g = 0.86\%$ and $\sigma_D = 12.02\%$, are taken as the sample mean and volatility of the real dividend growth rates from 1950 to 2002, provided by Robert Shiller. The two correlations for each scenario, $\rho_{D\mu}$, are chosen so that the volatility of the real stock return is approximately 14%. Results are reported only for scenarios (iii), (vi), and (ix), because it was not possible to generate stock return volatility of 14% in the other scenarios. The exogenous parameters are $\nu_\mu = 4\%$, $\bar{\mu} = 7\%$, $r = 1\%$, $\sqrt{\mathrm{Var}(R(1))} = 14\%$, $\alpha = 0$, and $\beta = 1$.

Scenarios	(iii)		**(vi)**		(ix)	
κ	0.02		**0.10**		0.50	
σ_μ	0.008		**0.018**		0.040	
$\rho_{P\mu}$	-0.90		**-0.90**		-0.90	
σ_P	0.144		**0.148**		0.155	
$\rho_{D\mu}$	0.679	-0.679	**0.655**	**-0.655**	0.610	-0.610
Panel A: Price Dividend Ratios						
$\mu_t = \bar{\mu} - \nu_\mu = 3\%$	31.39	64.01	**21.03**	**27.01**	16.91	19.00
$\mu_t = \bar{\mu} = 7\%$	17.05	27.54	**16.46**	**20.86**	15.80	17.76
$\mu_t = \bar{\mu} + \nu_\mu = 11\%$	10.79	13.86	**12.96**	**16.11**	14.74	16.56
Panel B: Dividend Yields						
$\mu_t = \bar{\mu} - \nu_\mu = 5\%$	3.2%	1.6%	**4.8%**	**3.7%**	5.9%	5.3%
$\mu_t = \bar{\mu} = 9\%$	5.9%	3.6%	**6.1%**	**4.8%**	6.3%	5.6%
$\mu_t = \bar{\mu} + \nu_\mu = 13\%$	9.3%	7.2%	**7.7%**	**6.2%**	6.8%	6.0%

Table 3a

Long Run Return Predictability and the Value of Dynamic versus Unconditional Strategies

This table reports the theoretical values under different scenarios of R^2 from regressions of long run returns on the value of μ at the beginning of the period. It also reports the ratios of the certainty equivalent wealth for an optimal dynamic strategy to the certainty equivalent wealth under an unconditional strategy for different horizons and initial values of μ_t. The exogenous parameters are $\nu_\mu = 4\%$, $\bar{\mu} = 7\%$, $r = 1\%$, $\sqrt{\mathrm{Var}(R(1))} = 14\%$, $\alpha = 0$, and $\beta = 1$. The risk aversion parameter is $\gamma = 5$.

Horizon	Scenarios	(i)	(ii)	(iii)	(iv)	(v)	**(vi)**	(vii)	(viii)	(ix)
	κ	0.02	0.02	0.02	0.10	0.10	**0.10**	0.50	0.50	0.50
	σ_μ	0.008	0.008	0.008	0.018	0.018	**0.018**	0.040	0.040	0.040
	$\rho_{P\mu}$	0.00	-0.50	-0.90	0.00	-0.50	**-0.90**	0.00	-0.50	-0.90
	σ_P	0.140	0.142	0.144	0.140	0.144	**0.148**	0.139	0.147	0.155
1 year	R^2	7.4%	7.4%	7.4%	6.9%	6.9%	**6.9%**	4.8%	4.8%	4.8%
	$\mu_t = 0.03$	1.01	1.01	1.01	1.01	1.01	**1.01**	1.01	1.01	1.01
	$\mu_t = 0.07$	1.00	1.00	1.00	1.00	1.00	**1.00**	1.00	1.00	1.00
	$\mu_t = 0.11$	1.01	1.01	1.01	1.01	1.01	**1.01**	1.01	1.01	1.01
5 years	R^2	26.5%	28.8%	31.0%	18.8%	22.4%	**26.8%**	4.6%	5.9%	7.8%
	$\mu_t = 0.03$	1.04	1.03	1.03	1.05	1.03	**1.02**	1.05	1.04	1.04
	$\mu_t = 0.07$	1.00	1.00	1.01	1.02	1.02	**1.03**	1.03	1.04	1.05
	$\mu_t = 0.11$	1.04	1.06	1.08	1.03	1.06	**1.11**	1.04	1.06	1.09
10 years	R^2	38.0%	44.1%	50.7%	20.4%	27.8%	**39.8%**	2.6%	3.6%	5.6%
	$\mu_t = 0.03$	1.10	1.06	1.04	1.11	1.07	**1.07**	1.09	1.09	1.13
	$\mu_t = 0.07$	1.02	1.02	1.04	1.05	1.06	**1.13**	1.08	1.09	1.15
	$\mu_t = 0.11$	1.06	1.14	1.25	1.06	1.14	**1.35**	1.08	1.11	1.20
20 years	R^2	45.6%	56.8%	71.1%	15.9%	24.6%	**46.1%**	1.3%	1.9%	3.3%
	$\mu_t = 0.03$	1.27	1.14	1.09	1.30	1.17	**1.29**	1.20	1.19	1.36
	$\mu_t = 0.07$	1.07	1.07	1.25	1.19	1.16	**1.57**	1.18	1.20	1.39
	$\mu_t = 0.11$	1.09	1.32	2.16	1.16	1.27	**2.25**	1.18	1.22	1.46
40 years	R^2	43.7%	58.0%	79.3%	8.8%	15.2%	**38.3%**	0.6%	1.0%	1.8%
	$\mu_t = 0.03$	2.17	1.46	1.42	1.90	1.43	**2.32**	1.43	1.45	2.00
	$\mu_t = 0.07$	1.50	1.27	2.43	1.70	1.42	**3.25**	1.41	1.45	2.05
	$\mu_t = 0.11$	1.29	1.58	9.02	1.62	1.56	**5.47**	1.41	1.48	2.14

Table 3b

Long Run Return Predictability and the Value of Dynamic versus Myopic Strategies

This table reports the theoretical values under different scenarios of R^2 from regressions of long run returns on the value of μ at the beginning of the period. It also reports the ratios of the certainty equivalent wealth for an optimal dynamic strategy to the certainty equivalent wealth under a myopic strategy for different horizons and initial values of μ_t. The exogenous parameters are $\nu_\mu = 4\%$, $\bar{\mu} = 7\%$, $r = 1\%$, $\sqrt{\mathrm{Var}(R(1))} = 14\%$, $\alpha = 0$, and $\beta = 1$. The risk aversion parameter is $\gamma = 5$.

Horizon	Scenarios	(i)	(ii)	(iii)	(iv)	(v)	**(vi)**	(vii)	(viii)	(ix)
	κ	0.02	0.02	0.02	0.10	0.10	**0.10**	0.50	0.50	0.50
	σ_μ	0.008	0.008	0.008	0.018	0.018	**0.018**	0.040	0.040	0.040
	$\rho_{P\mu}$	0.00	-0.50	-0.90	0.00	-0.50	**-0.90**	0.00	-0.50	-0.90
	σ_P	0.140	0.142	0.144	0.140	0.144	**0.148**	0.139	0.147	0.155
1 year	R^2	7.4%	7.4%	7.4%	6.9%	6.9%	**6.9%**	4.8%	4.8%	4.8%
	$\mu_t = 0.03$	1.00	1.00	1.00	1.00	1.00	**1.00**	1.00	1.00	1.00
	$\mu_t = 0.07$	1.00	1.00	1.00	1.00	1.00	**1.00**	1.00	1.00	1.00
	$\mu_t = 0.11$	1.00	1.00	1.00	1.00	1.00	**1.00**	1.00	1.00	1.00
5 years	R^2	26.5%	28.8%	31.0%	18.8%	22.4%	**26.8%**	4.6%	5.9%	7.8%
	$\mu_t = 0.03$	1.00	1.00	1.00	1.00	1.00	**1.00**	1.00	1.00	1.01
	$\mu_t = 0.07$	1.00	1.00	1.00	1.00	1.00	**1.01**	1.00	1.00	1.01
	$\mu_t = 0.11$	1.00	1.00	1.00	1.00	1.00	**1.02**	1.00	1.00	1.02
10 years	R^2	38.0%	44.1%	50.7%	20.4%	27.8%	**39.8%**	2.6%	3.6%	5.6%
	$\mu_t = 0.03$	1.00	1.00	1.00	1.00	1.00	**1.02**	1.00	1.01	1.04
	$\mu_t = 0.07$	1.00	1.00	1.02	1.00	1.01	**1.06**	1.00	1.01	1.05
	$\mu_t = 0.11$	1.00	1.01	1.04	1.00	1.02	**1.12**	1.00	1.01	1.06
20 years	R^2	45.6%	56.8%	71.1%	15.9%	24.6%	**46.1%**	1.3%	1.9%	3.3%
	$\mu_t = 0.03$	1.00	1.01	1.04	1.00	1.02	**1.20**	1.00	1.02	1.12
	$\mu_t = 0.07$	1.00	1.02	1.16	1.00	1.04	**1.38**	1.00	1.02	1.13
	$\mu_t = 0.11$	1.00	1.05	1.41	1.00	1.06	**1.66**	1.00	1.02	1.14
40 years	R^2	43.7%	58.0%	79.3%	8.8%	15.2%	**38.3%**	0.6%	1.0%	1.8%
	$\mu_t = 0.03$	1.00	1.04	1.35	1.00	1.08	**2.10**	1.00	1.05	1.31
	$\mu_t = 0.07$	1.00	1.09	2.32	1.00	1.11	**2.76**	1.00	1.05	1.33
	$\mu_t = 0.11$	1.00	1.20	5.56	1.00	1.14	**3.86**	1.00	1.05	1.34

Table 3c

Long Run Return Predictability and the Value of Dynamic versus Buy-and-Hold Strategies

This table reports the theoretical values under different scenarios of R^2 from regressions of long run returns on the value of μ at the beginning of the period. It also reports the ratios of the certainty equivalent wealth for an optimal dynamic strategy to the certainty equivalent wealth under an optimal buy-and-hold strategy for different horizons and initial values of μ_t. The exogenous parameters are $\nu_\mu = 4\%$, $\bar{\mu} = 7\%$, $r = 1\%$, $\sqrt{\mathrm{Var}(R(1))} = 14\%$, $\alpha = 0$, and $\beta = 1$. The risk aversion parameter is $\gamma = 5$.

Horizon	Scenarios	(i)	(ii)	(iii)	(iv)	(v)	**(vi)**	(vii)	(viii)	(ix)
	κ	0.02	0.02	0.02	0.10	0.10	**0.10**	0.50	0.50	0.50
	σ_μ	0.008	0.008	0.008	0.018	0.018	**0.018**	0.040	0.040	0.040
	$\rho_{P\mu}$	0.00	-0.50	-0.90	0.00	-0.50	**-0.90**	0.00	-0.50	-0.90
	σ_P	0.140	0.142	0.144	0.140	0.144	**0.148**	0.139	0.147	0.155
1 year	R^2	7.4%	7.4%	7.4%	6.9%	6.9%	**6.9%**	4.8%	4.8%	4.8%
	$\mu_t = 0.03$	1.00	1.00	1.00	1.00	1.00	**1.00**	1.01	1.01	1.01
	$\mu_t = 0.07$	1.01	1.01	1.01	1.01	1.01	**1.02**	1.01	1.01	1.02
	$\mu_t = 0.11$	1.02	1.02	1.02	1.02	1.02	**1.03**	1.03	1.03	1.03
5 years	R^2	26.5%	28.8%	31.0%	18.8%	22.4%	**26.8%**	4.6%	5.9%	7.8%
	$\mu_t = 0.03$	1.02	1.02	1.01	1.04	1.04	**1.05**	1.08	1.09	1.10
	$\mu_t = 0.07$	1.07	1.08	1.09	1.08	1.10	**1.12**	1.10	1.12	1.14
	$\mu_t = 0.11$	1.14	1.15	1.16	1.13	1.16	**1.18**	1.12	1.15	1.17
10 years	R^2	38.0%	44.1%	50.7%	20.4%	27.8%	**39.8%**	2.6%	3.6%	5.6%
	$\mu_t = 0.03$	1.04	1.05	1.05	1.10	1.13	**1.18**	1.19	1.24	1.32
	$\mu_t = 0.07$	1.16	1.19	1.24	1.17	1.23	**1.36**	1.21	1.27	1.36
	$\mu_t = 0.11$	1.29	1.35	1.43	1.26	1.36	**1.53**	1.24	1.30	1.42
20 years	R^2	45.6%	56.8%	71.1%	15.9%	24.6%	**46.1%**	1.3%	1.9%	3.3%
	$\mu_t = 0.03$	1.11	1.15	1.21	1.24	1.38	**1.76**	1.44	1.59	1.90
	$\mu_t = 0.07$	1.31	1.47	1.83	1.34	1.56	**2.27**	1.46	1.63	1.96
	$\mu_t = 0.11$	1.56	1.88	2.57	1.48	1.81	**2.92**	1.51	1.68	2.05
40 years	R^2	43.7%	58.0%	79.3%	8.8%	15.2%	**38.3%**	0.6%	1.0%	1.8%
	$\mu_t = 0.03$	1.27	1.47	2.11	1.58	2.09	**4.81**	2.11	2.65	3.97
	$\mu_t = 0.07$	1.55	2.12	5.24	1.70	2.40	**6.73**	2.15	2.71	4.07
	$\mu_t = 0.11$	2.00	3.28	15.66	1.92	2.91	**10.28**	2.22	2.81	4.25

Table 4

Parameter Estimates of the μ series

This table reports the parameter estimates associated with the instantaneous expected return, μ, series, which are derived from the long run discount rate k. The two real k series are calculated in Arnott and Bernstein (2002) (A&B) and Ilmanen (2002) (IL), while the two nominal k series are provided by Barclays Global Investors (BGI) and Wilshire Associates (WA). In each k series, we derive μ under two cases. In Case I, the dividend growth rate g is assumed to be a constant. In Case II, the dividend growth rate g is assumed to follow a mean-reverting process. When the real k from Ilmanen or A&Bis used as the long run discount rate, g is set to 0.86% in the first case, while $\bar{g}$ is set to 0.86% in the second case. When the nominal k is used, g is set to 4.82% in the first case, while $\bar{g}$ is set to 4.82% in the second case.

Parameters	κ_μ	σ_μ	$\bar{\mu}$	ν_μ	$\rho_{P\mu}$	σ_D	$\rho_{\mu D}$	κ_g	σ_g	$\rho_{\mu g}$	ρ_{Dg}
	Scenario (vi) Table 2										
	0.100	0.0180	0.090	0.0400	-0.900	0.1200	-0.879				
	Real Models										
A&B $\mu^{1,1}$	0.085	0.0173	0.047	0.0419	-0.977	0.0852	-0.126				
A&B $\mu^{1,2}$	0.083	0.0172	0.045	0.0423	-0.981	0.0776	-0.106	0.103	0.0090	0.328	-0.413
IL $\mu^{2,1}$	0.122	0.0196	0.066	0.0397	-0.884	0.0852	-0.117				
IL $\mu^{2,2}$	0.115	0.0224	0.066	0.0467	-0.885	0.0822	-0.066	0.025	0.0034	0.379	-0.055
	Nominal Models										
BGI $\mu^{3,1}$	0.091	0.0239	0.133	0.0560	-0.812	0.0859	0.234				
BGI $\mu^{3,2}$	0.085	0.0214	0.113	0.0519	-0.657	0.2076	0.249	0.209	0.0294	0.378	-0.326
WA $\mu^{4,1}$	0.122	0.0336	0.137	0.0680	-0.682	0.0859	0.095				
WA $\mu^{4,2}$	0.095	0.0266	0.111	0.0611	-0.714	0.2202	0.209	0.220	0.0339	0.402	-0.417

Table 5

Quarterly Return Prediction

This table reports the results of regressing real and nominal quarterly returns on the S&P500 stock index on forecasts of the return at the beginning of the quarter calculated from the estimated value of $\mu^{i,2}$ $(i = 1, 2, 3, 4)$ and the estimated parameters of the joint stochastic process using equation.

$$R(t, t + 0.25) = a_0 + a_1 \left[\frac{1.0 - e^{-\kappa/4}}{\kappa} \right] \mu_t^{i,2} + \epsilon_t, \quad i = 1, 2, 3, 4,$$

where $R(t, t + 0.25)$ is the one quarter real return on the S&P500 index in Panel A and is the corresponding nominal return in Panel B. In Panel A, the real return on the S&P 500 index is regressed on the estimated real A&B and IL $\mu^{i,2}$ $(i = 1, 2)$ series. In Panel B, *nominal* returns on the index is regressed on the estimated nominal BGI and WA $\mu^{i,2}$ $(i = 3, 4)$ series. The OLS $t-$ ratios are reported in parenthesis and the Newey-West adjusted $t-$ratios are in brackets.

	A. Real Return Predictive Regressions							
			$\mu^{1,2}$ as Predictor			$\mu^{2,2}$ as Predictor		
	Sample Period	Obs.	a_0	a_1	R^2(%)	a_0	a_1	R^2(%)
1.	1950.2-2002.2	209	0.005	0.874	1.43	0.009	0.701	1.12
			(0.46)	(1.74)		(0.86)	(1.53)	
			[0.43]	[1.60]		[0.79]	[1.27]	
2.	1950.2-1974.2, 1974.4-2000.2	208	0.005	0.981	1.94	0.008	0.800	1.57
			(0.43)	(2.02)		(0.85)	(1.81)	
			[0.38]	[1.83]		[0.75]	[1.50]	
3.	1950.2-1974.2	97	-0.019	2.157	4.37	-0.020	1.961	5.79
			(0.92)	(2.08)		(1.09)	(2.42)	
			[1.04]	[2.45]		[1.06]	[2.21]	
4.	1974.4-2002.2	111	0.011	0.708	1.35	0.017	0.494	0.70
			(0.82)	(1.22)		(1.40)	(0.88)	
			[1.07]	[1.86]		[2.36]	[1.80]	
	B. Nominal Return Predictive Regressions							
			$\mu^{3,2}$ as Predictor			$\mu^{4,2}$ as Predictor		
	Sample Period	Obs.	a_0	a_1	R^2	a_0	a_1	R^2
1.	1973.2-2002.2	117	0.003	1.026	2.89	0.006	0.924	2.74
			(0.20)	(1.86)		(0.37)	(1.80)	
			[0.18]	[1.69]		[0.34]	[1.66]	
2.	1973.2-1974.2, 1974.4-2000.2	116	0.008	0.953	2.75	0.010	0.852	2.57
			(0.48)	(1.80)		(0.67)	(1.73)	
			[0.41]	[1.60]		[0.54]	[1.42]	

Figure 1
Theoretical R^2 as a Function of the Horizon

The figure plots the theoretical regression R^2 as a function of the investment horizon τ, where τ varies from 0.08 (one month) to 20 years for the nine scenarios reported in Table 1. The exogenous parameters are $\nu_\mu = 4\%$, $\bar{\mu} = 9\%$, $r = 3\%$, $\sqrt{\mathrm{Var}(R(1))} = 14\%$, $\alpha = 0$, and $\beta = 1$.

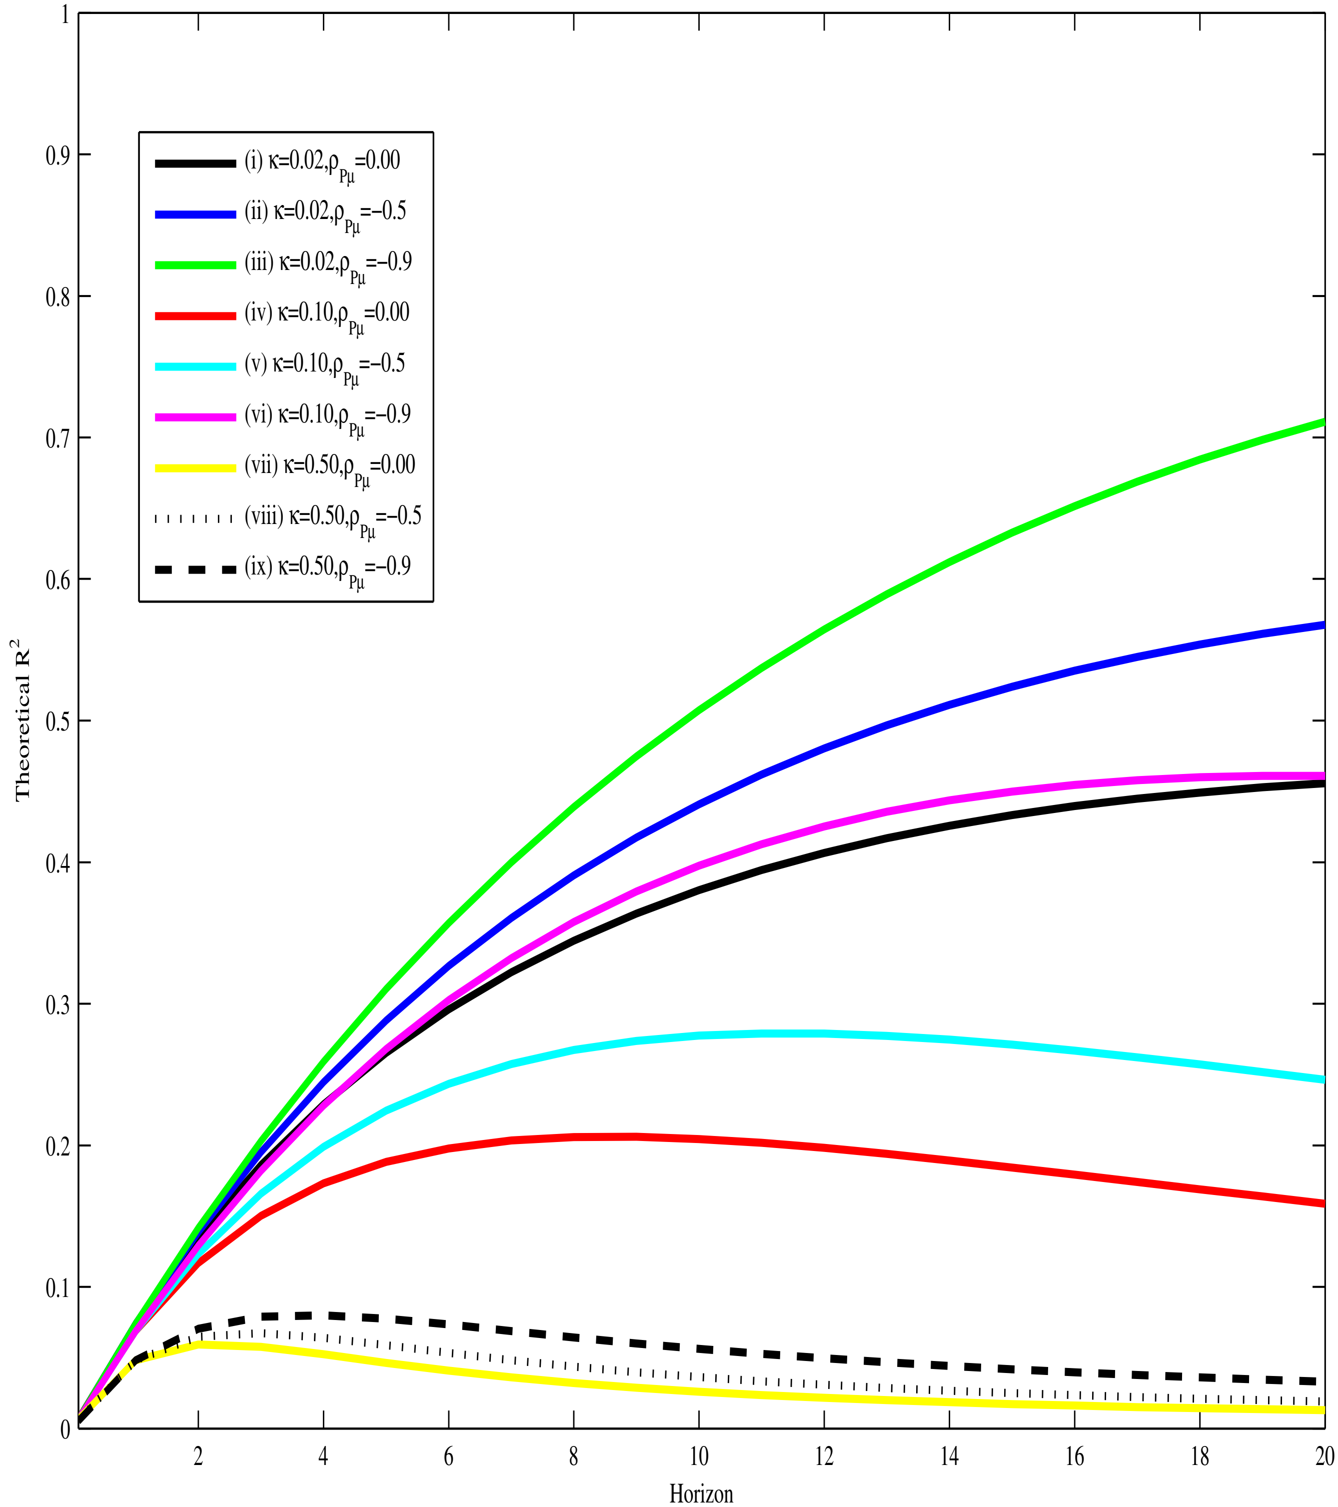

Figure 2
Theoretical Ten- and Twenty-Year R^2 Versus the One-Year R^2

The figure plots the theoretical regression R^2 at long horizon (ten and twenty years) as a function of the one-year R^2 for the nine scenarios reported in Table 1. The exogenous parameters are $\nu_\mu = 4\%$, $\bar{\mu} = 9\%$, $r = 3\%$, $\sqrt{\mathrm{Var}(R(1))} = 14\%$, $\alpha = 0$, and $\beta = 1$.

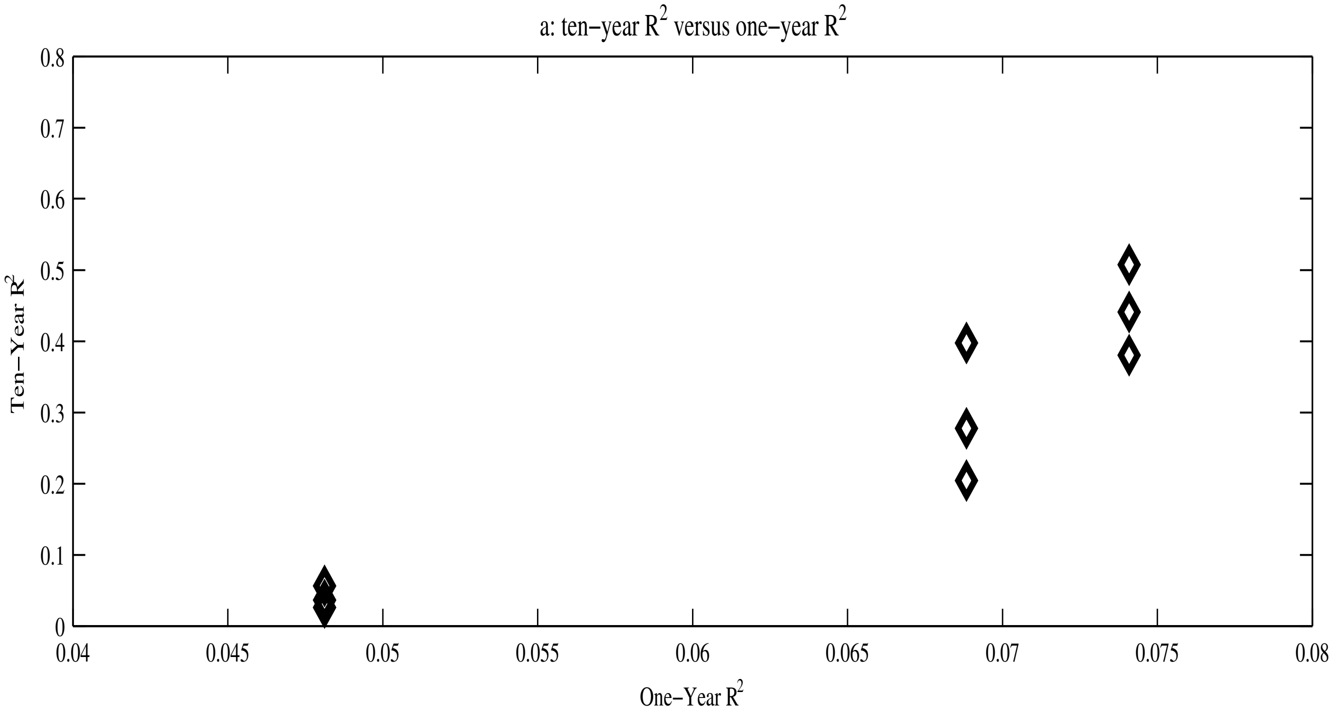

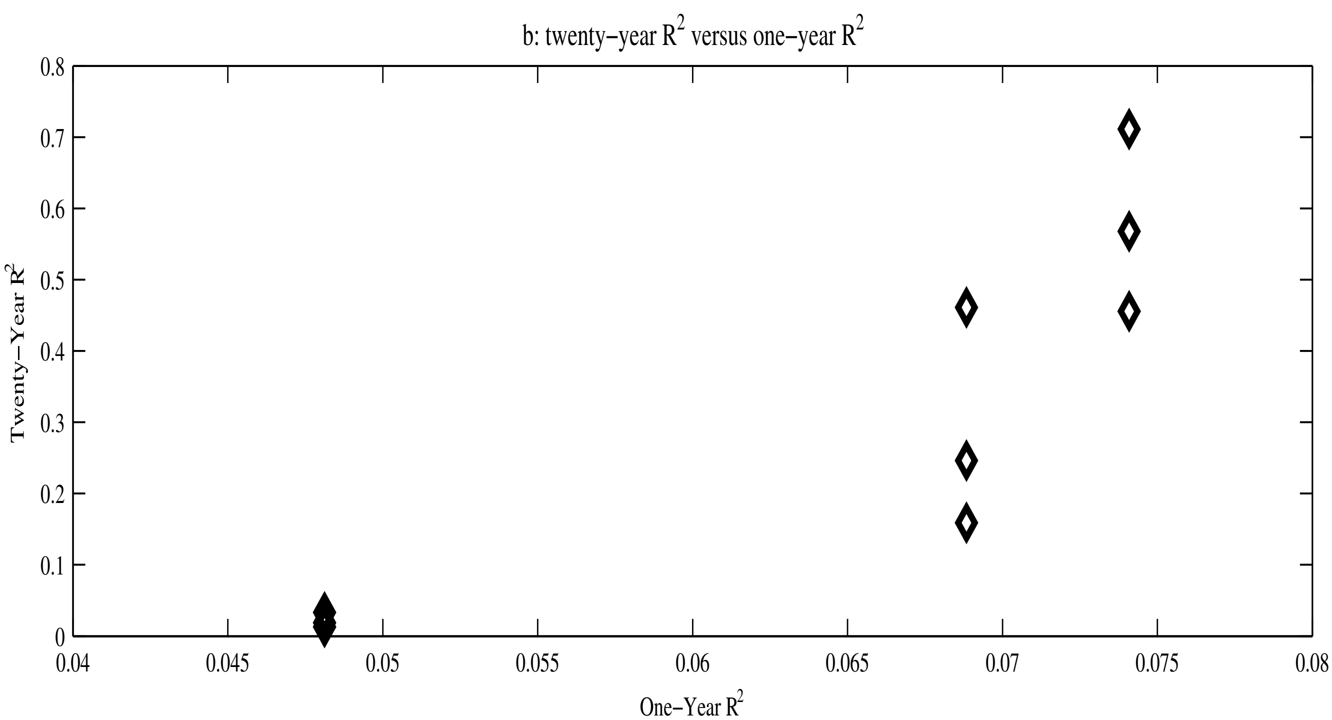

Figure 3

The Certainty Equivalent Wealth Ratio under the Optimal and the Unconditional Strategies for a Twenty-Year Horizon Versus the One-Year Predictive Regression R^2

The figure plots the certainty equivalent wealth ratios between the optimal and the unconditional strategies for a twenty-year horizon investor, $CEWR^{ou}$, as a function of the one-year R^2 for the nine scenarios reported in Table 1. The initial value of μ is set at respectively 5%, 9% and 13%. The exogenous parameters are $\nu_\mu = 4\%$, $\bar{\mu} = 9\%$, $r = 3\%$, $\sqrt{\mathrm{Var}(R(1))} = 14\%$, $\alpha = 0$, and $\beta = 1$.

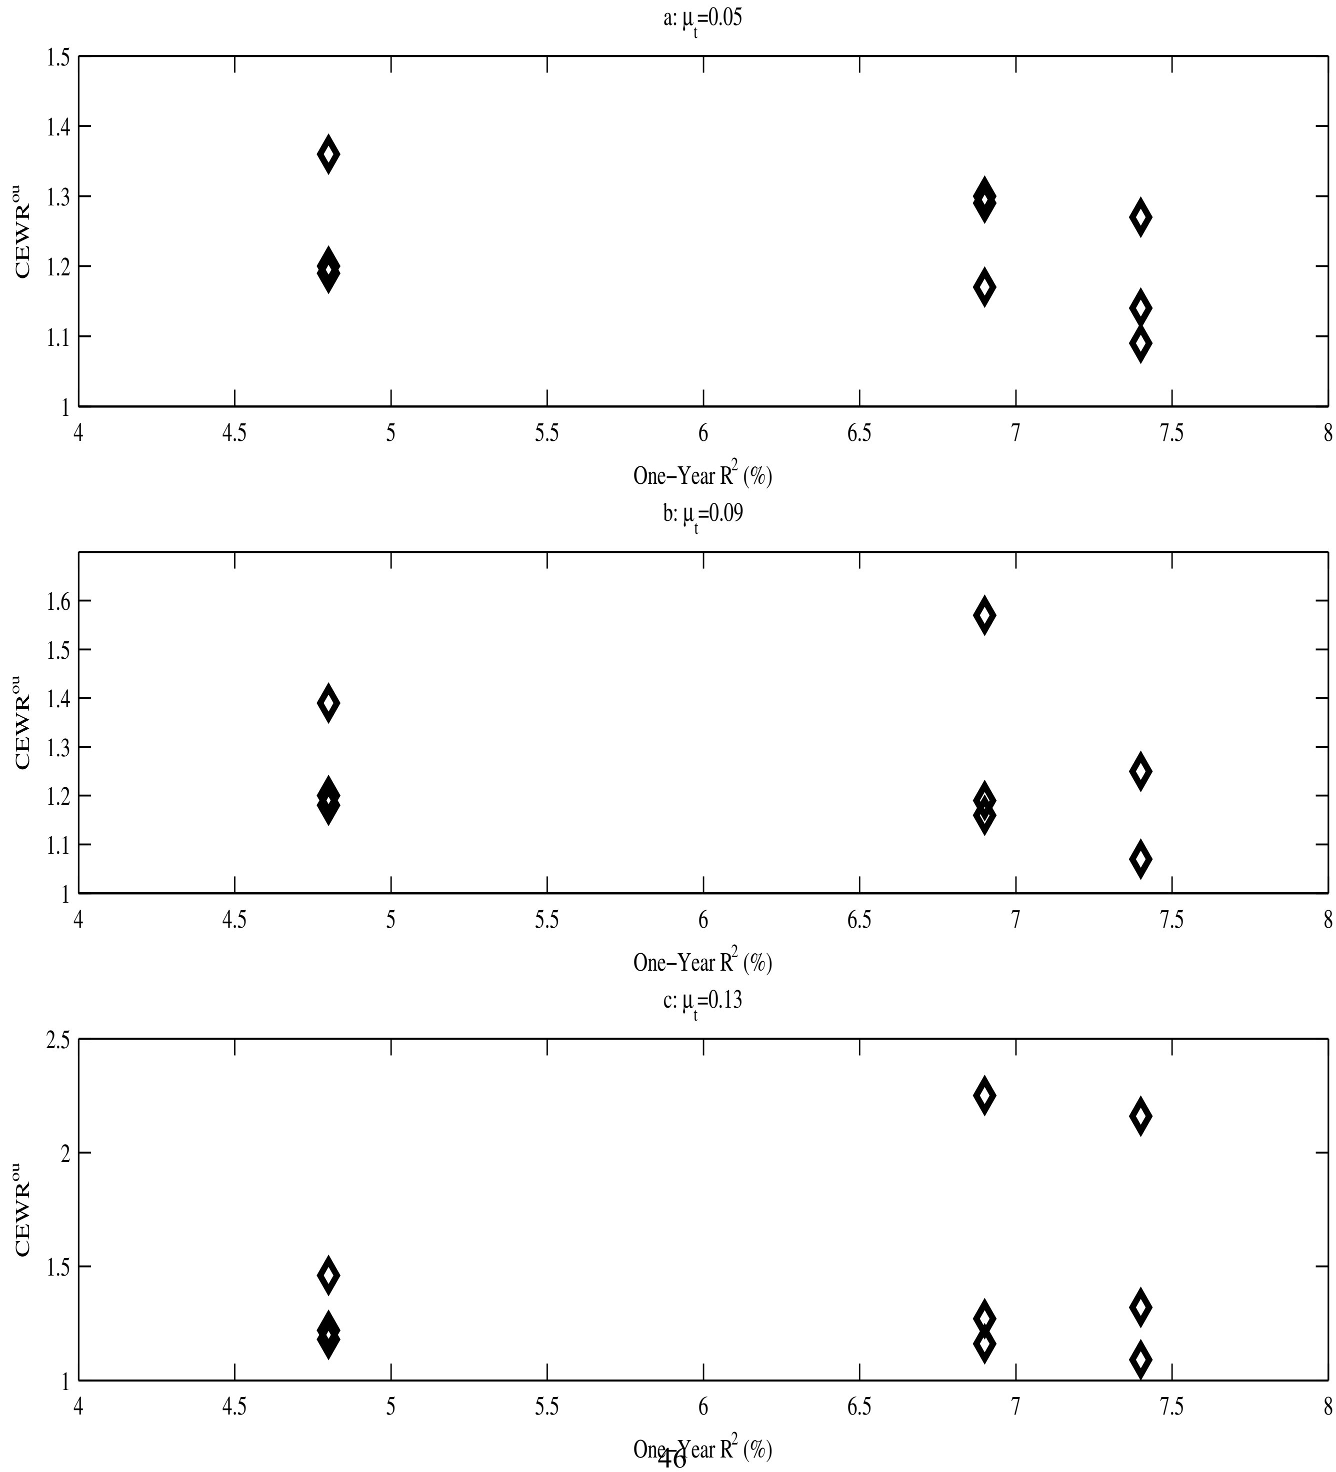

Figure 4
Estimates of the Real μ Series

Panel a plots the long run expected real return k^1 from Arnott and Bernstein (2002) (A&B) together with the estimated time series of $\mu^{1,1}$ and $\mu^{1,2}$, corresponding to the two cases of constant g and mean-reverting g. Panel b plots the long run expected real return k^2 from Ilmanen (2003) (IL) together with the estimated time series of $\mu^{2,1}$ and $\mu^{2,2}$, corresponding to the two cases of constant g and mean-reverting g. The period is from the 1st quarter of 1950 to the 2nd quarter of 2002 in both panels.

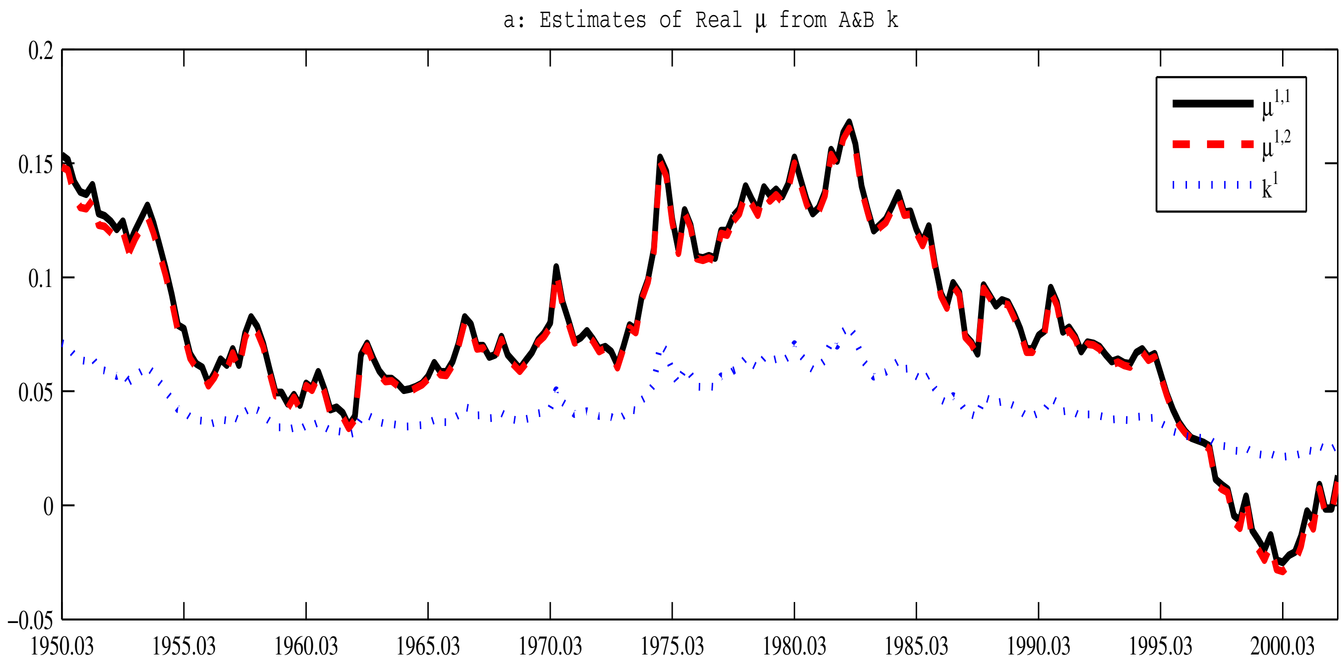

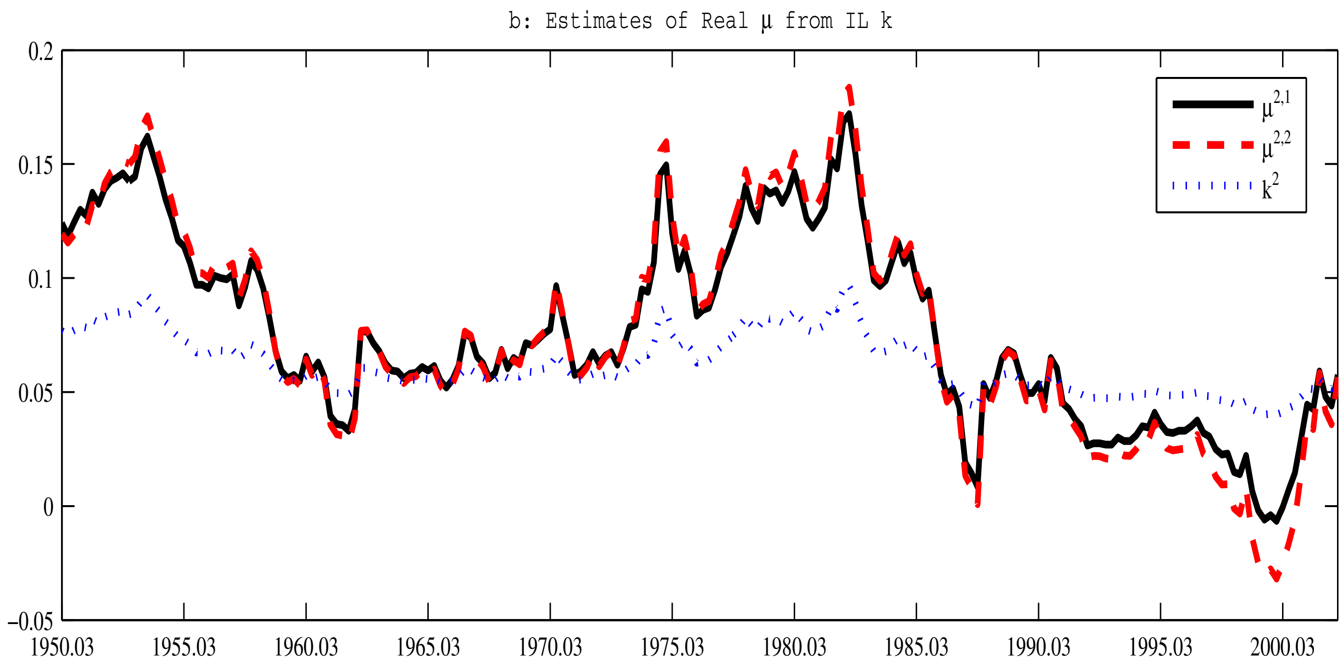

Figure 5

Estimates of the Nominal μ Series

Panel a plots the long run expected nominal return k^3 from Barclays Global Investors (BGI) together with the estimated time series of $\mu^{3,1}$ and $\mu^{3,2}$, corresponding to the two cases of constant g and mean-reverting g. Panel b plots the long run expected nominal return k^4 from Wilshire Associates (WA) together with the estimated time series of $\mu^{4,1}$ and $\mu^{4,2}$, corresponding to the two cases of constant g and mean-reverting g. The period in Panel a is from the 4th quarter of 1972 to the 1st quarter of 2002 while it is from the 1st quarter of 1973 to the 1st quarter of 2004.

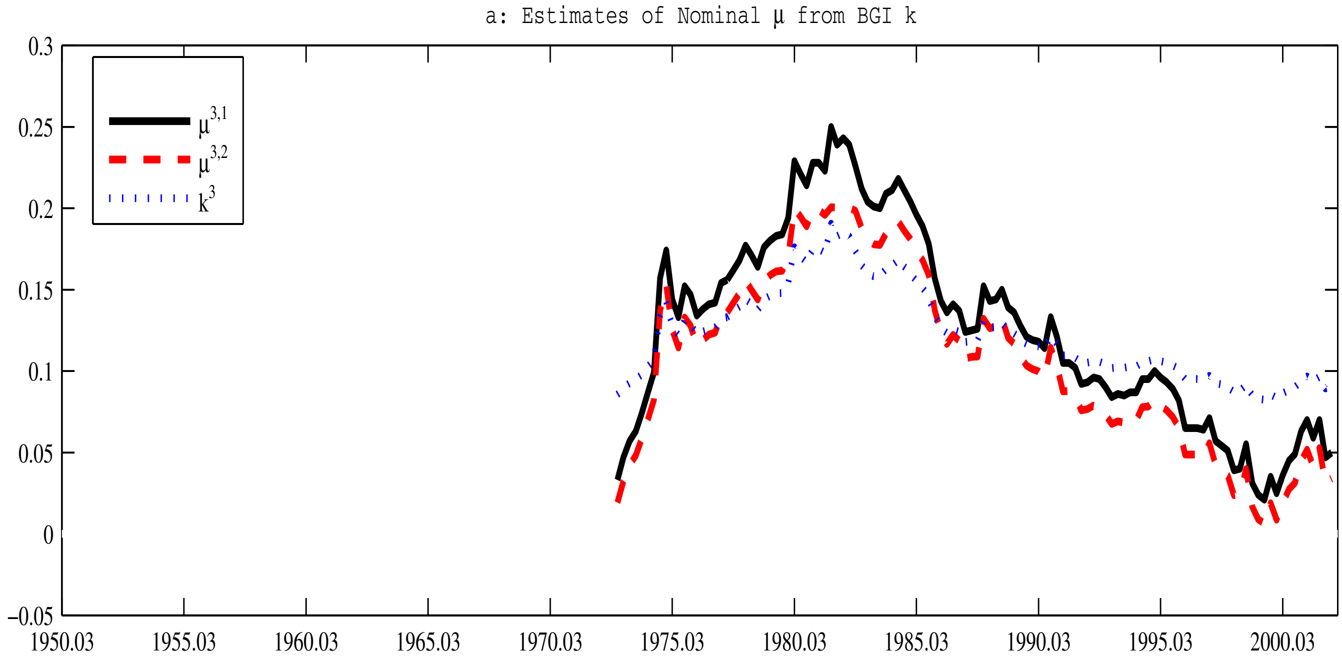

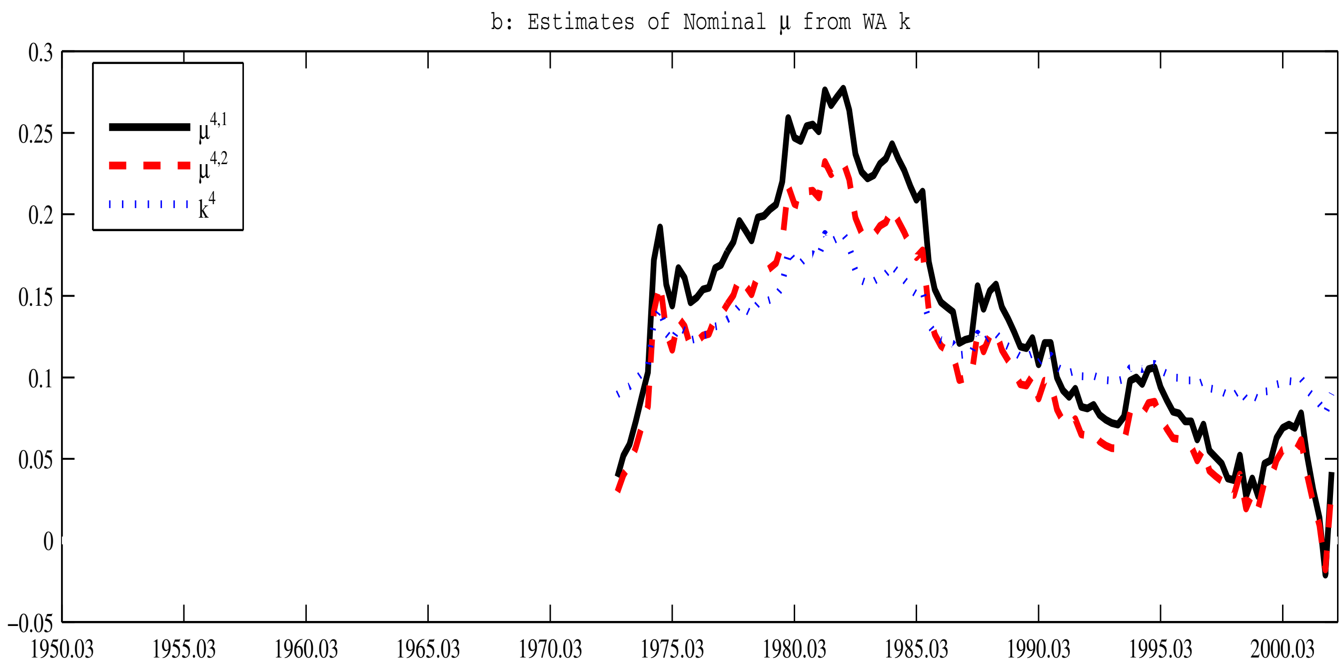

Figure 6

Cumulative Real Wealth under Optimal and Unconditional Strategies for a Long Horizon Investor ($\mu^{1,1}$ as the Predictor)

The figure plots the cumulative real wealth under the optimal and the unconditional strategy for long horizon investors with a risk aversion parameter $\gamma = 5$. The optimal strategy is based on the estimated real A&B $\mu^{1,1}$ series and its associated parameter estimates. Both the optimal and the unconditional strategies are constrained to have allocations between 0 and 1. In Panel a, the investment horizon is 20 years and the investor starts investing in 1950.1 with a terminal date in 1970.1. In Panel b, the investment horizon is also 20 years and the investor starts investing in 1970.1 with a terminal date in 1990.1. In Panel c, the investment horizon is 13 years and the investor starts investing in 1990.1 with a terminal date in 2002.2. In Panel d, the horizon is 53 years and the investor starts investing in 1950.1 with a terminal date in 2002.2.

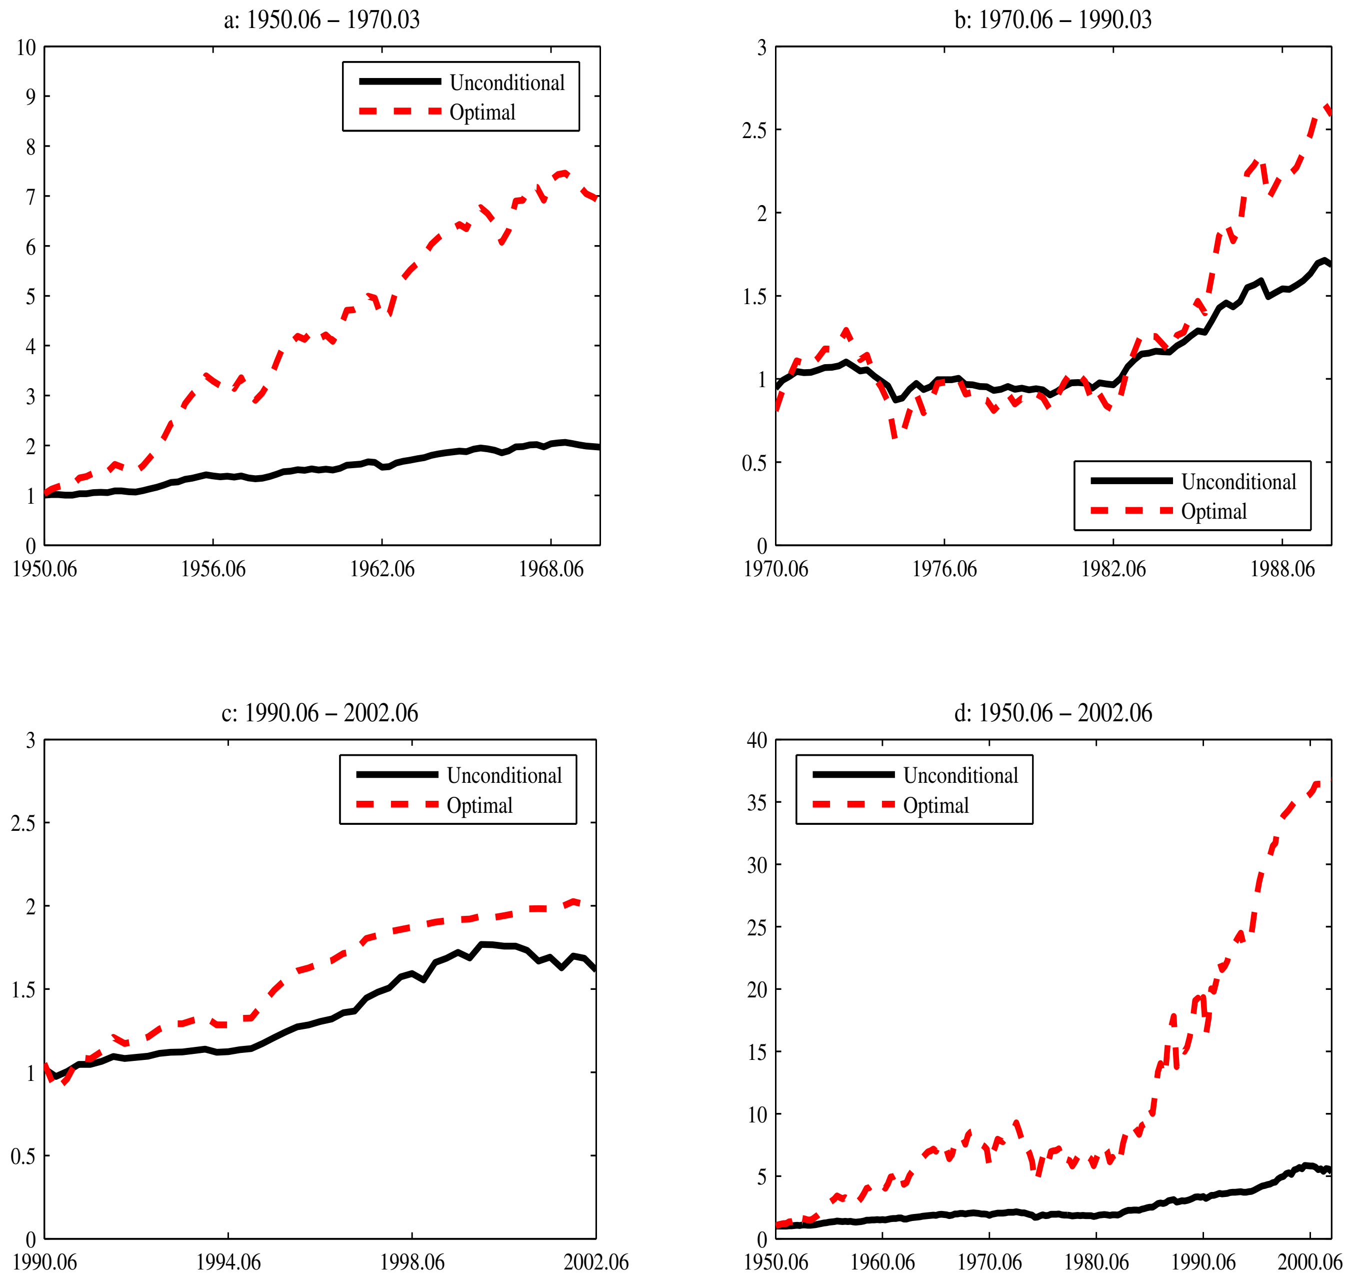

Figure 7
Cumulative Real Wealth under Optimal and Unconditional Strategies for a Long Horizon Investor ($\mu^{2,1}$ as the Predictor)

The figure plots the real wealth under the optimal and the unconditional strategy for long horizon investors with a risk aversion parameter $\gamma = 5$. The optimal strategy is based on the estimated real IL $\mu^{2,1}$ series and its associated parameter estimates. Both the optimal and the unconditional strategies are constrained to have allocations between 0 and 1. In Panel a, the investment horizon is 20 years and the investor starts investing in 1950.1 with a terminal date in 1970.1. In Panel b, the investment horizon is also 20 years and the investor starts investing in 1970.1 with a terminal date in 1990.1. In Panel c, the investment horizon is 13 years and the investor starts investing in 1990.1 with a terminal date in 2002.2. In Panel d, the horizon is 53 years and the investor starts investing in 1950.1 with a terminal date in 2002.2.

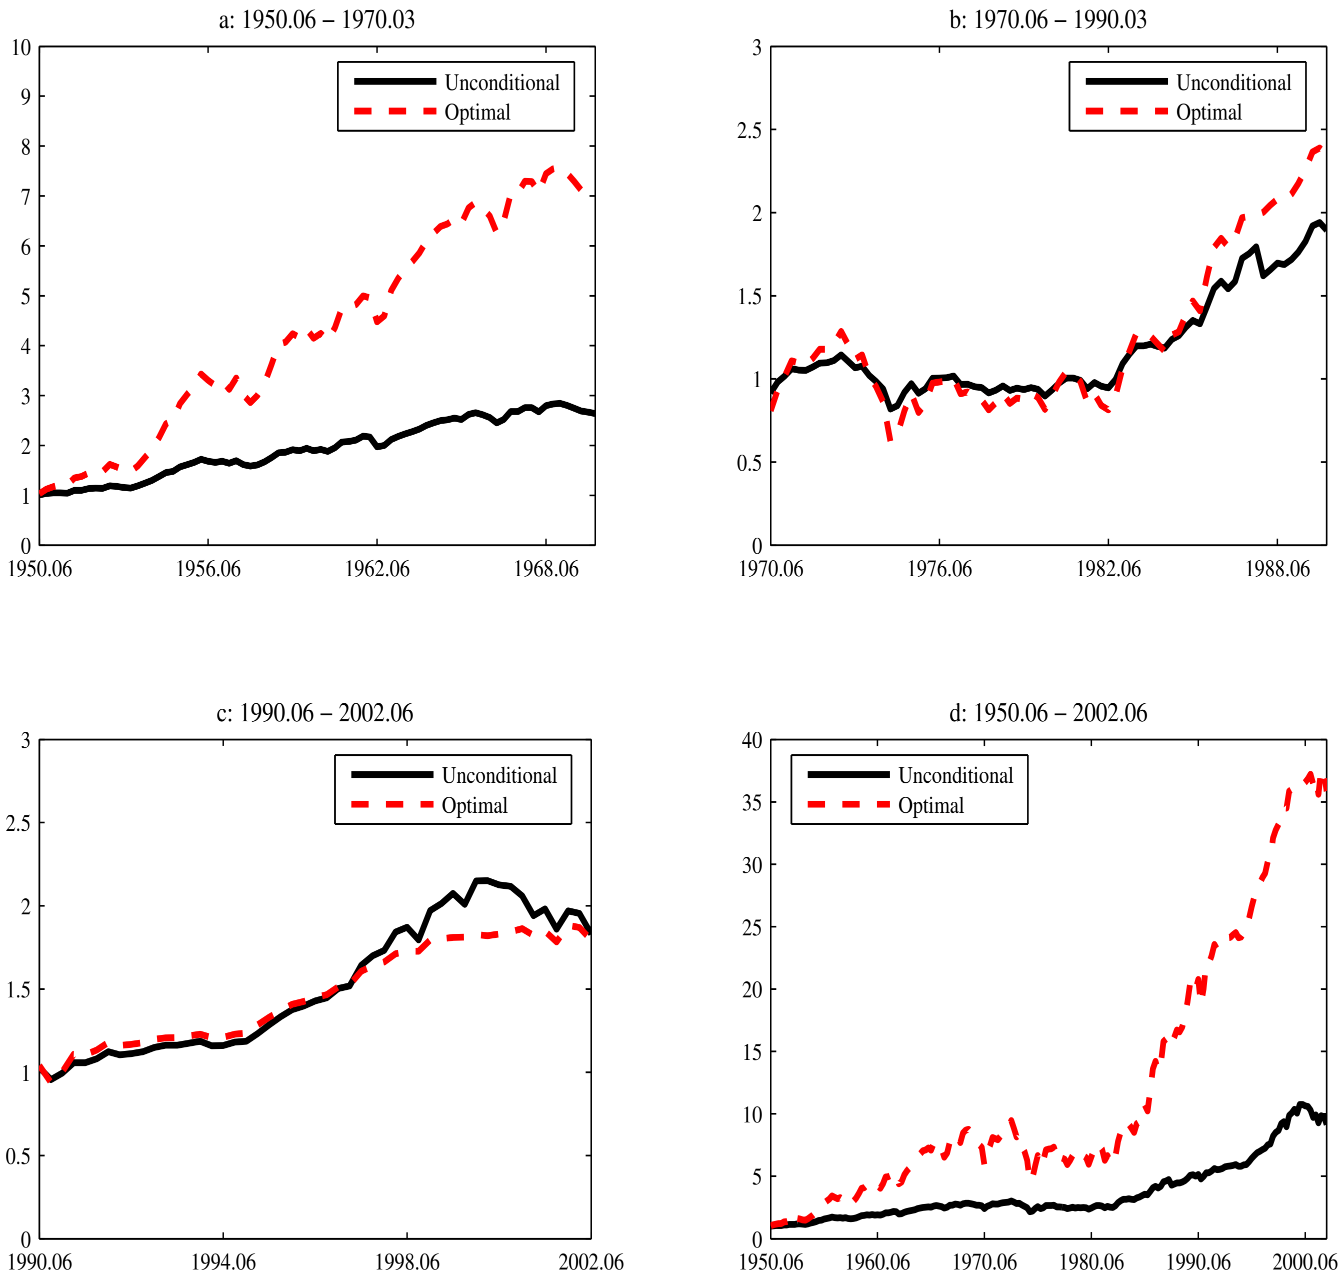

Figure 8
Cumulative Nominal Wealth under Optimal and Unconditional Strategies for a Long Horizon Investor
($\mu^{3,2}$ as the Predictor)

The figure plots the cumulative nominal wealth under the optimal and the unconditional strategy for long horizon investors with a risk aversion parameter $\gamma = 5$. The optimal strategy is based on the estimated nominal BGI $\mu^{3,2}$ series and its associated parameter estimates. Both the optimal and the unconditional strategies are constrained to have allocations between 0 and 1. In Panel a, the investment horizon is 20 years and the investor starts investing in 1972.4 with a terminal date in 1992.4. In Panel b, the investment horizon is around 10 years and the investor starts investing in 1992.4 with a terminal date in 2002.1. In Panel c, the investment horizon is about 30 years and the investor starts investing in 1972.4 with a terminal date in 2002.1.

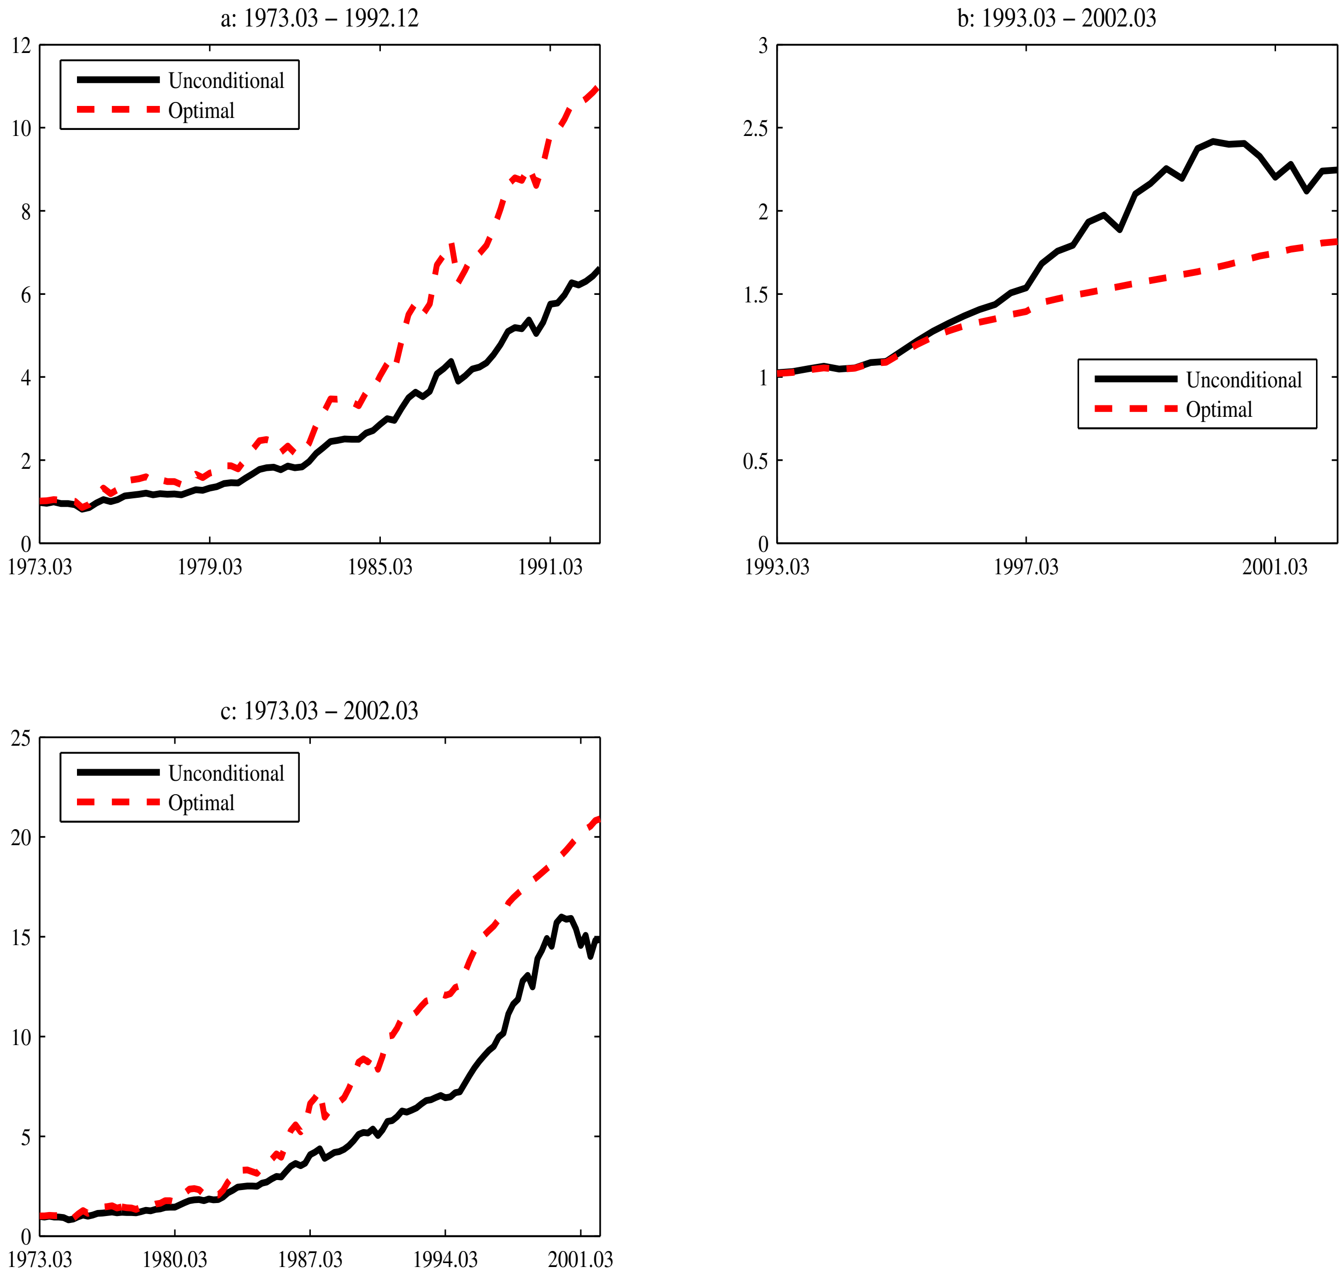

Figure 9

Cumulative Nominal Wealth under Optimal and Unconditional Strategies for a Long Horizon Investor

($\mu^{4,2}$ as the Predictor)

The figure plots the cumulative nominal wealth under the optimal and the unconditional strategy for long horizon investors with a risk aversion parameter $\gamma = 5$. The optimal strategy is based on the estimated nominal WA $\mu^{4,2}$ series and its parameter estimates. Both the optimal and the unconditional strategies are constrained to have allocations between 0 and 1. In Panel a, the investment horizon is 20 years and the investor starts investing in 1973.1 with a terminal date in 1993.1. In Panel b, the investment horizon is around 10 years and the investor starts investing in 1993.1 with a terminal date in 2002.2. In Panel c, the investment horizon is about 30 years and the investor starts investing in 1973.1 with a terminal date in 2002.2.

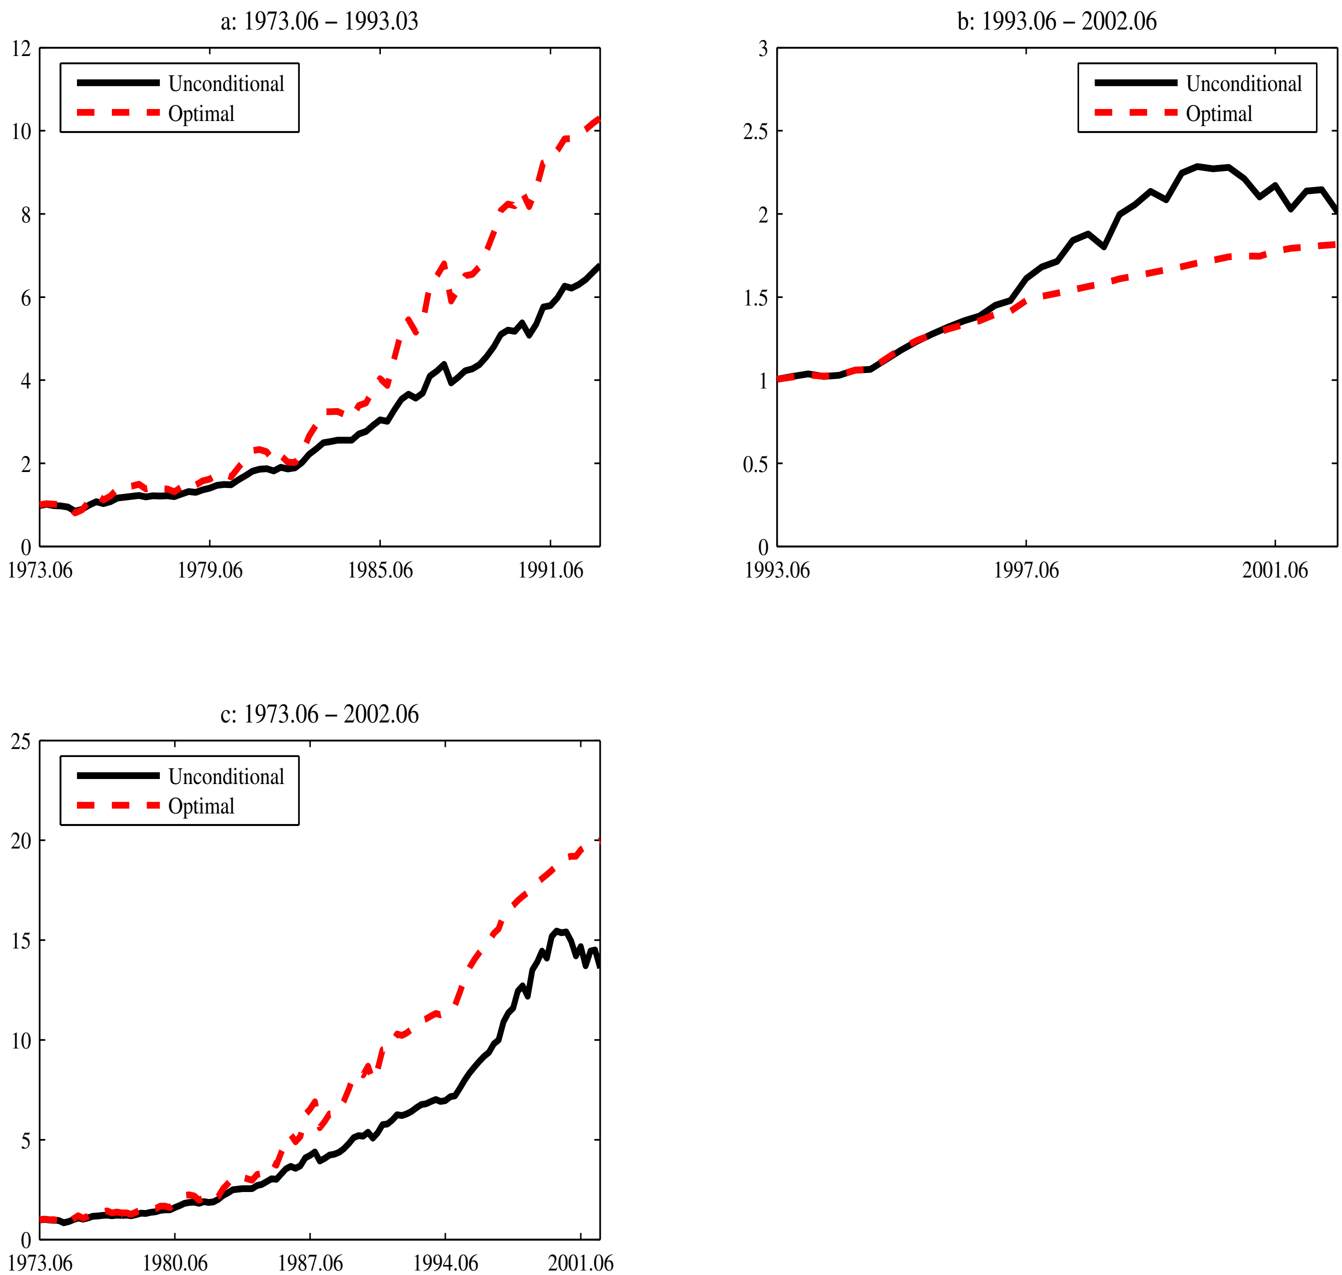

Figure 10

The Unconditional, the Myopic, and the Optimal Stock Allocations for a Long Horizon Investor ($\mu^{4,2}$ as the Predictor)

The figure plots the proportion of wealth allocated to stocks under the unconditional, the myopic, and the optimal strategies for long horizon investors with a risk aversion parameter $\gamma = 5$. The optimal strategy is based on the estimated nominal WA $\mu^{4,2}$ series and its parameter estimates. Both the optimal and the unconditional strategies are constrained to have allocations between 0 and 1. In Panel a, the investment horizon is 20 years and the investor starts investing in 1973.1 with a terminal date in 1993.1. In Panel b, the investment horizon is around 10 years and the investor starts investing in 1993.1 with a terminal date in 2002.2. In Panel c, the investment horizon is about 30 years and the investor starts investing in 1973.1 with a terminal date in 2002.2.