

Economic Implications of Bull and Bear Regimes in UK Stock Returns*

Massimo Guidolin Allan Timmermann
University of Virginia University of California San Diego

September 9, 2002

Abstract

This paper presents evidence of persistent ‘bull’ and ‘bear’ regimes in UK stock returns and considers their economic implications from the perspective of an investor’s portfolio decisions. We find that the perceived state probability has a large effect on the optimal allocation to stocks, particularly at short investment horizons. If ignored, the presence of such regimes gives rise to welfare costs that are substantial, particularly in the bear state where stock holdings should be significantly reduced. When we extend the return forecasting model to allow for predictability from the lagged dividend yield, we find that both dividend yields and regime switching have strong effects on the optimal asset allocation.

Key words: Optimal Asset Allocation, Regime Switching, Bull and Bear Markets, Model Specification.

*We would like to thank Giovanna Nicodano, Elisa Luciano and seminar participants at CERP, Moncalieri and Aarhus business school.

1. Introduction

Financial returns are difficult to predict and, for a long while, the absence of predictability served as one of the corner stones in financial economics. This proposition was largely supported by empirical studies. As recent as in the mid-seventies, the consensus among researchers was that, to a good approximation, returns in stock, bond and foreign exchange markets were unpredictable and prices were well characterized by a random walk. Following a string of papers that documented limited predictability of returns across different predictor variables, sample periods and asset classes, the earlier consensus has largely been replaced by a view that – although predictability may be over-stated because of data-snooping effects and small sample distortions – returns are predictable, particularly at longer horizons.¹

Predictability of returns does not, on its own, reject the notion that financial markets are efficient. In fact, since the predictable component in asset returns tends to be very small and uncertain, it is important to carefully consider how useful predictability really is to risk averse investors. Only recently have the economic implications of return predictability been explored by authors such as Barberis (2000), Brandt (1999), Campbell and Viceira (1999) and Kandel and Stambaugh (1996). These studies find that, faced with time-varying investment opportunities, it is optimal for investors to vary their stock holdings both as a function of a set of predictor variables and as a function of their investment horizon. Even though predictability of returns is generally weak from a statistical perspective, it is generally found to have strong effects on optimal portfolio holdings.

So far the literature has almost invariably explored the asset allocation implications of stock return predictability in the context of simple linear models designed to characterize predictability in the conditional mean of returns. However, for asset allocation purposes it is important to go beyond this and correctly model the full probability distribution of returns. Unless investors are assumed to have very restrictive preferences such as mean-variance utility, the calculation of expected utility and the derivation of optimal portfolio weights will reflect the full probability distribution of returns.

This paper finds strong evidence of nonlinearity in the process driving UK stock returns in the form of regime switching and considers its economic implications.

¹For a clear statement of the evidence and references to the literature, see Cochrane (2001).

We identify two states that can broadly be interpreted as a bull state that offers high mean returns and low volatility and a high volatility bear state. Specification tests that consider the full probability distribution of stock returns strongly reject single-state, linear models against the two-state alternative.

Since the risk-return trade-off on UK stocks varies substantially across the bull and bear states, their presence has the potential to significantly affect investors' optimal asset allocation. We consider the effect of regimes by studying an investor's decision between a broad portfolio of stocks (the FTSE All Share index) and T-bills. The presence of regimes gives rise to a wide variety of investment shapes linking optimal stock holdings to the investment horizon and generates very sensible patterns in optimal stock holdings.

Our paper makes several contributions to the existing literature. While bull and bear markets are part of financial folklore, their asset allocation implications have not previously been studied, certainly not in the rigorous framework that we consider.² Our model has a rich set of implications for optimal stock holdings as a function of the underlying dividend yield and the bull/bear regime probabilities. Our setup also allows us to compute the expected utility cost if regimes are ignored. This is a natural metric in an economic assessment of our findings and allows us to map the potential economic gains from considering regime switching. We find that the gains are large enough to be relevant to long-term investors such as pension funds and even to investors with shorter investment horizons.

We find that optimal stock holdings are strongly affected by investors' beliefs about the underlying state. A buy-and-hold investor who perceives a high probability of being in the moderately persistent bear state will invest very little in stocks in the short run. This investor will hold more in stocks at longer investment horizons as the likelihood of switching to the bull state grows. In contrast, in the highly persistent bull state, investors hold less in stocks the longer their investment horizon. This is because there is only a very small chance of leaving the bull state in the short run, while this probability grows as the investment horizon expands.

Evidence of two states remains strong when the return model is expanded to

²The only other paper that we are aware of that studies the effect of regimes on optimal asset allocation is Ang and Bekaert (2001). However, their paper studies international asset allocation and uses a very different methodology for model selection and asset allocation than that presented here.

include the lagged dividend yield as a predictor. In this extended model we find that both regime switching and predictability from the dividend yield significantly affect optimal stock holdings. Regime switching appears to capture a relatively fast moving mean-reverting component related to changes in volatility and its effect is strongest at relatively short investment horizons. In contrast, the dividend yield captures a slowly mean-reverting component in the return distribution that mostly affects long-run asset allocations.

Rebalancing of optimal portfolio holdings decrease the sensitivity of the optimal stock holding with respect to the investment horizon. This is to be expected since rebalancing means that the current position can be unwound if there is a change in the underlying state. Conversely, frequent rebalancing makes stock holdings more responsive to the current state. If investors find themselves in the bear state, they will reduce stock holdings to a level near zero in the expectation that they can raise them in a subsequent period. A long-term investor deprived of this possibility will not reduce stock holdings as aggressively and will instead base these on the average returns expected over the full investment horizon.

The outline of the paper is as follows. Section 2 presents estimation results for the regime switching model fitted to UK stock returns and provides a range of results from specification tests applied to the probability distribution of stock returns. Section 3 introduces the optimal asset allocation problem and reports empirical asset allocation results under a buy-and-hold investment scheme. Section 4 extends our empirical results to include predictability from the dividend yield. Section 5 considers rebalancing effects while Section 6 reports the outcome of calculations of welfare costs due to ignoring regimes and Section 7 concludes.

2. Regimes in UK Stock Returns

The possibility of predicting asset returns has fascinated generations of researchers in economics and finance and this question has spawned numerous empirical studies. One strand of the literature has adopted linear models and concentrated on documenting which state variables predict the conditional mean of stock returns. Examples of this approach include Campbell and Shiller (1988) and Fama and French (1988). For UK stock returns Clare, Thomas and Wickens (1994) find that the gilt-equity yield ratio has some predictive power, while Black and Fraser (1995) find that default- and term-premium variables have predictive power over returns.

Pesaran and Timmermann (2000) extend this list to find predictability of UK stock returns based on a variety of macroeconomic and financial variables.

Another strand of the literature on return predictability has broadened the scope by investigating the presence of regime dynamics in asset returns. Ang and Bekaert (2002), Driffill and Sola (1994), Gray (1996) and Hamilton (1988)) study regimes in interest rates while Ang and Bekaert (2001), Perez-Quiros and Timmermann (2000) and Turner, Starz and Nelson (1989) consider regimes in stock returns.

2.1. *The Model*

We will consider a framework that incorporates both the possibility of regimes as well as return predictability arising from other variables. As a starting point, suppose that the return or excess return on some risky asset, R_t , follows a Gaussian autoregressive process with mean, variance and autoregressive parameters that can vary across k regimes driven by a latent state variable, S_t :

$$R_t = \mu_{S_t} + \sum_{j=1}^k a_{j,S_t} r_{t-j} + \varepsilon_t. \quad (1)$$

Here S_t takes integer values between 1 and k , μ_{s_t} is the intercept in state S_t , a_{j,s_t} is the autoregressive coefficient at lag j in state S_t , and $\varepsilon_t \sim N(0, \sigma_{S_t}^2)$ is the return innovation which has mean zero and state-specific variance $\sigma_{s_t}^2$ and is assumed to be normally distributed. A linear model is obtained as a special case when $k = 1$.

Completing the model for returns requires specifying the process followed by the state variable, S_t . Following Hamilton (1989), we assume that S_t is driven by a first order, homogeneous Markov process with transition probability matrix \mathbf{P}

$$\mathbf{P}_{[i,j]} = \Pr(S_t = j | S_{t-1} = i) = p_{ij}, \quad i, j = 1, \dots, k. \quad (2)$$

The underlying state, S_t , is allowed to be unobserved. Predictability in this model has two sources. First, as long as at least two mean parameters vary across states, revisions in the perceived state probabilities give rise to time-varying expected returns. This effect is present even in the absence of autoregressive lags and also generates time-variations in the variance (volatility clustering), skew and kurtosis of returns, c.f. Timmermann (2000). Second, when present, the autoregressive terms in (1) directly imply some predictability.

To encompass multivariate models that incorporate a wider set of predictors, \mathbf{z}_t , (1) can be generalized to

$$\mathbf{y}_t = \boldsymbol{\mu}_{s_t} + \sum_{j=1}^p \mathbf{A}_{j,s_t} \mathbf{y}_{t-j} + \boldsymbol{\varepsilon}_t. \quad (3)$$

where $\mathbf{y}_t = (R_t \mathbf{z}_t)'$ and $\boldsymbol{\varepsilon}_t \sim N(\mathbf{0}, \Omega_{s_t})$. This model can capture richer patterns of predictability, combining nonlinear regime switching with linear predictability. The investor's information set at time t , \mathfrak{S}_t , is assumed to comprise the history of returns extended by the predictor variables, i.e. $\mathfrak{S}_t = \{R_j, \mathbf{z}_j\}_{j=1}^t$. Estimation of the parameters of the model, $\boldsymbol{\theta} = (\boldsymbol{\mu}_{s_t}, \mathbf{A}_{j,s_t}, \Omega_{s_t}, \mathbf{P})$ proceeds by maximizing the likelihood function through the EM algorithm, c.f. Hamilton (1990).

2.2. Data

Our data consists of monthly returns on the FTSE All Share stock market index, inclusive of dividends, and returns on 1-month T-bills. We use this data to model the return on stocks in excess of the T-bill rate. Section 4 of our analysis also uses the dividend yield computed as dividends over the preceding 12 months divided by the current stock price. The sample period is 1970:1 - 2000:12, a total of 372 monthly observations. All data was obtained from Datastream.

2.3. Model Selection Based on the Probability Distribution of Returns

Determining the number of states, k , and lags, p , in (1) can pose considerable difficulties, yet is clearly important to understanding the properties of the return distribution. While these design parameters are typically determined by means of statistical tests or information criteria, we will adopt an approach that is closer to the economic objectives of the modeling exercise.

Since we will be using the estimated models for asset allocation purposes, it is important to verify that they adequately capture the return distribution and are not misspecified. For this purpose we use the probability integral transform considered by Rosenblatt (1952) and recently used in economic analysis by Diebold et al (1998).

The probability integral transform or z -score is the probability of observing a value smaller than or equal to the realization r_{t+1} of returns under the null that the

model is correctly specified. Under the k -state mixture of normals, this is given by

$$\begin{aligned}
\Pr(R_{t+1} \leq r_{t+1} | \mathfrak{S}_t) &= \sum_{i=1}^k \Pr(R_{t+1} \leq r_{t+1} | s_{t+1} = i, \mathfrak{S}_t) \Pr(s_{t+1} = i | \mathfrak{S}_t) \\
&= \sum_{i=1}^k \Phi \left(\sigma_i^{-1} \left(r_{t+1} - \mu_i - \sum_{j=1}^p a_{j,i} r_{t+1-j} \right) \right) \Pr(s_{t+1} = i | \mathfrak{S}_t) \\
&\equiv z_{t+1}.
\end{aligned} \tag{4}$$

Here $\Phi(\cdot)$ is the cumulative density function of a standard normal variable. Provided that our model is correctly specified, z_{t+1} should be independently and identically distributed (IID) on the interval $[0, 1]$, with a uniform distribution c.f. Rosenblatt (1952). Based on this idea, Berkowitz (2001) proposes a likelihood-ratio test that inverts Φ to get a transformed z -score

$$z_{t+1}^* = \Phi^{-1}(z_{t+1}).$$

Under the null of a correctly specified model, z^* should be IID and normally distributed ($IIN(0, 1)$). This suggests conducting tests of normality based on moment conditions such as

$$\begin{aligned}
E[z_{t+1}^*] &= 0, \\
Var[z_{t+1}^*] &= 1, \\
Cov[z_{t+1}^*, z_t^*] &= 0, \\
Cov[(z_{t+1}^*)^2, (z_t^*)^2] &= 0, \\
Skewness[z_{t+1}^*] &= 0, \\
Kurtosis[z_{t+1}^*] &= 3.
\end{aligned} \tag{5}$$

We use a likelihood ratio test that focuses on a few salient moments of the return distribution. Suppose the log-likelihood function is evaluated under the null that $z_{t+1}^* \sim IIN(0, 1)$:

$$L_{IID\ N(0,1)} \equiv -\frac{T}{2} \ln(2\pi) - \sum_{t=1}^T \frac{(z_t^*)^2}{2}, \tag{6}$$

where T is the sample size. Under the alternative of a misspecified model, the log-likelihood function incorporates deviations from the null, $z_{t+1}^* \sim IIN(0, 1)$:

$$z_{t+1}^* = \mu + \sum_{j=1}^p \sum_{i=1}^l \rho_{ji} (z_{t+1-i}^*)^j + \sigma e_{t+1}, \tag{7}$$

where $e_{t+1} \sim IIN(0, 1)$. Obviously the null of a correct return model implies $p \times l + 2$ restrictions — $\mu = \rho_{ji} = 0$ ($j = 1, \dots, p$ and $i = 1, \dots, l$) and $\sigma = 1$ — in equation (7). Let $L(\hat{\mu}, \{\hat{\rho}_{ji}\}_{j=1}^p \{i=1}^l, \hat{\sigma})$ be the maximized log-likelihood obtained from (7). To test that the forecasting model (1) is correctly specified, we use the following test statistic

$$LR = -2 \left[L_{IIN(0,1)} - L(\hat{\mu}, \{\hat{\rho}_{ji}\}_{j=1}^p \{i=1}^l, \hat{\sigma}) \right] \sim \chi_{p \times l + 2}^2. \quad (8)$$

In addition to the standard Jarque-Bera (1980) test, we focus on three likelihood ratio tests proposed by Berkowitz (2001):

1. A test with $p = l = 0$ that only restricts the mean and variance of the distribution of the transformed z -scores: $LR_2 = -2 [L_{IIN(0,1)} - L(\hat{\mu}, \hat{\sigma})] \sim \chi_2^2$;
2. A test with $p = l = 1$ that also restricts the transformed z -scores to be serially uncorrelated: $LR_3 = -2 [L_{IIN(0,1)} - L(\hat{\mu}, \hat{\rho}_{11}, \hat{\sigma})] \sim \chi_3^2$;
3. A test with $p = l = 2$ that restricts the transformed z -scores as well as their squares to be serially uncorrelated: $LR_6 = -2 [L_{IIN(0,1)} - L(\hat{\mu}, \{\hat{\rho}_{ji}\}_{j=1}^2 \{i=1}^2, \hat{\sigma})] \sim \chi_6^2$.

Table 1 shows the outcome of the tests for a range of regime-switching models, including the special case of a linear, single-state model. Models without state-dependence in the volatility are strongly rejected by the first test suggesting that volatility definitely varies across regimes. The two-state model with state-dependent mean and variance passes all four tests. Since this is the most parsimonious model that appears to appropriately capture the distribution of stock returns, we use this model in our further analysis.³

2.4. Interpretation of the States

Using the two-state model fitted to monthly excess returns on the FTSE All Share index, Figure 1 plots the smoothed probability for state 1. This state picks up the

³Including a larger number of states in the model does not always lead to better performance of the misspecification tests. The reason is that the parameters of the larger models are not chosen to minimize the value of the probability integral test statistics.

very volatile markets in 1974-75 along with low return and high volatility episodes in 1976, 1981 and 1987.⁴ Table 2 confirms this interpretation by providing full-sample estimates for the two-state model. The first regime is a bull state where excess returns have a high mean (0.7% per month) and low volatility (4.5% per month). The second regime is a bear state with negative mean excess returns (-0.65% per month) and much higher volatility (15.2% per month).⁵ The parameters of the two-state model span those obtained for the single-state model: The estimated mean excess return in the linear model is 0.6% per month and the volatility is 6.2% per month.

While the mean excess return is statistically significantly different from zero in the bull state, it is not significantly different from zero in the bear state. This is likely to reflect the relatively small part of the sample spent in the bear state which leads to a larger standard error for the mean parameter in this state. In fact, at 0.986 and 0.844, the transition probability estimates show that the bull state is highly persistent with an average duration of 74 months while the expected duration of the bear state is only six months.

Though not visited frequently, the bear state is nevertheless important in determining both the expected value and the risk of stock returns and it clearly helps to capture outliers in the excess return distribution that cannot be accommodated by a Gaussian model. To demonstrate this point, Figure 2 plots the first four conditional moments as a function of the bull state probability using the parameter estimates reported in Table 2.

As the state probabilities change, the mean, volatility, skew and kurtosis of the distribution also change. The conditional mean increases linearly in the bull state probability, while the volatility decreases monotonically as a function of this probability. Increasing the probability of the bull state introduces a negative skew in the return distribution which peaks at -0.3 for a bull state probability of 0.9 only to increase to zero when the bull state probability is one. Kurtosis shows the opposite pattern. It increases from 3 (the kurtosis of the normal distribution), peaks just below 10 for a bull state probability of 0.9 and then drops to 3 again

⁴We also estimated a three-state model for monthly excess returns. The third state isolated the unusually large returns in January and February 1975 but did not include other parts of the sample.

⁵Although we label the two regimes ‘bull’ and ‘bear’ states, they could equally well be characterized as low and high volatility states.

when the bull state is known with certainty. Variations in the perceived state probabilities thus lead to large changes in the properties of the return distribution. As we shall see in the next section, this is key to understanding the optimal asset allocation results.

3. Optimal Stock Holdings

Using the two-state model for UK stock returns, this section studies the optimal stock holdings for a buy-and-hold investor with constant relative risk aversion preferences. In later sections we introduce predictability from the dividend yield and periodic rebalancing. Abstracting from these effects in the initial analysis simplifies the problem considerably and makes our results easier to understand. We first characterize the investor's optimization problem and then present empirical results.

3.1. The Investor's Optimal Asset Allocation Problem

Consider a buy-and-hold investor with unit wealth at time t and power utility defined over the level of wealth T periods from now, W_{t+T} :⁶

$$u(W_{t+T}) = \frac{W_{t+T}^{1-\gamma}}{1-\gamma}. \quad (9)$$

Here γ is the coefficient of relative risk aversion while T is the investment horizon. The investor is assumed to maximize expected utility at time t by allocating ω_t to stocks while $1 - \omega_t$ is invested in riskless, one-month T-bills.⁷ The problem solved by the investor is

$$\begin{aligned} \max_{\omega_t} \quad & E_t \left[\frac{W_{t+T}^{1-\gamma}}{1-\gamma} \right] \\ \text{s.t. } \quad & W_{t+T} = \{(1 - \omega_t) \exp(T r^f) + \omega_t \exp(R_{t+T}^s)\} \\ & \omega_t \in [0, 1] \end{aligned} \quad (10)$$

⁶Our analysis follows much of the existing literature by assuming power utility over final wealth, but the qualitative results appear to be robust to alternative functional forms.

⁷Following standard practice we assume that the risk-free rate is known and constant. This assumption means that we do not have to simultaneously model future T-bill rates and stock market returns. Because we model *excess* return on stocks, variations in T-bill rates are, however, likely to have a much smaller effect on optimal stock versus T-bill holdings since higher T-bill rates also increase stock returns computed as the excess return plus the T-bill rate.

where $R_{t+T}^s \equiv R_{t+1}^s + R_{t+2}^s + \dots + R_{t+T}^s$ is the continuously compounded stock return over the T -period investment horizon. The constraint $\omega_t \in [0, 1]$ rules out short-selling.

There is no closed-form solution to the optimal stock holdings under power utility, so we use Monte-Carlo methods to draw N time paths of T monthly excess returns from the regime switching model using the parameter estimates $\hat{\theta}_t = \{\hat{\mu}_t, \sigma_t^2, \hat{\mathbf{P}}_t\}$. These simulations account for the possibility of stochastic regime switching as governed by the transition matrix $\hat{\mathbf{P}}_t$. For a given value of ω_t , we can approximate the integral in the expected utility functional:

$$\hat{E}_t[U(W_{t+T})|\omega_t] = N^{-1} \sum_{n=1}^N \left\{ \frac{\left[(1 - \omega_t) \exp(Tr^f) + \exp\left(\sum_{i=1}^T R_{t+i,n}^s\right) \right]^{1-\gamma}}{1 - \gamma} \right\}. \quad (11)$$

Here $1 \leq n \leq N$ tracks the simulation number.⁸ A grid search across different values of ω_t determines the optimal stock holdings.

3.2. Empirical Results

Investors differ along several dimensions. They are likely to have different levels of risk aversion and different investment horizons. To assess the economic significance of regimes in stock returns we study optimal portfolio holdings by varying both the level of risk aversion (γ) and the time horizon (T). We start with the investment horizon. Speculators and some mutual funds may have relatively short investment horizons of up to one year, while pension funds have a longer investment horizon of, e.g., 10-20 years, depending on their liability structure.

Using the last half of our sample (1986-2000), Figure 3 shows the evolution in the real-time optimal stock holdings at four different investment horizons, namely a short horizon (1 month) two medium horizons (6 and 24 months) and a very long horizon (120 months). For comparison we also show the optimal allocation to stocks under no predictability, i.e. assuming that returns are identically and

⁸A large number of simulations is needed to adequately account for the occurrence of bear regimes with moderate steady-state probability. We experimented with different values of N between 5,000 and 50,000 and found that $N = 40,000$ is adequate to obtain sufficiently precise and stable solutions.

independently distributed. Following studies such as Barberis (2000) and Brandt (1999) we initially fix the coefficient of relative risk aversion at $\gamma = 5$.

To be more realistic and avoid a ‘benefit of hindsight’ bias, these results do not condition on end-of-sample parameter estimates or state probabilities and instead use the real-time recursive parameter estimates $\{\hat{\mu}_t, \hat{\sigma}_t^2, \hat{\mathbf{P}}_t\}_{t=1986:01}^{2000:12}$ and state probabilities $\{\hat{\pi}_t\}_{t=1986:01}^{2000:12}$. This explains why the optimal stock holdings change over time even under the no-predictability model.

The first conclusion emerging from Figure 3 is that optimal stock holdings are far more volatile at the short horizons than at the longer horizons. Throughout most of the sample a short-horizon portfolio objective implies a very high commitment to stocks. However, an investor with a 1-month horizon would have invested very little in stocks in the aftermath of October 1987, only to have raised stock holdings to 80% in early 1988. Despite this volatility in the optimal stock holdings, the short-sales constraint is never binding.

As the investment horizon grows the investor pays less attention to the current state and the optimal stock holdings become smoother over time. With the exception of a short period from the end of 1987 to the beginning of 1988, investors with a horizon of at least two years would have steadily invested between 20% and 40% of their wealth in the UK stock market. On the back of the high stock returns during the 1990s this percentage increased from about 30% to around 35% at the end of 2000.

It is interesting to compare the optimal stock holdings under this two-state model to those obtained under an assumption of no predictability ($k = 1$). The optimal stock holdings under no predictability are closest to those derived under a 1-month horizon when the probability of being in the bull state is high. It is easy to see why. The parameter estimates for the bull state are much closer to the full-sample values than the bear state parameter estimates, c.f. Table 2. Since the bull state is highly persistent, at short horizons where the current state matters the most, the optimal asset allocation will be similar in the bull state and under the no-predictability model. Of course, this conclusion does not carry over to the bear state since this state has a probability distribution for returns that is very different from the unconditional return distribution. It also does not hold at longer investment horizons since these put more weight on the return distribution in the bear state. Since the single-state model underestimates the probability of

months with large negative returns, the average stock holdings are roughly twice as large under no predictability as they are under regime switching and an investment horizon of two years or longer.

3.3. *State Beliefs and Horizon Effects*

Assuming a time-invariant investment opportunity set and power utility, in a classic paper Samuelson (1969) showed that optimal asset holdings are identical across different investment horizons. In our setup, the relationship between the investment horizon and the optimal stock holdings is considerably more complicated since the perceived investment opportunities change both with the underlying state probabilities and with the investment horizon.

Figure 4 shows the interaction between the perceived bull state probability and the investment horizon in determining optimal stock holdings. This figure tracks the effect on the optimal stock holdings of changing either the probability of being in the bull state or the investment horizon. The remaining parameters are set at their end-of-sample estimates.⁹

Holding the state probability fixed, the optimal stock holdings can be either constant, rising or declining as a function of the investment horizon. The intuition for these findings is as follows. Since bull markets are more persistent and likely in steady-state than bear markets, at long horizons stocks offer a high risk premium and relatively low volatility. At shorter investment horizons the risk-return trade-off offered by stocks changes significantly with the underlying state.

When the probability of being in a bear state is perceived to be high (as it was in the late eighties), stocks are unattractive to investors with a short horizon. These investors see a high chance of experiencing low and volatile stock returns in the event of several months spent in the bear market. Since the bear state is only moderately persistent, at longer horizons stocks again become attractive and the optimal allocation to stocks is an increasing function of T .

In contrast, when markets are dominated by uncertainty about the nature of the current regime, flat or even non-monotonic shapes are possible. For example, when the state beliefs are set at their steady-state values, the stock investment

⁹In this and all other exercises the bull state probability is the probability of *starting* from the bull state. It is *not* the probability of remaining in the bull state during the entire investment period.

schedule is flat and independent of T .

Finally, when the probability of being in a bull market is perceived to be high (as happened most of the time between 1989 and 2000), investors with a short horizon have strong incentives to aggressively buy stocks to exploit the very high persistence of bull markets. Although current beliefs may deem bear markets highly unlikely, as T grows the occurrence of bear market spells must be factored in. This explains the negative relation between horizon, T , and stock holdings.

If instead we fix the investment horizon and vary the state probability, the optimal allocation to stocks increases as the perceived bull state probability goes up. At short horizons the rise in stock holdings is very steep when the bull state probability increases, but it gets much flatter at the longest horizons considered here. This is due to the mean reversion in the state process, S_t , which implies that the current state matters less at long investment horizons.

3.4. *Effects of Risk Aversion*

So far we have fixed the coefficient of risk aversion at $\gamma = 5$. To study the effect on optimal stock holdings of different coefficients of relative risk aversion, we vary γ between 1 and 20. We show results for three different configurations of the initial state probabilities, $(\hat{\pi}_t, 1 - \hat{\pi}_t)$ corresponding to a bull state (1, 0), a bear state (0, 1) and a state with high uncertainty about the state (0.8, 0.2). We chose the latter probabilities to be close to the steady state probabilities and not too dissimilar to the state probabilities that were actually observed throughout the sample. For comparison we also report optimal stock holdings under no predictability. Figure 5 shows the outcome of this exercise using a short ($T = 1$) and a long ($T = 120$) horizon.

First consider the results in the left column of Figure 5 which assume a short investment horizon. As γ rises, the optimal stock holdings decline monotonically. In the bull state, at low levels of risk aversion ($\gamma \leq 4$) more than 90% of the portfolio is allocated to stocks, but this number declines gradually to around 30% for γ above 10. In the bear state stock holdings are at or below 40% for very small values of γ and quickly drop to less than 5% when $\gamma \geq 4$. Stock holdings are also very small for $\gamma \geq 4$ when there is high uncertainty about the current state.

At the longer horizon, the effect on stock holdings of raising γ appears to be even larger for values of γ between one and four. The optimal stock holdings go from

close to 100 percent to around 40% in the bull state and decline from 90% to 30% in the bear state or to 40% under high uncertainty. At the long horizon the two-state asset allocations are always very different from those under no predictability, whereas the two are very similar in the bull state when $T = 1$.

These findings show that the coefficient of risk aversion has a relatively large effect on the optimal level of stock holdings, but that most of the effect occurs at levels of risk aversion lower than our assumed value of $\gamma = 5$.

4. Stock Holdings Under Predictability from the Dividend Yield

Studies such as Barberis (2000), Brandt (1999), Campbell and Viceira (1999), Kandel and Stambaugh (1996), and Lynch (2001) have considered optimal asset allocation under predictability from the dividend yield so it is natural for us to extend our results to allow for predictability from this regressor. We do this in two steps. First, we consider the optimal asset allocation when the dividend yield is the only predictor, basing the results on a VAR model similar to that used in earlier studies. Having ensured comparability with existing results, we next introduce regimes and investigate the results in the context of the two-state bivariate regime switching model (3).

4.1. Predictability from the Dividend Yield

We use our more general model (3) to both consider a linear VAR and a regime-switching VAR with the dividend yield added as a predictor variable. Separate models fitted to the dividend yield and to excess returns suggested a two-state specification with almost identical state variables. We therefore continue to work with a two-state model extended to incorporate a lag to accommodate the strong evidence of a first-order autoregressive component in the yield.

Table 3 presents the full-sample parameter estimates for the bivariate model assuming either a linear or a two-state specification with an autoregressive component. In the linear model, the autoregressive terms are highly significant both in the equation for excess returns and in the equation for the dividend yield.

Turning to the two-state model, stock market volatility continues to be low in the first state (4.4% per month), while state 2 is a high volatility state (11.3% per month), albeit with less extreme volatility than we found in the univariate

model for excess returns. The dividend yield is also far more volatile in state 2 and frequently undergoes rapid moves in this state. The persistence of the dividend yield remains very high in the first state (0.98) but is somewhat lower in the second state (0.79).¹⁰

Interestingly, once regime-switching is accounted for, the dividend yield ceases to have predictive power over stock returns as shown by the insignificance of the coefficient on this regressor in the two-state model. There is no evidence of serial correlation in returns in the bull state, but the serial correlation is relatively strong and statistically significant in the bear state.

The lower graph in Figure 6 plots the smoothed state probabilities for the bivariate model fitted to stock returns and the dividend yield. Compared with Figure 1 there are now many more periods where the process is deemed to be in the high volatility bear state, such as in 1971, 1997 and, with less certainty, during some further months at the end of the sample.

4.2. Stock Holdings when the Dividend Yield is the only Predictor

Earlier papers on optimal stock holdings under predictability from the dividend yield all assume a single state. To compare our results to this literature and to disentangle the effect on stock holdings of predictability from the dividend yield and the presence of two regimes, we first constrain the general model (3) to $k = 1$, so that lagged values of the dividend yield is the only source of predictability and the model simplifies to a bivariate VAR:

$$\begin{pmatrix} r_t \\ dy_t \end{pmatrix} = \begin{pmatrix} \mu_r \\ \mu_{dy} \end{pmatrix} + \sum_{j=1}^p \begin{pmatrix} a_j^{r,r} & a_j^{r,dy} \\ a_j^{dy,r} & a_j^{dy,dy} \end{pmatrix} \begin{pmatrix} r_{t-j} \\ dy_{t-j} \end{pmatrix} + \begin{pmatrix} \varepsilon_{r,t} \\ \varepsilon_{dy,t} \end{pmatrix}. \quad (12)$$

We further constrain the model to match the assumptions in Barberis (2000) and Lynch (2001) by setting $p = 1$, and $a_j^{r,r} = a_j^{dy,r} = 0$ and estimating (12) by applying OLS equation by equation.

Figure 7 (upper window) shows the real-time optimal stock holdings at four different horizons. These holdings are very different from those obtained under the

¹⁰The strong negative simultaneous correlation between innovations to the dividend yield and stock returns reflects the fact that positive news that lead to higher stock prices (and hence positive returns) tend to simultaneously lower the dividend yield which has the stock price in the denominator.

two-state model in Figure 3. The short sales constraint is now binding in many periods. At the longest 10-year horizon, it is optimal to invest all money in stocks until 1997. The slow and persistent moves in the dividend yield are easy to detect in the stock holdings derived for the short and medium horizons: stock holdings rise after the October 1987 stock market crash and start coming down again after 1991 as the yield declines. At the shortest investment horizons, the model suggests going entirely out of stocks after 1993.

Figure 8 (upper window) shows the effect of the investment horizon on the real-time stock holdings during selected periods with different values of the yield. Optimal stock holdings vary significantly over the sample. Periods with low dividend yields (to the right of the figure) give rise to relatively flat horizon curves, while periods with high dividend yields (to the left of the figure) generate steeper investment curves. All investment curves are monotonically increasing.

4.3. *Regimes and Dividend Yield Effects*

We next consider the real-time optimal stock holdings in the general model that incorporates predictability from the dividend yield *and* allows for two states. The lower window in Figure 7 shows that introducing states in the extended model leads to significant changes in the optimal stock holdings. In stark contrast to the case with a single state, the long-horizon ($T = 120$) stock holdings are now very low between 1987 and 1990, high between 1990 and 1994 and high again around 1996 before they decline. Despite the extra uncertainty associated with the introduction of regimes, their effect is not to simply reduce stock holdings at all points in time. Rather, regimes reduce stock holdings at some times (near bear states) while increasing them at others.

The extended model with two regimes also leads to very different investment advice towards the end of the sample. Even at the short investment horizons, the model only suggests to exit stocks as late as in 1997. This is in contrast with the model that ignored the presence of regimes and suggested an exit from the stock market as early as 1993, at least for investors with short horizons.

Turning next to the optimal stock holdings as a function of the investment horizon, Figure 8 (lower window) shows that introducing regimes in the model with the dividend yield as a predictor variable leads to non-monotonic investment curves. For example, near bear states there were many occasions where the optimal

stock holding was high at the 1-month horizon, declined for horizons between 1 and 10 months and increased uniformly at longer horizons.

Figure 9 sheds light on the complex interaction between the perceived bull state probability and the dividend yield, using a short ($T = 1$) and a long investment horizon ($T = 120$). The perceived state probability is varied between zero and one while the dividend yield varies in the interval $[\overline{dy} - 2\hat{\sigma}_{dy}, \overline{dy} + 2\hat{\sigma}_{dy}]$, where \overline{dy} is the unconditional sample mean of the dividend yield (4.6%) and $\hat{\sigma}_{dy}$ is its standard deviation (1.34%).

At the short horizon the optimal allocation to stocks is zero for values of the dividend yield below 2.8%. It grows to 100% for values of the yield above 6.2%, so clearly the optimal allocation to stocks increases steeply in this range of yield values. Keeping the dividend yield fixed there is also a strong effect from changes in the state probability. For instance, when the dividend yield equals 4.5%, it is not optimal to hold stocks in the bear state, whereas the optimal allocation to stocks is 86% in the bull state.

At the long horizon the dividend yield has a strong effect on the optimal stock holdings which go from 34 to 100 percent as the yield moves from 2 to 4.2 percent. There is still an effect on stock holdings from the state probability, but this is clearly smaller than at the short horizon.¹¹

The intuition for these results is that both the state variable and the dividend yield track mean reverting components in stock returns. However, whereas the dividend yield variable captures a slow-moving mean reverting factor - as witnessed by the very high persistence of this variable - the regime variable captures mean reversion at a higher frequency. This means that the state probability has a stronger effect, the shorter the investment horizon while the dividend yield variable has its strongest effect at the longest horizons.

5. Portfolio Rebalancing

So far we have ignored the possibility of rebalancing. This may be realistic for an investor who faces high transaction costs and only gets to adjust the portfolio

¹¹It should also not be forgotten when interpreting these figures that whereas the dividend yield moves relatively smoothly and typically does not jump from one month to the next, in contrast the perceived state probability often undergoes sudden shifts, changing from zero to one over a span of one or two months.

weights infrequently. It also may be plausible for investors such as pension funds whose strategic asset allocation is constrained by trustees to lie within relatively narrow bands which precludes the fund manager from adjusting portfolio weights in the short run. For other investors this assumption is likely to be unrealistic so this section studies the effects of rebalancing.

Suppose that investors can adjust portfolio weights every $\varphi = \frac{T}{B}$ months at B equally spaced points $t, t + \frac{T}{B}, t + 2\frac{T}{B}, \dots, t + (B-1)\frac{T}{B}$ and let ω_b ($b = 0, 1, \dots, B-1$) be the portfolio weights on the risky assets at these rebalancing times. When $B = 1$, $\varphi = T$ and the investor simply implements a buy-and-hold strategy. Under power utility the derived utility of wealth conveniently simplifies to (c.f. Ingersoll (1987, page 242))

$$J(W_b, R_b^s, \boldsymbol{\theta}_b, \boldsymbol{\pi}_b, t_b) = \frac{W_b^{1-\gamma}}{1-\gamma} Q(R_b^s, \boldsymbol{\theta}_b, \boldsymbol{\pi}_b, t_b). \quad (13)$$

Investors are assumed to update their beliefs about the underlying state at each point in time using the formula (c.f. Hamilton (1989)):

$$\boldsymbol{\pi}_{t_b+1}(\hat{\boldsymbol{\theta}}_t) = \frac{\boldsymbol{\pi}_{t_b}(\hat{\boldsymbol{\theta}}_t) \hat{\mathbf{P}}_t^{\varphi(b+1)} \odot \boldsymbol{\eta}(R_{t_b+1}^s; \hat{\boldsymbol{\theta}}_t)}{(\boldsymbol{\pi}_{t_b}(\hat{\boldsymbol{\theta}}_t) \hat{\mathbf{P}}_t^{\varphi(b+1)} \odot \boldsymbol{\eta}(R_{t_b+1}^s; \hat{\boldsymbol{\theta}}_t))' \boldsymbol{\iota}_k}, \quad (14)$$

where $\boldsymbol{\eta}(R_{t_b+1}^s; \hat{\boldsymbol{\theta}}_t)$ is the predictive density of the stock return at time $t_b + 1$ and $\hat{\mathbf{P}}_t^{\varphi(b+1)}$ is the updated $\varphi(b+1)$ -step ahead transition probability matrix, $\hat{\mathbf{P}}_t^{\varphi(b+1)} \equiv \prod_{i=1}^{\varphi(b+1)} \hat{\mathbf{P}}_t$.¹² We choose ω_b to maximize

$$E_{t_b} \left[\left\{ (1 - \omega_b) \exp(\varphi r^f) + \omega_b \exp(R_{b+1}^s(s_b)) \right\}^{1-\gamma} Q(\boldsymbol{\pi}_{b+1}^j, t_{b+1}) \right]. \quad (15)$$

We solve this equation numerically by discretizing the domain of each of the state variables on G equi-distant points and using backward induction methods. The multiple integral defining the conditional expectation is again calculated by Monte Carlo methods. For each $\boldsymbol{\pi}_b = \boldsymbol{\pi}_b^j$, $j = 1, 2, \dots, G^{k-1}$ on the grid we draw in calendar time N samples of φ -period excess returns $\{R_{b+1,n}(S_b)\}_{n=1}^N$ from the regime switching model, where $R_{b+1,n}(S_b) \equiv \sum_{i=1}^{\varphi} R_{t_b+i,n}(S_b)$. The expectation (15) is then approximated as

$$N^{-1} \sum_{n=1}^N \left[\left\{ (1 - \omega_b) \exp(\varphi r^f) + \omega_b \exp(R_{b+1,n}^s(s_b)) \right\}^{1-\gamma} Q(\boldsymbol{\pi}_{b+1}^{(j,n)}, t_{b+1}) \right], \quad (16)$$

¹²The symbol \odot denotes element by element multiplication.

where $\pi_{b+1}^{(j,n)}$ denotes the element π_{b+1}^j on the grid used to discretize the state space. $G = 10$ delivers sufficient precision and at the same time keeps the computational burden at a feasible level.

Table 4 reports the outcome of this exercise. Again we study the optimal stock holdings under three scenarios for the state probabilities. For each value of $\hat{\pi}$ we consider a range of rebalancing frequencies, $\varphi = 1, 3, 6, 12, 24$, and T months and we report optimal weights for investment horizons $T = 1, 6, 12, 24, 60$, and 120 months. To save space we only report results for the two-state regime switching model without the dividend yield.

Several interesting results emerge. Our earlier finding that investment schedules (as a function of T) are downward sloping in the bull state and upward sloping in the bear state continues to be supported under rebalancing. However, investment schedules become flatter, the higher is φ .¹³ In the limit with rebalancing every month ($\varphi = 1$), the state probabilities and asset allocations are updated at the same frequency so the investment schedules are completely flat.

At all investment horizons the possibility of rebalancing makes an investor use information about the current state relatively aggressively for portfolio selection purposes. In bull states, $\hat{\omega}_t$ is uniformly increasing in φ . Indeed, while under buy-and hold the long-horizon optimal weight (less than 40%) is roughly half of its level when predictability is ignored, when $\varphi = 1$ the optimal stock holding is very close to the IID level (81%).

In bear states ω_t is a uniformly decreasing function of φ . Indeed, while under buy-and-hold the long-horizon optimal weight (32%) is still substantial, under monthly rebalancing, the optimal weight is very small (4%). Long-run buy-and-hold investors anticipate the end of the bear market and know that the bear state will occur infrequently in the long run. They are therefore willing to hold large investments in stocks even in the bear market. However the possibility of rebalancing makes it optimal even for long-run investors to drastically reduce the commitment to stocks in a bear state.

Under high uncertainty about the current state, stock holdings are at intermediate levels. The main difference is now that the investment curves are almost flat irrespectively of the investment horizon. Since investors are unsure of the current

¹³Of course, for $\varphi \geq T$, by construction the buy-and-hold and rebalancing results coincide. This explains why some entries are identical in the table.

state of the stock market, they are unwilling to take extreme positions that condition on very good (bull state) or very poor (bear state) prospects for the stock market.

We conclude from this analysis that the possibility of frequent rebalancing makes an investor use the information provided by the state probability estimates more aggressively. In the bear state, the higher is φ , the more drastic is the reduction in the optimal allocation to stocks. Conversely, at times when the market is perceived to be in the bull state, an investor increases the optimal weight in stocks more the higher the rebalancing frequency.

6. Welfare Costs of Ignoring Regimes

An economic assessment of the costs from ignoring regimes in stock returns is best conducted using a metric based on expected utility. In our setting this is equivalent to maximizing utility subject to constraining investors to choose at time t an optimal allocation ω_t^{IID} under the assumption that stock returns simply follow a normal distribution with mean vector $\hat{\mu}_t$, and variance $\hat{\sigma}_t^2$. In this case the portfolio choice and savings ratio are independent of the investment horizon and the value function for the constrained investor is

$$J_t^{IID} \equiv \frac{1}{1-\gamma} \sum_{b=0}^B \beta^b E_t [W_b^{1-\gamma}]$$

$$W_b = W_{b-1} [(1 - \omega_t^{IID}) \exp(\varphi r^f) + \omega_b \exp(R_{b+1} + \varphi r^f)].$$

The assumption of independent and identically distributed returns is a constrained case of the model with regime switching and predictor variables so, by construction, we have

$$J_t^{IID} \leq J(W_t, r_t, \mathbf{z}_t, \boldsymbol{\theta}_t, \pi_t, t).$$

We compute the increase in initial wealth η_t^{IID} an investor would require to derive the same level of expected utility from the constrained and unconstrained asset allocation problems. η_t^{IID} therefore solves the following equation:

$$(1 + \eta_t^{IID})^{1-\gamma} \sum_{b=0}^B \beta^b E_t [(W_b)^{1-\gamma}] = Q(R_b, \mathbf{z}_b, \boldsymbol{\theta}_b, \pi_b, t_b),$$

with solution

$$\eta_t^{IID} = \left\{ \frac{Q(R_b, \mathbf{z}_b, \boldsymbol{\theta}_b, \pi_b, t_b)}{\sum_{b=0}^B \beta^b E_t [(W_b)^{1-\gamma}]} \right\}^{\frac{1}{1-\gamma}} - 1. \quad (17)$$

This expression is relatively easy to calculate since $Q(r_b, z_b, \theta_b, \pi_b, t_b)$ is a by-product of the numerical solution to the investor's portfolio choice. $\sum_{b=0}^B \beta^b E_t [(W_b)^{1-\gamma}]$ can be computed through either Gaussian quadrature or Monte Carlo methods.

Table 5 reports implied welfare costs for the three configurations of the initial state probabilities considered in our earlier analysis. Again we consider rebalancing at different frequencies, φ , and investment horizons, T . The cost of ignoring regimes increases uniformly as a function of the investment horizon, T , but decreases in the rebalancing frequency, φ , provided that $\varphi > T$. Welfare costs increase in the investment horizon because the present value of having a suboptimal asset allocation is higher the longer this position is locked in. Conversely, the shorter the rebalancing frequency, φ , the more valuable it is for investors to use their information about the underlying state. For example, they can react more aggressively in reducing stock holdings in the short-lived bear state provided that rebalancing is possible.

For this reason the welfare costs are at their highest in the bear regime under frequent rebalancing both because the return distribution in this state is most different from the average return distribution and because this state is short-lived. At their highest level, welfare costs amount to 4% of the initial portfolio value at the longest investment horizon. Welfare costs are comparatively smaller in the bull regime, but clearly cannot be ignored even in this state where they peak at 3%.¹⁴

These costs are sufficiently large to be economically relevant to investors with a long horizon such as a pension fund. While we have ignored transaction costs, these can reasonably be assumed to be an order of magnitude lower than the estimated potential gains, partly because bear states do not occur all that often and also because a shift between T-bills and stocks can be inexpensively implemented through positions in futures contracts.

¹⁴We also computed the welfare costs under the real-time state probability and parameter estimates and found effects that were very similar to those reported in Table 5. For instance at the 120 month investment horizon with no rebalancing the average utility costs increased from 0.2 to around 0.8 percent when the initial bear state probability increased from zero to one.

7. Conclusion

This paper found clear empirical evidence of regimes with very different volatility and mean in UK stock returns. Our results suggest that a two-state specification with a highly persistent low-volatility regime mixed with a less persistent high volatility state capture important features of UK stock returns. Predictability from the dividend yield by no means subsumes the effects of predictability due to the presence of persistent bull and bear states in returns. However, the relative effects of the two factors very much depends on the assumed investment horizon. The shorter the investment horizon, the more important the perceived state probabilities are relative to variations in the dividend yield.

One way to summarize the economic difference between our two-state model and a model that assumes no predictability is by comparing the stock holdings associated with these models at the end of our sample (2000). At this point, the optimal stock holdings would have been close to twice as high under the single-state model compared to the two-state model that considers the possibility of a bad return state. Unsurprisingly in view of such differences, the expected utility costs arising from ignoring regimes can be quite significant.

We have not explored the potential equilibrium implications of time-variations in investors' asset allocation. Instead we considered the optimal asset allocation of a small investor whose actions do not affect equilibrium rates of return. However, our model is by no means inconsistent with equilibrium effects which could give rise to regimes in stock returns. Whitelaw (2001) constructs an equilibrium model where consumption growth follows a two-state process so that investors' intertemporal marginal rate of substitution also follows a regime process. This environment is consistent with the process for returns assumed in our paper and suggests that regime-switching in returns is consistent with equilibrium and need not violate efficiency in financial markets.

Another issue that we intend to address in future work is how investors could deal with estimation uncertainty. While our paper accounts for learning in the sense that investors were assumed to optimally update their state beliefs, we did not endow investors with prior beliefs over the current parameter values. Doing so would require a Bayesian analysis and thus goes well beyond the current paper.

References

- [1] Ang A., and G., Bekaert, 2001, "International Asset Allocation with Regime Shifts", forthcoming, *Review of Financial Studies*.
- [2] Ang, A. and G., Bekaert, 2002, "Regime Switches in Interest Rates", *Journal of Business and Economic Statistics*, 20, 163-182.
- [3] Barberis, N., 2000, "Investing for the Long Run When Returns Are Predictable", *Journal of Finance*, 55, 225-264.
- [4] Berkowitz, J., 2001, Testing Density Forecasts with Applications to Risk Management. *Journal of Business and Economic Statistics* 19, 465-474.
- [5] Black, A., and P., Fraser, 1995, "UK Stock Returns: Predictability and Business Conditions", *The Manchester School Supplement*, 85-102.
- [6] Brandt, M., 1999, "Estimating Portfolio and Consumption Choice: A Conditional Euler Equations Approach", *Journal of Finance*, 54, 1609-1645.
- [7] Campbell, J., and R., Shiller, 1988, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors", *Review of Financial Studies*, 1, 195-228.
- [8] Campbell, J., and L., Viceira, 1999, "Consumption and Portfolio Decisions when Expected Returns are Time Varying", *Quarterly Journal of Economics*, 114, 433-495.
- [9] Clare, A.D., S.H. Thomas and M.R. Wickens, 1994, "Is the Gilt-Equity Yield Ratio Useful for Predicting UK Stock Returns?", *Economic Journal*, 104, 303-15.
- [10] Diebold, F., Gunther, T., and A., Tay, 1998, "Evaluating Density Forecasts", *International Economic Review*, 39, 863-883.
- [11] Driffill, J. and M., Sola, 1994, "Testing the Term Structure of Interest Rates from a Stationary Switching Regime VAR", *Journal of Economic Dynamics and Control*, 18, 601-628.

- [12] Durlauf, S., 1991, "Spectral Based Testing of the Martingale Hypothesis", *Journal of Econometrics*, 50, 355-376.
- [13] Fama, E., and K., French, 1988, "Dividend Yields and Expected Stock Returns", *Journal of Financial Economics*, 22, 3-25.
- [14] Gray, S., 1996, "Modeling the Conditional Distribution of Interest Rates as Regime-Switching Process", *Journal of Financial Economics*, 42, 27-62.
- [15] Hamilton, J., 1988, Rational Expectations Econometric Analysis of Changes in Regime: An Investigation of the Term Structure of Interest Rates. *Journal of Economic Dynamics and Control* 12, 365-423.
- [16] Hamilton, J., 1989, "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle", *Econometrica*, 57, 357-384.
- [17] Ingersoll, J., 1987, *Theory of Financial Decision Making*, Rowman & Littlefield.
- [18] Jarque, C., and A., Bera, 1980, "Efficient Tests for Normality, Homoskedasticity, and Serial Independence of Regression Residuals", *Economics Letters*, 6, 255-259.
- [19] Kandel, S., and R., Stambaugh, 1996, "On the Predictability of Stock Returns: An Asset Allocation Perspective", *Journal of Finance*, 51, 385-424.
- [20] Lynch, A., 2001, "Portfolio Choice and Equity Characteristics: Characterizing the Hedging Demands Induced by Return Predictability", *Journal of Financial Economics*, 62, 67-130.
- [21] Perez-Quiros, G. and A., Timmermann, 2000, "Firm Size and Cyclical Variations in Stock Returns", *Journal of Finance*, 55, 1229-1262.
- [22] Pesaran, M.H. and A., Timmermann, 2000, "A Recursive Modeling Approach to Predicting UK Stock Returns", *Economic Journal* 110, 159-191.
- [23] Rosenblatt, M., 1952, "Remarks on a Multivariate Transformation", *Annals of Mathematical Statistics*, 23, 470-472.

- [24] Samuelson, P., 1969, "Lifetime Portfolio Selection by Dynamic Stochastic Programming", *Review of Economics and Statistics* 51, 239-246.
- [25] Timmermann, A., 2000, "Moments of Markov Switching Models", *Journal of Econometrics*, 96, 75-111.
- [26] Turner, C., R., Startz, and C., Nelson, 1989, "A Markov Model of Heteroskedasticity, Risk, and Learning in the Stock Market", *Journal of Financial Economics*, 25, 3-22.
- [27] Whitelaw, R.F., 2001, Stock Market Risk and Return: An Equilibrium Approach. *Review of Financial Studies* 13, 521-548.

Table 1

Specification Tests for Regime Switching Models

The table reports tests for the transformed z-scores generated by univariate regime-switching models

$$R_t = \mu_{s_t} + \sum_{j=1}^p a_{js_t} y_{t-j} + \sigma_{s_t} \varepsilon_t$$

where R_t is the excess return on the FTSE-All Share index. $\varepsilon_t \sim \text{I.I.D. } N(\mathbf{0},1)$ and s_t is governed by an unobservable, first-order Markov chain that can assume k distinct values (states). The sample period is 1970:01 – 2000:12. The tests are based on the principle that under the null of correct specification of the model, the probability integral transform of the one-step-ahead standardized forecast errors should follow an IID uniform distribution over the interval (0,1). A further Gaussian transform described in Berkowitz (2001) is applied to perform LR tests of the null that (under correct specification) the transformed z-scores are IIN(0,1) distributed. MSIAH(k,p) stands for Markov Switching Intercept Autoregressive Heteroskedasticity model with k states and p autoregressive lags. Simpler models with fewer letters do not allow for state dependency in the mean, autoregressive or volatility components.

Model	Number of parameters	Jarque-Bera test	LR ₂	LR ₃	LR ₆
Panel A – Univariate Models for Excess Stock Returns					
Base model: MSIA(1,0)					
MSIA(1,0)	2	7.193 (0.027)	189.323 (0.000)	203.027 (0.000)	468.945 (0.000)
MSIA(1,1)	3	2,657.1 (0.000)	0.096 (0.953)	2.987 (0.394)	18.602 (0.005)
Base model: MSIA(2,0)					
MSIA (2,0)	5	15.993 (0.000)	0.066 (0.968)	5.876 (0.118)	20.380 (0.002)
MSIA (2,1)	7	16.336 (0.000)	0.253 (0.881)	2.619 (0.454)	15.890 (0.014)
MSIH (2,0)	6	2.486 (0.289)	0.086 (0.958)	3.266 (0.352)	10.235 (0.115)
MSIAH (2,1)	8	3.393 (0.183)	0.097 (0.953)	2.061 (0.560)	13.099 (0.041)
MSIH (2,0) – AR(1)	7	3.364 (0.262)	0.084 (0.959)	2.003 (0.572)	12.748 (0.047)
Base model: MSIA(3,0)					
MSIA (3,0)	10	13,411.1 (0.000)	3.390 (0.184)	5.348 (0.148)	25.691 (0.000)
MSIH (3,0)	12	2.138 (0.343)	0.046 (0.977)	2.917 (0.405)	10.014 (0.124)
Base model: MSIA(4,0)					
MSIA (4,0)	17	55.872 (0.000)	0.244 (0.885)	3.432 (0.330)	14.516 (0.024)
MSIH (4,0)	20	0.531 (0.767)	0.005 (0.998)	4.002 (0.261)	12.632 (0.049)

Table 2**Regime Switching Model fitted to Stock Returns**

This table reports maximum likelihood estimates for a single state model and a two-state regime switching model fitted to monthly excess returns on the FTSE All Share index. The regime switching model takes the form:

$$R_t = \mu_{s_t} + \sigma_{s_t} \varepsilon_t$$

where $\varepsilon_t \sim N(\mathbf{0},1)$ is an unpredictable return innovation. The sample period is 1970:01 – 2000:12.

Panel A – Single State Model		
Mean excess return	0.603*	
Volatility	6.175	
Panel B – Two State Model		
Mean excess return		
Regime 1 (bull)	0.716**	
Regime 2 (bear)	-0.650	
Volatility		
Regime 1 (bull)	4.535	
Regime 2 (bear)	15.170	
Transition probabilities	Regime 1	Regime 2
Regime 1 (bull)	0.986	0.014
Regime 2 (bear)	0.156	0.844

* = significant at 5% level; ** = significant at 1% level

Table 3

Estimates for a Bivariate Regime Switching Model

This table reports maximum likelihood estimates for a bivariate VAR and a two-state regime switching model fitted to monthly excess returns and the dividend yield on the FTSE All Share index. The regime switching model takes the form

$$y_t = \mu_{s_t} + \sum_{j=1}^p a_{js_t} y_{t-j} + \sigma_{s_t} \varepsilon_t,$$

where y_t is a vector collecting the excess return and the dividend yield, μ_{s_t} is an intercept vector in state s_t , A_{1s_t} is a matrix of first-order autoregressive coefficients in state s_t and $\varepsilon_t = [\varepsilon_{1t} \ \varepsilon_{2t}]' \sim \text{I.I.D. } N(\mathbf{0}, \Omega_{s_t})$. s_t is governed by an unobservable, first-order Markov chain that can assume two distinct values. The data is monthly and covers the period 1970:01 – 2000:12. Panel A refers to the single state benchmark ($k = 1$) while panel B refers to the two-state model ($k = 2$). The values on the diagonals of the correlation matrices are volatilities, while off-diagonal terms are correlations.

Panel A – Single State Model		
	Excess stock returns	Dividend yield
Mean excess return	1.151**	0.073**
VAR(1) coefficients		
Excess stock returns	0.114**	0.826**
Dividend yield	-0.006**	0.958**
Correlations/Volatilities		
Excess stock returns	6.357	
Dividend yield	-0.923	0.386
Panel B – Two State Model		
	Excess stock returns	Dividend Yield
Intercepts		
Regime 1 (bull)	-1.221	0.081*
Regime 2 (bear)	-20.519**	1.246**
VAR(1) coefficients		
<i>Regime 1 (bull):</i>		
Excess stock returns	-0.001	0.404
Dividend yield	-0.0004*	0.980**
<i>Regime 2 (bear):</i>		
Excess stock returns	0.241*	3.676
Dividend Yield	-0.014**	0.788**
Correlations/Volatilities		
<i>Regime 1 (bull):</i>		
Excess stock returns	4.367	
Dividend yield	-0.939	0.197
<i>Regime 2 (bear):</i>		
Excess stock returns	11.349	
Dividend yield	-0.964	0.887
Transition probabilities	Regime 1	Regime 2
Regime 1 (bull)	0.970	0.030
Regime 2 (bear)	0.205	0.795

* = significant at 5% level; ** = significant at 1% level

Table 5

Expected Utility Costs

The table reports the expected utility cost due to being constrained to take portfolio decisions on the basis of a model that ignores regimes for an investor with power utility and a coefficient of relative risk aversion $\gamma = 5$. Excess returns on the FTSE-All Share index are assumed to be generated by a univariate two-state regime-switching model. Expected utility costs are identified with the compensatory variation, i.e. the percentage increase in initial wealth required by the investor in order to willingly ignore regimes. The table considers three possible values of the perceived probability of being in state 1, $\hat{\pi}_t = 1$ (bull state), $\hat{\pi}_t = 0$ (bear state) and $\hat{\pi}_t = 0.8$ (high uncertainty).

Rebalancing Frequency φ	Investment Horizon T (months)					
	T=1	T=6	T=12	T=24	T=60	T=120
Bull state						
$\varphi = T$ (buy-and-hold)	0.000	0.003	0.015	0.084	0.218	0.228
$\varphi = 24$ months	0.000	0.003	0.015	0.084	0.220	0.251
$\varphi = 12$ months	0.000	0.003	0.015	0.127	0.221	0.512
$\varphi = 6$ months	0.000	0.003	0.031	0.156	0.612	1.261
$\varphi = 3$ months	0.000	0.042	0.156	0.453	1.306	2.573
$\varphi = 1$ month	0.000	0.048	0.187	0.549	1.453	2.980
Bear state						
$\varphi = T$ (buy-and-hold)	0.028	0.118	0.227	0.421	0.692	0.764
$\varphi = 24$ months	0.028	0.118	0.227	0.421	0.435	0.910
$\varphi = 12$ months	0.028	0.118	0.227	0.341	0.674	1.299
$\varphi = 6$ months	0.028	0.118	0.198	0.524	1.114	2.382
$\varphi = 3$ months	0.028	0.108	0.327	0.895	2.016	3.169
$\varphi = 1$ month	0.028	0.413	0.824	1.207	2.161	4.026
High uncertainty						
$\varphi = T$ (buy-and-hold)	0.002	0.008	0.025	0.071	0.313	0.414
$\varphi = 24$ months	0.002	0.008	0.025	0.071	0.342	0.430
$\varphi = 12$ months	0.002	0.008	0.025	0.073	0.455	0.602
$\varphi = 6$ months	0.002	0.008	0.072	0.299	0.918	1.929
$\varphi = 3$ months	0.002	0.049	0.174	0.501	1.904	3.744
$\varphi = 1$ month	0.002	0.123	0.397	0.931	2.101	3.999

Figure 1
Smoothed Bear State Probability from the Two-State Model Fitted to
Stock (Excess) Returns

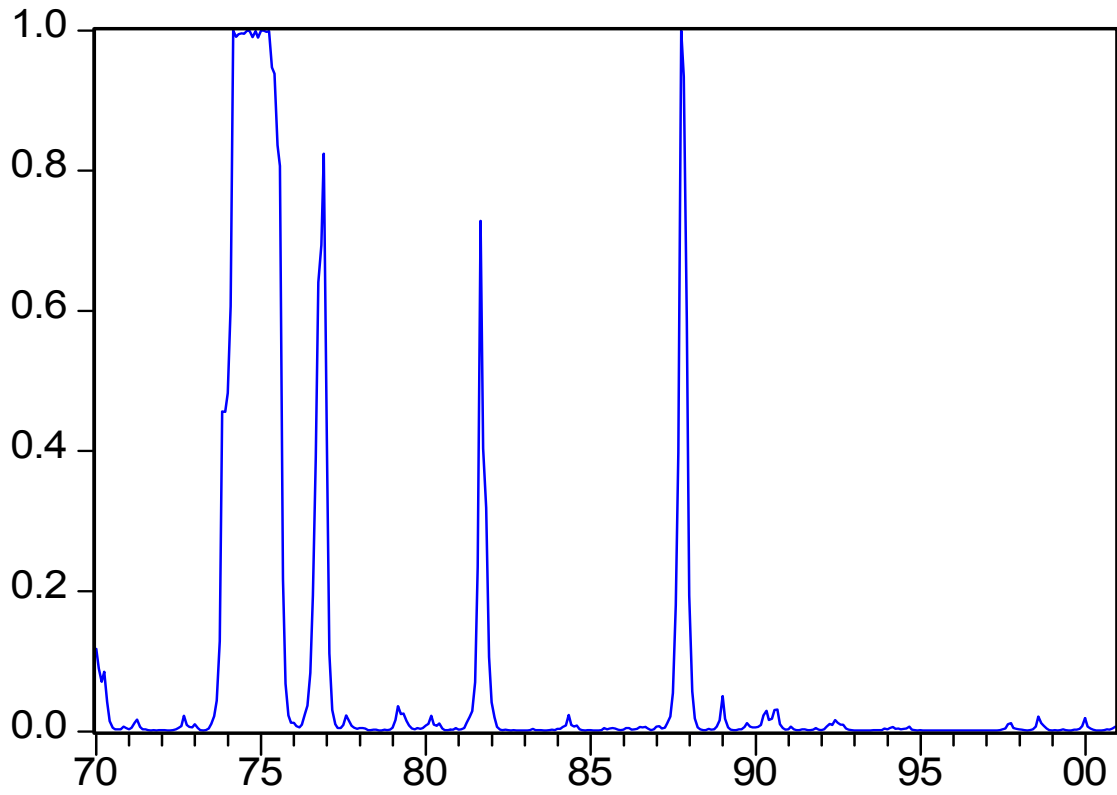


Figure 2

Moments of Stock Returns as a Function of the Probability of Bull State (1)

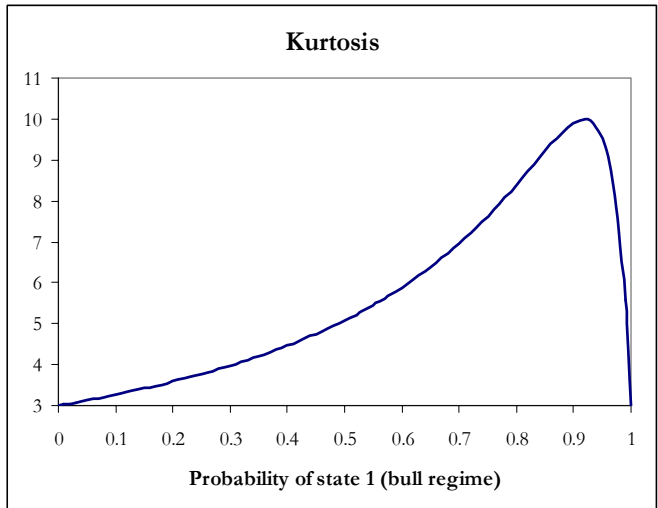
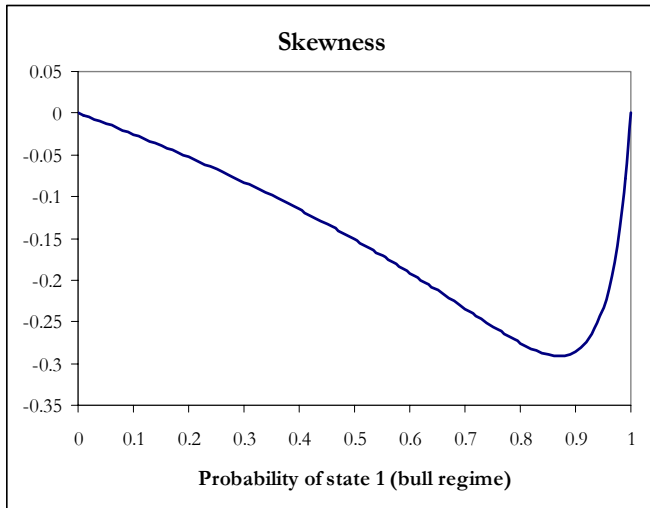
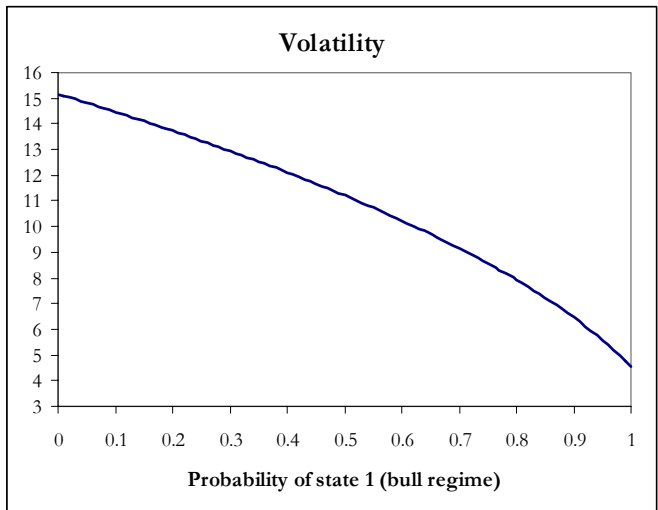
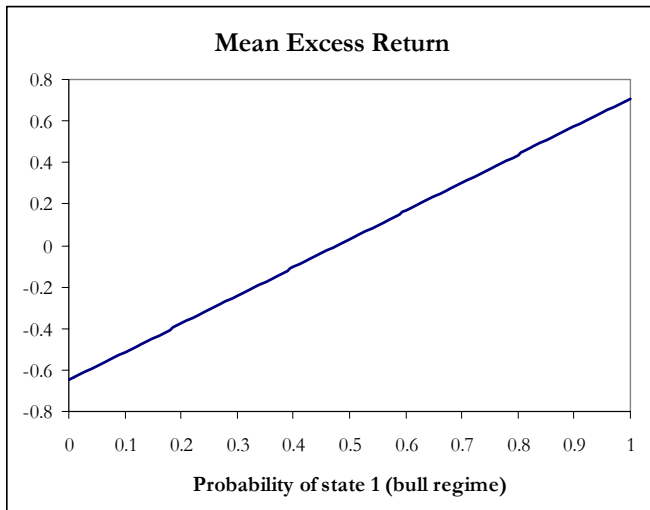


Figure 3

Real-Time Optimal Stock Holdings at Four Investment Horizons

This figure plots the time series of optimal stock holdings for an investor with power utility and coefficient of relative risk aversion $\gamma = 5$. All allocations are based on recursively updated real-time parameter estimates from the two-state model of excess returns with state-dependent mean and variance. We also show optimal holdings under the assumption of no predictability.

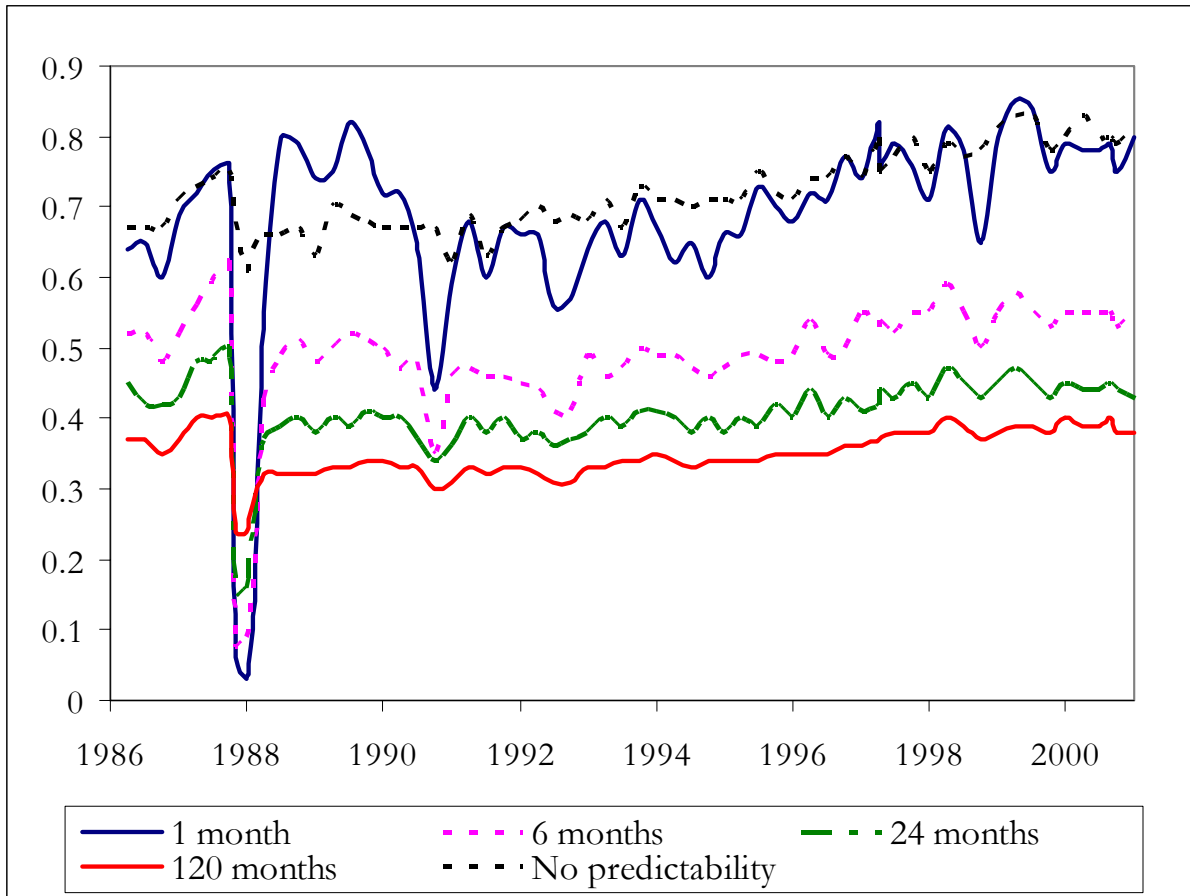


Figure 4

Probability Beliefs, Investment Horizon and Optimal Stock Holdings

This figure plots the optimal allocation to stocks as a function of the investment horizon and the perceived probability of a bull state for an investor with power utility and coefficient of relative risk aversion $\gamma = 5$.

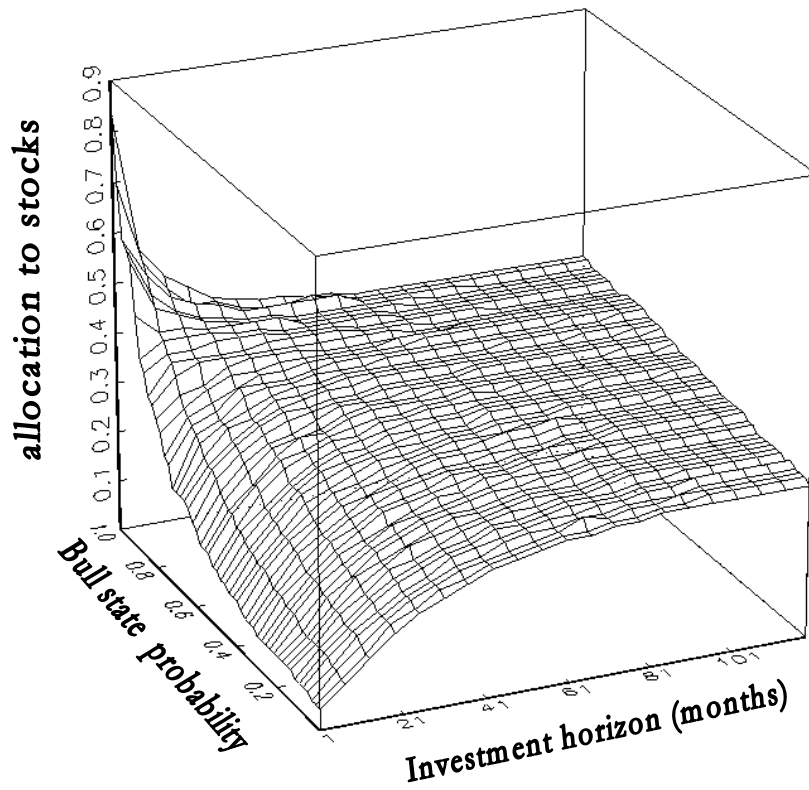


Figure 5

Effects of Relative Risk Aversion on Optimal Asset Allocation

The graphs plot optimal stock holdings as a function of the coefficient of relative risk aversion (γ) for an investor with power utility.

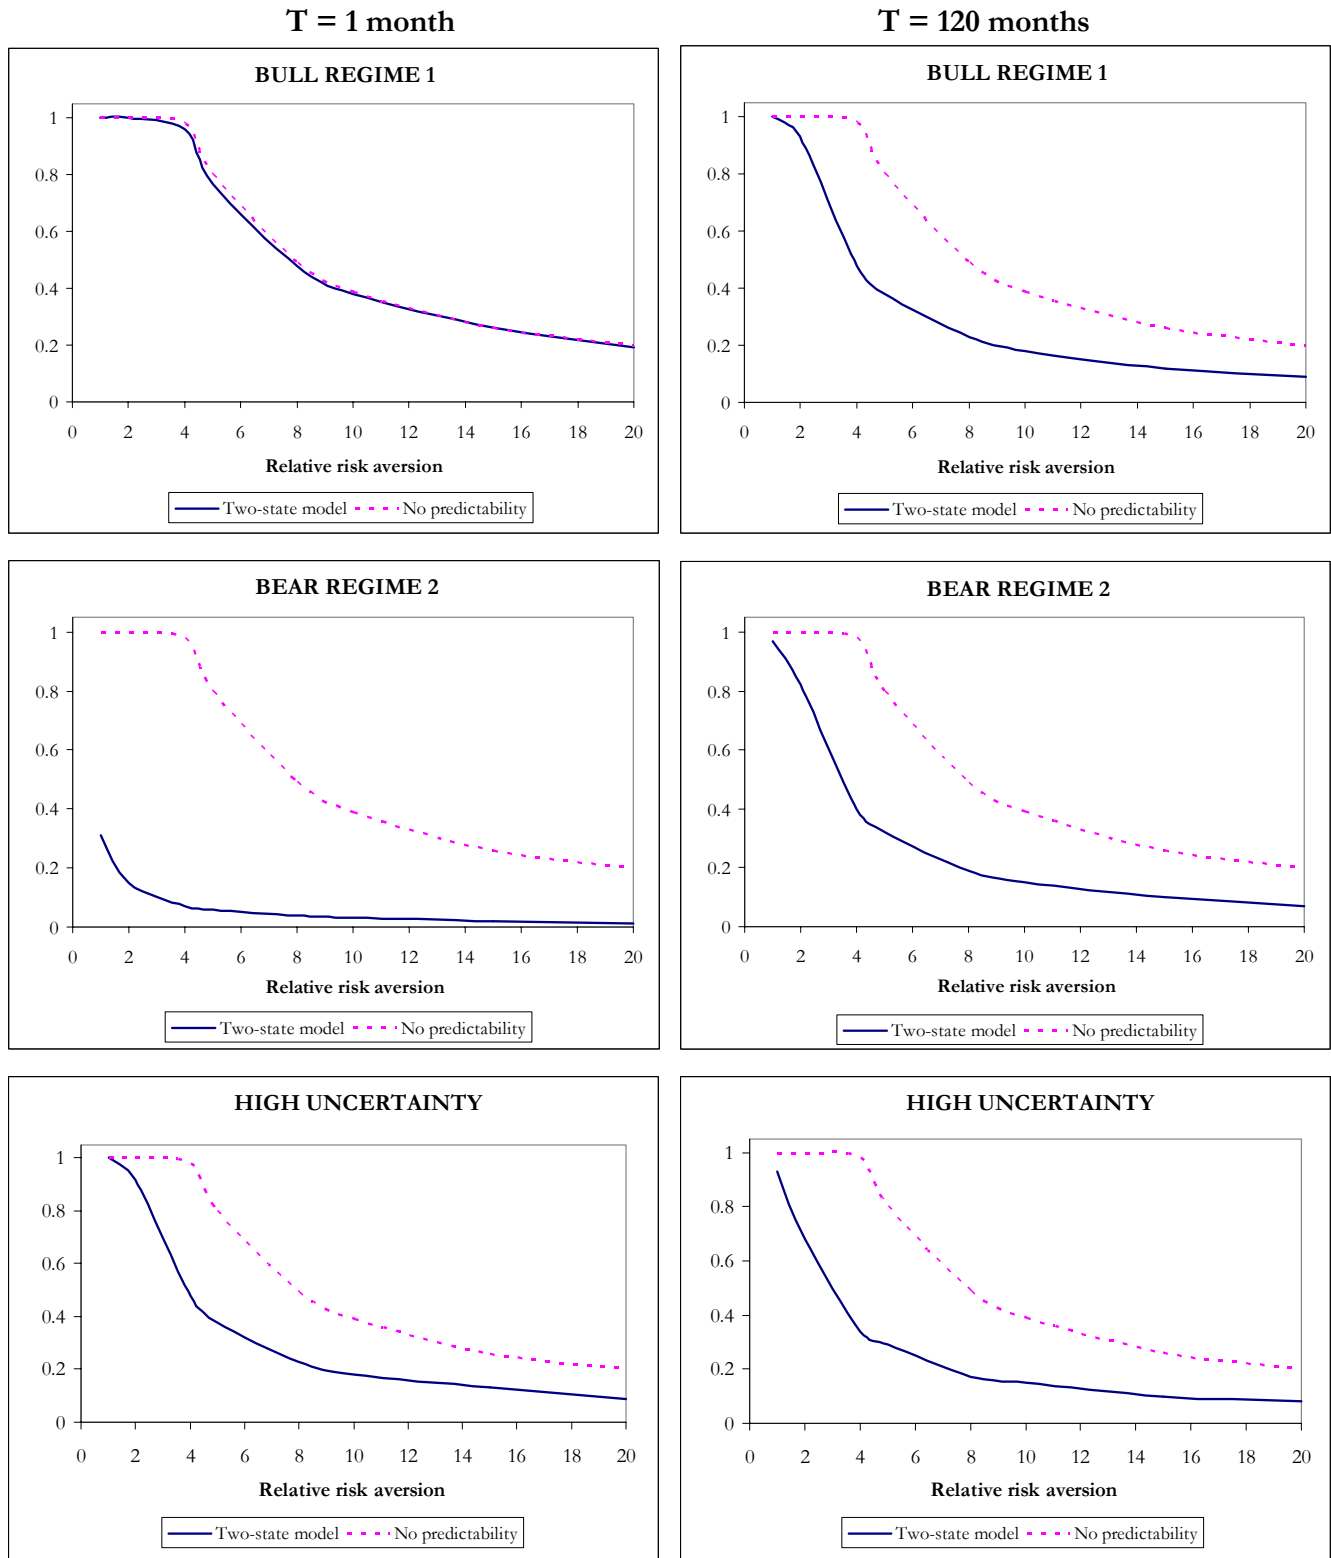


Figure 6

Smoothed Bear State Probability from the Bivariate Two-State Model Fitted to Excess Stock Returns and the Dividend Yield

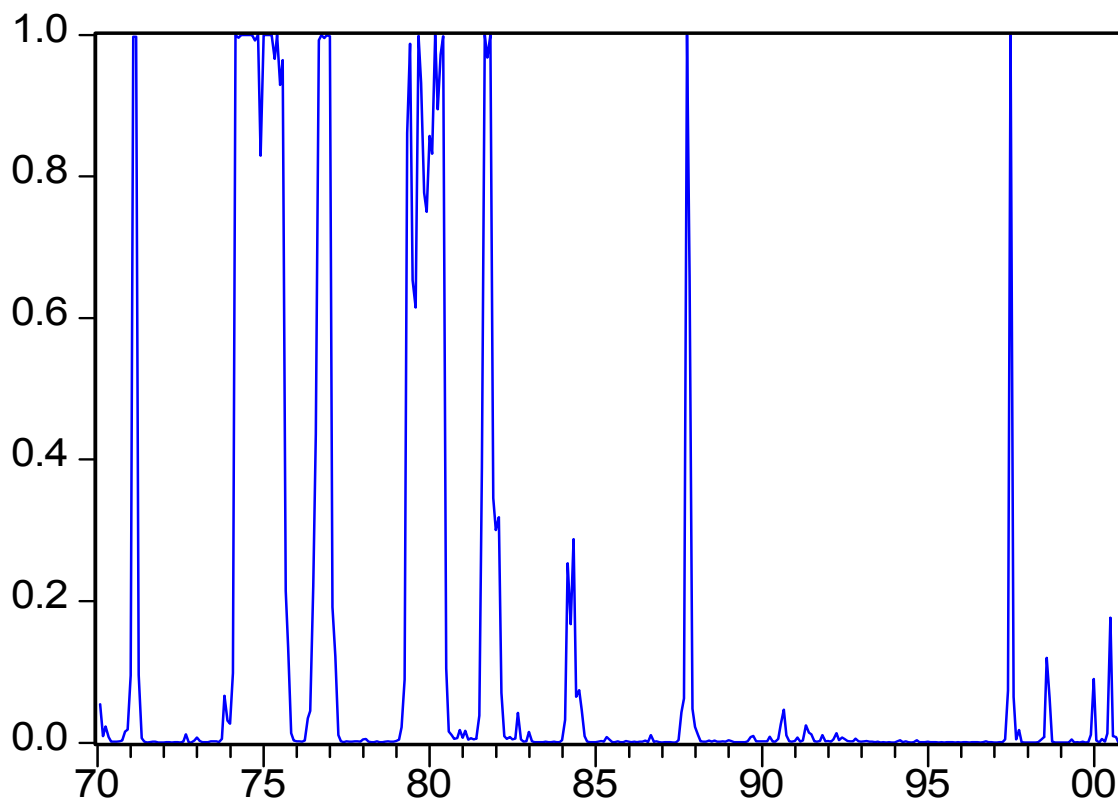


Figure 7

Real-Time Optimal Allocation to Stocks under Predictability from the Dividend Yield

This figure plots the time series of optimal stock holdings for an investor with power utility and coefficient of relative risk aversion $\gamma = 5$, in the presence of predictability from the dividend yield. All allocations are based on recursively updated real-time parameter estimates from k -regime VAR(1) models fitted to joint process for excess returns and dividend yields. The upper plot assumes $k = 2$, the bottom one $k = 1$ (Gaussian VAR(1)) when the matrix of autoregressive coefficients is restricted as in Barberis (2000). The restrictions imply no serial correlation in excess returns and no effects from lagged excess returns on the current dividend yield.

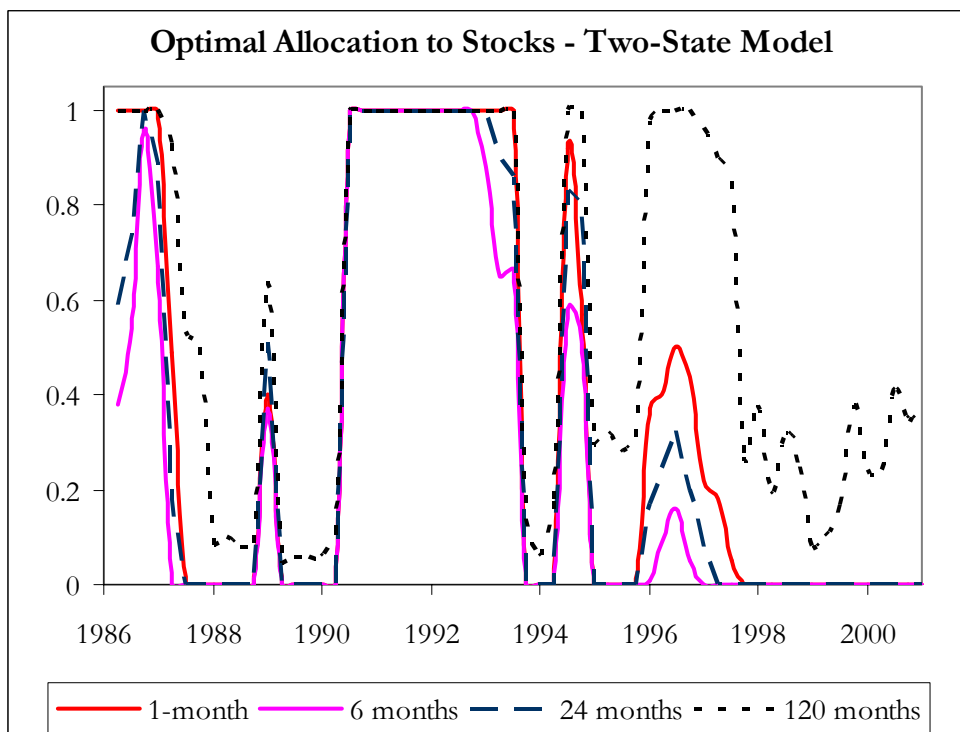
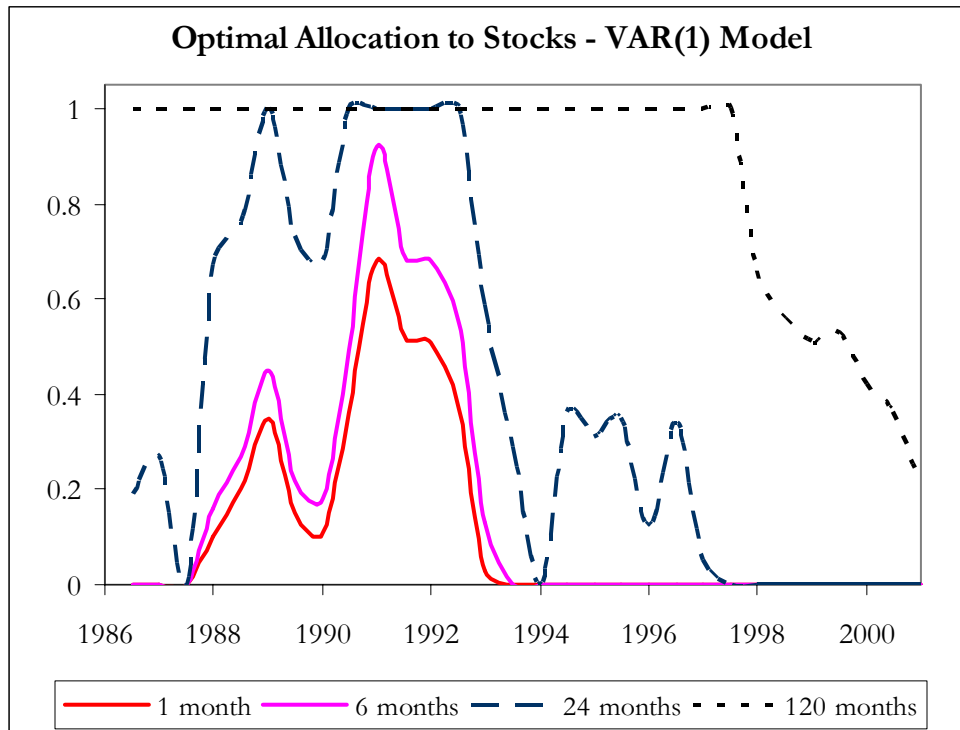


Figure 8

Optimal Stock Holdings For Different Values of the Dividend Yield

This figure plots the optimal allocation to stocks (as a proportion of the total portfolio) for various points in the sample selected to represent different values of the dividend yield. All allocations are based on recursively updated real-time parameter estimates from k -regime VAR(1) models fitted to the joint process for excess returns and dividend yields. The upper plot refers to $k = 2$, the bottom one to $k = 1$ (Gaussian VAR(1)) when the matrix of autoregressive coefficients is restricted to have zeros in the first column as in Barberis (2000). The restrictions imply no serial correlation in excess equity returns and no effects from lagged excess returns on the current dividend yield.

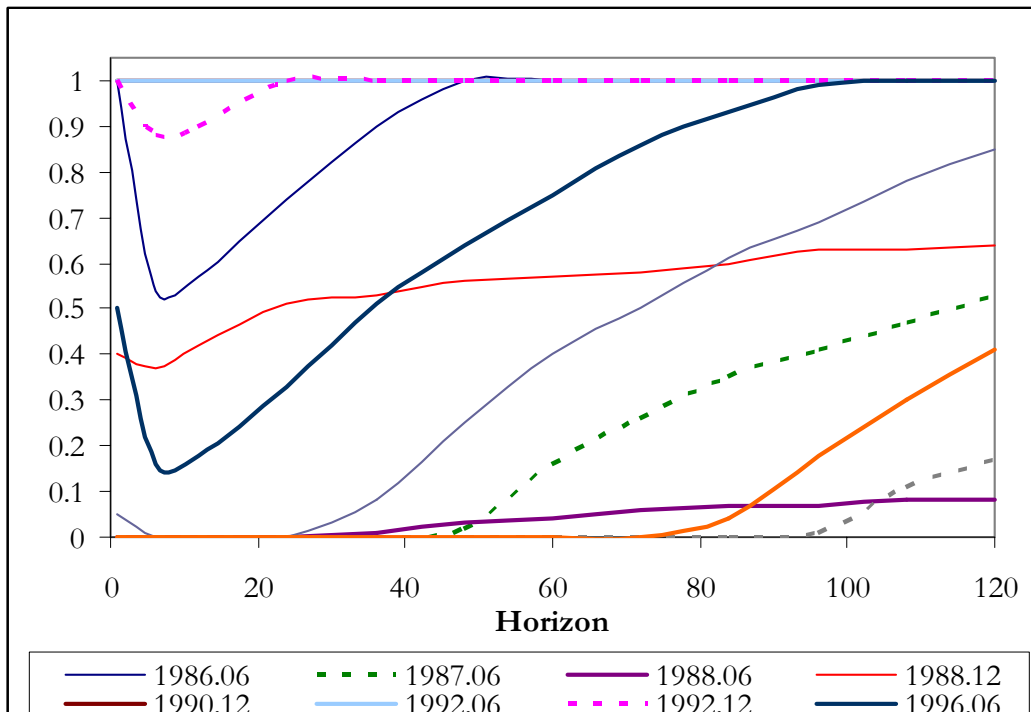
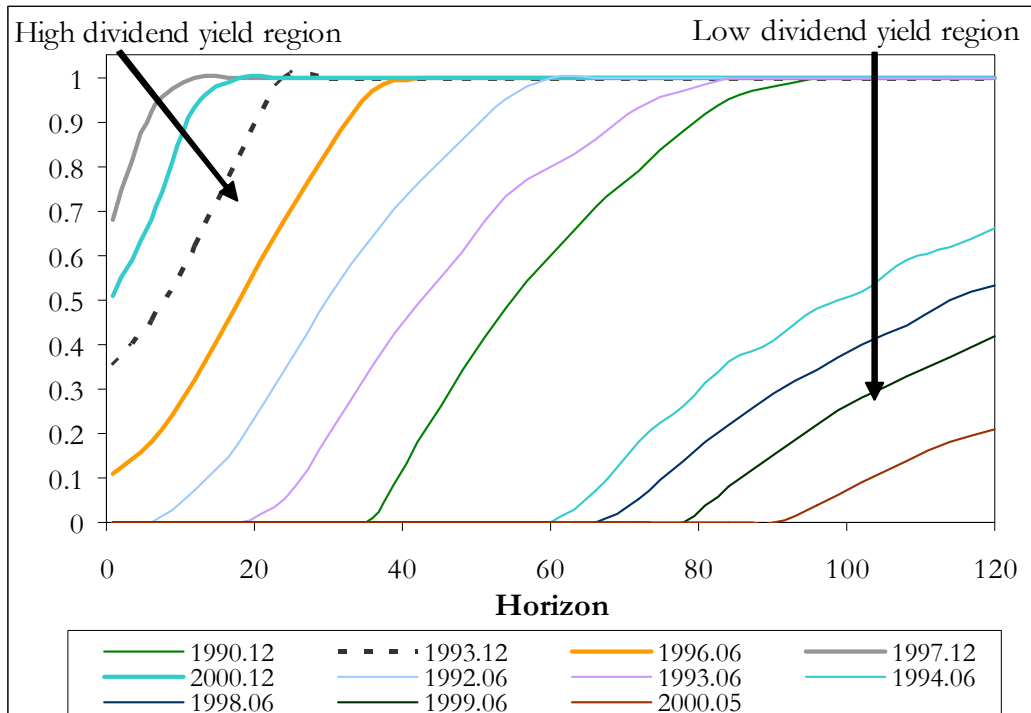


Figure 9

Probability Beliefs, Dividend Yields and Optimal Stock Holdings

This figure plots the optimal allocation to stocks as a function of the perceived probability of a bull state and the div

