Market Timing across Multiple Economic Regimes

Ron Guido[†], Joshua Pearl[‡] and Kathleen Walsh[‡]

ABSTRACT

This paper investigates whether a model which incorporates business cycle effects to infer the prevailing regime in US equity markets is beneficial to an investor who wishes to engage in market timing. The US equity premium is modelled as a Markov switching process where the regimes are dependent on economic variables. To characterise the economic regime, we employ the dimension reduction technique of Principal Components Analysis (PCA) to extract business cycle signals from a set of observed macroeconomic variables. By using these conditioning agents to infer the regime we find strong evidence of conditional normality and that the regime switching model for excess equity market returns provides a superior statistical fit. We then test a dynamic asset allocation strategy which invests in equity and cash on the basis of the predicted regimes. This timing strategy is shown to outperform a simple buy and hold strategy on a risk adjusted basis.

First Draft: 1 November 2006
Current Draft: 15 January 2007

[†] State Street Global Advisors, London UK

[‡] School of Banking and Finance, The University of New South Wales, Sydney, Australia.

Introduction

The financial idiom of 'bull' and 'bear' markets lends itself to the argument that there exist two distinct regimes within the equity sphere. This notion is considered here by modelling equity returns as a dual regime process and then attempts to identify these regimes ex-ante. From an investor's perspective, this presents a tantalising scenario. If we can formulate a strategy that involves the successful forecasting of regimes, it may be possible to reduce the volatility of the investment while simultaneously increasing return.

There has been overwhelming evidence of regime switching in equity returns (Turner, Startz and Nelson (1989), Hamilton and Susmel (1994) and Cai (1994). This technique is able to capture the effects of asymmetries, excess kurtosis and jumps of the return distribution. While the statistical properties give significant cause to model equity returns in a regime switching framework, the case is strengthened when you consider the regime switching properties of business cycles (see Hamilton (1989) and Filardo (1994)).

Previous research has shown that asset prices can be successfully represented by a mixture of distributions in a regime switching framework. The seminal study of Hamilton (1988) into the regime switching nature of interest rates provides a powerful precedent for studies that endeavour to model return series in a multi-regime framework. Evidence that financial return series are non-linear in nature has been provided by a number of preceding studies finding models that utilise a multiple regime methodology outperform their linear counterparts. Other studies have shown that the unconditional normal distribution is a poor representation for return series (see Peiro (1994), Zhou (1993) and Richardson and Smith (1993) offering implicit evidence as to why the dual regime model is superior. Despite the fact that the multi-regime framework tends to outperform the linear approach, the existing paradigm is to model equity returns in a single regime context. While the relative tractability of the single regime model is appealing, the efficiency and accuracy in comparison to the richer multi-regime technique is questionable.

A major contribution of this research is the use of economic variables to determine these regimes ex ante. There has been evidence of regime switching in macroeconomic variables such as inflation (Evans and Wachtel (1993) and Bhar and Hamori (2004)) and interest rates (Hamilton's (1988), Garcia and Perron (1996) and Gray (1996)). If any combination of these variables plays a deterministic role in the level of equity returns there is a valid argument as to why stock returns also exist in a multi-regime universe. We also consider economic variables that are likely to determine these regimes ex ante. These include economic variables that have been utilized in return forecasting such as inflation (Fama and Schwert (1977), Chen, Roll and Ross (1986)), interest rates (Folger, John and Tipton (1981), Fama and Schwert (1977) and Ferson (1989)), industrial production and consumption (Grauer, Litzenburger and Stehle (1976), Lettau and Ludvigson (2001) and Whitelaw (2000)) and dividend yield (Fama and French (1988) and Guidolin and Timmermann (2005)). Although a large number of studies have endeavoured to predict equity return levels, very few have used these to predict regimes.

This paper uses economic variables to predict regimes and then uses them to implement an asset allocation strategy. Ang and Bekaert (2004) is one of the few studies that considers asset allocation in the context of regime switching. They find that a strategy based upon this formulation yields greater raw and risk adjusted returns than alternate strategies. Instead of conditioning the asset allocation strategy on the switching nature of interest rates as in Ang and Bekaert (2004), we use the elements of the transition probability matrix of equity returns.

Given the set of financial and macroeconomic variables, we employ a principal components analysis (PCA) to condense all the information of the variables into an index. Despite its comparative advantages, the employment of this technique by studies within the equity sphere is limited. No prior study has allowed the switching nature of equity regimes to depend upon economic variables via a PCA. For the purposes of this study, the method is ideal. By aggregating all the information into an index, the model remains parsimonious and is expected to yield superior estimates. This is the first study that endeavours to condition the regimes on economic variables via the process of a PCA analysis, with the technique more commonly employed in inflation and business cycle analysis.

This study finds that, consistent with prior research, equity returns can be well represented by a mixture of normals and that such a representation is superior to modelling the return series in a single regime framework. The hypothesis that the series is drawn from a normal distribution is rejected, however when the data is split into two regimes, both of these regimes are found to be normally distributed. We find that strategic shifts in cash and equity on the basis of the regime switching probabilities offers investors a mechanism in which they can significantly reduce the overall risk of the portfolio. Using risk adjusted returns such an investment strategy is found to significantly outperform the market.

1. Methodology

1.1. The Statistical Structure

The focus of this study is to describe the evolution of equity returns whilst allowing for regime switching. The moments of the return series in any given period are governed by an unobservable latent state variable. By allowing multiple regimes one is essentially allowing the moments of the return series to be conditional on the prevailing regime. Specifically, the mean and variance of the series can be characterised as below:

$$E_t[r_{t+1}] = \mu_t = \begin{cases} \mu_1 & \text{if } S_t = 1 \\ \mu_2 & \text{if } S_t = 2 \end{cases} \tag{1}$$

$$E_t[r_{t+1} - \mu_t]^2 = \sigma_t = \begin{cases} \sigma_1 & \text{if } S_t = 1 \\ \sigma_2 & \text{if } S_t = 2. \end{cases} \tag{2}$$

The unobservable latent state variable S_t dictates the prevailing regime with its evolution over time following a first order Markov process. S_t is a stochastic (non-determinable) process in which all the observations are drawn from a discrete set. As a first order Markovian process, only the most recent realisation of the state variable has any effect on the distribution of the proceeding realisation. The description of the transition of the regimes from one period to the next will depend whether a constant transition probability matrix is being considered or whether the focus is on the time varying transition probability matrix. When the model considers constant transition probabilities, the process from S_t to S_{t+1} can be represented by the following Markov transition density matrix:

$$\Pi = \begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{bmatrix}. \tag{3}$$

where $p_{ij} = P(S_{t+1} = S_j \mid S_t = S_i)$ and $\sum_{j=1}^{2} p_{ij} = 1$ for all *i* with p_{ij} the probability of switching from a regime *i* period to a regime *j* period.

When time varying transition probabilities are used, the probability of the state depends upon the information available at the end of the preceding period. Under this scenario, the process from S_t to S_{t+1} can be represented by the transition density matrix:

$$P(S_{t+1}=s_{t+1} \mid S_t=s_t, X_t)=\begin{pmatrix} p_{11}(S_t,X_t) & p_{12}(S_t,X_t) \\ p_{21}(S_t,X_t) & p_{22}(S_t,X_t) \end{pmatrix}, \tag{4}$$

where $p_{ij}(S_t,X_t)=P(S_{t+1}=S_j \mid S_t=S_i,X_t)$ is the time varying probability of switching to regime *j* given the system was in regime *i*, and the information set prevailing at time t , denoted as X_t. In this study, we characterise X_t as the set of economic variables which jointly describe the economic state.

Given that S_t is a first order Markov process, the evolution of the return series can be described as $R_{\tau,t+1}=\mu(S_t)+\sigma(S_t)\ \varepsilon_t$ where $\mu(S_t)$ and $\sigma(S_t)$ are the state dependent mean and variance respectively and ε_t is a standard Gaussian process. Thus the return is measured with signal plus noise. The noise component depends upon both $\sigma(S_t)$ and ε_t, with the magnitude of the variation depending upon the volatility of the return which is state conditional. Given the nature of ε_t, the expected return can be characterised as $E[R_{\tau,t+1} \mid S_t=i]=\mu_i$. The state conditional expected value of the return, is equal to the first moment of the return series in the given regime. The observed return will deviate from the conditional mean of the regime governed by the state conditional volatility and the standard normal process ε_t. The greater the level of state dependent variance, the greater the amount that the observed return can deviate from expectations.

The dual regime model therefore allows the unconditional return series to behave as a mixture of normals. The conditional return series is a normal distribution whose moments depend upon the unobservable latent state variable. By modelling the series in such a fashion we can account for unconditional non-normality, skewness, excess kurtosis and heavy-tailed attributes of a given return series.

1.2. Estimation of the Model and Generation of the Regime Probabilities

Bayesian analysis is adopted in order to generate parameter estimates. The Bayesian methodology offers a number of benefits as it allow the user to modeller prior beliefs over the model and regimes. By combining the prior beliefs and the likelihood of the model given the data, the modeller arrives at the posterior distribution of the parameters; from which inference can then be performed.

In the following analysis two models shall be considered. An uninformative model which assumes we know nothing about the posterior distributions of the parameters in either regime. The second model places an informative prior on the first moment of the series in each regime. From here on we shall refer to these as the uninformative and informative models. The reason for employing informative priors in a market timing study is naturally and intuitively appealing. If the investor wishes to form a trading strategy that conditions upon whether the expected equity premium is positive or negative, the placing such priors to identify the regimes would be desirable.

Markov Chain Monte Carlo (MCMC) simulation is employed to estimate the posterior distributions of the state dependent parameters through an application of the Gibbs Sampler. For the purposes of the current model specification, we define the following quantities. Let $\theta = \left\{ \mu_i, \sigma_i^2, p_{ij} \right\}_{\forall i,j}$ represent the parameter vector for the model in question; $S = (S_1, ..., S_t)$ represents the vector of the latent state variables up to time that define the regime; $X = (X_1, ..., X_t)$ represents the observable vector economic variables which we will use to determine the regime and $Y = (Y_1, ..., Y_t)$ is the return series. In this context we seek to derive the joint posterior of the parameter vector and the latent state vector given the set of observable data: $f(S, \theta \mid Y, X)$.

According to Bayes theorem, the joint posterior can be considered as

$$f(S,\theta \mid Y,X) \propto f(Y,X \mid S,\theta) f(S \mid \theta) f(\theta) . \quad (5)$$

Given the model's analytical intractability, we employ the Gibbs sampling scheme to derive estimates for S and θ. Obtaining the Bayesian estimators for the model's parameters thus entails sampling from the set of full conditional posterior distributions. More detailed descriptions of this approach can be found in Casella and George (1992), Tanner (1996) and Chib and Greenberg (1996).

In order to generate the state variable when testing the asset allocation strategy, the methodology adopts an augmentation of the block sampling scheme developed Carter and Kohn (1994) and Chib (1995). So that there is no forward looking bias, the agent uses only the information in the data set up until point t when predicting the state of $t+1$. The following augmentation allows for the non-recursive estimate of the state variable. A Markov prior is imposed upon the indicator variable S_t and we generate S from Y, X and θ. For notational simplicity, the dependence upon θ is omitted and thus we are concerned with the generation of the state variable S from $\Pr(S \mid Y,X)$.

By noting that we can express the density function for the latent vector S as $\Pr(S \mid Y,X) = \Pr\{S_1 \mid Y_1\} \prod_{t=2}^{n} \Pr(S_t \mid Y_t, S_{t-1}, X_{t-1})$, we are able to derive a recursive scheme with which we can generate the state variable. We generate S by initially generating S_1 from $\Pr\{S_1 \mid Y_1\}$; for $t = 2:n$ in that order, S_t is generated from $\Pr(S_t \mid Y_t, S_{t-1}, X_{t-1})$. The reader is referred to Crater and Kohn (1994) for further details.

For the specification with time varying transition probabilities, the Carter and Kohn (1994) block sampling scheme is combined with the probit link function specification of Albert and Chib (1993b). This allows us to incorporate the effects of the economic variables on the transition probabilities of the state variable. The elements of the transition probability matrix are derived using estimates of a probit function whose parameters are themselves random variables. The transition matrix can be expressed as:

$$P(S_{t+1} = s_{t+1} \mid S_t = s_t, X_t) = \begin{pmatrix} p_{11}(S_t, X_t) & p_{12}(S_t, X_t) \\ p_{21}(S_t, X_t) & p_{22}(S_t, X_t) \end{pmatrix} \qquad (6)$$

We model each of the elements of the transition probability matrix as a function of the variables S and X. To specify the probit link function, consistent with Albert and Chib (1993b) and Filardo and Gordon (1998), we augment the parameter space by the inclusion of a latent variable, S_t^* which is positive when S_t = 1 and negative when S_t = 2; $P(S_t = 1) = P(S_t^* \geq 0)$. We model this latent variable to be a function of last period's regime and economic variables: $S_{t+1}^* = \beta_C + \beta_X' X_t + \beta_R R_t + u_t$ or more simply as $S_{t+1}^* = Z_t \beta + u_t$. R_t is a dummy variable which captures the effect the previous regime on the transition probability. When $S_t = 2$, R_t takes on the value of one and when $S_t = 1$, R_t takes on the value of zero. Without any loss of generality, to facilitate the derivation of the transition probability matrix it is assumed that $u_t \sim N(0,1)$.

We create the variable $S_{t+1}{}^*$ such that its distribution depends upon the value of S_{t+1}. A conditional truncated normal prior is placed on $S_{t+1}{}^*$ in order to limit the range of values it may take on. For $S_{t+1}{}^*$ we draw from the truncated normal bounded to be negative when $S_{t+1} = 1$ and a truncated normal bounded to be positive when $S_{t+1} = 2$:

$$\begin{aligned} S_{t+1}^{*} \mid S_{t+1} = 1, \beta &\sim TN_{(-\infty,0]}(Z_t\beta, 1) \\ S_{t+1}^{*} \mid S_{t+1} = 2, \beta &\sim TN_{(0,\infty)}(Z_t\beta, 1) \end{aligned} \quad . \tag{7}$$

After augmenting the parameter space with the latent variable S_{t+1}^{*} we are able to derive estimates of coefficient parameters. Placing a normal prior on β, $\beta \sim N(\alpha_0, A_0^{-1})$, the conditional posterior distribution reduces to $\beta \mid S^{*}_{n} \sim N(\hat{\beta}, (A^{1}_{0} + X'X)^{-1})$, where $\hat{\beta} = (A^{1}_{0} + X'X)^{-1}(A^{1}_{0}\alpha^{1}_{0} + X'S^{*}_{n})$. Once we have generated the state variable, the latent variable and the coefficient vector, we can generate the transition probabilities. Since, $\mathrm{P}(S_{t+1} = 2 \mid S_t = 2, X_t)$ = $P(S^{*}_{t+1} \leq 0 \mid S_t = 2, X_t)$ we therefore obtain

$$p_{22}(S_t, X_t) = \Phi(\beta_{\mathrm{C}} + \beta_{\mathrm{X}}{}' X_{\mathrm{t}} + \beta_{\mathrm{R}}), \tag{8}$$

Hence $p_{21}(S_t, X_t) = 1 - p_{22}(S_t, X_t) = \Phi(-(\beta_{\mathrm{C}} + \beta_{\mathrm{X}}{}' X_{\mathrm{t}} + \beta_{\mathrm{S}}))$.

The estimation of $\mathrm{P}(S_{t+1} = 1 \mid S_t = 1, X_t)$ is similarly achieved. Since $\mathrm{P}(S_{t+1} = 1 \mid S_t = 1, X_t)$ = $P(S^{*}_{t+1} \geq 0 \mid S_t = 1, X_t)$ then

$$p_{11}(S_t, X_t) = 1 - \Phi(\beta_{\mathrm{C}} + \beta_{\mathrm{X}}{}' X_{\mathrm{t}}), \tag{9}$$

which in turn yields an expression for $p_{12}(S_t, X_t) = 1 - p_{11}(S_t, X_t) = \Phi(\beta_C + \beta_X{}' X_t)$.

Given the nature of market timing it may be instructive to consider the impact the informative versus non-informative priors on the sign of the expected market premium: $\Pr(\mu_i \mid S_t)$. For example is S_t = 1, is characterised the positive equity premium state, our prior for μ_i could be such that the variable is restricted to be positive. Similarly, for S_t= 2, a negative equity premium state, the prior for $\Pr(\mu_i \mid S_t)$ could reflect that μ_i is strictly negative. In addition to diffuse

Normallly distributed priors for the state dependent expected returns, we suggest the following truncated normal distributions as the state dependent priors for the informative case:

$$\mu_1 \sim N_{tr}(0, \kappa^{-1})\mathbf{I}_{\mu_1 > 0}$$

$$\mu_2 \sim N_{tr}(0, \kappa^{-1})\mathbf{I}_{\mu_2 < 0}$$

where the variance κ^{-1} is chosen to be sufficiently large, to ensure μ_i has sufficient support. Thus the prior for μ_1 is a truncated normal with mean zero, variance κ^{-1} that is bound to be strictly positive, and the prior for μ_2 is truncated normal with mean zero, variance κ^{-1} that is bound to be strictly negative.

Having specified the form for the transition probabilities and sampling algorithm to generate the states along with the set of priors, we are able to generate inferences on the state dependent parameters by sampling from the set of full posterior densities. The full conditional densities for the parameters for static transition probability matrix case $\theta = \left\{ \mu_i, \sigma_i^2, p_{ij} \right\}_{\forall i,j}$ and for the dynamic transition probability matrix case $\theta = \left\{ \mu, \sigma^2, \beta \right\}$ are presented in the appendix.

In order to select the most appropriate specification for the data generating process, various models are compared on the basis of their Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC).

1.3. Construction of the Business Cycle Variable via PCA

We typically have a number of candidate economic signals that we could employ to determine the prevailing regime. These variables could include, interest rates and yield curve effects, consumption growth, industrial production growth, and aggregate dividend yields. Given a potentially large number signals, it would be desirable to reduce this number by capturing their common element which would hopefully capture the business cycle or economic impact on equity returns. In addition to considering the impact of individual signals, in particular the short term nominal interest rate on regime switching behaviour of market returns, the collection economic variables considered in this study are also transformed onto a new coordinate system via the process of principal components analysis. PCA effectively creates an index of information from the group of economic variables under consideration.

A new coordinate system is created such that the greatest variance by any projection of the data lies on the first coordinate (the first principal component), the second greatest variance on the second coordinate (the second principal component) and so on. It is the assumption that the lower co-ordinates contain the majority of the information of the dataset set. Each of the principal components are linear functions of the indicator variables.

PCA facilitates the estimation process through the reduction of the dimensions of the independent variables. For the purposes of this study, the first set of principal components that explain at least 65% of the variation of the group of economic variables are adopted as the transformed information (independent) variable(s).

1.4. Specification of the Market Timing Strategy

When testing whether a regime switching model can be employed to generate a profitable trading strategy, the following ex-ante strategy is employed by the agent. The agent estimates the model up to a certain point (time zero), which we define as the in-sample period. Using this period, the agent generates the parameters of the regime switching model which are used to determine the regime switching probabilities over the hold-out period. We consider two different strategies (A) and (B), both based upon the switching probabilities.

The strategies can be summarised as follows:

i) The agent has the data set that has prevailed up until time zero.

ii) They estimate the model, yielding $\hat{\beta}$ which dictates the transition probability matrix.

iii) The time varying transition probabilities gives the probability of what the next regime will be given the level of the economic variable and the regime in the prevailing period.

iv) (A) When $S_t = 1$, the positive mean return state, if $\Pr(S_{t+1} = 1 \mid S_t = 1, X_t) > 0.5$ the investor invests 100% of his wealth in equity in the following period and 0% in cash. Else if $\Pr(S_{t+1} = 1 \mid S_t = 1, X_t) < 0.5$ then the investor invests all in cash and nothing in equity.

When $S_t = 2$, the negative return state and $\Pr(S_{t+1} = 2 \mid S_t = 2, X_t) > 0.5$ then the investor invests solely in cash in the following period and nothing in stock. Else if $\Pr(S_{t+1} = 2 \mid S_t = 2, X_t) < 0.5$ the investor invests all in equity and nothing in cash.

(B) If $S_t = 1$, the positive mean return state, then the investor invests $\Pr(S_{t+1} = 1 \mid S_t = 1, X_t)$ of his wealth in equity in the following period and $1 - \Pr(S_{t+1} = 1 \mid S_t = 1, X_t)$ in cash. If $S_t = 2$, the negative mean return state, then the investor invests $\Pr(S_{t+1} = 2 \mid S_t = 2, X_t)$ of his wealth in cash in the following period and $1 - \Pr(S_{t+1} = 2 \mid S_t = 2, X_t)$ in stock.

v) At the end of every period, the agent updates their state probabilities, and hence their knowledge of the current regime and then invests according to (iv).

The performances of the strategies under regime switching are compared to the overall performance of the market. The Sharpe ratio is the adopted metric to generate the risk adjusted returns. The ratio is defined as excess return per unit of risk and has been adopted by a significant number of previous studies that compare performance based upon this benchmark.

We have prohibited the model from taking short positions. For example the analysis of Ang and Bekaert (2004) argued for the investor to short the Japanese equity market by up to 55%. While this scenario may be plausible for a select few, it is unfeasible and undesirable for the average investor.

2. Data

The model considers US monthly data from January 1959 to December 2005. The stock return series utilized is the annual rebalanced CRSP value weighted index of NYSE, AMEX and NASDAQ stocks available through WRDS.

To generate the excess return in order to derive the Sharpe ratio, the one month risk free rate is subtracted from the equity return series. The risk free rate is the Fama one month risk free average rate (average of the bid and ask rates). These variables were annualised by multiplying by a factor of 12 and were sourced from the CRSP database.

The level of the variables of short term interest rate, shape of the yield curve, inflation, growth in industrial production, dividend yield and consumption were all considered as possible predictors of equity return regimes. The affects of each of these variables on stock returns has a relevant grounding in economic and financial theory as previously discussed.

The short term interest rate was proxied by the artificial Fama three month risk free average rate. The dividend yield was derived by subtracting the return of the index excluding dividends from the return of the index including dividends.

The (shape of the) yield curve was calculated as the difference between the one month risk free rate and the 5 year yield. The one month risk free rate is the same as that used to calculate excess returns. The five year yield is calculated using Fama-Bliss artificial securities. Given the price, we can calculate the yield of the 5 year security using the following formula: $P = FV/(1+i)^5$. These variables were also sourced from the CRSP database.

The consumption variable is the seasonally adjusted percentage change in the real personal consumption expenditure of durable goods. The data was priced using the year 2000 dollars and extracted from the Federal Reserve Economic Data (FRED®) database.

Both the inflation and industrial production variables were sourced from Datastream. The inflation rate is the year on year percentage change of the USA Consumer Price Index. The industrial production variable is the percentage change of the USA Industrial Index (which uses the year 2000 as its base year). The variables are published on the 15th of each month and in order to bring them into line with the return series each must be adjusted. To derive values applicable to the beginning of each month, the average of the preceding and proceeding observations were taken. For example, to obtain the inflation rate at the 1st of March 2002, the average value of the observations at the 15th of February and the 15th of March 2002 was taken.

3. Results

The entire sample period is employed to ascertain whether the equity return series is well represented by a regime switching model using constant transition probabilities. We then test whether economic variables can predict future regimes via a dynamic transition probability matrix. The model is then calibrated over the in-sample period to obtain in-sample parameter estimates. From the perspective of the agent, we then test whether a profitable asset allocation strategy can be achieved by shifting in and out of equity on the basis of the predicted future regimes over the hold-out period.

The results are reported for the models that condition the transition probability matrix on the short term interest rate, the PCA exclusive dataset and a PCA inclusive dataset. The PCA exclusive dataset includes only the variables that are available at the time the agent predicts the proceeding regimes. These variables are the short term interest rate, shape of the yield curve, dividend yield and consumption. The PCA inclusive dataset includes all of the economic variables.

To estimate the parameters of the regime switching models, the number of iterations for the burn-in period is set to one thousand with ten thousand iterations used for the sample period.

3.1. Summary Statistics

The preliminary analysis into the distribution of the return series presented in Table 1 reveals that the property of normality for the equity return data vector is overwhelmingly rejected. Employing the Jarque-Bera and Lilliefors normality tests, the rejection of the hypothesis that the series is drawn from a normal distribution is uniform. The test statistics for each of the tests exceed the critical values when employing a 5% significance level.

[Insert Table 1 here]

A visual plot of the distribution provides similar evidence. Figure 1 plots a normal distribution that has the same moments as the equity return series over the histogram of the return series. The return series suffers from excess kurtosis, negative skewness and has heavy tails. From Table 1, the skewness and kurtosis statistics, defined as the third and fourth moments of the series, are significantly different from zero – the value expected under the normality assumption. The critical statistics are surpassed by factors of approximately 5 and 25 for the skewness and kurtosis statistics respectively.

[Insert Figure 1 here]

Consistent with Turner, Startz and Nelson (1989) we find strong evidence that equity returns exist in two distinct regimes characterised by a positive return regime with low volatility and a negative return regime with high volatility. When conducting normality tests on the entire series we reject the null hypothesis that the vector is normally distributed. However when the data vector is split into two regimes the null hypothesis that both of these regimes are normally distributed cannot be rejected. Table 2 presents the results of the normality tests for the regimes as specified by the informative and uninformative model employing static transition probabilities. Using a 5% significance level, we cannot reject the null hypothesis that the observations of regime one and regime two are normally distributed for both the uninformative and the informative model. Thus while we can reject the hypothesis that the return series is unconditionally normally distributed, strong evidence suggests that the series is conditionally normal, which is the manner in which we model the series.

[Insert Table 2 here]

3.2. Model Estimation Results

Table 3 provides the parameter estimates of the regime moments along with the Newy-West standard errors of the iterations which are used to correct for the heteroskedasticity and autocorrelation (Newey and West (1987)). The mean returns of the regimes are similar to that of the model that incorporates an informative prior on the mean of the distributions and that which adopts an uninformative prior. The dual regime model indicates that when the positive return regime prevails, the underlying volatility is much less in comparison to the negative return regime. For both the informative and the uninformative models the mean return of regime one is approximately 18% percent, while the volatility (square root of the variance) of this regime is around 39%. The mean return of regime two for the model that adopts an informative prior is -8.5% and -3.5% for the model that adopts an uninformative prior. The volatility of regime two for both the uninformative and informative model is approximately 71%. The fact that the uninformative model picks up a positive mean return and a negative mean return regime validates the adoption of an informative prior that factors in the belief that there exists two distinct regimes; one with a positive mean return and the other with a negative mean return.

[Insert Table 3 here]

The relevant mixture model is shown in Figure 2. The negative mean return distribution has a much larger variance than the positive mean return regime which exhibits a much greater level of peakedness. Thus regime one is a much more concentrated regime, with the majority of the observations drawn from values very close to the mean of the regime. Regime two is a much more disperse distribution, and although the regime picks up the very negative observations, it also tends to pick up some of the very large positive observations. This has some significant implications for the agent in the asset allocation strategy who wishes to exit the market when the model predicts a regime two period. While regime two is on average undesirable, there are periods when the agent will have preferred to have invested in equity in the regime two period – for example those periods where regime two picks up the large positive returns.

[Insert Figure 2 here]

From Table 3, the persistence of the positive regime is significant, at a level of approximately sixty-nine percent for the informative and sixty-six percent for the uninformative model. The persistence of the negative return regime is around thirty-three percent for the informative and thirty-six percent for the uninformative, suggesting that the duration of regime one is close to twice as persistent as regime two.

For the uninformative model the number of times the model allocates a positive regime to the state space is 451 in comparison to the informative model which allocates 422 periods. These numbers are significantly greater than the 346 periods that the value weighted return is greater than zero as provided in Table 1. This need not be a cause for concern. Given that the returns are characterised as a mixture of normals with variances of a non-trivial size, the two regimes will tend to have some overlapping area as seen in Figure 2. Consequently it may be possible for any one observation in the middle range of the two normals to be from either the positive or negative mean return regime. It would appear that the uninformative tends to be more optimistic in the sense that it is more likely to label a particular period as belonging to regime one as opposed to regime two. The informative model assigns 80% of the positive return periods as belonging to regime one and 33% of the negative return periods as belonging to regime two. The uninformative model on the other hand assigns 86% of the positive return periods to regime one and 30% of the negative return periods to regime two. This is by no means a detraction from the estimation process, and nor from the model. The purpose of the model is to fit a mixture model to the data series and assign each observation to one of the two regimes. It is not the purpose of the model to pick out every positive return period and label this period as belonging to regime one and pick out every negative return period and label this period as belonging to regime two.

The AIC measure as provided in Table 3 indicates that the informative model is preferable to the uninformative model, with these two superior to the single regime model. The model that characterises the return series as a single regime process picks up the return and variance essentially exactly. Though one would expect this to hold, if nothing else, the result gives one confidence that the methodology is operational and successful in its estimation.

[Insert Figure 3 here]

3.3. Full-Sample Market Timing Results

Figure 3 displays the growth of one dollar of wealth if the investor invested solely in equity, and the growth if the investor invested in equity when the model indicates the period is a regime one period and in cash when the model indicates a regime two period. If the investor invested in regimes according to the model that adopts an informative prior, the one dollar investment would have grown to over two hundred dollars. If the agent were instead to use the uninformative model the one dollar investment would have grown to nearly five hundred dollars. This compares to the growth of one dollar to one hundred and three dollars for the buy and hold equity strategy.

[Insert Table 4 here]

The results reported in Table 4 suggest that if one were able to predict the regimes as defined by the regime switching models, the potential gains are highly rewarding. Investing in equity in regime one periods and investing in cash in regime two periods would result in not only a greater raw return but also a much superior risk adjusted return. Thus if the investor was able to exit the market when the prevailing regime was regime two and invest in cash, and invest in equity when the prevailing regime was regime one, they would avoid both volatility and negative returns.

Table 4 reports that while the buy and hold strategy would have netted the investor an annualised average return of 11.07%, the return of the investment strategy based upon the regimes would have netted the investor a return of 12.03% based on the informative model and 13.86% if employing the uninformative model. The reduction of risk that the investor yields by investing in a low volatile positive mean return regime relative to the buy and hold position is significant. The average annualised volatility of the buy and hold position is 52.3%, 38.1% for the strategy that uses the uninformative model and just 35.4% for the strategy that invests in equity on the basis of the regimes according to the informative prior model. As a result, in terms risk adjusted return, the uninformative model offers a Sharpe ratio of 0.25, over twice as large than the passive buy and hold equity position of 0.11. The informative model achieves a risk adjusted return of 0.23, just shy of that achieved by the uninformative model.

For the following analysis that allows for a dynamic transition probability of the state variable, only the results for the informative prior are reported so to avoid repetition. The adoption of the informative prior allays fears of misidentification of the mean component of the regimes as the priors are asymmetric under this specification.

For the dynamic transition probability specification, recall that the transition probability matrix is augmented to factor in the effects business cycle conditioning agent, X_t, on the generation of the regime. This measure is achieved by characterising the transition matrix as a combination of probit functions, facilitated by the specification of Albert and Chib (1993b).

As described in the methodology section, PCA is used to extract business cycle effects from the collection of observed economic variables. To select the number of principal components, we use those that jointly explain at least 65% of the total variation in the group of variables. For the model that includes the entire set economic variables (the PCA inclusive dataset) this is achieved by using just the first principal component. For the model that includes all variables bar inflation and industrial production (the PCA exclusive dataset), this is achieved by employing the first and second principal components; see Table 5.

[Insert Table 5 here]

The following results are provided for the models that conditions the transition matrix on the interest rate variable, the PCA exclusive dataset and the PCA inclusive datasets.

[Insert Table 6 here]

Table 6 illustrates that consistent with the previous analysis, there exists two statistically significant regimes. These regimes can be characterised as one high volatile regime with a negative average mean return and another positive average mean return regime with a relatively smaller volatility. The mean of regime one as provided in Table 5 is approximately 18% and the mean of the regime two is around -3% for each of the models which incorporate the information of the economic variables. The volatility of regime one is around 36%. For regime two this figure is around 67%.

The parameters of the probit model indicate that being in a particular regime has an effect on whether or not that same regime prevails in the proceeding period. This can be seen through the fact that β_C is negative with a value of -0.86 and that β_R is positive with a value 1.60. Table 6 notes that the amount of information that the short term interest rate, the PCA exclusive dataset and the PCA inclusive dataset attribute to the generation of the state variable is modest. A positive (negative) coefficient reveals that the information variable has a positive (negative) effect on the persistence of regime two and a negative (positive) effect on the persistence of regime one (positive expected return/low volatility state). The interest rate tends to increase the persistence of regime two (negative expected return/high volatility state) and decrease the persistence of regime one. For the PCA exclusive dataset, the first and the second principal components have a negative effect on the persistence of regime two and a positive effect on the persistence of regime one. For the PCA inclusive dataset the first principal component has a negative affect on the persistence of regime two and a positive effect on the persistence of regime one. Whilst there is some dispersion in the posterior means of the parameters governing the transition probabilities, the conclusions provide some interesting economic insights that are consistent with previous studies; in particular; our finding with nominal interest rates are consistent with the findings of Ang and Beckaert (2004) and Fama and Schwert (1977) which suggest that when short term interest rates are low, subsequent equity returns are high. Furthermore, while not reported, inspection of Principal component loadings demonstrate that the PCA based variables are constructed with positive loadings on interest rates and inflation and negative loadings on consumption and industrial production growth. As such high levels in these variables are typically associated with downturn in economic activity. Therefore the results suggest that equity markets may be providing a hedge against consumption risk.

Both model selection criteria favour the PCA inclusive model relative to the PCA exclusive and the interest rate models. The difference between the models when using the selection criteria is minimal.

In relation to the frequency of the regimes under the different model specifications, the PCA inclusive model allocates 329 periods to the positive mean return regime period in comparison to 317 for the interest rate model and 314 for the model that conditions the transition probability matrix on the PCA exclusive dataset. One might like to consider the group of six variables (interest rate, yield curve, dividends, consumption, inflation and industrial production) as having collectively a more optimistic view of the equity return series.

The model that conditions the transition matrix on the PCA inclusive dataset allocates the same regime to a given period as does the Carter and Kohn (1994) methodology on average 62% of the time. The according values for the PCA exclusive and the interest rate models are 37% and 60% respectively. The differences in regime allocation between the dynamic transition models and the static transition model can be explained by the fact that each of the techniques fit different mixtures on the data depending upon the methodology and the information employed, with no model necessarily the ‘true’ model. However, it is reassuring that the model that introduces a fair amount of noise into the estimation technique assigns a significant portion of periods to the same regimes as the simple model which uses a smoother filter to generate the state variable.

[Insert Table 7 here]

As one can see from Table 7, on the basis of being able to predict the relevant regimes as assigned by the models, using the PCA inclusive dataset and the PCA exclusive dataset achieves slightly lower return than both the stock only strategy and the constant transition matrix model. The interest rate model outperforms the market on a raw return level. The interest rate, PCA exclusive and PCA inclusive models achieved returns of 11.09%, 9.94% and 9.53% respectively.

However, it is not just the first moment the investor cares about. The investor is ultimately interested in the tradeoff between risk and return as measured by the Sharpe ratio. Risk adjusted returns are far superior when employing the information from the economic variables, relative to the equity buy and hold and the static transition model. The Sharpe ratios of the interest rate, the PCA exclusive and the PCA inclusive models are 0.31, 0.27 and 0.22 respectively, relative to the buy and hold Sharp ratio of 0.11. The superior risk adjusted returns are achieved through a significant reduction in the volatility of the portfolio.

3.4. Out-of-Sample Market Timing Results

In the preceding analysis we have investigated the ability determine regimes and form trading strategies using the entire set of data. As econometricians this can be a useful exercise to understand the dynamics governing the market. Such analysis embeds within it a degree of look-ahead bias which can potentially distort any inferences as to the extent with which markets are efficient. Therefore in order to assess whether there is true market timing ability and hence the extent with which markets are efficient we seek to determine whether or not an agent can outperform the market on an ex-ante basis.

In order to test whether the agent can outperform the relevant benchmark, we calibrate the model using the first thirty years of data from 1959 to 1989, defining this period as the in-sample period. The in-sample period is used to obtain parameter estimates which are subsequently employed to predict regimes and thus allocate funds to the different asset classes of equity and cash in the hold-out sample which covers the period 1990 to 2005.

We assume that the agent knows which regime the current period is in but not what regime will prevail in the proceeding period. The approach is consistent with a vast number of preceding studies including that of Ang and Bekaert (2004) and Guidolin and Timmerman (2005). This need not be an excess assumption as one is still allowing for uncertainty in future regimes given all the information up to the current period.

The asset allocation of the following investment strategy is based upon the probability transition matrix, which in the context of the model, is used to predict the regime in the proceeding period, given all the information up until and including the current period. We consider the two different strategies (A) and (B), as previously defined.

[Insert Table 8 here]

Table 8 indicates that the mean and variance of the regimes as well as the dynamic transition probability matrix co-efficients are highly similar for the in-sample period estimates and the estimates based upon the entire sample period data. Explicitly, the means of regime one and regime two using the entire sample period are around 18% and -3% respectively for each of the three models. When using the in-sample period these values are approximately 16% and -3.5% for all three models. The volatility of regime one using the in-sample period is 36% (identical to the volatility of regime one for the entire sample period) and 70% in regime two (compared to 67% for the entire sample period). Given the noise of the parameter estimates the mean and variance of the regimes for the two sample periods are statistically indifferent.

Consistent with the study which analyses the entire data period, the parameters of the economic variables of the transition probabilities are of the same sign when the analysis is undertaken using just the in-sample period. The marginal affects of the economic variables on the persistence of the regimes is therefore the same. Furthermore, the parameter values are highly similar for the in-sample period and the analysis that uses the entire dataset. β_C is -0.87 for the in-sample period analysis, compared to -0.86 for the entire period analysis. The estimated β_R of 1.60 is identical for the in-sample and the full sample period analysis.

The model selection criterion tends to indicate that there is very little between the models**,** though as per the full sample analysis, the preferred model is the PCA inclusive model.

[Insert Table 9 here]

As Table 9 illustrates, the ex-ante active asset allocation strategies are able to achieve a significant amount of risk reduction through the process of withdrawing funds from equity and investing in cash when the model predicts a negative average return regime.

Comparing the active asset allocation strategies, the returns for the strategy (A) and strategy (B) for the different datasets are highly similar. The returns based upon the interest rate, the PCA exclusive dataset and the PCA inclusive dataset, are 5.64%, 6.31% and 9.31% respectively. For strategy (B), the returns are 5.91%, 6.31% and 8.02%. The best performing active asset allocation strategy is strategy (A) based upon the PCA inclusive dataset. This would have netted the agent an annualised average return of 9.31%. These figures compare to an annualised average return on the market of 11.24%.

One reason as to why the agent may underperform the market on the raw return level is that by endeavouring to invest just in regime one, they may miss out on those periods of regime two which are in fact large positive returns. While this may not be ideal they will be compensated by a large reduction in the volatility. This outcome would be more appealing the greater the risk aversion of the investor.

The volatility of the strategies that hold a portion of equity at all times, strategy (B), is 28.7% for the PCA inclusive model. For the interest rate model this is 31.6% and 28.1% for the PCA exclusive model. For the investment strategy (A), the volatility of the portfolio for the interest rate, PCA exclusive model and PCA inclusive model is 38.0%, 33.2% and 34.00% respectively. This is a large risk reduction relative to a passive equity position. The volatility of the buy and hold equity strategy is 51.90%. In other words, the active asset allocation strategy based on the switching probabilities has been successful to the point that it has been able to predict the low volatility regime and reduce the overall risk of the portfolio.

Despite the fact that the active strategies underperform the market on a raw return basis, a number of the active strategies outperform the market on a risk adjusted return level. The model which incorporates the PCA exclusive dataset to determine the transition probabilities and then invests on the basis of strategy (B), outperforms the passive buy and hold strategy.

The most outstanding performer is the PCA inclusive dataset which outperforms the passive buy and hold using risk adjusted returns for both strategy (A) and strategy (B). The annualised Sharpe ratios are 0.21 for asset allocation strategy (A) and 0.20 for asset allocation strategy (B). This compares to a Sharpe ratio of 0.14 for the market.

Though the agent outperforms the market using risk adjusted returns for the aforementioned models, we have thus far omitted explicit discussion of an important factor; transaction costs. The strategies that invest either all or nothing will tend to incur the greatest level of transaction costs, given the heavy shifts of weights across the asset classes. The strategies that make periodic shifts in asset classes but at all times hold a portion of wealth in equity will incur fewer transaction costs given that the largest shift in capital from one asset class to another is around 60%.

If the agent were to choose the PCA exclusive model and strategy (B), annual transactions would have to be essentially zero to stay on a par with the market in terms of risk adjusted return. For the PCA inclusive model that the agent would have chosen (on the basis of the preferred information criteria – AIC and BIC), annual transaction costs could erode return levels by around 1.5% annually for the strategy to achieve equality with the market's Sharpe Ratio. For the strategies based on this dataset, it may be possible to achieve a superior risk adjusted return to that of the market even when we take into account transaction costs.

A number of robustness checks have been performed. Specifically what is important to identify is whether or not the Markov Chain has converged to the invariant distribution. In order to ascertain whether this has been achieved, visual inspection of the plot of the iterations is performed every time the model is calibrated. Visual inspection of the iterations for the regime dependent variances in Figure 9 also serves as a check for the occurrence of label switching for a given parameter.

[Insert Figure 9 here]

Over the course of the analysis it has been assumed that the first and second moments of the return series are stable over time. Specifically, we assume that the in-sample parameters are the same as the parameters that are estimated using the entire sample period. This finding has been previously outlined and this in itself is suffice. As a further check the data period is split in half and calibrated over each of these two periods. The results of this robustness check allays any fears that the moments are trending or changing significantly over time (see Table 10).

[Insert Table 10 here]

4. Conclusion

This paper tests whether equity returns are well modelled by a mixture of normal distributions, whether one can use economic variables to predict these return distributions and whether a profitable trading strategy can be generated on the basis of successful prediction of the state variable.

We find strong evidence that equity returns exist in a dual regime framework and that these regimes can be characterised by one positive return low volatile regime and one negative return highly volatile regime. The results indicate that if an agent were able to identify these returns ex-ante they would enjoy much superior raw and risk adjusted returns relative to the agent who buys and holds equity.

When we sought to identify these regimes ex-ante we found that the dominant determinant of future states was today’s state; the effects of the economic information variables on the state variable were found to be marginal. However when investigating whether or not the economic signals could be employed to predict future states and generate a market trading strategy; we find that such information can be used successfully to outperform the appropriate benchmark. We find that the dynamic asset allocation strategy based upon the time-varying transition probabilities which explicitly condition on business cycle variables clearly outperforms simple market buy and hold strategies on risk adjusted terms.

These preliminary suggests that even after controlling for look-ahead bias due to potential data mining, there still appears to be information in business cycle variables to predict equity returns. These results have important links to recent studies such as Avramov and Chordia (2006) that cite the predictability of cross sectional equity market returns when conditioning on business cycle variables. Taken together these results suggest that inefficiencies in the market may be due market participants failing to incorporate the regime dependent structure of the economy. More research is required to further explore and explain these results.

Appendix: Full Conditional Densities for Various Model Parameters

Having considered in Section 1 the sampling density for the vector of parameters governing the transition probabilities in the dynamic case, in this appendix we present the sampling densities for the transition probabilities under the static case. Let $\Pi_i \equiv (p_{i1},.., p_{iK})$ represent the ith row of the transition matrix Π; the vector of state transition probabilities given $S_t = i$. By construction, these probabilities must sum to unity. The full conditional distribution for Π_i can then be expressed by Bayes rule as $\Pr(\Pi_i \mid Y, S, \theta_{-\Pi_i}) = \Pr(\Pi_i \mid S) \propto \Pr(S \mid \Pi_i)\Pr(\Pi_i)$. Given that S_t evolves according to a first order Markov process, the joint likelihood for S can be expressed as a Dirichlet process. By adopting conjugate priors the posterior denisty too will be Dirichlet, and hence parameters for Π_i can be jointly sampled from the following Dirichlet distribution:

$$\Pi_i \mid S^n \sim \text{Dir}(d_{i1}, d_{i2},\dots, d_{ik}) \qquad \text{(A1)}$$

where $d_{ij} = n_{ij} + u_{ij}$, n_{ij} represents the number of transitions from state i to state j : $n_{ij} = \sum_{t=2}^{n} \mathbf{I}_{it-1}\mathbf{I}_{jt}$, and u_{ij} are the hyperparameters of the Dirichlet prior, where $\mathbf{I}_{it}$ equals one when S_t equals i, and zero otherwise.

Turning to the posterior densities of the state dependent expected returns variances, we first consider the uninformative prior case. Given that the joint sampling conditional density or conditional likelihood for the returns Y_t is Gaussian, using uninformative conjugate priors for $\mu_i : N(\xi, \kappa^{-1})$ and $\sigma_i^2 :\ \text{IG}(\alpha_i, \beta_i)$, and applying Bayes rule, it is straightforward to construct the full conditional densities for μ_i and σ_i^2 :

$$\mu_i \mid Y, S, \sigma_i^2 \sim N(BA^{-1}, A^{-1}) \qquad \text{(A2)}$$

where $\mathrm{A} \equiv \sigma_i^{-2}\sum_{t=1}^{T}\mathbf{I}_{it} + \kappa$, $\mathrm{B} \equiv \sigma_i^{-2}\sum_{t=1}^{T}X_t\mathbf{I}_{it} + \kappa\xi$, N is a Normal Distribution; and

$$\sigma_i^2 \mid S^T, X_T, \mu_i \sim \mathrm{IG}\left(\frac{1}{2}\sum_{t=1}^{T}\mathbf{I}_{it} + \alpha_i, S\right), \tag{A3}$$

where $S = \beta_i + 0.5\sum_{t=1}^{T}\frac{(Y_t - \mu_i)^2}{T}\mathbf{I}_{it}$ and IG is an Inverse Gamma.

For the informative prior case on the expected returns, combining the truncated normal prior defined in the paper with the likelihood function of the data yields the following posteriors densities:

$$\mu_1 \mid Y, S, \sigma_1^2 \sim N_{tr}(B_1 A_1^{-1}, A_1^{-1})\mathbf{I}_{\mu_1 > 0} \tag{A4}$$

where $\mathrm{A}_1 \equiv \sigma_1^{-2}\sum_{t=1}^{T}\mathbf{I}_{1t} + \kappa$, $\mathrm{B} \equiv \sigma_1^{-2}\sum_{t=1}^{T}Y_t\mathbf{I}_{1t}$, and

$$\mu_2 \mid Y, S, \sigma_1^2 \sim N_{tr}(B_2 A_2^{-1}, A_2^{-1})\mathbf{I}_{\mu_2 > 0} \tag{A5}$$

where $\mathrm{A}_2 \equiv \sigma_2^{-2}\sum_{t=1}^{T}\mathbf{I}_{2t} + \kappa$, $\mathrm{B} \equiv \sigma_2^{-2}\sum_{t=1}^{T}Y_t\mathbf{I}_{2t}$.

Table I **Characterisation of Equity Return Series**

Table 1 reports the basic statistics for the annualised monthly equity return series covering the period 1959 to 2005. The statistics include the mean and variance of the return series, the number of months the return series yields a positive return and the number of months the series yields a negative return. The table also provides the test statistics (critical values provided in brackets) for the Jarque-Bera (JB) and Lilliefors (LL) normality tests and for the skewness and kurtosis tests. For the hypothesis tests a significance level of 5% is employed.

Mean	Variance	Pos. Periods	Neg. Periods	JB T-stat	LL T-Stat	Skewness T-Stat	Kurtosis T-Stat
0.111	0.274	346	217	109.12	0.0406	-0.468	4.968
				(5.992)	(0.0373)	(0.103)	(0.206)

Table II **Conditional Normality Test**

Table 2 reports the normality tests for the regimes for the annualised monthly equity return series covering the period 1959 to 2005. Normality tests are conducted for both regime one and regime two and for both the informative and uninformative models using the Carter and Kohn (1994) static transition matrix. The table provides the test statistics (critical values provided in brackets) for the Jarque-Bera (JB), and Lilliefors (LL) normality tests. A significance level of 5% is employed.

	Informative Model		**Uninformative Model**	
	JB T-Stat	LL T-Stat	JB T-Stat	LL T-Stat
Regime one	1.439	0.0320	0.289	0.0269
	(5.992)	(0.0431)	(5.992)	(0.0417)
Regime two	0.897	0.0455	0.987	0.0364
	(5.992)	(0.0746)	(5.992)	(0.0837)

Table III Results of Static Transition Model

Table 3 presents the results derived from the calibration of the Carter and Kohn (1994) methodology using annualised monthly equity return series covering the period 1959 to 2005. The table provides the parameters for the model that adopts a single regime, a model that assumes a dual regime with an informative prior and a model that assumes a dual regime with an uninformative prior. The table provides the mean and variance of the two regimes, regime dependent probabilities (measuring the persistence of the regimes), the Akaike Information Criterion and the frequency of the regimes in terms of the number of times the model attributes an observation to belonging to the regime. For all of the estimated parameters, the Newey-West standard errors are provided below.

Regime	Unconditional Model	Informative Model	Uninformative Model
A. Regime Dependent Mean Returns			
Regime one	0.111	0.185	0.183
	0.0171	0.039	0.037
Regime two		-0.0854	-0.0349
		0.084	0.090
B. Regime Dependent Variance			
Regime one	0.274	0.154	0.149
	0.0162	0.0267	0.0257
Regime two		0.522	0.502
		0.109	0.104
C. Regime Dependent Probabilities			
Regime one	1	0.699	0.662
		0.118	0.113
Regime two	-	0.336	0.366
		0.114	0.108
D. Model Selection Criteria			
AIC	881.60	717.98	723.75
	7.76	27.12	27.40
E. Frequency of Regimes			
Regime one	583	422	451
Regime two	-	141	112

Table IV Percentage of Variation explained by the Principal Components

Table 4 shows how much variation of the entire dataset is explained by the relevant principal components. The exclusive dataset only considers those variables that are available at the end of each month, namely: dividend yield, consumption, shape of the yield curve and the (3 month) interest rate. The inclusive dataset includes all variables: inflation, industrial production, dividend yield, consumption, shape of the yield curve and the interest rate. The dataset covers the period from 1959-2005.

	Exclusive Data Set	Inclusive Data Set
1^{st} PC	43.00%	69.22%
2^{nd} PC	36.77%	14.90%
3^{rd} PC	15.42%	8.78%
4^{th} PC	4.82%	3.73%
5^{th} PC		2.28%
6^{th} PC		1.09%

Table V Static Transition Model Performance 1959-2005

Table 5 presents the mean return, volatility and Sharpe ratio of the various asset allocation strategies for models calibrated using the annualised monthly equity return series covering the period 1959 to 2005. Specifically, these are the simple buy and hold, the model that adopts the Carter and Kohn (1994) methodology and assumes an uninformative prior and the model that assumes an informative prior. The agent invests in equity if the model brands a given period as belonging to regime one and in cash if the model brands the observation as a regime two observation.

Measure	Stock	Informative Prior	Uninformative Prior
Mean Return	11.07%	12.03%	13.86%
Standard Deviation	52.30%	35.40%	38.10%
Sharpe Ratio	0.111	0.230	0.253

Table VI Results of Dynamic Transition Model

Table 6 presents the results derived from the calibration of the model that augments the Albert and Chib (1993b) probit methodology with the Carter and Kohn (1994) methodology using the non-recursive filter to generate the state variable. The input data includes the annualised monthly equity return series as well as inflation, industrial production, dividend yield, consumption, shape of the yield curve and the (3 month) interest rate, covering the period from 1959 to 2005. For the various models, the transition probability matrix depends upon the interest rate, the PCA exclusive dataset (PCA that only considers the variables of dividend yield, consumption, shape of the yield curve and the interest rate) and the PCA inclusive dataset (all six economic variables). The table provides the mean and variance of the two regimes, the parameters that dictate the transition probability matrix, the Akaike and Bayesian Information Criterion, and the frequency of the regimes in terms of the number of times the model attributes an observation to belonging to a particular regime. For all of the estimated parameters, the Newey-West standard errors are provided below.

Regime	Interest Rate	PCA - Exclusive	PCA - Inclusive
A. Regime Dependent Mean Returns			
Regime one	0.183	0.183	0.184
	0.0645	0.0579	0.0622
Regime two	-0.0290	-0.0299	-0.0299
	0.0516	0.0522	0.0516
B. Regime Dependent Variance			
Regime one	0.127	0.126	0.126
	0.0389	0.0381	0.0394
Regime two	0.449	0.443	0.450
	0.0804	0.0776	0.0809
C. Probability Transition Matrix			
β_c	-0.858	-0.858	-0.856
	0.191	0.135	0.135
β_{z1}	0.00683	-0.00362	-000711
	2.525	1.885	0.418
β_{z2}		-0.0115	
		2.012	
β_R	1.603	1.603	1.603
	0.0567	0.0500	0.0457
D. Model Selection Criteria			
AIC	717.73	718.49	717.14
	44.32	36.21	38.03
BIC	730.40	743.82	729.81
	44.33	36.21	38.03
E. Frequency of Regimes			
Regime one	317	314	329
Regime two	246	249	234

Table VII Dynamic Transition Model Performance 1959-2005

Table 7 presents the excess return, volatility and Sharpe ratio of the various asset allocation strategies for models calibrated using the annualised monthly equity return series covering the period from 1959 to 2005. Specifically, these models include the simple buy and hold, the model that adopts a time varying transition probability matrix that depends upon the interest rate, the PCA exclusive dataset and the PCA inclusive dataset. The methodology employs a non-recursive filter and adopts an informative prior. The agent invests in equity if the model brands a given period as belonging to regime one and in cash if the model brands the observation as a regime two observation.

Measure	Stock	Interest Rate	PCA - exclusive	PCA - inclusive
Mean Excess	11.07%	11.09%	9.94%	9.53%
Volatility	52.30%	27.09%	25.98%	28.72%
Sharpe Ratio	0.111	0.308	0.274	0.222

Table VIII In-Sample Period Results for Dynamic Transition Model

Table 8 presents the results derived from the calibration of the model that augments the Albert and Chib (1993b) probit methodology with the Carter and Kohn (1994) methodology using the non-recursive filter to generate the state variable. The input data includes the annualised monthly equity return series as well as inflation, industrial production, dividend yield, consumption, shape of the yield curve and the (3 month) interest rate, covering the period from 1959 to 1989. For the various models, the transition probability matrix depends upon the interest rate, the PCA exclusive dataset (PCA that only considers the variables of dividend yield, consumption, shape of the yield curve and the interest rate) and the PCA inclusive dataset (all six economic variables). The table provides the mean and variance of the two regimes, the parameters that dictate the transition probability matrix, the Akaike and Bayesian Information Criterion and the frequency of the regimes in terms of the number of times the model attributes an observation to belonging to the regime. For all of the estimated parameters, the Newey-West standard errors are provided below.

Regime	Interest Rate	PCA - exclusive	PCA - inclusive
A. Regime Dependent Mean Returns			
Regime one	0.161	0.161	0.163
	0.0845	0.0715	0.0637
Regime two	-0.0358	-0.0351	-0.0350
	0.0675	0.0538	0.0640
B. Regime Dependent Variance			
Regime one	0.131	0.131	0.129
	0.0481	0.0488	0.0481
Regime two	0.484	0.485	0.485
	0.154	0.143	0.157
C. Probability Transition Matrix			
β_c	-0.872	-0.871	-0.867
	0.215	0.188	0.217
β_{z1}	0.0516	-0.00216	-0.00730
	3.215	2.719	0.515
β_{z2}		-0.0419	
		2.822	
β_R	1.605	1.606	1.605
	0.0707	0.0779	0.0774
D. Model Selection Criteria			
AIC	490.41	490.42	487.88
	34.02	30.48	32.51
BIC	502.26	514.11	499.73
	34.02	30.48	32.51
E. Frequency of Regimes			
Regime one	216	228	212
Regime two	157	145	161

Table IX Hold-Out Period Model Performance

Table 9 presents the mean return, volatility and Sharpe ratio of the various asset allocation strategies for the dynamic transition models whose parameters were generated using the in-sample period. We test whether the agent can beat the market using these parameters to predict the regime that prevails in future periods over the hold-out period 1990-2005. The agent adopts the strategies of (A) and (B) as defined in Section 1. The agent invests the larger portion of their wealth in equity if the model predicts the future as a regime one period and in cash if the model predicts a regime two period.

Measure	Stock	Interest Rate (A)	PCA - Exc (A)	PCA - Inc (A)	Interest Rate (B)	PCA - Exc (B)	PCA - Inc (B)
Return	11.24%	5.64%	6.31%	9.31%	5.91%	6.31%	8.02%
Volatility	50.89%	38.02%	33.21%	33.98	31.56%	28.14%	28.70%
Sharpe	0.144	0.0856	0.123	0.206	0.125	0.144	0.200

Table X Full Sample and Sub Sample Conditional Normality Tests

Table 10 reports the normality tests for the regimes for the annualised monthly equity return series covering the period 1959 to 2005. Normality tests are conducted for both regime one and regime two for the informative model which employ a dynamic transition matrix. The table provides the test statistics (critical values provided in brackets) for the Jarque-Bera (JB) and Lilliefors (LL) normality tests. For the hypothesis tests a significance level of 5% is employed.

Panel A: Full Sample 1959 to 2005

	Interest Rate		PCA Exclusive		PCA Inclusive	
	JB T-stat	LL T-stat	JB T-stat	LL T-stat	JB T-stat	LL T-stat
Regime One	1.307 (5.992)	0.0282 (0.0498)	1.0152 (5.992)	0.0347 (0.0500)	6.651 (5.992)	0.0269 (0.0488)
Regime Two	3.727 (5.992)	0.0393 (0.0565)	3.257 (5.992)	0.0497 (0.0561)	4.991 (5.992)	0.0337 (0.0579)

Panel B: Sub-Sample 1959 to 1989

	Return Series		Interest Rate		PCA Exclusive		PCA Inclusive	
	JB Tstat	LL Tstat	JB Tstat	LL Tstat	JB Tstat	LL Tstat	JB Tstat	LL Tstat
Regime One	0.0511 (0.0459)	96.06 (5.992)	0.7860 (5.992)	0.0423 (0.0603)	0.791 (5.992)	0.0378 (0.0587)	0.132 (5.992)	0.0254 (0.0609)
Regime Two			2.515 (5.992)	0.0492 (0.0707)	3.616 (5.992)	0.0489 (0.0736)	11.14 (5.992)	0.0482 (0.0698)

Figure I **Histogram of Equity Mean Return Series**

Figure 1 plots a normal distribution with the same moments as the return series, over the histogram of the equity return series. The dataset covers the period from 1959-2005.

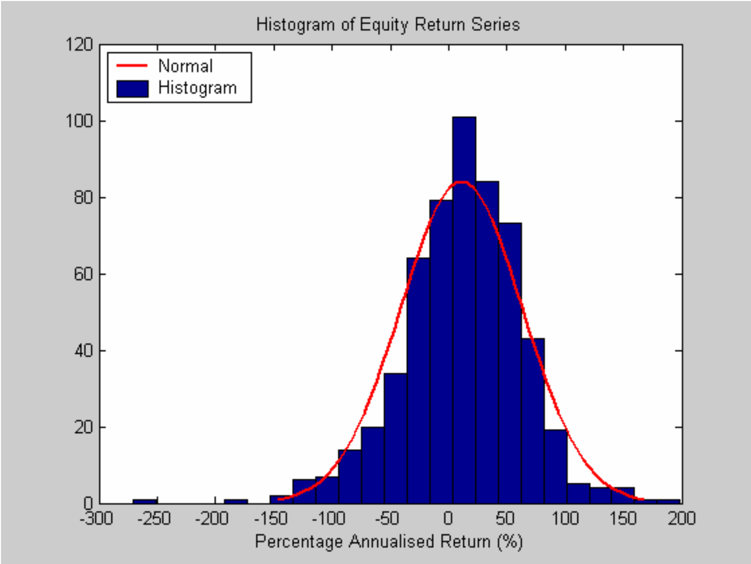

Figure II **Plot of Posterior Densities for the Regime Dependent Expected Excess Returns**

Figure 2 plots the distributions that the Carter and Kohn (1994) informative model fits to the annualised return series data. The distribution in blue is the posterior density for the expected return for regime 1, while the distribution in red is the posterior density for the expected returns in regime 2. Regime 1 is characterised as having posterior mean for expected returns that is positive, while Regime 2 is characterised as having a posterior mean for expected returns that is negative. The dataset covers the period from 1959-2005.

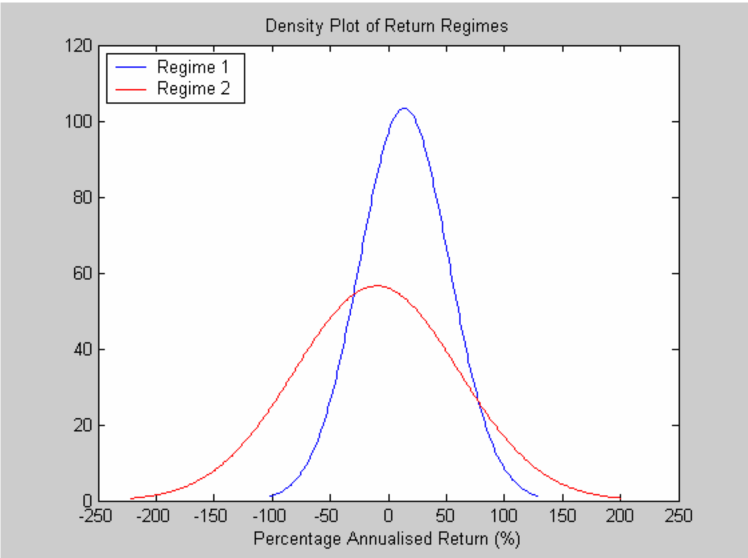

Figure III Plot of Growth of $1.00 of Wealth over the Period 1959-2005

Figure 3 plots the growth of $1 of wealth over the period of 1959-2005. The different scenarios include a stock buy and hold, and a strategy based upon buying stock if the model indicates a regime one period and holding cash if the model indicates a regime two period. The regimes are defined as per the output of the Carter and Kohn (1994) methodology, for the models that adopt both an informative and an uninformative prior.

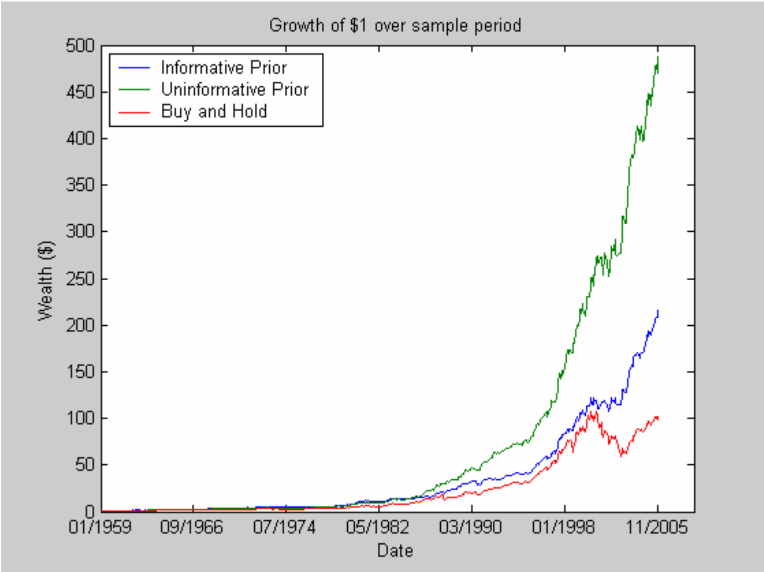

Figure IV Iterates of Variance for the Dynamic Transition Model

Figure 4 plots the iterates from the MCMC simulation for the posterior densities of the variance for the two regimes using the dynamic transition probability specification with informative priors. The iterations show rapid mixing to the posterior distribution, using a warm-up period of 1000 iterations and a sampling period of 10000. We note that the two regimes produce distinct (non-overlapping) posterior densities for the variance under each regime.

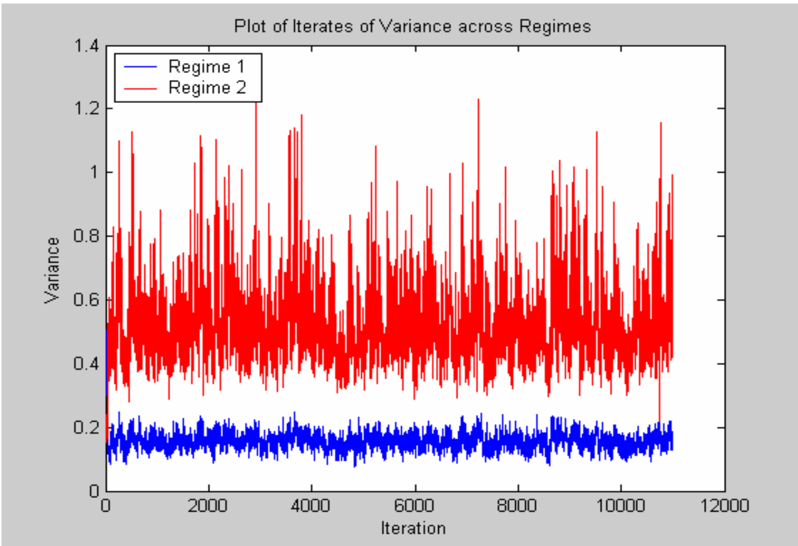

References

Albert, J. and S. Chib (1993a), 'Bayes Inference via Gibbs sampling of Autoregressive Time Series Subject to Markov Mean and Variance Shifts', *Journal of Business and Economic Statistics*, 11(1), 1-15.

Albert, J. and S. Chib (1993b), 'Bayesian Analysis of Binary and Polychotomous Response Data', *Journal of the American Statistical Association*, 88, 669-679.

Ang, A., and G. Bekaert (2004), 'How Regimes Affect Asset Allocation', *Financial Analysts Journal*

Avramov, D. and T. Chordia (2006), 'Asset Pricing Models and Financial Market Anomalies,' *Review of Financial Studies*, 19, 1001-1040.

Bhar, R., and S. Hamori (2004), 'Link between Inflation and Inflation Uncertainty: Evidence from G7 Countries', *Empirical Economics*, 29 (4), 825-853.

Cai, J. (1994), 'A Markov model of switching-regime ARCH, *Journal of Business and Economic Statistics*, 12, 309-316.

Carter, C. K., and R. Kohn (1994), 'On Gibbs Sampling for State Space Models', *Biometrika*, 81, 541-553.

Casella, G., and E. I. George (1992), 'Explaining the Gibbs Sampler', *The American Statistician*, 46, 167-190.

Chen, N., R. Roll and S. A. Ross (1986), 'Economic Forces and the Stock Market', *The Journal of Business*, 59(3), 383-403.

Chib, S. (1995), 'Marginal Likelihood from the Gibbs Output', *Journal of the American Statistical Association*, 95, 957-70.

Chib, S., and E. Greenberg (1996), 'Markov Chain Monte Carlo Simulations in Econometrics', *Econometric Theory*, 12, 409-431.

Evans, M. and P. Wachtel (1993), 'Inflation Regimes and the Sources of Inflation Uncertainty', *Journal of Money, Credit and Banking*, 25(3), 475-511.

Fama, Eugene F. and Kenneth R. French, 1988, Dividend yields and expected stock returns, *Journal of Financial Economics*, 22, 3-25.

Fama, E., and W. Schwert (1977), 'Asset returns and inflation', *Journal of Financial Economics*, 5, 115-146.

Ferson, W. E. (1989), 'Changes in Expected Security Returns, Risk, and the Level of Interest Rates', *Journal of Finance*, 44(5), 1191-1217.

Filardo, A. J. (1994), 'Business Cycle Phases and their Transition Dynamics', *Journal of Business and Economic Statistics*, 12, 299-308.

Filardo, A. J. and S. F. Gordon (1998), 'Business Cycle Durations', *Journal of Econometrics*, 85, 99-123.

Folger, H. R., K. John and J. Tipton (1981), ‘Three Factors, Interest Rate Differentials and Stock Groups’, *Journal of Finance*, 36(2), 323-335.

Garcia, R., and P. Perron (1996), ‘An analysis of the real interest rate under regime shifts’, *Review of Economics and Statistics*, 78, 111-125.

Grauer, F. L. A., R. H. Litzenburger and R. E. Stehle (1976), ‘Sharing Rules and Equilibrium in an International Capital Market Under Uncertainty’, *Journal of Financial Economic*, 3, 236-256.

Gray, S. F. (1996), ‘Modelling the Conditional Distribution of Interest Rates as a Regime-Switching Process’, *Journal of Financial Economics*, 42, 27-62.

Guidolin, M., and A. Timmermann (2005), ‘Economic Implications of Bull and Bear Regimes in UK Stock and Bond Returns’, *The Economic Journal*, 115, 111-143.

Hamilton, J. D. (1988), ‘Rational Expectations Econometric Analysis of Changes in the Regime: An Investigation into the Term Structure of Interest Rates’, *Journal of Economic Dynamics and Control*, 12, 385-423

Hamilton, J. D. (1989), ‘A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle’, *Econometrica*, 57, 357-384.

Hamilton, J. D. and R. Susmel (1994), ‘Autoregressive Conditional Heteroskedasticity and Changes in Regime’, *Journal of Econometrics*, 64, 307-333.

Lettau, L., and S. Ludvigson (2001), ‘Consumption, Aggregate Wealth, and Expected Stock Returns’, *Journal of Finance*, 56(3), 815-849.

Marcellino, M., J. H. Stock and M. W. Watson (2003), ‘Macroeconomic forecasting in the Euro area: Country specific versus area-wide information’, *European Economic Review*, 47(1), 1-18.

Newey, W. & K. West (1987), ‘A simple positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix’, *Econometrica*, 55, 703-708.

Peiro, A. 1994, ‘The Distribution of Stock Returns: International Evidence’, *Applied Financial Economics*, 4(6), 431-439.

Richardson, M. and T. Smith (1993), ‘A Test For Multivariate Normality in Stock Returns’, *Journal of Business*, 66(2), 295-321.

Stock, J. H. and M. W. Watson (2002), ‘Forecasting using principal components from a large number of predictors’, Journal of the American Statistical Association, 97, 1167-1179.

Tanner, M. A. (1996), *Tools for Statistical Inference*, 3rd edn. New York: Springer.

Turner, C. M., R. Startz and C. R. Nelson (1989), ‘A Markov Model of Heteroskedasticity, Risk and Learning in the Stock Market’, *Journal of Financial Economics*, 25, 3-22.

Whitelaw, R. F. (2000), ‘Stock Market Risk and Return: An Equilibrium Approach’, *Review of Financial Studies*, 13(3), 521-547.

Zhou, G. (1993), ‘Asset Pricing Tests Under Alternative Distributions’, *Journal of Finance*, 48(5), 1927-1942.