Accepted Manuscript

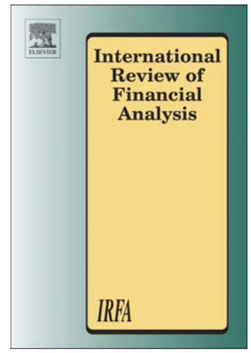

Equity Premium Estimates from Economic Fundamentals under Structural Breaks

Simon C. Smith

PII: S1057-5219(17)30053-4
DOI: doi: 10.1016/j.irfa.2017.04.011
Reference: FINANA 1102

To appear in: *International Review of Financial Analysis*

Received date: 5 September 2016
Revised date: 27 April 2017
Accepted date: 28 April 2017

Equity Premium Estimates from Economic Fundamentals under Structural Breaks

Simon C. Smith[a,*]

[a] *Lancaster University, Lancaster, LA1 4YX, UK*

Abstract

This article compares three estimates of the conditional equity premium using dividend and earnings growth rates to measure the expected rate of capital gain. The premia are estimated using a theory-informed Bayesian model that admits structural breaks. The equity premium fell from 8.16% in 1951 to 1.15% in 1985. Approximately half of this decline was reversion of a high conditional premium to the long run mean and the remainder resulted from a decline in the expected stock return. The decline in the expected stock return was largely driven by the Fed Accord (1951) and the Fed's 'monetarist policy experiment' (1979-1982).

Keywords: Equity premium, Structural Break, Bayesian analysis

JEL classifications: G10, C11, C15

*Corresponding author: Simon C. Smith. The comments of Michael Brennan, George Bulkley, David Leslie, Mark Shackleton, Evarist Stoja, and Ian Tonks have been helpful. Brian Lucey (the editor) and an anonymous referee get special thanks. All errors are my own.

Email address: `s.smith3@lancaster.ac.uk` (Simon C. Smith)

1. Introduction

The equity premium is one of the most important quantities in finance playing a central role in portfolio allocation decisions, evaluation of fund manager performance, and Social Security investment decisions. Estimates of the equity premium vary considerably depending upon the sample and data set under consideration. Estimates of the premium from realised returns on a market portfolio in excess of the risk-free rate however are almost universally found to be in the region of 7 percent, a value far too high to reconcile with the observed volatility of consumption (Mehra and Prescott 1985).

This article compares estimates of the conditional equity premium from 1871 to 2013 using dividend and earnings growth rates to measure the expected rate of capital gain (Fama and French 2002). These estimates of the path of the equity premium are compared with the estimate from realised excess returns. Estimating the path of the premium from these three series will make several contributions to the literature. First, direct evidence will be provided to determine which estimated path is likely to be closest to the path of the true equity premium. Second, the estimates allow for permanent shifts in the premium known as structural breaks therefore the timing of breaks will provide valuable information about which key economic events are driving the premium. Third, estimating the magnitude by which the premium shifted between two given dates enables us to try to disentangle the exact amount of the shift that is attributable to different factors such as, for instance, mean reversion from a high or low conditional premium to its unconditional mean or a permanent shift in the expected stock return.

The approach of Fama and French (2002) is straightforward. Assuming the stock price is cointegrated with dividends and earnings the growth rates of dividends and earnings measure the expected rate of capital gain given a sufficiently long sample. They focus on the unconditional equity premium and find that while the estimates from fundamentals and realised excess returns are similar for the period 1871-1950 they diverge thereafter. Fama and French (2002) (see also Claus and Thomas 2001) report equity premium estimates from fundamentals far below the average stock return for the period 1951 to 2000 (2.55% and 4.32% compared with 7.43%) and provide direct evidence that the estimates from fundamentals are closer to the true expected return. Pastor and Stambaugh (2001) estimate the path of the equity premium using realised excess returns and Blanchard et al. (1993) estimate the path of the expected stock return, but neither article compares estimates of the path of the premium.

Since Fama and French (2002) focus on the unconditional equity premium they assume that the dividend- and earnings-price ratios are stationary throughout the sample. Such an assumption seems strong, however, especially given their long sample from 1871 to 2000. Fama and French (2002) note that it is reasonable to expect the two price ratios to be subject to structural breaks as a result of changing government policy or technological innovation, for example. Given the increasing evidence that the equity premium is subject to breaks (see e.g. Pastor and Stambaugh 2001; Kim, Morley and Nelson 2005; and Rapach, Strauss and Zhou 2010) it is important to allow for breaks when estimating the path of the equity premium over such a long period. Fama and French (2002) note that their approach is still valid in the presence of breaks as long as the dividend- and earnings-price ratios mean-revert within regimes which is weaker than their assumption that the price-ratios are stationary throughout the entire sample. If the break dates can be estimated their approach can be used to estimate the premium between breaks thereby giving the path of the premium over the entire sample. It is therefore necessary to formally model the underlying breakpoint process to enable estimation of the path of the equity premium from the three return series.

One attractive model of the equity premium that allows for multiple structural breaks is introduced by Pastor and Stambaugh (2001). This Bayesian model allows for multiple breaks occurring at unknown times enabling the uncertainty surrounding the timing of breaks to be incorporated into the estimates of the equity premium. The more popular frequentist approach, however, estimates the timing of the breaks and subsequently estimates the parameters assuming these estimates are the true locations.[1] In finite samples such an approach can compromise inference because it ignores estimation risk. The model is motivated by economic theory through the specification of the prior distributions. First, it is assumed a priori that the equity premium is unlikely to be subject to breaks of a large magnitude. Second, the positive relation between the premium and volatility within each regime, the 'price of risk', is specified as a flexible prior distribution and the strength of this relation is adjusted to evaluate the effect on inferences. Third, a hierarchical prior on the equity premium allows the premium in a given regime to be estimated as a combination of the data observed in that regime and the value of the estimated premium in previous regimes. Finally, the equity premium is assumed to be positive since risk-averse investors must be compensated for being exposed to greater risk (Merton 1980).

One limitation of this attractive model is that the number of breaks must be fixed in advance due to Pastor and Stambaugh (2001)'s use of Chib (1998)'s algorithm to estimate the breaks.[2] This article uses the reversible jump Markov chain Monte Carlo (RJMCMC) algorithm of Green (1995) (see also Bulkley, Leslie and Smith 2015) to determine the number of breaks endogenously. This approach enables the uncertainty surrounding both the number and timing of breaks to be incorporated into estimates of the equity premium and thereby enhances Pastor and Stambaugh (2001)'s model even further.

[1] A handful of frequentist changepoint studies include Bai (1997), Liu, Wu and Zidek (1997), Bai and Perron (1998), Bai and Perron (2003), and Chen and Hong (2012).

[2] Other Bayesian changepoint studies include Carlin, Gelfand and Smith (1992), Inclan (1993), Kim and Nelson (1999), Pesaran, Pettenuzzo and Timmermann (2006), Koop and Potter (2007), Maheu and Gordon (2008), Giordani and Kohn (2008), and Song (2014).

The equity premium estimates from realised returns and fundamentals follow a similar path until 1950 but diverge thereafter. The estimates from fundamentals fall from 8.16% in October 1951 to 1.15% in March 1985 while the estimates from realised returns fluctuate between 6% and 7% during this period. Relative to the findings of Fama and French (2002) this article provides two key insights. First, our methodology can disentangle how much of the decline was anticipated and how much was not. Approximately half (3.68%) of the 7.01% decline in the conditional premium from 1951 to 1985 was anticipated and reflects reversion of a high conditional premium to the long run mean of 4.48%. The remaining 3.33% decline was unanticipated and, having explored a range of plausible causes, we conclude this was likely the result of a decline in the expected stock return.

The second key insight our methodology provides is the exact dates at which this decline occurred. Specifically we find that the decline in the expected stock return was driven largely by two breaks that occurred in 1951 and 1979 corresponding to the Treasury Fed Accord of 1951 and the Fed's 'monetarist policy experiment' between 1979 and 1982. This suggests that monetary policy was the main driver of the expected stock return over the past century. Monetary policy is an intuitive candidate for shifts in the expected stock return. For example, the tightening of the money supply in 1979 to combat the high inflation experienced in the 1970s is likely to have reduced future economic growth prospects and this may have caused agents to revise downwards their expectations of stock returns.

Three pieces of direct evidence from Sharpe ratios, book-to-market ratios, income return on investment, and the precision of the estimates suggest that the estimates from fundamentals for the period 1951 to 2013 are likely to be closer to the true equity premium than is the realised return estimate. Specifically the standard error of the estimates from fundamentals is lower than that from realised returns suggesting the estimates are more precise. Second, the Sharpe ratio calculated using fundamentals falls from the first subsample, 1871-1950, to the second, 1951-2013. Sharpe ratios are related to aggregate risk aversion and since the first subsample contains by far the most volatile period of the entire sample it is unsurprising that the Sharpe ratio falls from the first subsample to the second. The Sharpe ratio calculated using realised returns however increases in the second subsample which seems implausible. Third, and most important, the book-to-market ratio for the period 1961-2013 averages 0.79 and is only greater than 1 for five years. A book-to-market ratio of less than 1 suggests the average corporate investment during this period was profitable. Therefore one would expect the average return on investment for the period to exceed the cost of equity capital (characterised by the expected stock return.) This relation holds for the equity premium estimates from fundamentals. The estimate from realised returns however is greater than the average return on investment suggesting the average corporate investment over the period 1961 to 2013 was not profitable. This seems difficult to reconcile with an average book-to-market ratio of less than 1.

The remainder of the paper is structured as follows. Section 2 sets out the methodology. Section

3 presents the empirical analysis. Section 4 concludes.

2. Methodology

2.1. Fama and French's Approach

The rate of return on equities comprises the dividend yield and the rate of capital gain

$$R_t = D_t/P_{t-1} + GP_t, \tag{1}$$

in which D_t denotes the dividend for year t, P_{t-1} denotes the price at the end of year $t-1$, and $GP_t = (P_t - P_{t-1})/P_{t-1}$ denotes the rate of capital gain.

Fama and French (2002) note that if the dividend-price ratio, D_t/P_t, is stationary then given a long sample period the compound dividend growth rate will approach the compound rate of capital gain. An alternative expected return estimate is therefore

$$RD_t = D_t/P_{t-1} + GD_t, \tag{2}$$

in which $GD_t = (D_t - D_{t-1})/D_{t-1}$ denotes the dividend growth rate. Let us refer to (2) as the dividend growth model.

The reasoning that applies to (2) can be applied to any variable that is cointegrated with the stock price. Assuming the earnings-price ratio, Y_t/P_t, is stationary then given a long sample period the compound earnings growth rate will approach the compound rate of capital gain. Another expected return estimate is therefore

$$RY_t = D_t/P_{t-1} + GY_t, \tag{3}$$

in which $GY_t = (Y_t - Y_{t-1})/Y_{t-1}$ is the earnings growth rate. Let us refer to (3) as the earnings growth model.

Fama and French (2002) are interested in the unconditional expected return and therefore assume the dividend- and earnings-price ratio are stationary throughout the sample, that is the price-ratios have constant unconditional means. It is unreasonable to assume the two price ratios are random-walk processes that can drift off to infinity but occasional permanent shifts (or structural breaks) in the ratios is reasonable.

Fama and French (2002) note that their approach remains valid in the presence of structural breaks as long as the dividend- and earnings-price ratios mean-revert within regimes, that is, the price ratios have constant means within regimes. Their approach in this setting however would estimate the overall average of the regime-specific expected returns. The aim of this article is to estimate the regime-specific expected returns themselves. We therefore assume that the price ratios mean-revert within regimes. In this case equations (1), (2), and (3) hold on average within regimes, and our method estimates the regime-specific expected returns.

It is therefore necessary to formally model the breakpoint process when estimating the equity premium using the realised excess return series and the dividend and earnings growth models. Pastor and Stambaugh (2001) introduce an attractive economically motivated model of the equity premium that allows for multiple stochastic structural breaks.

2.2. Model

The model of Pastor and Stambaugh (2001) is now introduced. Their approach is Bayesian enabling economic theory to motivate the prior distributions, while allowing the uncertainty surrounding the timing of breaks to be incorporated into the equity premium estimates. Let θ, x, K, and q denote the parameter vector, data, number of changepoints, and a vector of changepoint locations. A frequentist approach involves a two-stage procedure. First, the changepoint locations $\hat{q}$ are estimated and subsequently the parameters are estimated assuming that the estimated changepoint locations are the true locations. Such an assumption ignores estimation risk and could therefore compromise parameter estimates in finite samples. Pastor and Stambaugh (2001) point out that a Bayesian approach approximates the posterior distribution of the changepoint locations $p(q \mid x)$ enabling the changepoint locations q to be marginalised when estimating the parameters

$$p(\theta \mid x, K) = \int_q p(\theta \mid x, q, K) p(q \mid x, K) dq. \tag{4}$$

Pastor and Stambaugh (2001) use Chib (1998)'s algorithm to estimate the breaks. This procedure however requires the number of breaks to be pre-specified by the user. The first modification relative to Pastor and Stambaugh (2001) is to estimate the breaks using a different algorithm to allow the number of breaks to be determined endogenously. The algorithm of Chib (1998) is replaced with the RJMCMC approach of Green (1995) as proposed by Bulkley, Leslie and Smith (2015). The methodology applied in this paper is fully developed by Bulkley, Leslie and Smith (2015). In addition to approximating the posterior distribution of the break locations this algorithm also approximates the posterior distribution of the number of changepoints $p(K \mid x)$ enhancing the model of Pastor and Stambaugh (2001) even further by allowing the uncertainty surrounding the number of changepoints to also be incorporated into parameter estimates

$$p(\theta \mid x) = \int_K \int_q p(\theta \mid x, q, K) p(q \mid x, K) p(K \mid x) dq dK. \tag{5}$$

The second modification relative to Pastor and Stambaugh (2001) is the removal of the transition regimes (TRs) that are short periods during which the distribution of excess returns are gradually changing. Specifically prices are likely to contemporaneously shift in the opposite direction to shifts in the premium during a TR. The TRs are removed because they arise with realised returns but not with expected returns that are estimated by the dividend and earnings growth models. For direct comparison it is preferable to use the same model to estimate the three equity premia paths and therefore a framework which removes the TRs and assumes breaks occur instantly is implemented. It is assumed that the dividend- and earnings-price ratios mean-revert within regimes which is a milder assumption than the one made by Fama and French (2002).

Since each break occurs abruptly a break is characterised by a single changepoint. The period between two successive changepoints is referred to as a regime. Within each regime the equity premium estimate is informed by the corresponding volatility estimate through the 'price of risk', the positive relation between the premium and volatility. The strength of this relation is specified through a prior distribution while the estimated price of risk in previous regimes informs the current estimate. Since large shifts in the equity premium are assumed to be unlikely a prior distribution is specified placing most of the prior mass on relatively small shifts in the level of the premium. Risk-averse investors must be compensated for being exposed to greater risk and therefore the equity premium is truncated at zero (Merton 1980).

The stochastic framework and prior distributions of the model are now formally described.

2.2.1. Stochastic Framework

Let x_t denote the excess return at time t and $x = (x_1, \ldots, x_T)$ denote the vector of T observations of excess returns in the sample. Given K structural breaks the sample is divided into $K+1$ regimes. A changepoint vector denoted $q = (q_1, \ldots, q_K)$ is characterised by the final time point in each regime.[3] The final regime continues through the end of the sample. The first time point in the ith regime is $q_{i-1}+1$ and the final time point is q_i. The duration of regime i is denoted $l_i = q_i - q_{i-1}$, and the time span of the ith regime is

$$i = q_{i-1} + 1, \ldots, q_i, \qquad i = 1, \ldots, K+1. \tag{6}$$

Within each regime the excess returns are normally distributed with mean μ_i and variance σ_i^2

$$x_t \sim \mathrm{N}(\mu_i, \sigma_i^2), \qquad t \in i, \qquad i = 1, \ldots, K+1. \tag{7}$$

[3] Assume $q_0 = 0$ for convenience.

Let $\mu = (\mu_1, \ldots, \mu_{K+1})$ and $\sigma = (\sigma_1, \ldots, \sigma_{K+1})$ denote the equity premiums and volatilities in the $K+1$ regimes. The likelihood function is

$$p(x \mid \mu, \sigma, q, K) \propto \left(\prod_{i=1}^{K+1} \frac{1}{\sigma_i^{l_i}} \right) \exp \left\{ -\frac{1}{2} \sum_{i=1}^{K+1} \sum_{t=q_{i-1}+1}^{q_i} \frac{(x_t - \mu_i)^2}{\sigma_i^2} \right\}. \tag{8}$$

2.2.2. Prior Beliefs

A Bayesian framework enables economic theory to inform the model through the specification of prior distributions. Prior information is combined with information in the data transmitted through the likelihood function and inference is performed on the resulting posterior distribution. The prior specifications follow Pastor and Stambaugh (2001) and therefore the choices are not discussed but simply noted for completeness.

2.2.3. Beliefs About Regime Durations

The duration of the ith regime follows a geometric distribution given the corresponding nontransition probability $p_{i,i}$

$$p(l_i \mid p_{i,i}) = p_{i,i}^{l_i - 1}(1 - p_{i,i}), \qquad i = 1, \ldots, K. \tag{9}$$

The probability of observing a length l_{K+1} between the final changepoint q_K and the end of the sample T is

$$p(l_{K+1} \mid p_{K+1,K+1}) = p_{i,i}^{l_i}. \tag{10}$$

A beta distribution is placed over $p_{i,i}$ for $i = 1, \ldots, K+1$

$$p(p_{i,i}) = \frac{\Gamma(a_i + c_i)}{\Gamma(a_i)\Gamma(c_i)} p_{i,i}^{a_i - 1}(1 - p_{i,i})^{c_i - 1}, \tag{11}$$

in which a_i=224 and c_i=3 following Pastor and Stambaugh (2001) to give a prior expected regime duration of approximately ten years. The unconditional prior mean duration of the ith regime is

$$E(l_i) = \frac{a_i + c_i - 1}{c_i - 1}, \tag{12}$$

giving a prior expected regime duration of 113 months.

Since a dimension-switching RJMCMC algorithm is used rather than the Hierarchical Hidden Markov model of Chib (1998) the $p_{i,i}$s can be marginalised for simplicity giving a marginal prior on the regime durations, $p(l_i)$, for $i = 1, \ldots, K$ of

$$p(l_i) = \int_0^1 p(l_i \mid p_{i,i}) p(p_{i,i}) dp_{i,i},$$

$$p(l_i) = \frac{c_i \Gamma(a_i + c_i) \Gamma(l_i + a_i - 1)}{\Gamma(a_i) \Gamma(a_i + l_i + c_i)}.$$

The marginal prior on the final regime length, l_{K+1} is

$$p(l_{K+1}) = \frac{\Gamma(a_{K+1} + c_{K+1}) \Gamma(a_{K+1} + l_{K+1})}{\Gamma(a_{K+1}) \Gamma(a_{K+1} + c_{K+1} + l_{K+1})}. \tag{13}$$

Since the $p_{i,i}$s have been marginalised and are no longer needed they are discarded.

To combat overfitting a geometric model probability is specified that places more weight on fewer regimes

$$p(K) = (1 - \delta)^{K-1} \delta, \qquad k_{max} + 1 > K > 0, \tag{14}$$

in which δ is set equal to 0.9 following Bulkley, Leslie and Smith (2015) and k_{max} denotes a pre-specified upper limit on the number of breaks (set equal to T). The value of 0.9 for δ is chosen so that the effective prior on the number of structural breaks corresponds to the findings of Pastor and Stambaugh (2001). To evaluate the effective prior on the number of breaks resulting from the combination of the prior on the regime durations and the prior model probability the model was run without data. The effective prior on the number of breaks occurring in the sample 1871 to 2013 for the realised excess return and the dividend growth models gives a mode of 16 breaks, corresponding to a break occurring approximately every ten years with a spread from 1 to 35 breaks. The effective prior on the number of breaks occurring over the sample 1951 to 2013 for the earnings growth model has a mode of seven breaks corresponding to a break occurring approximately every ten years, with a spread ranging from one to 20 breaks.

2.2.4. Beliefs About the Premium's Association with Volatility

Within each regime the prior relation between the premium and volatility is

$$\mu_i = \gamma \psi_i \sigma_i^2, \qquad i = 1, \ldots, K + 1, \tag{15}$$

in which $\gamma > 0$ and $\psi = (\psi_1, \ldots, \psi_{K+1})$. The prior on the aggregate price of risk γ follows a gamma distribution

$$p(\gamma) \propto \gamma^{a_\gamma - 1} \exp\left\{ -\frac{\gamma}{b_\gamma} \right\}, \qquad \gamma > 0, \tag{16}$$

in which a_γ and b_γ are specified using an empirical Bayes approach that follows Pastor and Stambaugh (2001).[4]

The regime-specific price of risk ψ_i also follows a gamma distribution

$$p(\psi_i) \propto \psi_i^{(\nu/2)-1} \exp\left\{-\frac{\psi_i \nu}{2}\right\}, \qquad \psi_i > 0, \qquad i = 1, \ldots, K+1 \tag{17}$$

in which ν is specified to be equal to ten following Pastor and Stambaugh (2001). Such a specification places a moderate prior on the regime-specific price of risk.

2.2.5. Beliefs About Magnitudes of Changes in the Premium

The vector of mean excess returns follows a hierarchical multivariate normal distribution with mean $\overline{\mu}$ and covariance matrix V_μ^{-1}

$$\begin{aligned} p(\mu \mid \overline{\mu}, K) &\propto \exp\left\{-\frac{1}{2}(\mu - \overline{\mu}\imath)' V_\mu^{-1} (\mu - \overline{\mu}\imath)\right\}, && \mu > 0, \\ p(\overline{\mu} \mid K) &\propto 1, && \overline{\mu} > 0, \end{aligned}$$

in which $\imath$ denotes a $(K+1) \times 1$ vector of ones. The prior on $\overline{\mu}$ is uniform and truncated at zero. The $(K+1) \times (K+1)$ covariance matrix is specified as

$$V_\mu^{-1} = \sigma_\mu^2 \begin{pmatrix} 1 & -\rho & 0 & \cdots & 0 \\ -\rho & 1+\rho^2 & -\rho & \cdots & 0 \\ 0 & -\rho & 1+\rho^2 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & \cdots & 1+\rho^2 & -\rho & 0 \\ 0 & \cdots & -\rho & 1+\rho^2 & -\rho \\ 0 & \cdots & 0 & -\rho & 1 \end{pmatrix}, \tag{18}$$

in which σ_μ^2 denotes the unconditional prior variance of μ_i. The vector of mean excess returns μ is specified as a conditionally independent AR(1) process. The value of ρ is set equal to 0.9 following Bulkley, Leslie and Smith (2015), and σ_μ^2 is set equal to 0.001 to specify a small prior probability of a large shift in the level of the premium. All specified hyperparameter values are displayed in Table 1.

Let the parameter vector, θ, consist of γ, ψ, μ, and $\overline{\mu}$. All the priors except those on μ and $\overline{\mu}$ are assumed to be independent. Using (15) to substitute for the elements of σ^2 the joint posterior

[4] The values of a_γ and b_γ are data-dependent and hence are different from Pastor and Stambaugh (2001), and vary among the realised return, dividend and earning growth estimates.

Table 1

Hyperparameter	Real	Div	Earn
a_i	224	224	224
c_i	3	3	3
a_γ	37.1	203.79	8.46
b_γ	0.1	0.12	0.23
ν	10	10	10
σ_μ	0.001	0.001	0.001
ρ	0.9	0.9	0.9
δ	0.9	0.9	0.9

Table 1: Hyperparameter Values. This table displays the specified hyperparameter values for each of the priors of the model used to estimate the equity premium with realised excess returns (Real), the dividend growth model (Div), and the earnings growth model (Earn). Recall from Section (2.2.5) that the specification of the covariance matrix (and hence σ_μ and ρ) differs from Pastor and Stambaugh (2001). The specification of ρ and δ follows Bulkley, Leslie and Smith (2015). The remaining hyperparameter values are specified in the same manner as in Pastor and Stambaugh (2001) but the data-dependent values of a_γ and b_γ differ because different data sets are used.

distribution is

$$p(\theta \mid x) \propto \left(\prod_{i=1}^{K+1} p(l_i)\right) p(\gamma) \left(\prod_{i=1}^{K+1} p(\psi_i)\right) p(\mu \mid \overline{\mu}, K) p(\overline{\mu} \mid K) p(K) \\ \times \left(\prod_{i=1}^{K+1} \left(\frac{\psi_i \gamma}{\mu_i}\right)^{l_i/2}\right) \exp\left\{-\frac{1}{2}\sum_{i=1}^{K+1} \sum_{t=q_{i-1}+1}^{q_i} \frac{(x_t-\mu_i)^2}{\mu_i}\gamma\psi_i\right\}. \tag{19}$$

The model is estimated in the same manner as Bulkley, Leslie and Smith (2015). The specific details of the estimation procedure, explained in the Appendix, are different from Bulkley, Leslie and Smith (2015) due to the removal of the TRs from the model.

3. Empirical Analysis

3.1. The Conditional Expected Stock Return

This article uses monthly returns because the annual return series contains fewer than 150 observations which is insufficient for the data-driven model estimated here. The market portfolio is the S&P 500. The Consumer Price Index is the deflator and is sourced from Robert Shiller's website. The data on earnings, dividends, prices, and returns for the period 1871-1926 are from Shiller and for the period 1927-2013 are from the Center for Research in Security Prices (CRSP). Due to concerns about the quality of the S&P earnings data during the beginning of the sample (Shiller 1992) the equity premium is estimated using the earnings growth model only for the period 1951 to 2013 following Fama and French (2002). Real returns are used in the empirical analysis.

The rate of return on a risk-free instrument for 1920-2013 is the Treasury-bill rate available from CRSP. However the Treasury-bill rate did not exist prior to 1920 and nor did any other short-term risk-free debt. Therefore the risk-free rate used from 1871-1919 is the series that Welch and Goyal (2008) estimate using Commercial paper rates for New York City available from the NBER Macrohistory database. The primary focus of this paper is to compare the equity premium estimates from the three competing models and so the risk-free rate is simply an additive constant that does not affect inferences.

Book data for all North American companies available on Compustat since 1961 is used. Book value is equal to total assets. An aggregate book value is calculated as the mean of all firms in each year, and this is divided by the market value to get the book-to-market ratio. Following Azar et al. (2015) market value is calculated as book value (total assets) + market value of equity (equal to stock price at fiscal year close times the number of common shares outstanding minus the book value of common equity). Linear interpolation is used to derive monthly values from annual observations.

3.1.1. Structural Breaks

Extremely large shifts in the equity premium are unlikely and it is reasonable that there is at least some positive relation between the equity premium and volatility. The equity premium estimates from the three different models presented in the main empirical analysis use such moderately informative prior beliefs about these two characteristics. Recall from (17) that the price of risk, the relation between the mean excess return and volatility, is characterised through the specification of the hyperparameter ν in the prior on ψ_i. These hyperparameter values are specified in accordance with Pastor and Stambaugh (2001). Specifically ν is set equal to ten specifying a moderate relation between the mean excess return and volatility. Pastor and Stambaugh (2001) state that this specification implies there is a ten percent probability that in any regime the price of risk is less than half its sample value, and a ten percent probability that the price of risk is more than 1.6 times its overall sample value. The degree by which the equity premium can shift from one regime to the next is specified by σ_μ which is set equal to 0.001 thereby specifying that a priori large shifts in the equity premium are unlikely.

Figures 1a and 1b display the posterior distributions of the number and timing of structural breaks when estimating the equity premium using the realised excess return series. The posterior mode number of breaks is 5 with a probability of 37%. Such a mode corresponds to a break occurring approximately every thirty years. The posterior has a spread ranging from 3 breaks to 15 with more than 50 percent of the mass lying between 5 and 6 breaks. The equity premium appears to be most unstable around the late 1920s, the early 1930s, the early 1940s, the early 1990s, and the late 2000s. These unstable periods correspond to major global events, namely, the Wall Street Crash and the Great Depression, World War II, the recession of the early 1990s, and the recent financial crisis.

This suggests that major global events are driving permanent shifts in the underlying distribution of the equity premium.

Figures 2a and 2b show that the estimated break dates correspond to times at which the dividend- and earnings-price ratios are changing direction. It is interesting that the break dates are similar for the dividend and earnings growth models.

3.1.2. The Equity Premium

Figure 3a plots the path of the equity premium from 1871 through 2013. The solid line plots the equity premium estimated with the dividend growth model while the dashed and dotted lines plot the equity premium estimated with the earnings growth model and realised excess returns. The solid vertical line marks the time at which the equity premium is first estimated with the earnings growth model, January 1951.

From 1871 to the 1940s the equity premium estimates from realised excess returns and the dividend growth model follow a similar path. There is strong evidence of a break in the late 1920s corresponding to the Wall Street Crash. Figure 3b plots the posterior standard deviations from all three models and shows that the new regime is characterised by higher volatility. Figure 4a graphs the paths of the conditional Sharpe Ratios computed using each of the three equity premium estimates. Sharpe ratios correspond to aggregate risk aversion and therefore the rising Sharpe ratios in the new regime suggest risk aversion in the market increased sharply. The higher volatility and increased aggregate risk aversion of this regime drove up the equity premium to 8%, its highest value of the sample.

While the equity premium estimates from the competing models follow a similar path for the first part of the sample, Figure 3a shows how their paths diverge thereafter. The dividend growth estimate falls more sharply in the 1950s and estimates from fundamentals remain below the realised return estimate for the remainder of the sample. The divergence in the estimates from fundamentals and the realised return estimate is concentrated around the early 1950s and the early 1980s. This suggests the stock market underwent periods of secular growth. Such growth cannot continue indefinitely and a break around the late 1990s, corresponding to the dotcom bubble and its subsequent bursting, caused the paths of realised and expected returns to converge.

3.1.3. Evaluating the Expected Return Estimates from 1951 to 2013

Since the paths of the equity premium estimates diverge in the latter part of the sample it is necessary to evaluate which estimate is most likely to be closest to the path of the true equity premium. Three pieces of evidence are provided that suggest the equity premium estimates from fundamentals are likely to be closer to the true premium than is the estimate from realised returns.

First, the expected return estimates from fundamentals are more precise. Figure 3b shows that the average of the posterior standard deviations for the entire sample 1871-2013 using the realised return and dividend growth models are 1.47% and 0.49%.[5] Using these values the standard errors for the realised return and dividend growth models are 0.036% and 0.012%. The estimate from the dividend growth model has a lower standard error and is hence more precise corroborating the findings of Claus and Thomas (2001) and Fama and French (2002).

Second, Table 2 displays the means of the Sharpe ratios calculated using the equity premium estimate from each of the three models for the entire sample, 1871-2013, and the two subsamples, 1871-1950 and 1951-2013. Let us compare the Sharpe Ratios from the dividend growth and realised excess return models since they are the only models for which we have Sharpe Ratios in both subsamples. Figure 3b shows that the first subsample, 1871-1950, contains the most volatile period of the entire sample corresponding to the Wall Street Crash. One might therefore expect aggregate risk aversion to be higher in the first subsample. The Sharpe ratio computed using the dividend growth model is in line with this hypothesis falling from 0.38 in the first subsample to 0.25 in the second. The Sharpe ratio computed using the realised excess return however rises from 0.45 to 0.54 from the first subsample to the second. It seems unlikely that aggregate risk aversion rose in the second subsample given that the first subsample contains the most volatile period of the entire sample. Furthermore the Sharpe ratio computed using the earnings growth model in the second subsample averages 0.17 far closer to the Sharpe ratio from the dividend growth model.

Table 2: Sharpe Ratios and Expected Returns

	SD	SY	SR	$E(R_Y)$	$E(R_D)$	$E(R)$	$A(Y_t/B_{t-1})$
1871-2013	0.33		0.49				
1871-1950	0.38		0.45				
1951-2013	0.25	0.17	0.54				
1961-2013				3.9%	5.4%	7.4%	6.2%

Table 2: Sharpe ratios, expected stock returns, and the average income return on book value. This table displays the means of the Sharpe ratios calculated from the equity premium estimates using the competing models for the entire sample and two subsamples, 1871-1950 and 1951-2013. SD, SY, and SR denote the Sharpe ratios for the dividend growth, earnings growth, and realised excess return models, respectively. The numerator for each Sharpe ratio is the average of the conditional equity premium estimate from the corresponding model over the given sample. The denominator is common across all three models and is computed as the average of the posterior standard deviation when estimating the stock return (using the same return series as the realised excess return model without subtracting the risk-free rate). $E(R)$, $E(R_D)$, and $E(R_Y)$ denote the average of the conditional expected stock return estimated using the realised return series, dividend growth, and earning growth models. The expected stock returns are estimated using all three models without subtracting the risk-free rate from each of the series. The average income return on book value is denoted $A(Y_t/B_{t-1}) = 1 + y_t/b_{t-1} * L_{t-1}/L_t - 1$. Nominal earnings for month t are denoted y_t, and nominal book values at the end of month t are denoted b_t. The rate of inflation for month t is denoted $L_t/L_{t-1} - 1$ where L_t denotes the price level at the end of month t. All variables except for the Sharpe ratios are multiplied by 1200 to express them as annualised percentages.

Third, and most important, investment theory and fundamentals suggest that the estimates

[5] Since the earnings growth model is only estimated from 1951 onwards we ignore it in this comparison.

from fundamentals are likely to be closer to the true equity premium than the realised excess return estimate. Figure 4b graphs the book-to-market ratio from 1961 to 2013. The book-to-market ratio for the period 1961 to 2013 averages 0.79 and is only greater than 1 for five years from 1974 to 1978, while never being greater than 1.19. A market value greater than the book value suggests that on average the cost of capital will be lower than the expected return on investment.

Following Fama and French (2002) the average income return on book equity is a measure of the expected return on the share of assets allocated to equity assuming an investment at time t generates a series of revenues at times $t+1$, $t+2$, ..., N, with a constant expected value. If, as a book-to-market ratio below 1 suggests, the expected return on investment is greater than the cost of capital then the average income return on book equity should exceed the cost of equity characterised by the expected stock return.

$$A(Y_t/B_{t-1}) > E(R). \tag{20}$$

The expected stock return estimates for the three models are obtained by re-estimating each model using the total return series without subtracting the risk-free rate. Table 2 displays the expected stock return estimates using the three models alongside the average income return on book value, $A(Y_t/B_{t-1})$, for the period 1961 to 2013. The $A(Y_t/B_{t-1})$ value is 6.2%, and the expected stock return estimates from the realised return, dividend growth, and earning growth models for the 1961-2013 subsample are 7.4%, 5.4%, and 3.9%, respectively. The expected stock return estimates using fundamentals are below the $A(Y_t/B_{t-1})$ value for this period. The expected stock return however using realised returns is above the $A(Y_t/B_{t-1})$ value implying that the average investment during this period was unprofitable which is not consistent with a book-to-market ratio of less than 1.

3.2. Unexpected Capital Gains

Fama and French (2002) note that there are three main explanations for why the average stock return is greater than expected for the period 1951-2013. First, dividends and earnings growth are unexpectedly high for the period 1951-2013. Second, the post-2013 growth rates of dividends and earnings are higher than expected. Third, the expected stock return is lower than expected at 2013. We now address these three possibilities in turn.

3.2.1. Is Dividend Growth for 1951-2013 Unexpectedly High?

As Fama and French (2002) note unexpectedly high growth in dividends and earnings for the period 1951-2013 cannot explain why the expected return estimates from fundamentals are so much lower than that produced by the realised excess return. If dividends and earnings growth were higher than

Table 3: Means of the Growth Rates of Dividends and Earnings

	GD	GY
1951-1960	0.68	-0.23
1961-1970	2.76	1.66
1971-1980	-0.75	2.95
1981-1990	2.29	-0.58
1991-2000	0.89	6.07
2001-2010	3.61	10.72
2011-2013	9.40	6.64

Table 3: Dividend and earnings growth rates. This table displays the average growth rates of dividends and earnings for each decade over the sample 1951 to 2013. Dividend and earnings growth rates are expressed as annualised percentages and have therefore been multiplied by 1200. The average growth rate of dividends for a given subsample is denoted GD $= (d_t/d_{t-1}) * (L_{t-1}/L_t) - 1$. The average growth rate of earnings for a given subsample is denoted GY $= (y_t/y_{t-1}) * (L_{t-1}/L_t) - 1$. Nominal dividends and earnings for month t are denoted d_t and y_t. The inflation rate for month t is $(L_{t-1}/L_t) - 1$ where L_t denotes the price level at the end of month t.

expected for the period unexpected capital gains are produced. Unexpected growth in dividends and earnings for the period, however, does not affect the 1951 or the 2013 dividend- and earnings-price ratios. For instance, assume that the price ratios are the same in 1951 and 2013 then the total growth in dividends and earnings for the period will equal the total growth in prices and all three models will produce similar expected return estimates.

3.2.2. Are post-2013 Expected Earnings and Dividend Growth Rates Unusually High?

There seems to be little evidence that long term expected growth rates of dividends and earnings are unusually high at the end of the sample period, 2013. Table 3 shows in fact that the average growth rate of real dividends falls from 1.72 percent for 1951-1970 to 1.51 percent for the period 1971-2010.

Following Fama and French (2002) the regression analysis in Table 4 presents further evidence on the optimal forecast of post-2013 real dividend growth rates. Forecasts of the dividend growth rates are made at the one and two year horizons. For the subsample 1876-1950 there is a moderate amount of predictability at the one year horizon. The lagged dividend yield and the lagged dividend-price ratio appear to have some predictive power of dividend growth. The regression explains 12 percent of the variation in the growth rate of dividends. Extending the forecast horizon to two years causes all forecast power to dissipate.

For the second subsample, 1951-2013, there is moderate predictability at the one year horizon. Many of the regressors have at least some predictive power of the dividend growth rate and the regression explains 14 percent of the variation in dividend growth. Increasing the forecast horizon to two years only the dividend-yield and dividend-price ratio are now significant and the regression captures just 8 percent of the variation in dividend growth. Without presenting the results extending the horizon further causes forecast power to dissipate corroborating the results of Fama and French

(2002). In sum, there is some forecast power of dividend growth rates at the one and two year horizons but this dissipates when the forecast horizon is extended. The finding that dividend growth is largely unpredictable over the past six decades is in line with previous studies (Campbell 1991; Cochrane 1994; Campbell and Shiller 1998; Fama and French 2002).

If dividend growth is essentially unpredictable over the sample 1951-2013 then the best forecast in 2013 of future long term dividend growth is the historical average growth rate. The historical average implies mean-reverting behaviour in the dividend growth rate and thus there is no evidence to suggest that post-2013 expected dividend growth rates are unusually high.

Turning to earnings growth it appears to have no obvious trend from 1951 to 2013. The growth rate is low at the beginning of the period averaging -0.23 percent per year from 1951 to 1960. It then rises averaging 2.95 percent per year during 1971 to 1980. It averages its lowest value of the sample in the following decade (-0.58%) before rising again to average its highest value of 10.72% from 2001 to 2010. Earnings growth therefore appears highly volatile with no clear trend.

Table 5 presents regressions that evaluate the predictability of earnings growth at the one and two year horizons. At the one year horizon there is moderate forecast power of earnings growth with several regressors found to be significant and 11 percent of the variation in earnings growth is explained by the regression. All forecast power appears to dissipate at the two year horizon however. Only two variables are now significant (lagged values of earnings growth) and the regression captures just one percent of the variation in earnings growth.

In short, there is moderate predictability of earnings and dividend growth rates one year ahead but this dissipates rapidly as the forecast horizon is extended suggesting the historical average offers the best forecast of future growth rates. There is no evidence here to suggest that expected post-2013 dividend and earnings growth rates are unusually high.

It should be noted that a large literature exists on forecasting the equity premium. For instance Welch and Goyal (2008) argued that many of the popular predictive variables such as the dividend yield and the short rate are unable to deliver superior out-of-sample forecasts of the equity premium relative to the prevailing mean which is a ‘no predictability’ benchmark model. Campbell and Thompson (2008), Pettenuzzo et al. (2014), and Li and Tsiakas (2016) find that imposing economic constraints on the equity premium can improve predictability, while Rapach et al. (2010) report that combining forecasts can lead to out-of-sample gains relative to the prevailing mean. Neely et al. (2014) show that combining predictive variables with technical indicators offers strong predictability.

3.2.3. Do Expected Stock Returns Fall During the 1951 to 2013 Period?

The final possible explanation is that the expected stock return is unexpectedly low in 2013. The dividend-price ratio falls from 7.44 percent at the end of 1950 to 1.94 percent at the end of 2013

Table 4: Regressions to Forecast Real Dividend Growth Rates, GD_t, at 1 and 2 Year Horizons

Forecast GD_t at 1 Year Horizon														
1876- 1950	Int	d_{t-12}/y_{t-12}	d_{t-12}/p_{t-12}	GD_{t-12}	GD_{t-24}	GD_{t-36}	GD_{t-48}	GD_{t-60}	R_{t-12}	R_{t-24}	R_{t-36}	R_{t-48}	R_{t-60}	R^2
$Coef$	36.82	-18.55	-4.27	-0.56	-0.33	0.01	0.00	0.03	0.039	-0.01	-0.00	-0.02	-0.01	0.12
$t-Stat$	9.93	-4.47	-7.36	-1.68	-1.01	0.35	0.05	0.98	3.00	-0.54	-0.85	-0.28	-1.49	
1951- 2013														
$Coef$	5.47	-2.99	-0.64	0.09	-0.17	-0.14	0.03	0.07	0.03	-0.00	-0.01	0.00	-0.00	0.14
$t-Stat$	5.62	-2.82	-2.82	2.35	-5.10	-4.16	0.87	2.25	4.77	-0.15	-0.94	0.09	-0.50	
Forecast GD_t at 2 Year Horizon														
1876- 1950	Int	d_{t-24}/y_{t-24}	d_{t-24}/p_{t-24}	GD_{t-12}	GD_{t-24}	GD_{t-36}	GD_{t-48}	GD_{t-60}	R_{t-12}	R_{t-24}	R_{t-36}	R_{t-48}	R_{t-60}	R^2
$Coef$	8.51	-2.62	-0.99		-0.06	-0.00	-0.01	0.00		0.01	-0.01	-0.00	-0.02	-0.00
$t-Stat$	2.18	-0.59	-1.60		-1.80	-0.04	-0.31	0.06		0.78	-0.40	-0.19	-1.29	
1951- 2013														
$Coef$	1.95	1.93	-0.22		-0.14	-0.17	0.00	0.07		0.00	-0.01	0.00	-0.00	0.08
$t-Stat$	1.78	-0.97	-3.86		-5.07	0.03	1.87	0.61		-0.95	0.01	-0.33	-0.63	

Table 4: Regression forecasts of dividend growth at the one and two year horizons. This table displays regression forecasts of dividend growth at the one and two year horizons. The variables are constructed following Fama and French (2002). The nominal stock price at the end of month t is denoted p_t, and nominal dividends and earnings for month t are denoted d_t and y_t. L_t denotes the level of prices at the end of month t. Real growth of dividends $GD_t = (d_t/d_{t-1}) * (L_{t-1}/L_t) - 1$. The realised real return on the aggregate market portfolio (S&P500) for month t is denoted R_t. The R^2 is the adjusted R-squared. The regression intercept is denoted Int and $t - Stat$ denotes the t-Statistic. All variables are expressed as annualised percentages and are therefore multiplied by 1200.

Table 5: Regressions to Forecast Real Earnings Growth Rates, GY_t, at the 1 and 2 Year Horizons

Forecasts of GY_t at 1 Year Horizon 1951 - 2013	Int	Y_{t-12}/B_{t-24}	d_{t-12}/y_{t-12}	y_{t-12}/p_{t-12}	GY_{t-12}	GY_{t-24}	GY_{t-36}	GY_{t-48}	GY_{t-60}	R_{t-12}	R_{t-24}	R_{t-36}	R_{t-48}	R_{t-60}	R^2
$Coef$	0.01	0.00	0.03	-0.00	-0.23	-0.13	-0.09	-0.87	-0.32	-0.00	-0.00	-0.00	-0.00	-0.00	0.11
$t-Stat$	0.82	0.11	3.80	-2.13	-6.19	-3.45	-2.53	-2.09	-4.02	-2.94	-0.88	-0.29	-0.86	0.12	

Forecasts of GY_t at 1 Year Horizon 1951 - 2013	Int	Y_{t-24}/B_{t-36}	d_{t-24}/y_{t-24}	y_{t-24}/p_{t-24}	GY_{t-12}	GY_{t-24}	GY_{t-36}	GY_{t-48}	GY_{t-60}	R_{t-12}	R_{t-24}	R_{t-36}	R_{t-48}	R_{t-60}	R^2
$Coef$	0.01	-0.00	0.00	-0.00		-0.12	-0.07	-0.06	-0.27		-0.00	-0.00	-0.00	0.00	0.01
$t-Stat$	1.41	-0.30	0.32	-1.12		-2.92	-1.56	-1.25	-3.10		-0.60	-0.03	-0.62	0.33	

Table 5: Regression forecasts of earnings growth. This table displays regression forecasts of earnings growth at the one and two year horizons. The R-squared value is the adjusted R squared. The variables are constructed following Fama and French (2002). The nominal stock price and book value at the end of month t are denoted p_t and b_t, and nominal dividends and earnings for month t are denoted d_t and y_t. L_t denotes the level of prices at the end of month t. Annual observations on book value are linearly interpolated to obtain a monthly series. Real growth of earnings $GY_t = (y_t/y_{t-1}) * (L_{t-1}/L_t) - 1$. The realised real return on the aggregate market portfolio (S&P500) for month t is denoted R_t. The R^2 is the adjusted R-squared. The regression intercept is denoted Int and $t - Stat$ denotes the t-Statistic. All variables are expressed as annualised percentages and are therefore multiplied by 1200.

so the growth in stock price is 3.84 times the growth in dividends from 1950 to 2013. Similarly the earnings-price ratio falls from 14.38 percent at the end of 1950 to 5.54 percent at the end of 2013 so the growth in stock price is 2.59 times the growth in dividends from 1950 to 2013. The dividend- and earnings-price ratios are driven by expected future returns and expected dividend and earnings growth. The previous regression analysis suggests that expected future dividend and earnings growth rates are not unusually high in 2013. Therefore the high capital gain relative to dividend and earnings growth for the period 1951 to 2013 is largely driven by a falling expected stock return. Fama and French (2002) note that some of this decline is probably anticipated, caused by a reversion to the unconditional mean given the high conditional return in 1950. The remaining decline is the result of a decline in expected returns to values far below the mean.

A key contribution of this article relative to Fama and French (2002) is disentangling the amount of the decline in the premium that is caused by mean reversion and declining expected stock returns, respectively. Approximately half (3.68%) of the 7.01% decline in the conditional premium from October 1951 to March 1985 was anticipated and reflects reversion of a high conditional premium to the long run mean of 4.48% (see Figure 5). The remaining 3.33% decline was unanticipated and resulted from a decline in the expected stock return.

The second key insight our methodology provides relative to Fama and French (2002) is the exact dates at which this decline occurred. We find that the decline in the expected stock return was driven largely by two breaks that occurred in 1951 and 1979 corresponding to the Treasury Fed Accord of 1951 and the Fed's 'monetarist policy experiment' between 1979 and 1982. This suggests that monetary policy was the main driver of the expected stock return over the past century.

4. Conclusions

This article provides a long term comparison of three competing estimates of the path of the equity premium from 1871 to 2013. In addition to the realised excess return series two additional series are constructed using dividend and earnings growth rates to measure the expected rate of capital gain (Fama and French 2002). Permanent shifts in the dividend- and earnings-price ratios are likely given the long sample and therefore it is important to formally model the underlying breakpoint process when estimating the path of the premium from the three return series. The economically-motivated Bayesian model of Pastor and Stambaugh (2001), with the removal of the TRs, is used to estimate the premium. The Bayesian RJMCMC approach proposed by Bulkley, Leslie and Smith (2015) is adopted to estimate the structural breaks because it can determine the number of breaks endogenously rather than following Pastor and Stambaugh (2001)'s use of Chib (1998)'s algorithm which requires the number of breaks to be pre-specified by the user. The model therefore enables the uncertainty surrounding the timing and the number of breaks to be incorporated into the premium estimates. It allows pre-break data to inform the current estimate while ruling out nonpositive premium estimates on theoretical grounds. The specification of the priors enable large shifts in the premium across regimes to be unlikely, while placing a prior positive relation between the premium and volatility.

From 1871 to the late 1920s the estimates from fundamentals and realised excess returns are similar. They both suggest the premium was stable at around 4 and 6 percent, respectively. A structural break occurred in the late 1920s, however, corresponding to the Wall Street Crash. The new regime is characterised by increased volatility and aggregate risk aversion driving the premium up to 8 percent. Upon exiting the high volatility regime of the 1930s corresponding to the Wall Street Crash, the estimate from fundamentals suggests that aggregate risk aversion and the equity premium return to their pre-crash levels. The realised return estimate exhibits a slow decline and fluctuates between 6 and 7 percent from 1951 to 2013. Risk aversion remains at a high level for this period following the Crash. The dividend growth estimate however declines sharply following a break in 1951 corresponding to the Treasury Fed Accord, reaching approximately 3 percent by the late 1950s. Another break in 1979 corresponding to the change in the Fed's operating procedures causes the premium to fall to approximately 1 percent. Finally, a break in 2000 corresponding to the dotcom bubble causes the premium to gradually rise reaching its unconditional mean of 4.35% by the end of the sample, still far below the realised return estimate of approximately 6.5%. The earnings growth estimate follows a similar path to the dividend growth estimate.

Three pieces of evidence suggest the estimate from fundamentals are likely to be closer to the path of the true equity premium. First, the estimates from fundamentals have lower standard errors and therefore are estimated more precisely. Second, Sharpe ratios are related to aggregate risk aversion. Splitting the sample into two periods 1871-1950 and 1951-2013 one might expect the average Sharpe ratio to fall across these periods since the first subsample contains the most volatile period of the entire sample around the Wall Street Crash. The Sharpe ratio calculated from the dividend growth model falls across subsamples however the Sharpe ratio calculated from realised excess returns actually increases which seems implausible. Third, and most important, the expected stock return estimated from realised returns is greater than the average income return on investment for the period 1961 to 2013 suggesting the average corporate investment for the period was unprofitable. This is difficult to reconcile with an average book-to-market ratio of 0.79 for the period. The expected stock returns from fundamentals however are less than the average income return on investment and are therefore more consistent with the observed book-to-market ratio.

The main findings therefore are that the premium fell from 8% in 1951 to 3.5% in 1957 due to a structural break corresponding to the Treasury Fed Accord in 1951. Of this 4.5% decline approximately 3.5% of it was anticipated as a reversion from a high conditional premium to the unconditional mean and the remaining 1% decline was the result of a decline in the expected stock return. The premium then fell from approximately 3% following a structural break in 1979 corresponding to a change in the Fed's operating procedures to its lowest value of the sample of roughly 1% in 1982. All of this 2% decline was due to a decline in the expected stock return. Finally the premium reverted from this low conditional value to its unconditional mean of 4.35% by 2013 following a break in 2000 corresponding to the dotcom bubble.

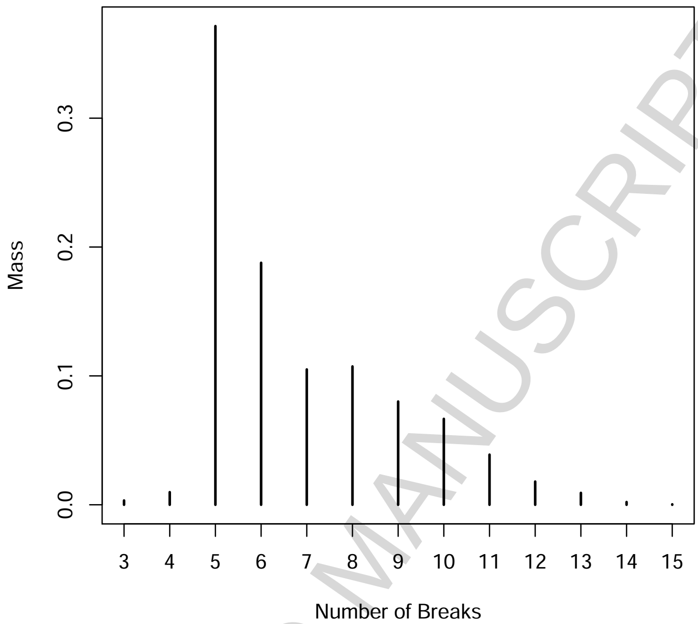

(a) Posterior Model Probability

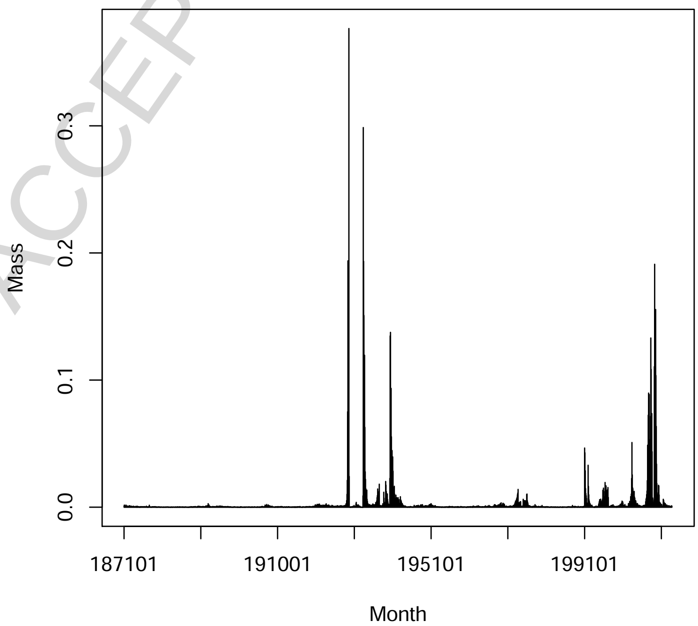

(b) Posterior Break Dates

Figure 1: This Figure displays the posterior model probabilities (top panel) and the posterior break dates (bottom panel) for the sample 1871 to 2013 when estimating the equity premium using the realised excess return series and the hyperparameters set out in Table 1.

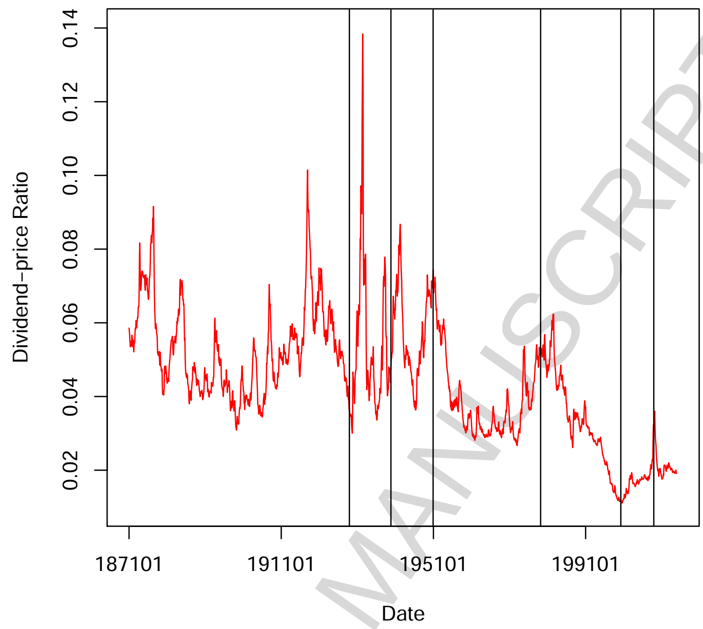

(a) Dividend-price ratio

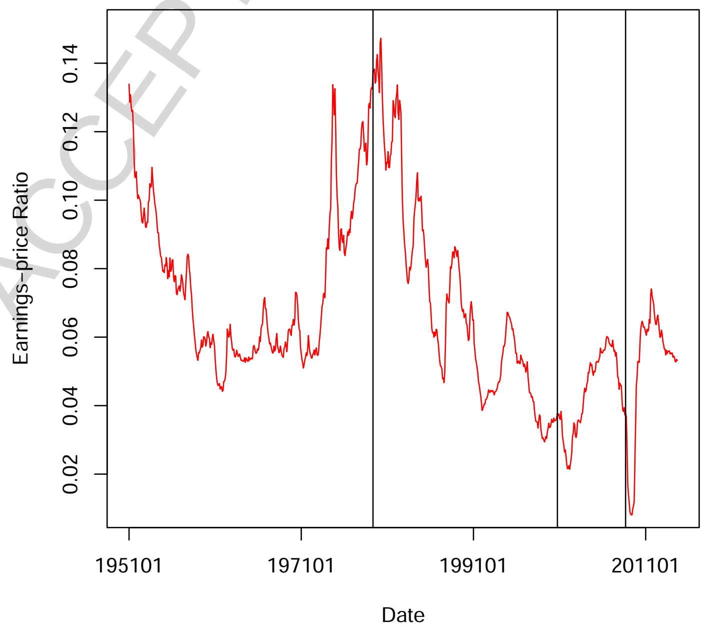

(b) Earnings-price ratio

Figure 2: This figure graphs the dividend-price ratio from 1871 to 2013 (top panel) and the earnings-price ratio from 1952 to 2013 (bottom panel). The solid vertical lines mark the posterior break location modes when estimating the equity premium using the respective models. The breaks occur at February 1929, December 1939, March 1951, April 1979, May 2000, and January 2009 for the dividend-price ratio, and at May 1979, October 2000, and September 2008.

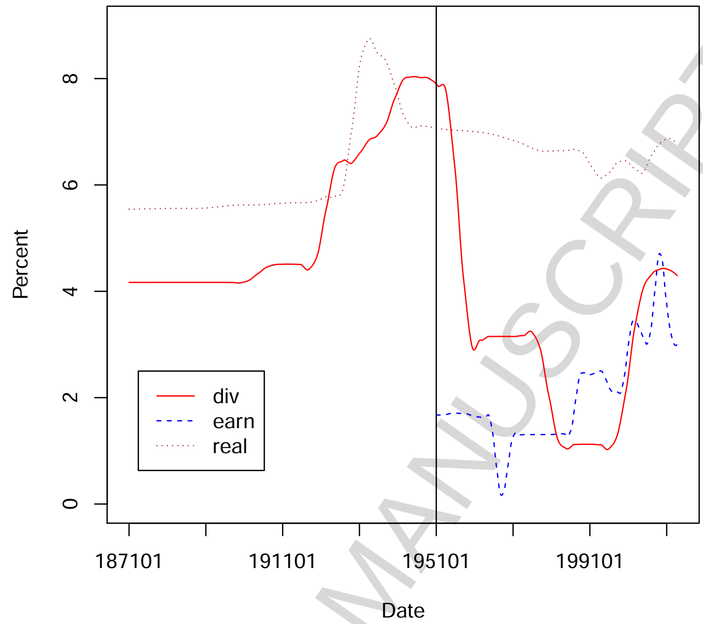

(a) Equity Premium

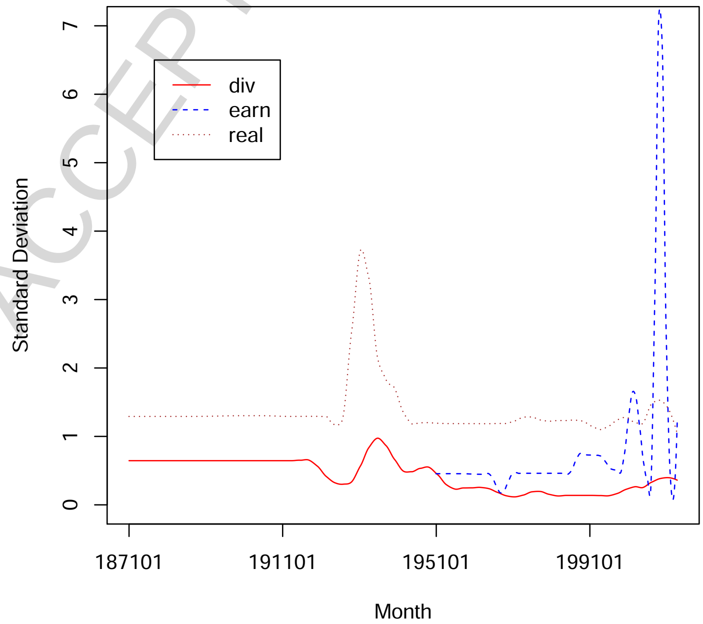

(b) Standard Deviation

Figure 3: This figure graphs the path of the equity premium (top panel) and corresponding standard deviation (bottom panel) from 1871 to 2013 estimated using all three models. Values are expressed as annualised percentages. Specifically the solid, dotted, and dashed lines graph the equity premium and standard deviation estimated using the dividend growth, realised excess returns, and earnings growth models, respectively. The vertical line marks the date at which the first estimate is made using the earnings growth model, January 1951.

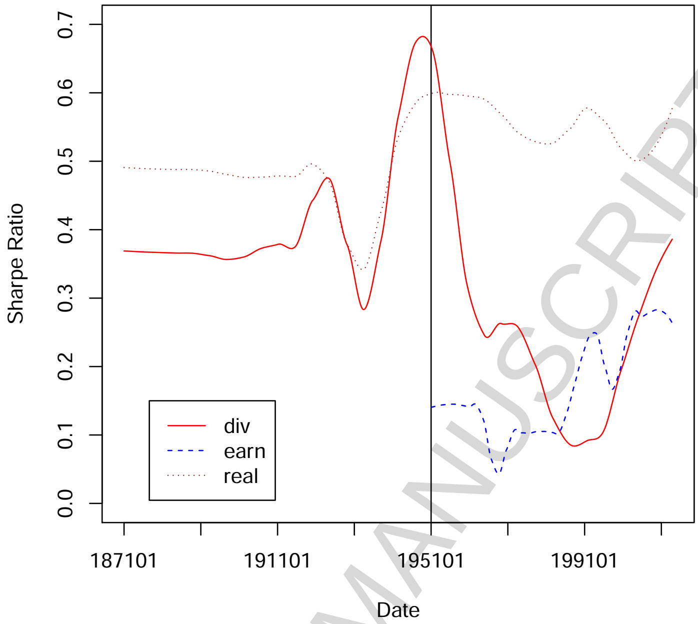

(a) Sharpe Ratios

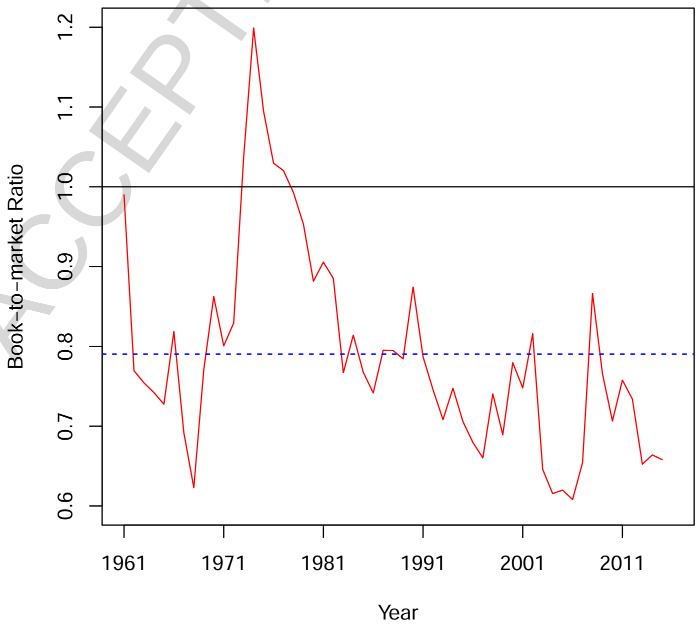

(b) Book-to-market Ratio

Figure 4: This figure graphs the path of the Sharpe ratio (top panel) calculated using each of the three equity premium estimates. The numerator of each Sharpe ratio at a given date is the equity premium estimate from the corresponding model. The denominator is the posterior standard deviation when estimating the path of the stock return (using the same return series as for the realised excess return model without subtracting the risk-free rate). Specifically the solid, dotted, and dashed lines graph the conditional Sharpe Ratios using the dividend growth, realised excess returns, and earnings growth models, respectively. The vertical line marks the date at which the first estimate is made using the earnings growth model, namely January 1951. The bottom panel graphs the book-to-market ratio from 1961 through 2013. The ratio is above 1 for only a few years of the sample, 1974 to 1978, and has a mean of 0.79.

Figure 5

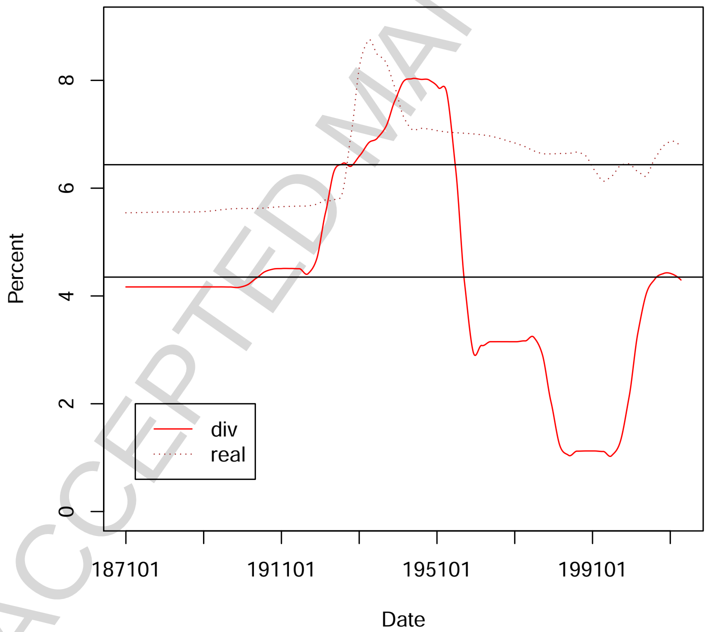

Figure 5: This figure graphs the path of the equity premium from 1871 through 2013. The solid and dashed lines correspond to the dividend growth and realised excess return models. The higher and lower solid vertical lines mark the average of the equity premium estimates for the whole sample from the realised excess return and dividend growth models.

Appendix A. Estimating the Model

Appendix A.1. Updating the Parameter Vector

Conditional on the number of breaks K and their locations q the parameter vector θ is updated using the Metropolis Hastings and Gibbs steps. The components of θ are drawn in turn from their full conditionals. Specifically, every component of θ is drawn using the Gibbs step except for μ which is drawn using the one-at-a-time Metropolis Hastings step. The proposal distribution for μ is a simple normal distribution truncated at zero with mean equal to the current value of μ and variance tuned to give the desired acceptance ratio. Performing the change of variables $\lambda \equiv 1/\gamma$ and $\phi_k^2 \equiv 1/\psi_k$ as in Pastor and Stambaugh (2001) the full conditionals of the components of θ are

$$\overline{\mu} \mid \cdot \sim \mathrm{N}\left(\frac{\imath' V_\mu^{-1}\mu}{\imath V_\mu^{-1}\imath}, \frac{1}{\imath' V_\mu^{-1}\imath}\right), \qquad \overline{\mu} > 0, \tag{Appendix A.1}$$

$$\phi_i^2 \mid \cdot \sim \frac{\nu + \sum\limits_{t=q_{i-1}+1}^{q_i} \frac{(x_t-\mu_i)^2}{\lambda\mu_i}}{\chi^2_{\nu+l_i}}, \qquad i = 1, \ldots, K+1, \tag{Appendix A.2}$$

$$\lambda \mid \cdot \sim \frac{\frac{2}{b_\gamma} + \sum\limits_{i=1}^{K+1} \sum\limits_{t=q_{i-1}+1}^{q_i} \frac{(x_t-\mu_i)^2}{\phi_i^2\mu_i}}{\chi^2_{2a_\gamma + \sum_{i=1}^{K+1} l_i}}, \tag{Appendix A.3}$$

(Appendix A.4)

$$p(\mu \mid \cdot) \propto \left(\prod_{i=1}^{K+1} \mu_i^{(-l_{i/2})}\right) \times \exp\left\{-\frac{1}{2}\left[(\mu - \overline{\mu}\imath)' V_\mu^{-1}(\mu - \overline{\mu}\imath) + \sum_{i=1}^{K+1} \sum_{t=q_{i-1}+1}^{q_i} \frac{(x_t-\mu_i)^2}{\lambda\phi_i^2\mu_i}\right]\right\}, \qquad \mu > 0. \tag{Appendix A.5}$$

The sampler is run for 50,000 iterations discarding the first 5,000 and thinning the remaining draws at an interval of 5 to combat autocorrelation. For each sampler two chains are run with different seeds and starting values to check convergence.

Appendix A.2. Updating the Number of Changepoints

This article develops the RJMCMC algorithm of Green (1995) (see also Bulkley, Leslie and Smith 2015) that enables a Bayesian model to jump between different dimensions and thereby determine the dimensionality of the problem in the process. In the current context this involves jumping between models with different numbers of changepoints K and thereby approximating the posterior distribution of the number of changepoints $p(K \mid x)$. Drawing this posterior enables K to be marginalised when estimating θ and therefore the uncertainty surrounding the number of changepoints to be incorporated into the parameter estimates.

On each iteration of the sampling scheme the RJMCMC algorithm attempts with equal probability to either add a break (birth move) or remove a break (death move). Letting b_k and d_k denote the probabilities of proposing a birth and death move then $b_k = d_k = 0.5$. However if $K = 0$ then $d_0 = 0$ and if $K = k_{max}$ then $b_{k_{max}} = 0$. Corresponding parameters are then proposed for the new regime(s). The acceptance probability that ensures detailed balance is maintained across the entire parameter space is calculated and the proposed move is accepted with this probability. If the move is accepted the new changepoints and parameters are updated else they are discarded and the original ones are maintained.

Birth Move: A birth move attempts to increase K to $K+1$ by proposing a new changepoint. The new changepoint denoted q_{i^*} is proposed uniformly from the sample

$$q_{i^*} \sim U[1, T]. \tag{Appendix A.6}$$

The move is immediately rejected if an existing changepoint is proposed, that is if $q_{i^*} \in q$. Let i^c denote the existing regime in which the proposed changepoint falls. Assuming $q_{i^*} \notin q$ the existing regime has been split into two new regimes denoted i^* and i^*+1. There are now two new regimes where once there was one so two new mean excess return values and two new regime-specific price of risk values, to compute the variance in the two new regimes, must be proposed.

Let the proposed mean excess return values in the two new regimes be denoted μ_{i^*} and μ_{i^*+1}. They are proposed from a truncated normal distribution with mean $\overline{x}$ set equal to the mean of the corresponding excess return series

$$\mu_{i^*}, \mu_{i^*+1} \sim \mathrm{N}(\overline{x}, \sigma^2_{\mu^*}), \qquad \mu_{i^*}, \mu_{i^*+1} > 0, \tag{Appendix A.7}$$

in which $\overline{x} = \sum_{t=1}^{T} x_t/T$. The values of $\overline{x}$ for the realised excess return, dividend growth, and earnings growth model are 0.00486, 0.00347, and 0.00488, respectively, while $\sigma^2_{\mu^*}$ is set equal to 0.005 for all three models. New $\phi^2_{i^*}$ and $\phi^2_{i^*+1}$ values are proposed from their full conditional and the variances for the two new regimes $\sigma^2_{i^*}$ and $\sigma^2_{i^*+1}$ are calculated using (15).

This completes the set of new parameters required for the birth move. The death move must now be specified before calculating the acceptance probability.

Death Move: The death move proposes to remove an existing changepoint and thus reduce K to $K-1$. The changepoint to remove denoted q_{i^c} is proposed uniformly from the set of existing changepoints q

$$q_{i^c} \sim U[q]. \tag{Appendix A.8}$$

The move attempts to replace the two regimes that are separated by q_{i^c} denoted i^c and i^c+1 with a single regime denoted i^*. The length of i^* is set equal to the sum of the lengths of the two existing regimes

$$l_{i^*} = l_{i^c} + l_{i^c+1}. \tag{Appendix A.9}$$

The existing values of μ_{i^c}, μ_{i^c+1}, $\phi^2_{i^c}$, and $\phi^2_{i^c+1}$ are discarded. The mean excess return value for the new regime μ_{i^*} is proposed using equation (Appendix A.7) while the new $\phi^2_{i^*}$ value is proposed from its full conditional enabling the variance of the new regime σ^2_{i*} to be calculated using (15). Table A.6 displays the parameters that need to be proposed for the birth and death moves.

Table A.6: Parameters to propose

Birth move
μ_{i^*}
μ_{i^*+1}
$\phi^2_{i^*}$
$\phi^2_{i^*+1}$
Death move
μ_{i^*}
$\phi^2_{i^*}$

Table A.6: Parameters to propose. This table displays all the parameters that need to be proposed for the birth and death moves.

Acceptance probability: The proposed move must be accepted with a probability that ensures detailed balance is maintained across the entire parameter space and this is set out in the remainder of this section.

Consider a move from K to K^* proposed with probability $j(K, K^*)$. Let θ_K and $\theta^*_{K^*}$ denote the respective parameter vectors for models K and K^*. The acceptance probability is

$$\max\left\{1, \frac{p(K^*)}{p(K)} \frac{p(\theta^*_{K^*} \mid K^*)}{p(\theta_K \mid K)} \frac{p(x \mid \theta^*_{K^*}, K^*)}{p(x \mid \theta_K, K)} \frac{j(K^*, K)}{j(K, K^*)} \frac{q(\theta^*_{K^*}, \theta_K)}{q(\theta_K, \theta^*_{K^*})}\right\} \tag{Appendix A.10}$$

in which the density of θ_K when moving from $\theta^*_{K^*}$ is denoted $q(\theta^*_{K^*}, \theta_K)$ and likewise the density of $\theta^*_{K^*}$ when moving from θ_K is denoted $q(\theta_K, \theta^*_K{}^*)$.

All the elements of θ_K that do not correspond to the regime that is being split in a birth move (or the regimes that are being merged in a death move) remain unchanged in $\theta^*_{K^*}$. The new elements of $\theta^*_{K^*}$ are proposed either from an independent distribution such as the normal distribution for $\mu^*_{i^*}$ and $\mu^*_{i^*+1}$ or from full conditionals such as $\phi^2_{i^*}$ and $\phi^2_{i^*+1}$. Parameters proposed from full conditionals cancel in the acceptance probability resulting in fewer effective dimensions thereby increasing the probability of a proposed move being accepted.

Since b_k= d_k=0.5, for each proposed move $j(K,K^*) = j(K^*,K)$=0.5 except at K=0 and $K = k_{max}$ where we may have $j = 0$ or $j = 1$. Only two proposal densities need to be considered because a move always either introduces or removes a single changepoint. When a birth move proposes to introduce a changepoint the proposal density is

$$q(\theta_K, \theta^*_{K+1}) = \frac{1}{T} f_{\bar{x},\sigma^2_{\mu^*}}(\mu^*_{i^*}) f_{\bar{x},\sigma^2_{\mu^*}}(\mu^*_{i^*+1}) p(\phi^{2*}_{i^*} \mid \cdots) p(\phi^{2*}_{i^*+1} \mid \cdots) \qquad \text{(Appendix A.11)}$$

in which the $(T)^{-1}$ corresponds to the uniform selection of any given time point in the sample at which to introduce the new changepoint, $f_{\bar{x},\sigma^2_{\mu^*}}$ denotes the density of a zero-truncated normal random variable with mean $\bar{x}$ and variance $\sigma^2_{\mu^*}$, and $p(\phi^{2*}_{i^*} \mid \cdots)$ is the full conditional of $\phi^{2*}_{i^*}$.

When a death move proposes to remove an existing changepoint the proposal density is

$$q(\theta_{K+1}, \theta^*_K) = \frac{1}{K+1} f_{\bar{x},\sigma^2_{\mu^*}}(\mu_{i^*}) p(\phi^{2*}_{i^*} \mid \cdots)$$

in which $(K+1)^{-1}$ corresponds to the uniform sampling of one of the $K+1$ existing changepoints to remove.

Appendix A.3. Updating the Changepoint Locations

Since the RJ step provides the sampler with a global movement enabling the introduction or removal of a changepoint at any location all that is required to keep the changepoint locations true is a simple local adjustment. This local adjustment is provided by a simple random-walk Metropolis Hastings step. A proposal is made to perturb the location of each changepoint q_i for $i = 1, \ldots, K$ by a number τ sampled uniformly from the interval $[-\tau, \tau]$ whereby τ is tuned to achieve the desired acceptance rate. The proposed perturbation to each changepoint location is accepted with probability $\min(1, \alpha)$, in which $\alpha = p(\hat{\theta} \mid x)/p(\theta \mid x)$ using (19), and $p(\hat{\theta} \mid x)$ denotes the posterior of θ calculated using the perturbed changepoint locations $\hat{q}$.

References

Azar, J., J.-F. Kagy and M. C. Schmalz, "Can Changes in the Cost of Carry Explain the Dynamics of Corporate Cash Holdings," *Review of Financial Studies, forthcoming* (2015).

Bai, J., "Estimating multiple breaks one at a time," *Econometric Theory* 13 (1997), 315–352.

Bai, J. and P. Perron, "Estimating and testing linear models with multiple structural changes," *Econometrica* 66 (1998), 47–78.

———, "Computation and analysis of multiple structural change models," *Journal of Applied Econometrics* 18 (2003), 1–22.

Blanchard, O. J., R. Shiller and J. J. Siegel, "Movements in the equity premium," *Brookings Papers on Economic Activity* (1993), 75–138.

Bulkley, G., D. S. Leslie and S. C. Smith, "Forecasting the Equity Premium in the Prsence of Structural Breaks," (2015).

Campbell, J. and S. Thompson, "Predicting excess stock returns out of sample: Can anything beat the historical average?," *Review of Financial Studies* 21 (2008), 1509–1531.

Campbell, J. Y., "A variance decomposition for stock returns," *Economic Journal* 101 (1991), 157–179.

Campbell, J. Y. and R. J. Shiller, "Valuation ratios and the long-run stock market outlook," *The Journal of Portfolio Management* 24 (1998), 11–26.

Carlin, B., A. Gelfand and A. Smith, "Hierarchical Bayesian analysis of changepoint problems," *Applied Statistics* 41 (1992), 389–405.

Chen, B. and Y. Hong, "Testing for smooth structural changes in time series models via nonparametric regression," *Econometrica* 80 (2012), 1157–1183.

Chib, S., "Estimation and comparison of multiple change-point models," *Journal of Econometrics* 86 (1998), 221–241.

Claus, J. and J. Thomas, "Equity premia as low as three percent? Evidence from analysts' earnings forecasts for domestic and international stock markets," *Journal of Finance* (2001), 1629–1666.

Cochrane, J. H., "Permanent and transitory components of GNP and stock prices," *The Quarterly Journal of Economics* (1994), 241–265.

Fama, E. and K. French, "The equity premium," *Journal of Finance* 57 (2002), 637–659.

Giordani, P. and R. Kohn, "Efficient Bayesian inference for multiple change-point and mixture innovation models," *Journal of Business & Economic Statistics* 26 (2008), 66–77.

Green, P., "Reversible jump Markov chain Monte Carlo computation and Bayesian model determination," *Biometrika* 82 (1995), 711–732.

Inclan, C., "Detection of multiple changes of variance using posterior odds," *Journal of Business & Economic Statistics* 11 (1993), 289–300.

Kim, C., J. Morley and C. Nelson, "The structural break in the equity premium," *Journal of Business & Economic Statistics* 23 (2005).

Kim, C. and C. Nelson, "Has the US economy become more stable? A Bayesian approach based on a Markov-switching model of the business cycle," *Review of Economics and Statistics* 81 (1999), 608–616.

Koop, G. and S. Potter, "Estimation and forecasting in models with multiple breaks," *Review of Economic Studies* 74 (2007), 763–789.

Li, J. and I. Tsiakas, "Equity premium prediction: The role of economic and statistical constraints," *Journal of Financial Markets* (2016).

Liu, J., S. Wu and J. Zidek, "On segmented multivariate regression," *Statistica Sinica* 7 (1997), 497–525.

Maheu, J. M. and S. Gordon, "Learning, forecasting and structural breaks," *Journal of Applied Econometrics* 23 (2008), 553–583.

Mehra, R. and E. C. Prescott, "The equity premium: A puzzle," *Journal of Monetary Economics* 15 (1985), 145–161.

Merton, R., "On estimating the expected return on the market: An exploratory investigation," *Journal of Financial Economics* 8 (1980), 323–361.

Neely, C. J., D. E. Rapach, J. Tu and G. Zhou, "Forecasting the equity risk premium: the role of technical indicators," *Management Science* 60 (2014), 1772–1791.

Pastor, L. and R. Stambaugh, "The Equity Premium and Structural Breaks," *Journal of Finance* 56 (2001), 1207–1239.

Pesaran, M., D. Pettenuzzo and A. Timmermann, "Forecasting time series subject to multiple structural breaks," *Review of Economic Studies* 73 (2006), 1057–1084.

PETTENUZZO, D., A. TIMMERMANN AND R. VALKANOV, "Forecasting stock returns under economic constraints," *Journal of Financial Economics* 114 (2014), 517–553.

RAPACH, D., J. STRAUSS AND G. ZHOU, "Out-of-sample equity premium prediction: Combination forecasts and links to the real economy," *Review of Financial Studies* 23 (2010), 821–862.

SHILLER, R. J., *Market volatility* (1992).

SONG, Y., "Modelling regime switching and structural breaks with an infinite hidden Markov model," *Journal of Applied Econometrics* 29 (2014), 825–842.

WELCH, I. AND A. GOYAL, "A comprehensive look at the empirical performance of equity premium prediction," *Review of Financial Studies* 21 (2008), 1455–1508.

Highlights

- The equity premium has undergone five structural breaks over the period 1871–2013
- Estimates from fundamentals are below those from realised returns for 1951–2013
- Valuation theory suggests fundamentals estimates are closer to true equity premium
- Unexpected capital gain result of changes in Feds operating procedures (1951 and 1979)