

A Bayesian solution to the equity premium puzzle¹.

A. Jobert, A. Platania, & L. C. G. Rogers

Statistical Laboratory, University of Cambridge

First draft: March 2005. This draft: March 2006.

Abstract

This paper describes a Bayesian solution to the *equity premium puzzle*, that is, the inability of standard intertemporal economic models to account for the magnitude of the observed excess return earned by a risky security over the return on T-bills. We follow convention and assume a single representative agent, but the main difference is that we suppose that the agent is not certain about the parameters of the dividend process, modelling this uncertainty by a prior distribution, and making inferences in a Bayesian fashion. The price of the stock is still the NPV of future dividends, but the agent is now averaging not only over the possible future paths of the dividend process, but also over the parameters that govern its dynamics. We then use particle filtering to work out the posterior distribution of the parameters of the problem, and find a striking conclusion; coefficients of relative risk aversion lie in the interval (1,2) with high probability - in other words, there is no equity premium puzzle.

JEL Classification: G10, G12.

Keywords: asset pricing, equity risk premium, risk-free rate premium, Bayesian approach.

Introduction.

The average annual return (i.e. inflation-adjusted) on the US stock market for the past 110 years has been about 7.9%, whereas in the same period the real return on a relatively riskless Treasury bill was roughly 1%. The difference between these two returns (i.e. 6.9%) is known as the equity premium. If standard intertemporal economic theory is consistent with our usual conception of risk according to which, on average, stocks should return more than bonds, it has been shown by Mehra and Prescott (Mehra and Prescott, 1985) to fail to capture the

¹We thank participants at the Cambridge-Princeton workshop, September 2005, and seminar participants in Sheffield and Cambridge, for their comments and support. We thank particularly Bill Janeway, Hashem Pesaran, Jose Scheinkman, Chris Sims and Seth Stafford.

magnitude of the equity premium. The observed 6.9 % premium was indeed much greater than could be explained by neoclassical representative agent paradigms² as a simple premium for bearing risk. The large equity premium and the low risk-free rate gave rise to the so-called “equity and risk-free rate puzzles”.

In their original analysis, Mehra and Prescott take a discrete-time model with a CRRA representative agent and a consumption growth rate assumed to be controlled by a two-state Markov chain. They show that the difference in the covariances of these returns with consumption growth is only large enough to explain the difference in the average returns if the typical investor is implausibly averse to risk. This is the *equity premium puzzle*: stocks are not sufficiently riskier than Treasury bills to explain the spread in their returns. As for the *risk-free rate puzzle* (Weil, 1989), it simply indicates that we do not know why people save even when bond returns are low³.

The equity premium puzzle is exhibited by Mehra and Prescott (Mehra and Prescott, 1985) in the simple case of lognormal growth rates of dividends⁴, where a simple closed-form solution is available for the equity premium, as the product of the coefficient of relative risk-aversion and the variance of the growth rate of consumption. But the latter variance is estimated to be 0.00125, which explains why unless the coefficient of relative of risk-aversion is taken to be implausibly large⁵, a high equity premium is impossible.

Over the last fifteen years, a huge literature has grown, presenting various attempts to explain the equity premium puzzle; excellent surveys are provided by Kocherlakota (Kocherlakota, 1996) and Mehra & Prescott (Mehra and Prescott, 2003). Without pretending to be an exhaustive list, different approaches to the problem have included: changing the law of the

²See among others (Lucas, 1978), (Breedon, 1979). A key idea underlying such models is that consumption today and consumption in some future period are treated as different goods, whose relative prices are equal to people’s willingness to substitute between these goods. In this type of model, a security’s risk can be measured using the covariance of its return with per capita consumption. Representative agent models, which incorporate the Lucas-Breedon paradigm for explaining asset-return differentials, have played a crucial role in our understanding and intuition of modern macroeconomics and the inability of such models to fit financial market data on stock returns has posed therefore a great challenge to the entire economics community.

³Although Treasury bills offer only a low rate of return, individuals defer consumption (i.e. save) at a sufficiently fast rate to generate average per capita consumption growth of around 2% per year.

⁴Dividends are assumed equal to consumption, by market clearing. This assumption is relaxed in various subsequent papers.

⁵Fischer Black proposed that a coefficient of risk-aversion taken to be equal to 55 would solve the puzzle; empirical studies (see among others (Arrow, 1971; Friend and Blume, 1975; Kydland and Prescott, 1982), however, provide evidence for such a coefficient being no more than 10.

dividends⁶, changing the preference structure⁷, introducing incomplete markets⁸, introducing hidden variables⁹, introducing market segmentation¹⁰, survivorship bias¹¹, borrowing constraints¹², liquidity premium¹³, taxation effects¹⁴, leverage effects¹⁵, to list but a few of the relevant papers on the equity premium puzzle. Some of these alternatives are more successful than others, and a detailed comparison and discussion of the literature is beyond the scope of this paper, but it is the verdict of Kocherlakota (Kocherlakota, 1996) and Mehra & Prescott (Mehra and Prescott, 2003) that there is not yet an explanation that is entirely satisfactory. Maybe more than one of the effects studied should be included, but we remark that the simplicity of the basic Lucas-Breeden representative agent analysis is strongly appealing, and the more complicated the alternative story, the less appealing it becomes.

What we offer in this paper is another possible explanation of the equity premium puzzle. In contrast to all the other explanations that we are familiar with, we shall take the basic Lucas-Breeden model, and change *nothing!* The process for the log dividends will still be a random walk, even a random walk with IID Gaussian increments, preferences will not be changed, consumption will equal dividend at all dates. To explain what will be different, let us first consider a simple but extremely illuminating example.

Example 1 (The 20's example). *Suppose you observe daily prices for T years of a stock with an annual rate of return of 20%, and an annual volatility of 20%. You want to observe for long enough so that your 95% (= $\frac{19}{20}$) confidence interval estimates of the parameters are good to 1 in 20 (that is your confidence interval is $\pm 1\%$).*

- (i) *How large must T be to give this level of precision in the estimate of the volatility?*
- (ii) *How large must T be to give this level of precision in the estimate of the rate of return?*

The answers are: (i) about 11, (ii) about 1550!!

Thus uncertainty regarding the rate of return is enormous, and any analysis which over-

⁶Cecchetti, Lam & Mark (Cecchetti et al., 1993), Rietz (Rietz, 1988), Mehra & Prescott (Mehra and Prescott, 1988), Tsionas (Tsionas, 2005), Barro (Barro, 2005).

⁷Campbell & Cochrane (Campbell and Cochrane, 1999), Constantinides (Constantinides, 1990), Epstein & Zin (Epstein and Zin, 1991), Ferson & Constantinides (Ferson and Constantinides, 1991)

⁸Constantinides & Duffie (Constantinides and Duffie, 1996), Weil (Weil, 1989), Heaton & Lucas (Heaton and Lucas, 1996), Lucas (Lucas, 1994), Mankiw (Mankiw, 1986), Mehra & Prescott (Mehra and Prescott, 1985).

⁹Brennan & Xia (Brennan and Xia, 2001), Veronesi (Veronesi, 1999), (Veronesi, 2000)

¹⁰Mankiw & Zeldes (Mankiw and Zeldes, 1991)

¹¹Brown, Goetzmann & Ross (Brown et al., 1995)

¹²Constantinides, Donaldson & Mehra (Constantinides et al., 2002)

¹³Bansal & Coleman (Bansal and Coleman, 1996)

¹⁴McGrattan & Prescott (McGrattan and Prescott, 2001)

¹⁵Kandel & Stambaugh (Kandel and Stambaugh, 1991), Benninga & Protopapadakis (Benninga and Protopapadakis, 1990)

looks this is unlikely to sit well with reality. We therefore assume alternatively that both the mean and the volatility of the growth rate are *unknown* to the representative agent, who is endowed with a prior distribution about these two parameters. Observing market rates of return, he can update his posterior beliefs about these parameters in a Bayesian way. Such a Bayesian procedure is significantly different from the literature described above, where the risk premium and the risk-free rate are effectively calibrated by plugging into the model prices the sample mean and the sample variance of past growth rates. We are inserting the full posterior distribution of growth parameters into our model prices, not just their point estimates¹⁶.

The hidden variable stories of Brennan & Xia (Brennan and Xia, 2001) and Veronesi (Veronesi, 1999), (Veronesi, 2000) have something of the flavour of our approach, but are very different in fundamental ways. Firstly, the returns in their models are not IID, but are instead stationary Markov processes; and secondly, the agent is supposed to know exactly what the law of the returns process is. This second feature seems to us to be hard to justify in view of the 20's example. In the language of the filtering community, (Brennan and Xia, 2001), (Veronesi, 1999), (Veronesi, 2000) are about *filtering* an unknown signal with known dynamics, whereas what we are doing is *estimating* the dynamics of a more simple model.

Thus we shall analyse the price process from a Bayesian perspective. One issue, highlighted by (Geweke, 2001), is that obvious choices of priors can lead to divergent expressions for the stock price in the standard Lucas tree model. We use a methodological innovation to deal with this. If we propose a Gamma-Gaussian prior for the precision and mean of the (Gaussian) log returns, then we shall obtain divergent expressions. But this is because this particular prior proposes parameter combinations which imply infinite asset price. We deal with this by inserting an additional factor into the prior density which eliminates the (absurd) parameter combinations which would lead to infinite stock price; the simple prior-to-posterior updating is not affected, but the divergences are gone.

When it comes to fitting the model to the data, we have to make inferences about the parameters of the problem, which here are the parameters of the representative agent's preferences, and of his prior density. At this point, we could apply conventional likelihood-based methods, but the functional forms of the stock and bond prices as functions of parameters and observations are not particularly simple, so we have resorted to computational Bayesian techniques. We chose to apply particle filtering techniques to calibrate our model to the growth rate of real consumption and real returns on the US Treasury bill and the SP500 index. This Bayesian approach is shown to account for the equity premium puzzle.

¹⁶A recent preprint of Weitzman (Weitzman, 2005) proposes a Bayesian approach similar to ours.

1 The basic model.

Let us consider the case of a frictionless economy that has a single representative agent, ordering its preferences over random consumption paths by

$$E_0 \left[\sum_{t=0}^{\infty} \beta^t U(c_t) \right], \quad (1)$$

where $0 < \beta < 1$ is a subjective discount factor, E_t denotes the expectation conditional on information available at time t , U is an increasing, continuously differentiable concave utility function and c_t represents per-capita consumption. The utility function is taken to be of constant relative risk-aversion (CRRA) form, namely

$$U(c) = \frac{c^{1-R}}{1-R}, \quad (2)$$

where $0 < R < \infty$ denotes the coefficient of relative risk-aversion, $R \neq 1$.

In this economy, the outputs y_t in period t (period t dividends assumed to be produced by one productive unit) are assumed to equal the consumption c_t in period t (so that market clears). One share with price S_t is competitively traded and is viewed as a claim on the dividend process y_t . Since the agent's time- t marginal utility is equal to $U'(c_t) = U'(y_t) = y_t^{-R}$, his state-price density at time t is given by $\zeta_t \equiv \beta^t y_t^{-R}$ and it follows from the usual marginal pricing story that equity is priced at time t by

$$\begin{aligned} S_t &= E_t \left[\sum_{j \geq 0} \beta^j \left(\frac{y_{t+j}}{y_t} \right)^{-R} y_{t+j} \right] \\ &= y_t E_t \left[\sum_{j \geq 0} \beta^j \left(\frac{y_{t+j}}{y_t} \right)^{1-R} \right]. \end{aligned} \quad (3)$$

The time t -price of a one-year bond which pays 1 at maturity is given by

$$B_t = \beta E_t \left(\frac{y_{t+1}}{y_t} \right)^{-R}. \quad (4)$$

The prices of stock and bond depend on the postulated dynamics of y , and on the preferences of the representative agent, embodied in the parameters β, R . Now in the original study of Mehra & Prescott, it is assumed that the agent knows the law of y with certainty, but in view of the 20's example, this assumption is hard to defend. Abandoning it leads us to propose that the agent knows a parametric form for the dynamics of y , but does not know the parameters; the agent will carry out a Bayesian inference on the parameters, and the expectations (3), (4) will be averages over the parameters of averages over future paths of y . As we shall see, this change of viewpoint substantially alters the problem.

2 The Bayesian representative agent.

We suppose that the log returns $\xi_t \equiv \log(y_t/y_{t-1})$ are independent random variables with common $N(\mu, \tau^{-1})$ distribution, where the agent does not know the parameters μ, τ , but takes a prior density $\pi_0(\mu, \tau)$ for them, and makes Bayesian inferences about the parameters in the light of the observations. The approach we are going to use could be described as ‘doubly Bayesian’. This is because firstly the agent acts as a Bayesian; he knows the parameters (β, R) of his preferences, but he does not know the true parameters (μ, τ) of the dynamics of y . He proposes a prior for (μ, τ) , and updates his beliefs with the observations, and computes at each time t what S_t and B_t should be, using (3), (4). The prices he comes up with will be complicated functions of his preference parameters (β, R) , the parameters specifying his prior, and the observations, but in what follows we will make this functional form quite explicit.

The second Bayesian level of thinking going on is that of the econometrician, who observes the same market data as the representative agent, but does *not* know the preference parameters (β, R) , or the parameters the agent used to specify his prior; the econometrician proposes a prior for these (to him) unknown parameters, and uses Bayesian methods to make inferences about them.

The particular form of the prior density π_0 that the representative agent uses is

$$\pi_0(\mu, \tau) \propto \varphi(\mu, \tau) \tau^{\alpha_0 - 1} \exp\left(-\frac{1}{2} K_0 \tau \mu^2 - b_0 \tau\right) \sqrt{\tau}, \quad (5)$$

where

$$\varphi(\mu, \tau) \equiv e^{-c/2\tau^2} (1 - \beta e^{-\nu\mu + \nu^2/2\tau})^+, \quad (6)$$

and $\nu \equiv R - 1$ and the constant c is positive. We shall assume that $R > 1$; this is not essential but makes various subsequent statements cleaner. In any case, the equity premium puzzle arises because of the apparently large values of R needed, so excluding small values is unimportant. The term $\exp(-c/2\tau^2)$ in the prefactor is included to guarantee the convergence of various integrals to be detailed later.

The form of (5) is mainly quite conventional; the law of τ is Gamma distributed, and given τ , μ is normally distributed with zero mean and precision¹⁷ τ - or at least that is what the joint law would be were it not for the prefactor $\varphi(\mu, \tau)$. The reason for this prefactor becomes clear when we consider the expression (3) for the share price. Using the assumed structure of y , we have

$$\begin{aligned} S_t &= y_t E_t \left[\sum_{j \geq 0} \beta^j \left(\frac{y_{t+j}}{y_t} \right)^{-\nu} \right] \\ &= y_t E_t \left[\sum_{j \geq 0} \beta^j \exp\left(-\nu \sum_{r=t+1}^{t+j} \xi_r\right) \right] \\ &= y_t E_t \left[\sum_{j \geq 0} \beta^j \exp\left(-\nu \mu j + j\nu^2/2\tau\right) \right]. \end{aligned}$$

¹⁷The precision is the reciprocal of the variance.

This is the reason for the prefactor; if $\beta \exp(-\nu\mu + \nu^2/2\tau) \geq 1$ then the sum inside the last expectation is infinite¹⁸! We must therefore ensure that the prior (and hence the posteriors) give no mass to this region, and the prefactor in (5) is one way to ensure this. Developing the expression for S_t one line further shows why this choice of prefactor is particularly advantageous:

$$S_t = y_t E_t \left[\frac{1}{1 - \beta e^{-\nu\mu + \nu^2/2\tau}} \right]. \quad (7)$$

The point of the choice of the prefactor φ is that it will give a cancellation in the evaluation of the stock price, resulting in considerable simplification.

Let us record the analogous result to (7) for bond prices, referring to (4):

$$B_t = \beta E_t \exp(-R\mu + R^2/2\tau). \quad (8)$$

In order to compute the expectations in (3) and (4), we now need to find the posterior distribution of (μ, τ) given (ξ_1, \dots, ξ_t) , denoted by $\pi_t(\mu, \tau | \xi_1, \dots, \xi_t)$. After observing ξ_1, \dots, ξ_t (assumed to be IID), the posterior density will be

$$\begin{aligned} \pi_t(\mu, \tau | \xi_1, \dots, \xi_t) &\propto \varphi(\mu, \tau) \exp \left[-\frac{1}{2} K_0 \tau \mu^2 - b_0 \tau \right] \tau^{\alpha_0 - 1} \exp \left[-\frac{\tau}{2} \sum_{k=1}^t (\xi_k - \mu)^2 \right] \sqrt{\tau} \tau^{t/2} \\ &\propto \varphi(\mu, \tau) \exp \left[-\frac{1}{2} K_t \tau (\mu - m_t)^2 \right] \tau^{\alpha_t - 1} \exp[-b_t \tau] \sqrt{\frac{K_t \tau}{2\pi}} \end{aligned} \quad (9)$$

where

$$\begin{aligned} K_t &= K_0 + t, \\ m_t &= \frac{t \bar{\xi}_t}{K_t}, \\ \alpha_t &= \alpha_0 + t/2, \\ b_t &= b_0 + \frac{1}{2} S_{\xi\xi}(t) + \frac{1}{2} t \bar{\xi}_t^2 K_0 / K_t \end{aligned}$$

where $\bar{\xi}_t = \frac{1}{t} \sum_{k=1}^t \xi_k$ and $S_{\xi\xi}(t) = \sum_{k=1}^t (\xi_k - \bar{\xi}_t)^2$. As is well known, the Gamma-Gaussian family updates nicely under Gaussian observations, and we see that the posterior density π_t is of the same structural form as the prior π_0 .

For the calculation of share and bond prices, it is helpful to introduce the function

$$F_t(\lambda, \rho) \equiv \int_0^\infty d\tau \int_{\frac{1}{\nu} \log \beta + \frac{\nu}{2\tau}} d\mu \frac{\pi_t(\mu, \tau)}{1 - \beta e^{-\nu\mu + \nu^2/2\tau}} \exp(-\lambda\mu + \lambda^2/2\tau - \rho\mu + \rho^2/2\tau) \quad (10)$$

The integral defining F_t is guaranteed to converge, because of the term $\exp(-c/\tau^2)$ which controls the integrand near $\tau = 0$. Further analysis of F is required before we can implement

¹⁸This divergence problem has also been observed by Geweke (Geweke, 2001).

the computation; we defer this to the Appendix. For now, what we need is that in terms of F_t , the expressions (3) and (4) become respectively

$$S_t = y_t \frac{F_t(0, 0)}{F_t(0, 0) - \beta F_t(\nu, 0)} \equiv f_S(t, \Xi_t, \theta), \quad (11)$$

$$B_t = \beta \frac{F_t(0, R) - \beta F_t(\nu, R)}{F_t(0, 0) - \beta F_t(\nu, 0)} \equiv f_B(t, \Xi_t, \theta), \quad (12)$$

where we let $\Xi_t = (\xi_0, \dots, \xi_t)^T$ denote the vector of values of dividend growth, and let θ denote the parameter vector $\theta = (\beta, R, \alpha_0, b_0, K_0, v_1, v_2, v_3)^T$, where the parameters v_1, v_2, v_3 will be explained in the next section. In order to calculate the prices of stock and bond, the representative agent needs to know \mathbf{y}_t ; his preference parameters (β, R) ; and the parameters (α_0, b_0, K_0) of his prior¹⁹.

3 Calibration.

We used the same data as Mehra and Prescott in their original analysis (Mehra and Prescott, 1985). Their original dataset was used to generate three time series, which we will be calibrating the model to, namely the growth rate of real consumption, the real return on a risk-free asset and the real return on the SP500. The observations are described in Figure 1.

In terms of our existing notation, the observation $Y_t \equiv (Y_t^1, Y_t^2, Y_t^3)^T$ is just

$$\begin{pmatrix} Y_t^1 \\ Y_t^2 \\ Y_t^3 \end{pmatrix} = \begin{pmatrix} \xi_t \\ -\log B_t \\ \log(S_t/S_{t-1}) \end{pmatrix} + \begin{pmatrix} \varepsilon_t^1 \\ \varepsilon_t^2 \\ \varepsilon_t^3 \end{pmatrix} \quad (13)$$

where the ε_t^i are independent zero-mean Gaussians, with variance of ε_t^i equal to v_i . We require there to be the possibility that there is some noise in the observations, as it would in general be impossible to fit the observed values exactly.

We briefly explain the main ideas of particle filtering, how they are modified for the present example, and how the methodology is applied. Excellent surveys are to be found in (Arumpalam et al., 2001), (Doucet et al., 2001), (Crisan and Doucet, 2002) that explain the methodology in greater detail.

A discrete-time Markov process $(X_t)_{t \geq 0}$ with *known* transition density $p(\cdot|\cdot)$ is imperfectly observed through $(Y_t)_{t \geq 0}$, where the density of Y_t given $\mathbf{X}_t \equiv (X_0, \dots, X_t)^T$ is $f(\cdot|\mathbf{X}_t)$. The posterior density π_t of \mathbf{X}_t given $\mathbf{Y}_t \equiv (Y_0, \dots, Y_t)^T$ is approximated by

$$\hat{\pi}_t = \sum_{i=1}^N w_t^i \delta_{x_t^i}, \quad (14)$$

¹⁹In principle, he needs to know c as well, but since this is just needed to make integrals converge we shall propose a value for this and hold it fixed, for no reason other than to lighten the calculations. We shall show in the Appendix that the choice of c has negligible effect on the outcome.

an atomic probability measure concentrated on the ‘particles’ $(x_t^i)_{1 \leq i \leq N}$ at time t . The exact Bayesian updating

$$\begin{aligned} P(X_{t+1} \in A, Y_{t+1} \in B | \mathbf{Y}_t) &= \int \pi_t(dx) \int_A dx' \int_B dy p(x'|x) f(y|x') \\ P(X_{t+1} \in A | \mathbf{Y}_{t+1}) &= \frac{\int \pi_t(dx) \int_A dx' p(x'|x) f(Y_{t+1}|x')}{\int \pi_t(dx) \int dx' p(x'|x) f(Y_{t+1}|x')} \end{aligned}$$

gets approximated by selecting a ‘descendent’ x_{t+1}^i of x_t^i according to the density $\tilde{p}(\cdot | x_t^i, y_{t+1})$, and recomputing the weights

$$w_{t+1}^i \propto w_t^i \frac{f(Y_{t+1}|x_{t+1}^i) p(x_{t+1}^i|x_t^i)}{\tilde{p}(x_{t+1}^i|x_t^i, y_{t+1})} \quad (15)$$

The most natural choice²⁰ for \tilde{p} is to take $\tilde{p}(x'|x, y) = p(x'|x)$, but the added flexibility is frequently helpful; see, for example, (Crisan and Doucet, 2002).

Applying this algorithm without further modification will in practice lead to posteriors $\hat{\pi}_t$ with virtually all weight concentrated on a small number of particles - impoverishment. To counteract this, a resampling step is generally used²¹; once the new particles (x_{t+1}^i) and weights (w_{t+1}^i) have been found, we select a sample of size N from the set $\{x_{t+1}^i : i = 1, \dots, N\}$ according to a multinomial distribution²² with probabilities (w_{t+1}^i) . The criterion used to determine when to resample was as follows. Given the N particles, we set n to be the integer part of $N^{0.4}$, and then sorted the particles into decreasing order of w_t^i . If

$$\sum_{i=1}^n w_t^i < RSLEVEL,$$

where $RSLEVEL \in [0, 1]$ is some pre-assigned value, then we do multinomial resampling, otherwise we do not resample. A value for $RSLEVEL$ of 0 means there is no resampling, a value of 1 means that there is always resampling, and we tried various values on different runs of the particle filter.

In our application, the Markov process is

$$X_t = (t, \xi_t, \bar{\xi}_t, S_{\xi\xi}(t), \theta)^T,$$

with known transitions. We initialise the population of particles by choosing the values of ξ_0 uniformly from the interval $[-0.2, 0.2]$, and then setting $\bar{\xi}_0 = \xi_0$, and $S_{\xi\xi}(0) = 0$. Next the value of β is chosen according to a $B(30.7, 2.6355)$ distribution²³. Next, values of R are

²⁰... which we shall use here

²¹The details may be varied; for example, resampling need not be done every period, or it may only be done among particles with weight less than some low threshold value.

²²This is the most obvious resampling scheme, but others can be used: see (Crisan and Doucet, 2002).

²³The mean look-ahead time of an agent with a given value of β is $(1 - \beta)^{-1}$. The particular parameters chosen for the Beta distribution were selected to make sure that the probability that $P[(1 - \beta)^{-1} < 6] = 0.05$ and $P[(1 - \beta)^{-1} > 50] = 0.05$. It seems unlikely that an agent would be looking ahead much more than 50 years, or much less than 6 years.

generated in the form $R = 1 + Z_1 + Z_2$, where Z_1, Z_2 are independent exponential variables with means which varied from one run to the next, but were in the range $[1, 2]$; we do not expect values of R to be higher than 5, and this choice of initial distribution makes such values possible but unlikely. We let initial values of α_0 be uniform on $[0.02, 25]$, initial values of b_0 be uniform on $[0.01, 2]$, and initial values of K_0 be uniform²⁴ on $[5, 80]$. Finally, we generate values for v_i as $(Z_1 + Z_2)$, where Z_i are independent exponentials, whose means vary from run to run, but are typically chosen so that the prior mean of v_i is about half the residual sum-of-squares of the observed Y^i ; in basis points (10^{-4}) the residual sums-of-squares of the data were (12.86, 35.04, 273.63).

We now detail the procedure used to generate the moves of the particles. Given the value $X_t = x$ of the Markov process at time t , we simulate the value X_{t+1} by firstly simulating a pair (μ, τ) according to the Gamma-Gaussian density proportional to

$$\exp \left[-\frac{1}{2} K_t \tau (\mu - m_t)^2 \right] \tau^{\alpha_t - 1} \exp [-b_t \tau] \sqrt{K_t \tau};$$

see (9). Then we perform rejection sampling, using the prefactor φ ; we compute $\varphi(\mu, \tau)$, and accept (μ, τ) if $\varphi(\mu, \tau) < U$, with U some independent $U[0, 1]$ variable, otherwise we propose another (μ, τ) and repeat until there is an acceptance. Though the principle is clear, we had to be very careful to carry out this rejection sampling efficiently to ensure rapid acceptance; we explain the details of this in the Appendix. Once we have accepted a pair (μ, τ) , we simulate $\xi_{t+1} \sim N(\mu, \tau^{-1})$, and update $(\bar{\xi}_t, S_{\xi\xi}(t))$ using the new value ξ_{t+1} .

One further modification of the particle-filtering algorithm is needed, and this concerns θ . Under the Markovian law of X , the θ -components *never change*. An undesirable consequence of this is that once the collection of particles at time 0 has been chosen, the set of possible θ -values will never change, and so at all later time the particle posterior for θ will be restricted to this set. If this set did not put enough points in the places where the data is telling us θ should be, the particle-filter estimates of θ will be poor.

To counteract this, we propose to replace (14) with

$$\hat{\pi}_t = \sum_{i=1}^N w_t^i q_\varepsilon(x_t^i, \cdot) \quad (16)$$

where q_ε is a Markov transition function creating a small perturbation of the initial state. This is in some sense a *simulated annealing* step. Of course, this introduces a bias (we are now approximating $\pi_t q_\varepsilon$, not π_t), but it is essential to do this, particularly as the dimension of θ is quite large. In the absence of any observations, this would cause the parameter values to evolve according to some (slow) Markov process, which is not what constant parameters should do! Something has to be done to correct this; (Liu and West, 2001), for example, propose that there should be some shrinkage of the parameter values to the population mean. What we

²⁴The value of K_0 can be thought of as the factor by which the annual precision of ξ values gets scaled; in effect, if we were to observe for M years, and use that data to form our initial estimate of the precision, then the precision would be $M\tau$. It seems reasonable then that the Bayesian agent would be basing his prior for τ on somewhere between 5 and 80 years of data.

have done is firstly to use only random perturbations which keep the mean unchanged; once we have gone through the algorithm with these random perturbations, we can then approximate the posterior discovered in this way to choose an initial set of particles which put most θ -values in the most likely places, and then run the analysis again, but this time permitting no change in the parameter values.

Once again, the exact details of how the perturbations (16) were performed are given in the Appendix.

4 Results.

We present here the results of the analysis. We performed a number of different runs, with various parameters set to different values, to test the sensitivity of the conclusions to the possible inputs, and to find situations where the fit was good. A fit is good if we find that the posterior distribution of the v_i is small, ideally a lot smaller than the values obtained by assuming the Y_i are IID; this means that we are relying little on the observational errors to reconcile the model values and the observed values. In each run, we fixed the value of c ; its numerical value was in the range 5×10^5 to 10^7 . Other parameters were chosen as described in the previous section. Many different runs were made to build up a picture; the resulting posteriors were generally quite similar.

In Figure 2, we present the posterior densities of the eight parameters β , R , α_0 , b_0 , K_0 , v_1 , v_2 , v_3 . These results came out of a run with 25000 particles, $RSLEVEL = 0.5$, $c = 10^7$, prior mean for R equal to 3, prior means for the v_i set at 5, 10 and 100 bp, and the parameters ε_i used for the shaking (see the Appendix) set at 10^{-4} .

The posteriors of the agent's preference parameters, β and R , are presented in Figure 3 and 4. The interquartile ranges [94.79%, 94.95%] for β , [1.62, 1.71] for R , are completely plausible values for these parameters.

The posteriors for the v_i are informative, and deserve further comment, as they allow us to assess the effectiveness of the model in explaining the data. If we simply supposed that the Y_t were independent Gaussian variables, then fitting the mean would leave a residual sum-of-squares vector of 10^{-4} times (12.86, 35.04, 273.62). From the posterior plots for the v_i , we see that the residual sum-of-squares from fitting the model is of the order of 10^{-4} times (1, 24, 200). The residual sum-of-squares of all three is substantially reduced; the quality of the fit is therefore improved. We should not however expect much better than this from fitting a model that is time-homogeneous, because a casual inspection of the data suggests that there have been substantial variations in the volatility of the three time series over the 90-year span of the data.

5 Conclusions.

The equity premium puzzle arises from a failure to account for the natural uncertainty about structural growth parameters and especially the rate of return, whose point estimation is well known to be extremely imprecise. Inserting the full posterior distribution of growth parameters (and not just their point estimates) into our formulae for equity and bond prices leads us to a standard Bayesian calibration procedure based upon particle filtering techniques. The observed equity premium is shown to be reconciled with the neoclassical paradigm of a representative agent, to whom both the mean and the volatility of the growth rate are unknown. No implausibly large values for the coefficient of relative risk-aversion are needed and the equity premium puzzle is therefore solved in a fully Bayesian way.

Table 1: Summary values of the posterior.

Parameter	25 %ile	Median	75 %ile
β	0.9479	0.9487	0.9497
R	1.6263	1.6618	1.7153
α_0	2.5736	2.7362	2.8738
b_0	0.07169	0.07715	0.08211
K_0	65.219	69.441	75.013
v_1	0.94 bp	1.03 bp	1.09 bp
v_2	24.40 bp	25.67 bp	27.24 bp
v_3	188.43 bp	200.67 bp	210.54 bp

A Appendix

A.1 Analysis of F .

For the calculation of share and bond prices, we introduced the function F defined as an integral, and were able to express the prices in terms of it. In order to implement the algorithm effectively, we need to make sure that the evaluation of F is done as rapidly as possible consistently with accuracy. For this, we need to understand what part of the (infinite) region of integration really matters, and restrict function calls to that region.

In order to understand this, we develop the integral defining F :

$$\begin{aligned}
F_t(\lambda, \rho) &\equiv \int_0^\infty d\tau \int_{\frac{1}{\nu} \log \beta + \frac{\nu}{2\tau}}^\infty d\mu \frac{\pi_t(\mu, \tau)}{1 - \beta e^{-\nu\mu + \nu^2/2\tau}} \exp(-\lambda\mu + \lambda^2/2\tau - \rho\mu + \rho^2/2\tau) \\
&= \int_0^\infty d\tau \tau^{\alpha_t-1} \exp\left[-b_t\tau + (\lambda^2 + \rho^2)/2\tau - \frac{c}{2\tau^2}\right] \\
&\quad \int_{\frac{1}{\nu} \log \beta + \frac{\nu}{2\tau}}^\infty d\mu \exp\left\{-\frac{1}{2} K_t\tau(\mu - m_t)^2 - (\lambda + \rho)\mu\right\} \sqrt{\frac{K_t\tau}{2\pi}} \\
&= \int_0^\infty d\tau \tau^{\alpha_t-1} \exp\left[-b_t\tau + (\lambda^2 + \rho^2)/2\tau - \frac{c}{2\tau^2}\right] e^{-(\lambda+\rho)m_t} \\
&\quad \int_{\frac{1}{\nu} \log \beta + \frac{\nu}{2\tau}}^\infty d\mu \exp\left\{-\frac{1}{2} K_t\tau\left(\mu - m_t + \frac{\lambda + \rho}{K_t\tau}\right)^2 + \frac{1}{2} \frac{(\lambda + \rho)^2}{K_t\tau}\right\} \sqrt{\frac{K_t\tau}{2\pi}} \\
&= e^{-(\lambda+\rho)m_t} \int_0^\infty d\tau \tau^{\alpha_t-1} \exp\left[-b_t\tau + (\lambda^2 + \rho^2)/2\tau + \frac{1}{2} \frac{(\lambda + \rho)^2}{K_t\tau} - \frac{c}{2\tau^2}\right] \\
&\quad \bar{\Phi}\left(\left(\frac{1}{\nu} \log \beta - m_t + \frac{\nu}{2\tau} + \frac{\lambda + \rho}{K_t\tau}\right) \sqrt{K_t\tau}\right) \\
&= e^{-(\lambda+\rho)m_t} \int_0^\infty dz z^{-\alpha_t-1} \exp\left[-\frac{b_t}{z} + \frac{\lambda^2 + \rho^2}{2} z + \frac{(\lambda + \rho)^2}{2K_t} z - \frac{cz^2}{2}\right] \\
&\quad \bar{\Phi}\left(\left(\frac{1}{\nu} \log \beta - m_t + \left(\frac{\nu}{2} + \frac{\lambda + \rho}{K_t}\right)z\right) \sqrt{\frac{K_t}{z}}\right).
\end{aligned}$$

Here, as usual, $\bar{\Phi}$ is the tail of the standard Gaussian distribution. Notice that although F_t is not available in closed form, its evaluation requires only the computation of a single univariate integral, so this can be done rapidly numerically. The integral defining F_t is guaranteed to converge, because of the term $\exp(-c/\tau^2)$ which controls the integrand near $\tau = 0$. However, the exact choice of c is relatively unimportant, as we see when we study the integrand. It is

well known that the function $\tilde{\Phi}(x) \equiv e^{x^2/2}\bar{\Phi}(x)$ varies as $x^{-1}(2\pi)^{-1/2}$ as $x \rightarrow \infty$, while $\bar{\Phi}(x)$ is $O(1)$ for negative values of x .

Simplifying the notation by dropping the t -subscripts, and setting $a_0 = \nu^{-1} \log \beta - m$, $a_1 = \nu/2 + (\lambda + \rho)/K$, $D = (\lambda^2 + \rho^2)/2 + (\lambda + \rho)^2/2K$, we see that the integral is more simply expressed as

$$I \equiv \int_0^\infty z^{-\alpha-1} \exp \left[-\frac{b}{z} + Dz - \frac{cz^2}{2} \right] \bar{\Phi}((a_0 + a_1z)\sqrt{K/z}) dz. \quad (17)$$

Using the asymptotic for $\bar{\Phi}$, the integrand (for large z) is in effect the exponential of

$$-\frac{b}{z} + Dz - \frac{cz^2}{2} - \frac{1}{2}(a_0 + a_1z)^2 \frac{K}{z}.$$

Now as t gets larger, $K_t = K_0 + t$ gets larger, and so we see that we do not in fact *need* the convergence term $-cz^2/2$, once t has got so large that $\gamma^2 K_t/2 > D_t$.

Numerical evaluation of the integral (17) requires some care. The first point that can be problematic is that the integrand may be very small everywhere, so a numerical integration routine quickly decides that the integral is zero; but this is no use for (11), (12). It is clear what must be done; the integrands must be scaled appropriately. In order to do this, we have to know (at least approximately) what the maximum value of the integrand is. For this, we notice that (with $Z \equiv (a_0 + a_1z)\sqrt{K/z}$) we have approximately either $\bar{\Phi}(Z) = \exp(-Z^2/2)/Z\sqrt{2\pi}$ or $\bar{\Phi}(Z) = 1$, the first eventuality obtaining if Z is reasonably large. Since Z gets large if either z or z^{-1} is large, we have three regimes to consider, where $s = 0, 1, -1$, and the log of the integrand is approximated by

$$-(1 + \alpha + s/2) \log(z) - b/z + Dz - cz^2/2 - \frac{s^2 K}{2z} (a_0 + a_1z)^2. \quad (18)$$

Here, $s = 0$ corresponds to $\bar{\Phi} \simeq 1$, $s = 1$ corresponds to the case of large Z , large z , and the case $s = -1$ corresponds to the case of large Z , small z . To find the maximum of the function (18), we differentiate to obtain the cubic equation

$$cz^3 - (D - s^2 K a_1^2/2) z^2 + (1 + \alpha + s/2)z - (b + s^2 K a_0^2/2) = 0.$$

The calculation now finds the real positive roots of these three cubics, evaluates the log integrand at those points, and picks the place where the integrand is maximal.

A.2 Details of the rejection sampling.

The simulation methodology requires us to sample the pair (μ, τ) from the posterior (9) before simulating the next ξ from a $N(\mu, \tau^{-1})$ distribution. This is not quite as easy as it might at first sight appear; rejection-sampling is the natural methodology to use here, but we have to be careful not to pick a proposal distribution that will result in billions of proposals before an acceptance.

The first step is to sample a value of τ according to a density which is proportional to²⁵

$$\tau^{\alpha-1} \exp(-c/2\tau^2 - b\tau). \quad (19)$$

What we shall do is sample a value of $Z \equiv 1/\tau$ from a density proportional to

$$e^{\varphi_0(z)} = z^{-1-\alpha} \exp(-cz^2/2 - b/z). \quad (20)$$

We shall do this by proposing a value for Z from a $N(a, c^{-1})$ distribution, and then doing rejection sampling. This requires us to find a constant B such that

$$e^{\varphi_0(z)} \leq B \exp(-c(z-a)^2/2) \equiv B e^{\varphi_1(z)}$$

for all z , and then accepting the proposed value of Z with probability $\exp(\varphi_0(z) - \varphi_1(z))/B$. Clearly, the smaller the value of B , the better the chances of acceptance. The question we have to answer is how to choose a . Holding a fixed, and maximising $\varphi_0(z) - \varphi_1(z)$ over z leads to the equation²⁶

$$acz^2 + Az - b = 0 \quad (21)$$

to determine the value z_* where the maximum is attained. The difference $\Delta(z_*) \equiv \varphi_0(z_*) - \varphi_1(z_*)$ must now be minimised over choice of a . We may equivalently think of it as a minimisation over z_* , where a is expressed as a function of z_* using (21). Solving $\Delta'(z) = 0$ leads to a quartic equation, one of whose roots is $z = 2b/A$; this root corresponds to a value $a = -A^2/4bc$, with negative second derivative, so Δ achieves a local maximum there. Factoring out this root, the equation $\Delta'(z) = 0$ becomes the cubic equation

$$cz^3 + Az - b = 0, \quad (22)$$

which has a unique positive root given by Cardan's formula:

$$z_* = Y/6c - 2A/Y, \quad (23)$$

where

$$Y \equiv \left(c^2(108b + 12\sqrt{3/c}\sqrt{4A^3 + 27cb^2}) \right)^{1/3}. \quad (24)$$

The second derivative Δ'' is positive at z_* , so we have a minimum.

Having found τ , the second stage of the simulation is to generate a value of μ from a density proportional to

$$(1 - \beta e^{-\nu\mu + \nu^2/2\tau})^+ \exp\left[-\frac{1}{2} K\tau(\mu - m)^2\right]. \quad (25)$$

If we define

$$b^* \equiv \frac{1}{\nu} \log \beta + \frac{\nu}{2\tau},$$

²⁵For this discussion, we omit the subscript t .

²⁶For brevity, we set $1 + \alpha \equiv A$ for this discussion.

then the prefactor in (25) is positive for $\mu > b^*$. Two cases then arise, the first (simpler) being when $m > b^*$, for then we just propose a value of μ from a $N(m, (K\tau)^{-1})$ distribution, and accept with probability $(1 - \beta e^{-\nu\mu + \nu^2/2\tau})^+$. More interesting is the situation where $b^* > m$, for then we have to sample out in the tail of a Gaussian density, and this requires some ingenuity. What we shall do is to propose a value of μ from the density

$$(\mu - m)\sqrt{K\tau} \exp\left[-\frac{1}{2}K\tau(\mu - m)^2\right], \quad (26)$$

and then accept with probability

$$\frac{(1 - \beta e^{-\nu\mu + \nu^2/2\tau})^+}{\gamma(\mu - m)} \equiv \frac{(1 - e^{-\nu(\mu - b^*)})^+}{\gamma(\mu - m)}, \quad (27)$$

where $\gamma = \min\{\nu, (b^* - m)^{-1}\}$ is chosen so that the acceptance probability is always at most 1. How do we simulate a random variable X from density (26)? Notice that if X has the density (26), then $Y \equiv (X - m)^2 K\tau/2$ has a standard exponential law, so what we shall do is generate a standard exponential variable η , set²⁷ $Y = \eta + (b^* - m)^2 K\tau/2$, and then recover our proposed value for μ as

$$\mu = m + \sqrt{\frac{2Y}{K\tau}}.$$

A.3 Perturbing the parameters.

The perturbation of the different components of our vector $\theta = (\beta, R, \alpha_0, b_0, K_0, v_1, v_2, v_3)^T$ of parameters has to be handled in different ways.

- (i) Given the current value of β , the perturbed value β' is chosen from a $B(a, b)$ distribution, where $a = \beta/\varepsilon_\beta$, $b = (1 - \beta)/\varepsilon_\beta$, where ε_β is a small parameter of the calibration. The mean of β' is β , the variance is $ab/(a + b)^2(a + b + 1)$, which is of the order of ε_β .
- (ii) The current value of R gets perturbed to $R' = 1 + Z(R - 1)$, where the random variable Z is expressed as

$$Z = \frac{X}{1 - X} \quad (28)$$

for some $B(a, a + 1)$ random variable X . It is easily verified that Z has mean 1 and variance $2/(a - 1)$. We set $a = 1/\varepsilon_R$, where ε_R is a small parameter of the calibration.

- (iii) For the remaining components θ_i of the parameter vector, the perturbation is done in the same way, namely to set $\theta'_i = Z\theta_i$, where the random variable Z is expressed as (28) for some $B(a, a + 1)$ random variable X . We set $a = 1/\varepsilon_i$, where ε_i is a small parameter of the calibration.

²⁷We are only going to accept a value for μ if it exceeds b^* , equivalently, if Y exceeds $(b^* - m)^2 K\tau/2$, so we shall condition Y to exceed that value; of course, the overshoot of Y over that value is again standard exponential.

References

- Arrow, K. (1971). Essays in the theory of risk-bearing. *North-Holland, Amsterdam*.
- Arumpalam, S., Maskell, S., Gordon, N., and Clapp, T. (2001). A tutorial on particle filters for on-line non-linear/non-gaussian bayesian tracking. *IEEE Transactions on Signal Processing*, XX:100–117.
- Bansal, R. and Coleman, J. W. (1996). A monetary explanation of the equity premium, term premium and risk free rate puzzles. *Journal of Political Economy*, 104:1135–1171.
- Barro, R. J. (2005). Rare events and the equity premium. Harvard University preprint.
- Benninga, S. and Protopapadakis, A. (1990). Leverage, time preference and the ‘equity premium puzzle’ . *Journal of Monetary Economics*, 25:49–58.
- Breeden, D. (1979). An intertemporal asset pricing model with stochastic consumption and investment opportunities. *Journal of Financial Economics*, 7:265–296.
- Brennan, M. J. and Xia, Y. (2001). Stock price volatility and equity premium. *Journal of Monetary Economics*, 47:249–283.
- Brown, S., Goetzmann, W., and Ross, S. (1995). Survival. *Journal of Finance*, 50:853–873.
- Campbell, J. Y. and Cochrane, J. H. (1999). By force of habit: a consumption-based explanation of aggregate stock market behaviour. *Journal of Political Economy*, 107:205–251.
- Cecchetti, S. G., Lam, P.-S., and Mark, N. C. (1993). The equity premium and the risk-free rate: matching the moments. *Journal of Monetary Economics*, 31:21–45.
- Constantinides, G. (1990). Habit formation: a resolution of the equity premium puzzle. *Journal of Political Economy*, 98:519–543.
- Constantinides, G. and Duffie, D. (1996). Asset pricing with heterogeneous consumers. *Journal of Political Economy*, 104:219–240.
- Constantinides, G. M., Donaldson, J. B., and Mehra, R. (2002). Junior can’t borrow: a new perspective on the equity premium puzzle. *Quarterly Journal of Economics*, 118:269–296.
- Crisan, D. and Doucet, A. (2002). A survey of convergence results on particle filtering methods for practitioners. *IEEE Transactions on Signal Processing*, 50:736–746.
- Doucet, A., De Freitas, N., and Gordon, N., editors (2001). *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York.
- Epstein, L. and Zin, S. (1991). Substitution, risk aversion and the temporal behavior of consumption and asset returns: an empirical analysis. *Journal of Political Economy*, 99:263–286.
- Ferson, W. E. and Constantinides, G. M. (1991). Habit persistence and durability in aggregate consumption. *Journal of Financial Economics*, 29:199–240.
- Friend, I. and Blume, M. (1975). The demand for risky assets. *American Economic Review*, 65:900–922.

- Geweke, J. (2001). A note on some limitations of crra utility. *Economics Letters*, 71:341–345.
- Grossman, S. J. and Shiller, R. J. (1981). The determinants of the variability of stock market prices. *The American Economic Review*, 71:222–227.
- Heaton, J. and Lucas, D. J. (1996). Evaluating the effects of incomplete markets on risk sharing and asset pricing. *Journal of Political Economy*, 104:443–487.
- Kandel, S. and Stambaugh, R. F. (1991). Asset returns and intertemporal preferences. *Journal of Monetary Economics*, 27:39–71.
- Kocherlakota, N. (1996). The Equity Premium: It’s Still a Puzzle. *Journal of Economic Literature*, 34:42–71.
- Kydland, F. and Prescott, E. (1982). Time to build and aggregate fluctuations. *Econometrica*, 50:1345–1370.
- Liu, J. and West, M. (2001). Combined parameter and state estimation in simulation-based filtering. *in Sequential Monte Carlo Methods in Practice*.
- Lucas, D. J. (1994). Asset pricing with undiversifiable risk and short sales constraints: deepening the equity premium puzzle. *Journal of Monetary Economics*, 34:325–341.
- Lucas, R. (1978). Asset prices in an exchange economy. *Econometrica*, 46:1429–1445.
- Mankiw, N. G. (1986). The equity premium and the concentration of aggregate stocks. *Journal of Financial Economics*, 17:211–219.
- Mankiw, N. G. and Zeldes, S. P. (1991). The consumption of stockholders and nonstockholders. *Journal of Financial Economics*, 29:97–112.
- McGrattan, E. R. and Prescott, E. C. (2001). Taxes, regulations, and asset prices. Technical Report 8623, National Bureau for Economic Research.
- Mehra, R. and Prescott, E. (1985). The Equity Premium: A Puzzle. *Journal of Monetary Economics*, 15:145–161.
- Mehra, R. and Prescott, E. C. (1988). The equity premium: a solution? *Journal of Monetary Economics*, 22:133–136.
- Mehra, R. and Prescott, E. C. (2003). The equity premium in retrospect. *Handbook of the Economics of Finance*.
- Rietz, T. A. (1988). The equity risk premium: a solution. *Journal of Monetary Economics*, 22:117–131.
- Tsionas, E. G. (2005). Likelihood evidence on the asset returns puzzle. *Review of Economic Studies*, 72:917–946.
- Veronesi, P. (1999). Stock market overreaction to bad news in good times: a rational expectations equilibrium model. *Review of Financial Studies*, 12:975–1007.
- Veronesi, P. (2000). How does information quality affect stock returns? *Journal of Finance*, 55:807–837.

- Weil, P. (1989). The Equity Premium Puzzle and the Risk-Free Rate Puzzle. *Journal of Monetary Economics*, 24:401–421.
- Weitzman, M. (2005). A unified Bayesian theory of equity puzzles. *Preprint, Harvard University*.

Data from Mehra & Prescott

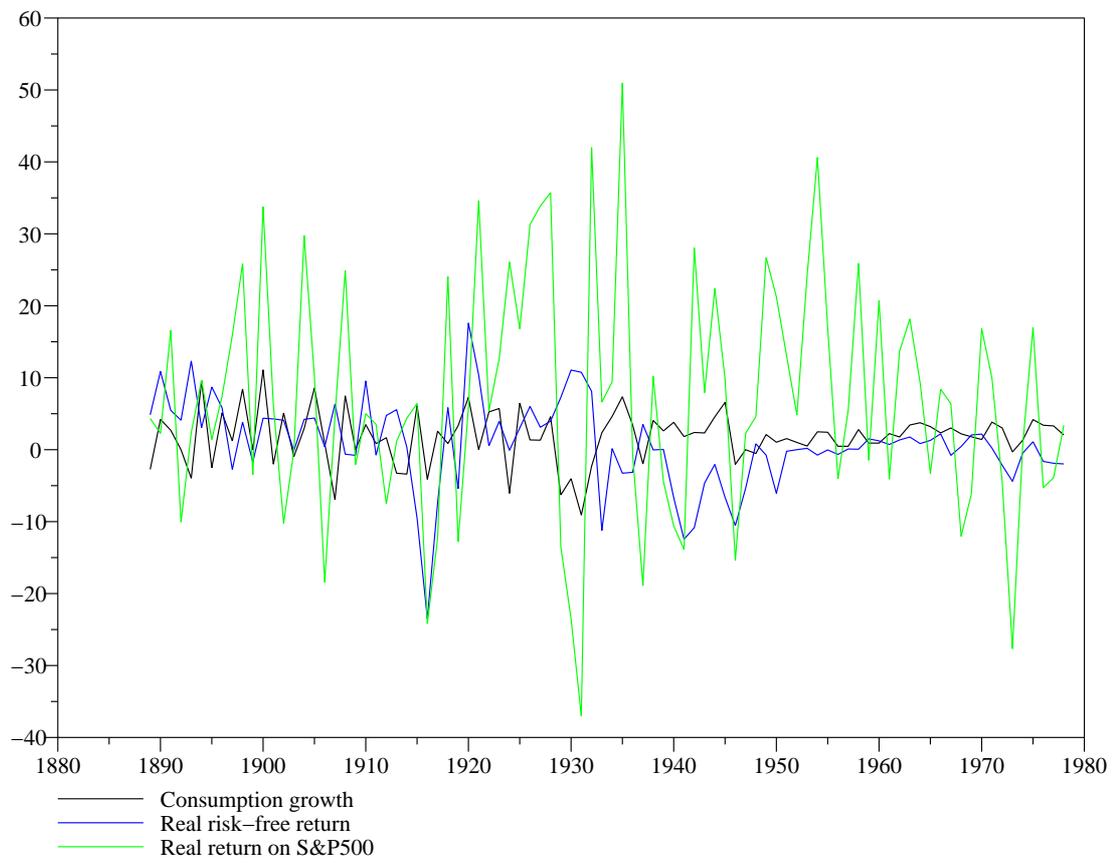
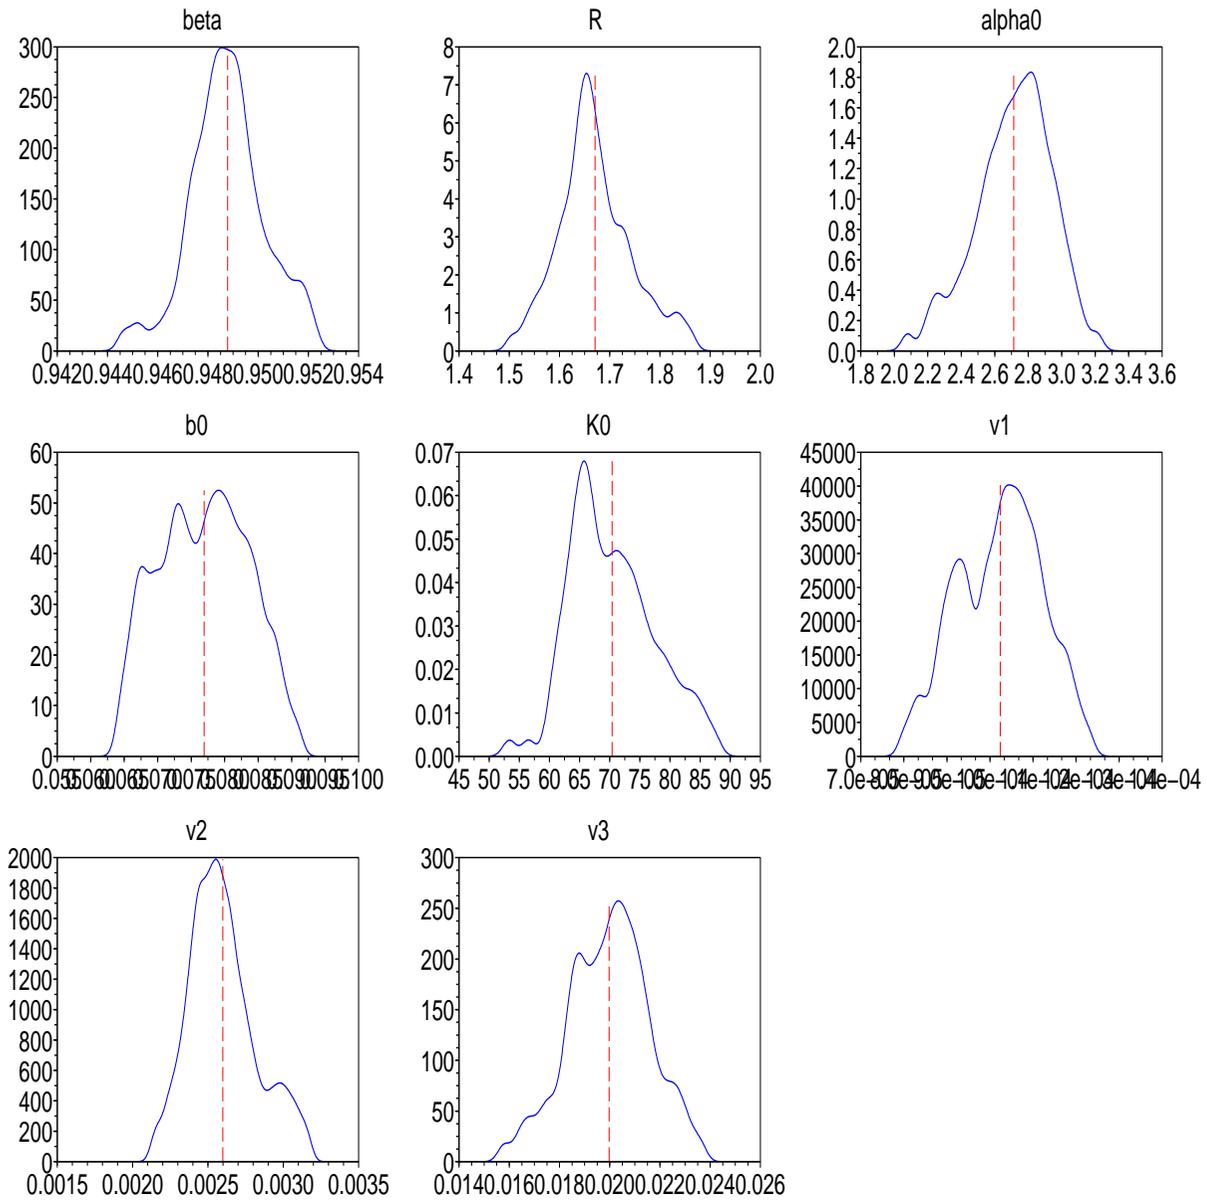


Figure 1: Growth rate of real consumption, real return on US Treasury Bill and real return on the SP500 index over the period 1889-1978 (all expressed in %).

Figure 2: Posterior density for all the parameters



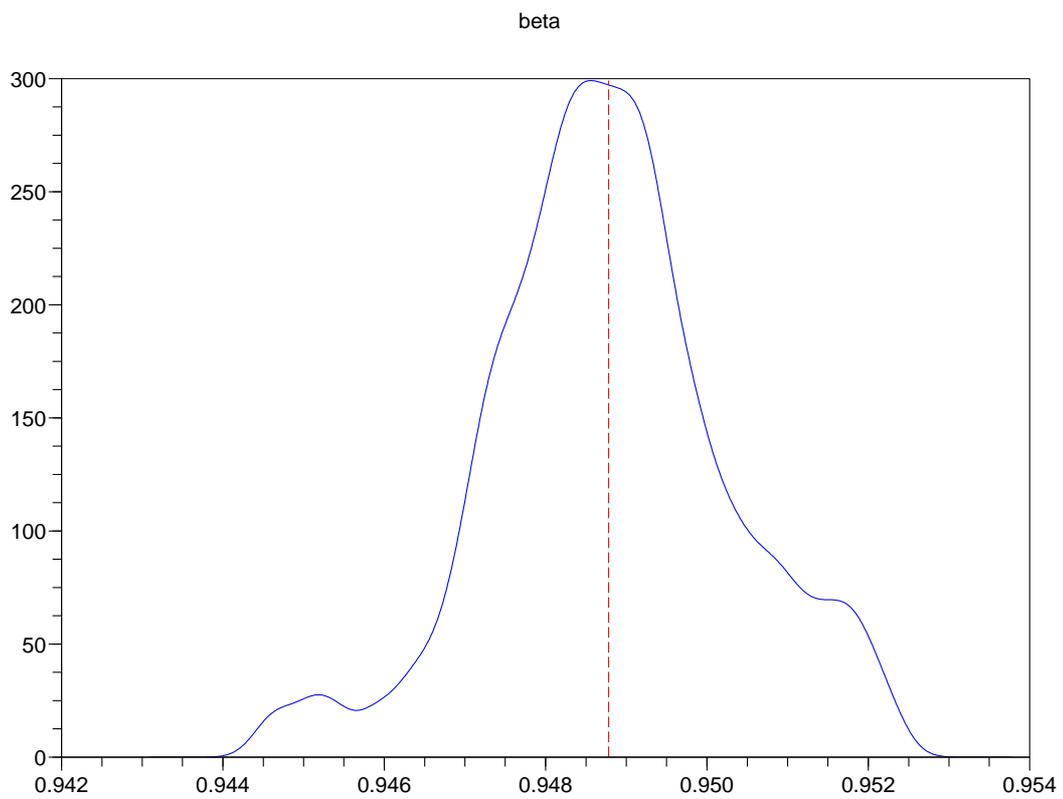


Figure 3: Posterior density for β

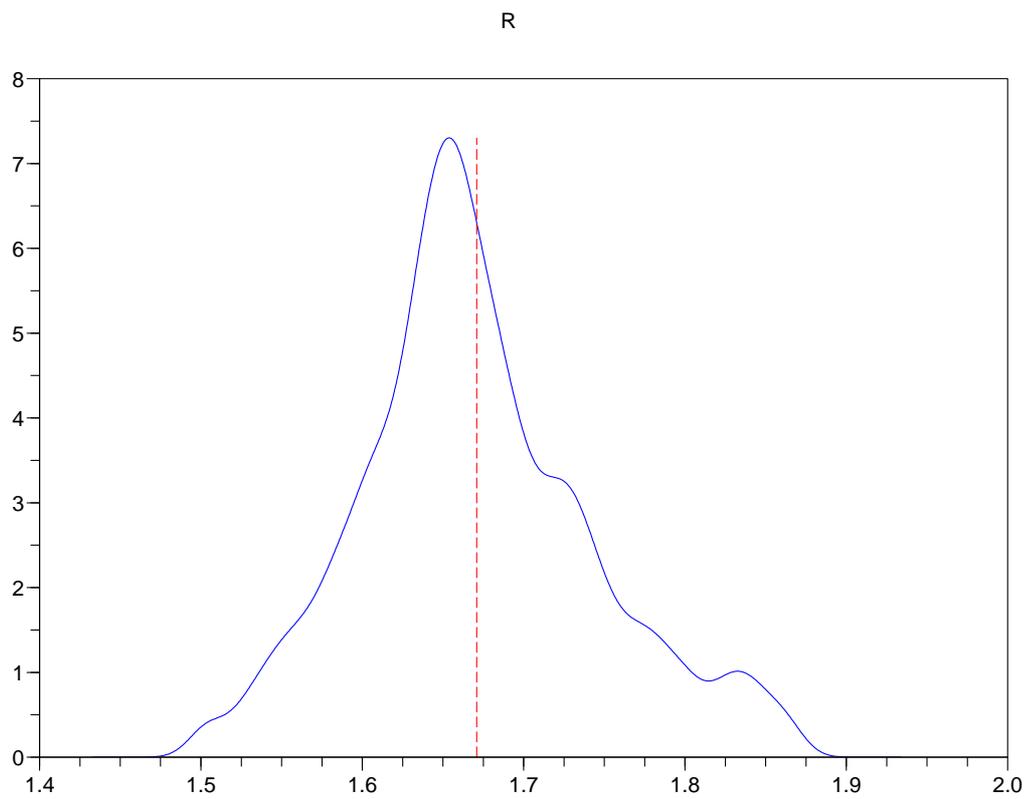


Figure 4: Posterior density for R