

A Bayesian solution to the equity premium puzzle.

A. Jobert
Statistical Laboratory
Wilberforce Road
Cambridge CB3 0WB, GB
A.Jobert@statslab.cam.ac.uk

A. Platania
Statistical Laboratory
Wilberforce Road
Cambridge CB3 0WB, GB
A.Platania@statslab.cam.ac.uk

L.C.G. Rogers
Statistical Laboratory
Wilberforce Road
Cambridge CB3 0WB, GB
L.C.G.Rogers@statslab.cam.ac.uk
phone: +44 1223 766806

May 9, 2005

Abstract

This paper describes a Bayesian solution to the equity premium puzzle, the inability of standard intertemporal economic models to account for the magnitude of the observed historical US equity premium, i.e. the return earned by a risky security in excess of that earned by a relatively risk-free US Treasury bill. A single-representative agent is assumed to have a prior distribution about the rate of return and the variance of the growth rate of the underlying economy. Standard Bayesian calculations are carried out, which yield to an appealing solution to the equity puzzle formulated by Mehra and Prescott.

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 (?) to fail to capture the magnitude of the equity premium. The observed 6.9 % premium was indeed much greater than could be explained by neoclassical representative agent paradigms ¹ 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* (?), it simply indicates that we do not know why people save even when bond returns are low ².

The equity premium puzzle is illustrated by Mehra and Prescott (?) in the simple case of lognormal growth rates of consumption and dividends, 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 large ³, a high equity premium is impossible.

Over the last fifteen years, a huge literature has been attempting to solve for both the equity premium and the risk-free rate puzzles. Kocherlakota (?) provides an excellent survey of those various attempts. A first attempt to solve equity puzzles advocated the use of alter-

¹See among others (?), (?). 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-Breeden 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.

²Although 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.

³Fischer Black proposed that a coefficient of risk-aversion taken to be equal to 55 would solve the puzzle; empirical studies (see among others (?; ?; ?), however, provide evidence for such a coefficient being no more than 10.

native preference structures. We will sketch here very briefly the main ideas underlying the latter research, since they already give a pretty good picture of the scope of the solutions proposed till recently: although there are several possible explanations for the low level of Treasury returns, the large equity premium is still an open puzzle.

An important restriction imposed by CRRA preferences is that the coefficient of relative risk-aversion is intrinsically linked to the elasticity of intertemporal substitution (one is the reciprocal of the other). This implies that if an individual is averse to variation of consumption in different states at a particular point in time (i.e. he dislikes risk), he will be averse to consumption variation over time (i.e. he will dislike growth). Given that the large equity premium implies that investors in the Mehra-Prescott setup are highly risk-averse, CRRA preferences would then imply that they want to smooth consumption over time and have little incentive to save. The demand for bonds is low and the risk-free rate derived from this model can be therefore easily understood to be too high. Epstein and Zin (?) consider a recursive utility structure which enables to disentangle the coefficient of relative risk-aversion and the elasticity of intertemporal substitution. This can potentially resolve the risk-free rate puzzle.

Constantinides (?) modified standard preferences by incorporating habit formation and allowing utility to be affected not only by current consumption but also by past consumption. The induced aversion to consumption risk increases the demand for bonds and reduces thereby the risk-free rate. Such an approach can therefore eliminate the risk-free rate puzzle. Some other work includes market incompleteness (?) or market imperfections such as transaction costs (?). As explained by Mehra and Prescott (?), none of the latter attempts provides a satisfactory answer to the equity premium puzzle.

Let us now motivate our Bayesian approach to these equity puzzles. In the Mehra-Prescott setup, the agent is assumed to know with certainty the distributional properties of the growth rate of dividends, but this assumption would be hard to defend in the face of the known imprecision in estimating rates of return. An example makes the point.

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 overlooks 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 estimates.

While this study was in progress, we learned of a paper by Weitzman (?) proposing a similar explanation of the equity puzzles, based upon a Bayesian approach. He is taking for his growth parameters a normal-gamma prior, which has to be truncated at low precision in order for the expectations defining the equity to be finite. He then argues that some very similar values of δ can be used to solve both the equity premium and the risk-free rate puzzles, at least in mean.

We have attempted to reproduce Weitzman's calculations, but this is made difficult by the fact that he does not state certain numerical values required. His method for dealing with the problem of small precisions (large variances) is to take a prior for the precision which is zero to the left of some cut-off value δ . He argues correctly that as $\delta \downarrow 0$ the equity premium increases to infinity, so any given value could be achieved; he then claims that very similar values of δ can be found which match empirically observed risk premia and risk-free rates. However, taking⁴ $m = 10$ and $R = 2$, and using the data from Mehra & Prescott, we find that his function $\Pi(\delta)$ for the equity premium takes the numerical value 0.0024345 for δ in the range $[0.003, 300]$, and yet at 0.0028 its value is in excess of 10^{304} ! Now the 'consensus' value 0.0024345 is much too low for the equity premium values observed in practice, which Weitzman quotes as 0.045, and yet in some small range for δ this desired value will be achieved! What Weitzman is doing is to rely on the numerical explosion of Π to match point estimates of the equity premium and risk-free rate; it is therefore not so surprising that similar values of δ will match the values estimated from the data, as this explosion will occur in very much the same place for all the functions he is considering. The numerical analysis in this region is actually quite delicate, and we believe that it would be very hard to compute any of these values reliably. It seems therefore that the analysis in Weitzman's paper is not conclusive.

Moreover, the methodology proposed, to appeal to a Bayesian argument to explain why the equity premium puzzle might not exist, but then to try to match point estimates in the style of a frequentist, is a curious hybrid; having taken up Bayesian arguments, why not follow them through to their conclusion? This is what we do in this paper.

In any case, the emphasis Weitzman attaches to the truncation level δ is rather unnatural to us. We should be finding that the action focuses on the mean and not the volatility of the rate of return, since we know that the mean is the parameter whose point estimate is most imprecise.

We are applying 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 both the equity premium and the risk-free rate puzzles.

⁴We use Weitzman's notation for this discussion

1 A Bayesian approach to the equity puzzle.

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

$$\mathbf{E}_0 \left[\sum_{t=0}^{\infty} \beta^t U(c_t) \right], \quad (1)$$

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), 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 equity share with price S_t is competitively traded and is viewed as a claim on the dividend process y_t . Let us assume that the logarithmic growth rates of the underlying economy, i.e. $\xi_t = \log(y_t) - \log(y_{t-1})$, are independent and identically distributed as $\mathcal{N}(\mu, \sigma^2) \equiv \mathcal{N}(\mu, \frac{1}{\tau})$. 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

$$S_t = \mathbf{E}_t \left[\sum_{j \geq 0} \beta^j \left(\frac{y_{t+j}}{y_t} \right)^{-R} y_{t+j} \right] \quad (3)$$

$$= y_t \mathbf{E}_t \left[\sum_{j \geq 0} \beta^j \left(\frac{y_{t+j}}{y_t} \right)^{1-R} \right] \quad (4)$$

$$= y_t + y_t \mathbf{E}_t \left[\sum_{j \geq 1} \beta^j \exp \left((1-R) \sum_{k=t+1}^{t+j} \xi_k \right) \right] \quad (5)$$

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

$$\beta \mathbf{E}_t \left(\frac{y_{t+1}}{y_t} \right)^{-R} = \beta \mathbf{E}_t \exp(-R\xi_{t+1}). \quad (6)$$

Now suppose that the agent has a prior a prior distribution for (μ, τ) with density proportional to

$$g(\tau) \exp \left[-\frac{1}{2} \tau \mu^2 - b\tau \right] \tau^{\alpha-1}$$

where b, α are positive parameters and $g(\tau)$ is a factor in the prior needed to ensure convergence of certain integrals. Weitzman uses $g(\tau) = \mathbf{I}_{[\delta, \infty)}(\tau)$ for some positive δ which needs to be 'tuned'; we use $g(\tau) = \exp(-\frac{c}{\tau^2})$ for some positive c which is fitted from data. In order to compute the expectations in (5) and (6), we now need to find the posterior distribution of (μ, τ) given (ξ_1, \dots, ξ_t) , namely $\pi(\mu, \tau | \xi_1, \dots, \xi_t)$. After observing ξ_1, \dots, ξ_t (assumed to be IID), the posterior density will be

$$\pi(\mu, \tau | \xi_1, \dots, \xi_t) \propto g(\tau) \exp\left[-\frac{1}{2}\tau\mu^2 - b\tau\right] \tau^{\alpha-1} \left(\frac{\tau}{2\pi}\right)^{\frac{t}{2}} \exp\left[-\frac{\tau}{2} \sum_{k=1}^t (\xi_k - \mu)^2\right] \quad (7)$$

$$= g(\tau) \sqrt{\tau} \exp\left[-\frac{1}{2}K'\tau(\mu - m)^2\right] \tau^{\alpha'-1} \exp[-b'\tau] \quad (8)$$

where $m = \frac{t\bar{\xi}_t}{t+1}$, $K' = 1 + t$, $\alpha' = \alpha + \frac{t}{2} - \frac{1}{2}$, $b' = b + \frac{1}{2}S_{\xi\xi} + \frac{1}{2}\frac{t}{t+1}\bar{\xi}_t^2$, $\bar{\xi}_t = \frac{1}{t} \sum_{k=1}^t \xi_k$ and $S_{\xi\xi} = \sum_{k=1}^t (\xi_k - \bar{\xi}_t)^2$.

To price equity according to (5), we must therefore compute

$$\mathbf{E}_t \left[\sum_{j \geq 1} \beta^j \exp\left((1-R) \sum_{k=t+1}^{t+j} \xi_k\right)\right] \quad (9)$$

$$= \mathbf{E}_t \left[\sum_{j \geq 1} \beta^j \exp\left((1-R)j\mu + \frac{1}{2}j \frac{(1-R)^2}{\tau}\right)\right] \quad (10)$$

$$= \mathbf{E}_t \left[\sum_{j \geq 1} \beta^j \exp\left((1-R)jm + \frac{1}{2}j^2 \frac{(1-R)^2}{K'\tau} + \frac{1}{2}j \frac{(1-R)^2}{\tau}\right)\right]. \quad (11)$$

For any $\tau > 0$ the series is divergent and the above expectation is infinite ... To get round the exploding infinities, we suppose that the valuation of equity is performed by a representative agent with finite rather than infinite horizon. The use of an infinite horizon is merely a mathematical convenience, to be discarded if it becomes mathematically inconvenient! We check our conclusions against different values of N , and find them virtually unaltered. The expression for the stock price therefore becomes

$$S_t = y_t \mathbf{E}_t \left[\sum_{j=0}^N \beta^j \left(\frac{y_{t+j}}{y_t}\right)^{1-R} \right] = y_t + y_t \sum_{j=1}^N \beta^j \mathbf{E}_t \left[\exp\left(\sum_{k=t+1}^{t+j} (1-R)\xi_k\right)\right]. \quad (12)$$

Using the posterior derived above,

$$\mathbf{E}_t \left[\exp\left(\sum_{k=t+1}^{t+j} (1-R)\xi_k\right)\right] = \mathbf{E}_t \left[\exp\left((1-R)jm + \frac{1}{2}j^2 \frac{(1-R)^2}{K'\tau} + \frac{1}{2}j \frac{(1-R)^2}{\tau}\right)\right] \quad (13)$$

$$\propto \exp((1-R)jm) \int_0^\infty \tau^{\alpha'-1} g(\tau) \exp\left[-b'\tau + \frac{c_j}{\tau}\right] d\tau. \quad (14)$$

where $c_j \equiv \frac{(1-R)^2}{2}j(1 + \frac{j}{R})$. To determine the normalising constant, note that if $R = 1$ this expression should be 1, so we get

$$\mathbf{E}_t \left[\exp \left(\sum_{k=t+1}^{t+j} (1-R)\xi_k \right) \right] = \exp((1-R)jm) \frac{\int_0^\infty \tau^{\alpha'-1} g(\tau) \exp[-b'\tau + \frac{c_j}{\tau}] d\tau}{\int_0^\infty \tau^{\alpha'-1} g(\tau) \exp[-b'\tau] d\tau}. \quad (15)$$

Summing over j gets us to the year- t stock price,

$$S_t = y_t \sum_{j=0}^N \beta^j \exp((1-R)jm) \frac{\int_0^\infty \tau^{\alpha'-1} g(\tau) \exp[-b'\tau + \frac{c_j}{\tau}] d\tau}{\int_0^\infty \tau^{\alpha'-1} g(\tau) \exp[-b'\tau] d\tau} \quad (16)$$

$$\equiv y_t F(t, \bar{\xi}_t, S_{\xi\xi}(t); \theta) \equiv y_t \exp(\psi_t(\theta)). \quad (17)$$

One point is worth emphasizing here. If the law of the ξ_t was known with certainty, then

$$\frac{S_t}{S_{t-1}} = \frac{y_t}{y_{t-1}} = \exp(\xi_t). \quad (18)$$

However, when there is uncertainty in the law of ξ_t , we have to modify the expression, in this case to

$$\frac{S_t}{S_{t-1}} = \frac{y_t}{y_{t-1}} \exp(\psi_t(\theta) - \psi_{t-1}(\theta)) = \exp(\xi_t + \psi_t(\theta) - \psi_{t-1}(\theta)). \quad (19)$$

This is a point that Weitzman seems to have overlooked. For the bond, we need to compute

$$\mathbf{E}_t \left[\left(\frac{y_{t+1}}{y_t} \right)^{-R} \right] = \mathbf{E}_t \exp(-R\xi_{t+1}), \quad (20)$$

which is like the preceding but with $j = 1$ and $1-R$ replaced by $-R$. Thus, if $\tilde{c} \equiv \frac{R^2}{2} (1 + \frac{1}{R})$, we get

$$\mathbf{E}_t \left[\left(\frac{y_{t+1}}{y_t} \right)^{-R} \right] = \exp(-Rm) \frac{\int_0^\infty \tau^{\alpha'-1} g(\tau) \exp[-b'\tau + \frac{\tilde{c}}{\tau}] d\tau}{\int_0^\infty \tau^{\alpha'-1} g(\tau) \exp[-b'\tau] d\tau}. \quad (21)$$

So the risk-free rate prevailing from time t to time $t+1$ is

$$\rho_t = m_t R - \log(\beta) + \log \left(\frac{\int_0^\infty \tau^{\alpha'-1} g(\tau) \exp[-b'\tau] d\tau}{\int_0^\infty \tau^{\alpha'-1} g(\tau) \exp[-b'\tau + \frac{\tilde{c}}{\tau}] d\tau} \right) \equiv G(t, \bar{\xi}_t, S_{\xi\xi}(t); \theta) \quad (22)$$

where $\theta^T \equiv (R, \alpha, b, c, \beta, v_1, v_2)$ is the parameter vector and v_i are the variances of the components of the observations. In the following section, the latter integrals (convergent for any positive values of our parameters) will be evaluated numerically using Gaussian quadrature ⁵.

⁵We used the 'qagiu' adaptive integration routine on infinite intervals from the GNU scientific library.

2 Calibration.

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

The observation $Y_t \equiv (Y_t^1, Y_t^2, Y_t^3)^T$ is just

$$(\xi_t, G(t, \bar{\xi}_t, S_{\xi\xi}(t); \theta), \xi_t + \psi_t(\theta) - \psi_{t-1}(\theta))^T \quad (23)$$

plus some noise.

The hidden state to be filtered in this example is just the parameter θ , which we shall allow to evolve (a little) according to the transition density ϕ . We refer the interested reader to (?; ?; ?) for a more detailed discussion of particle filtering techniques. If our approximation at time n to the posterior is

$$\sum_{i=1}^{N_p} w_n^i \delta_{z_n^i} \quad (24)$$

then the approximation to the posterior at time $n + 1$ will be proportional to

$$x \mapsto \sum_{i=1}^{N_p} w_n^i \phi(z_n^i, x) f(Y_{n+1}|x) \quad (25)$$

where we suppose that

$$f(Y_{n+1}|x) = \exp \left[-\frac{1}{2} \sum_{j=1}^2 \frac{(Y_{n+1}^j - \eta^j)^2}{v_j} \right] (2\pi)^{-1} (v_1 v_2)^{-\frac{1}{2}} \quad (26)$$

for concreteness (the noises do not have to be Gaussian of course). Here,

$$\eta = \begin{pmatrix} G(n+1, \bar{\xi}_{n+1}, S_{\xi\xi}(n+1); x) \\ \xi_{n+1} + \psi_{n+1}(\theta) - \psi_n(\theta) \end{pmatrix} \quad (27)$$

is the mean of the observations given the parameters.

Our main results can now be summarised by Figures 2 and 3, representing respectively the posterior distributions for the coefficient of relative risk-aversion R and the subjective discount factor β , for four different values of N .

Note firstly that the particular value of N does not really alter significantly our conclusions. N can be taken large enough to approximate realistically the finite time horizon of an individual. Most importantly, the posterior distribution for R tells us that it lies between 1 and

5 with a 95 % probability and a peak around 2.5. This is consistent with our intuition of what a plausible coefficient of risk-aversion should be equal to. There is therefore no need to take an implausibly large coefficient of risk-aversion to account for the historically observed equity premium and the puzzle as initially formulated by Mehra and Prescott is explained. As expected, the posterior discount factor can additionally be seen to be close to 1.

We finally provide in Figure 4 the plots for the posteriors of τ and μ . Not surprisingly, the posterior for the mean rate of return remains significantly more spread than the one we obtain for the precision.

3 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.

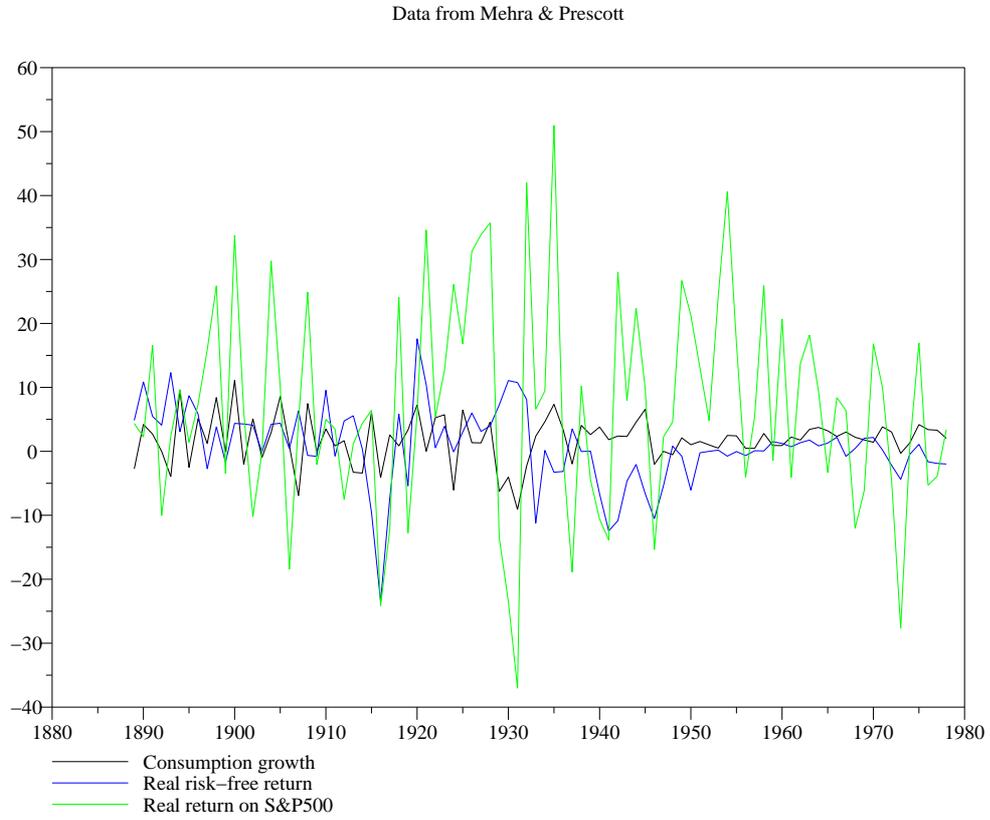


Figure 1: Growth rate of real consumption, real return on US Treasury Bill and real return on the SP500 index over the period 1889-1978 (all expressed in %).

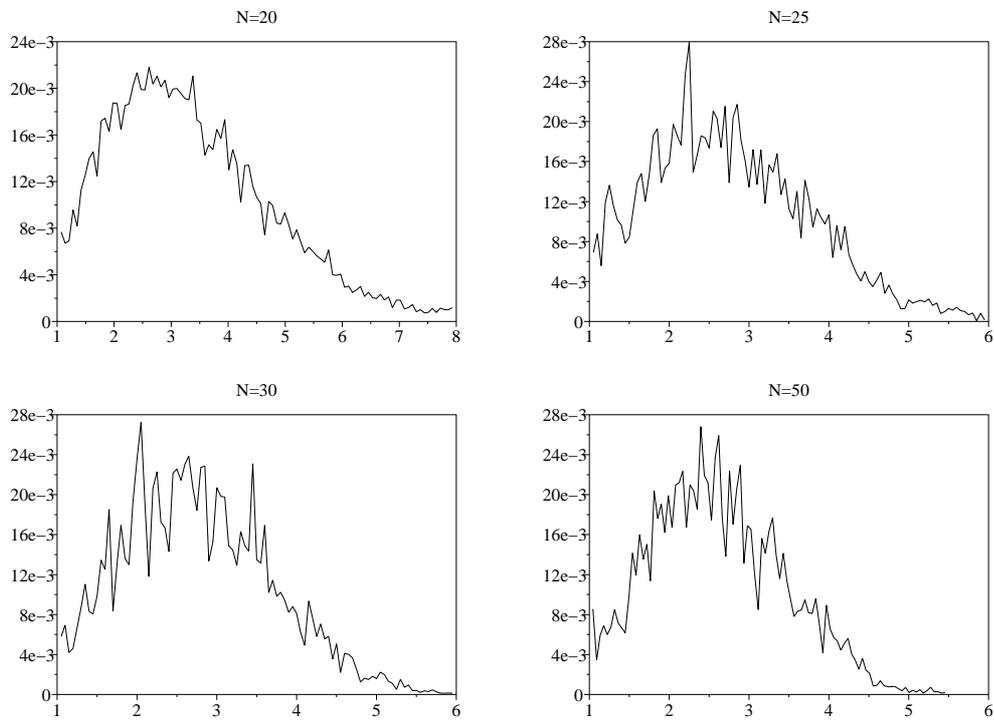


Figure 2: Posterior distribution of the coefficient of relative risk-aversion for different time-horizons.

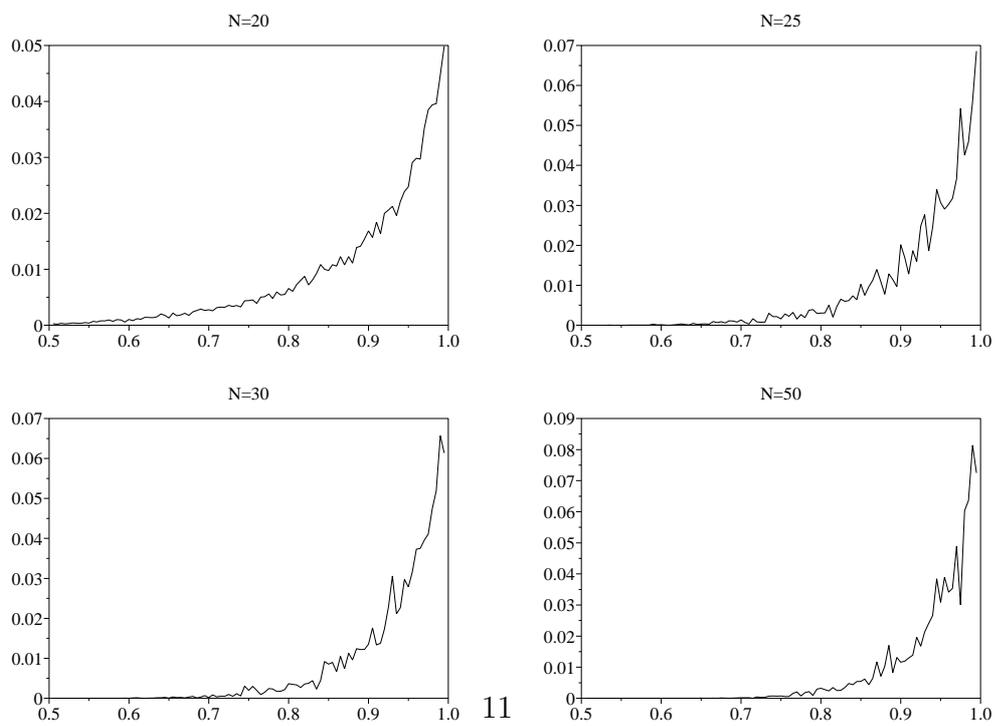


Figure 3: Posterior distribution of the subjective discount factor for different time-horizons.

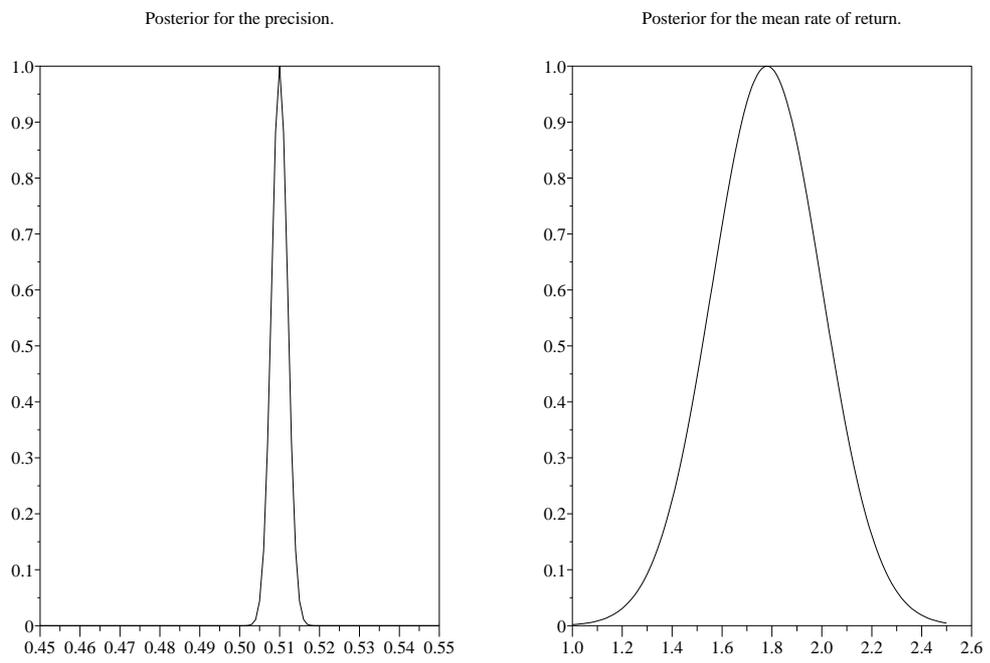


Figure 4: Posterior distributions for the precision and the mean of the logarithmic growth rates.