EFA Eastern Finance Association

The Financial Review 45 (2010) 167–203

The Financial Review

Predictability in Consumption Growth and Equity Returns: A Bayesian Investigation

Alex Paseka*

University of Manitoba

George Theocharides

Sungkyunkwan University

Abstract

We use a Bayesian method to estimate a consumption-based asset pricing model featuring long-run risks. Although the model is generally consistent with consumption and dividend growth moments in annual data, the conditional mean of consumption growth (a latent process) is not persistent enough to satisfy the restriction that the price-dividend ratio be an affine function of the latent process. The model also requires relatively high intertemporal elasticity of substitution to match the low volatility of the risk-free return. These two restrictions lead to the equity volatility puzzle. The model accounts for only 50% of the total variation in asset returns.

Keywords: consumption-based asset pricing, consumption growth predictability, return predictability, equity volatility puzzle

JEL Classifications: C11, C13, C22, G12

**Corresponding author:* University of Manitoba, I.H. Asper School of Business, 476-181 Freedman Crescent, Winnipeg MB, Canada R3T 5V4; Phone: (204) 474-8353; Fax: (204) 474-7545; E-mail: paseka@cc.umanitoba.ca.

We thank Fuyu Angela Yang and seminar participants at the 2007 Athens Institute for Education and Research Finance Conference, the 2007 FMA Conference, and the 2007 International Congress on Modeling and Simulation, Christchurch, New Zealand for helpful comments and discussions. We thank also Arnie Cowan (a former editor) and two anonymous referees for constructive comments and suggestions. Any errors are our responsibility.

1. Introduction

Since the "equity premium puzzle" paper by Mehra and Prescott (1985), researchers have discovered a number of other related anomalies that are difficult to explain within the context of rational asset pricing models. Such anomalies include the "risk-free rate puzzle" (Weil, 1989) that accounts for the low real return on T bills, or the "equity volatility puzzle" (Campbell, 2003) that arises due to the excessively high unconditional volatility of real stock returns relative to the unconditional volatility of real consumption growth. One of the responses to the aforementioned puzzles is the long-run risks (LRR) model proposed by Bansal and Yaron (2004). Their model is based on a small, but highly persistent expected growth rate component in the consumption process, as well as fluctuating economic uncertainty. Using calibration and simulation exercises, they are able to show that the above two channels can help explain the equity premium, the risk-free rate, and the volatility puzzle with a relatively low risk aversion. In more recent years, the LRR model has been extensively used to account for several asset pricing phenomena.

Although the assumed long-run properties of consumption growth could help capture the empirical puzzles theoretically, it is not clear whether shocks to consumption growth are persistent enough in the data to explain the puzzles. The main focus of our study, in the context of the Bansal and Yaron (2004) model, is the empirical assessment of how important consumption persistence is to the resolution of the financial puzzles. Furthermore, although calibration and simulation exercises can be very useful, consumption-based models need to be also formally estimated and tested as advocated by Cochrane (2006). To this end, we estimate a consumption-based model based on Bansal and Yaron (2004) using the fully Bayesian Markov Chain Monte Carlo (MCMC) method. Intuitively, the MCMC method generates samples of parameters and state variables from their full conditional densities. Under suitable conditions, these samples converge to draws from the corresponding marginal densities. With empirical marginal densities at hand, an inference is natural and easy to implement. Using this technique we find that the persistence of the latent conditional mean of consumption growth is only moderate in the data leading to underprediction of equity volatility, overprediction of relative risk aversion, and lack of predictability in equity returns under the model. We restrict our attention to the first channel proposed by Bansal and Yaron (2004), the fluctuating expected growth rates (Case I in their model), and test it on the time series of aggregate stock market data and interest rates.

Recent empirical work in this direction is scarce, and it mainly focuses on simulation and calibration (Bansal and Yaron, 2004) or variants of the generalized method of moments (GMM) estimation (Bansal, Gallant, and Tauchen, 2007; Bansal, Kiku, and Yaron, 2007; Constantinides and Ghosh, 2008). We choose the Bayesian method as it has several advantages over classical methods. Due to its conditional nature, Bayesian MCMC approach allows us to avoid explicit computation of data likelihood marginalized over the state variable (expected consumption growth in the Bansal and Yaron, 2004, model). For the same reason, we do not require any optimization procedure for estimation. We can infer the entire time series of the latent state variable without parameter estimation risk, which may be especially problematic in problems involving latent state processes. Under Bayesian MCMC estimation, the marginal posterior integrates out the parameters as opposed to classical methods, which obtain estimates conditional on parameters (see Johannes and Polson, 2009, for application of Bayesian methods in finance). We conduct a simulation study to show that our parameter estimates using this method are reliable for a typical sample of consumption, dividend and asset return data.

We estimate the model using equity returns, consumption growth, dividend growth, and interest rate data for the post-depression period 1934–2005 and the post-war period 1948–2005. Although some parameter estimates slightly differ across the two periods, our main conclusions are unaffected by our choices of the sample periods. For example, the estimate of the latent state variable persistence is robust across the two sample periods and different prior specifications. The posterior mean of the persistence parameter in 1934–2005, 0.372, is only 4% higher than that for 1948–2005, 0.357. We find that, consistent with the model, consumption growth contains a small predictable latent component that captures its conditional expectation. Our results suggest that the persistence of the latent mean consumption growth is only moderate, corresponding to half life of consumption growth shocks of 1.3 years for the 1934–2005 sample period. At the same time, the solution to the Bansal and Yaron (2004) model with the homoskedastic consumption growth implies that price-dividend ratios are deterministic affine functions of the latent process and, therefore, must be perfectly correlated with it. In fact, under the model, price-dividend ratio autocorrelations must have the same decay rates as those of consumption growth. The last property is at odds with high persistence of price-dividend ratios. Our estimate of the half life of price-dividend shocks of approximately 13 years implies that the model's affine restriction is not supported by the data. Put differently, despite evidence of persistence of shocks to the mean consumption growth, the persistence is not large enough to make long-term consumption volatility similar to the observed volatility of aggregate wealth. As wealth and consumption are related through a budget constraint, the lack of high persistence in conditional mean of consumption growth is an indirect evidence for mean-reverting wealth (as advocated by Campbell and Cochrane, 1999) rather than an upward trend in long-term consumption volatility.

We also argue that the lack of persistence in the mean of consumption growth results in the equity volatility puzzle. We show that in the Bansal and Yaron (2004) model equity volatility consists of three components, two of which, the unconditional variance of conditional mean of consumption growth and the conditional variance of price-dividend ratio, are strongly positively related to the persistence parameter. Moderate values of the parameter (in the range 0.3–0.4) render these components too small to match equity volatility. The remaining component is simply equal to the unconditional volatility of the risk-free rate of return and inversely proportional to the square of the intertemporal elasticity of substitution (IES). According to our estimates, explaining the remaining 3.7% of the 17% sample standard deviation (SD) of equity returns not explained by the first two components requires a value of 7.7 for the inverse of the IES. The high value of IES is, however, at odds with the low volatility of the interest rates. As a consequence, the model explains only 50% of the total return volatility. This result is robust across different subsamples of the data.

We also find that the equity beta on consumption growth is small, 1.9 in our 1934–2005 sample (even smaller at 0.63 in the 1948–2005 sample). This outcome is mainly due to only moderate persistence of the mean of consumption growth because the equity beta is highly sensitive to the value of the persistence parameter. Yet, the model can still explain the equity premium. However, explaining the equity premium requires unusually high price of consumption risk and, consequently, high risk aversion. The real unconditional equity premium of 7% per year is about 424 times our estimate of the annual variance of conditional consumption growth. Therefore, our estimate of the equity beta of 1.9 implies a price of consumption risk of 223 in annual data. Such high price of consumption risk, in turn, leads to high risk aversion needed to match the equity premium. Our estimate of the relative risk aversion is around 140.

2. Related literature

The traditional consumption capital-asset pricing model (CCAPM) (Rubinstein, 1976; Lucas, 1978; Breeden, 1979) implies that the risk premium on an asset should be the compensation for consumption fluctuations of the representative household. In other words, the covariance between the returns on an asset and consumption growth should determine the risk premium. Despite the elegance and appeal of the above model, it unfortunately performs poorly when confronted with data (Mankiw and Shapiro, 1986; Breeden, Gibbons, and Litzenberger, 1989).

One of the responses to the failures of the original CCAPM is to modify the representative investor's utility preferences to include habit formation. Sundaresan (1989), Constantinides (1990), and Heaton (1995) are some of the earlier papers that pursue this direction. More recently, Campbell and Cochrane (1999) show that a version of a habit formation model can explain a number of the asset pricing puzzles, albeit with a high risk aversion. The key assumption of their model is that consumption growth innovations are identically and independently distributed (i.i.d.). This assumption is consistent with other empirical findings of little or no predictability in consumption growth (Hall, 1988; Campbell and Mankiw, 1989; Cochrane, 1994; Lettau and Ludvigson, 2001).

Theoretically, aggregate wealth and consumption are linked through the budget constraint and must have the same long-term volatility. Yet, consumption is very smooth, and wealth is volatile in the short run. Models with i.i.d. consumption growth must imply that the volatility of wealth decreases with the horizon to reach the low volatility of consumption growth.

An alternative solution is to model a small, predictable, and persistent shock component in consumption growth process so that its annualized volatility increases with the horizon to the levels of the volatility of the wealth growth. This assumption is the center of the LRR framework presented in Bansal and Yaron (2004).[1] Combining the effect of the expected growth channel with time-varying economic uncertainty, they show via a calibration exercise that the model can potentially go a long way in capturing the existing puzzles.

LRR model has been used to account for a number of asset pricing phenomena in the domestic equity, bond, or other markets. Bansal, Dittmar, and Lundblad (2005) show that the long-run covariance of consumption growth and dividends can be used to explain the cross-sectional differences in risk premia. In a related paper, Bansal, Dittmar, and Kiku (2009) examine the importance of a cointegrating relation between consumption and dividends. They show that this relation helps to capture consumption risk at all investment horizons, but is especially important at long horizons. Parker and Julliard (2005) evaluate the performance of the LRR model using a cross-section of returns. In their model, consumption growth is measured over a long horizon rather than over the length of one quarter. Kiku (2006) finds that the model can account for the value premium.

Within the context of the bond market, Piazzesi and Schneider (2006) use this framework to study bond risk premia. In their model, inflation can carry bad news signals regarding long-run consumption growth and make nominal yield curves slope upward. Eraker (2008) provides a general affine equilibrium model that builds on that of Bansal and Yaron (2004). Hasseltoft (2007) extends the LRR model to capture both the cyclical properties of interest rates as well as the level and volatility of equity returns.

With respect to other markets, Bansal and Shaliastovich (2007) use the LRR model to explain asset pricing anomalies found in the equity, bond, and currency markets. Colacito and Croce (2006) use the LRR model in a two-country setup to account for the poor correlation of consumption growth between countries, while the stochastic discount factors move in opposite directions.[2]

However, despite the popularity of this model, to the best of our knowledge, little work has been done so far to estimate it. Exceptions are a few recent papers, Bansal, Gallant, and Tauchen (2007), Bansal, Kiku, and Yaron (2007), and Constantinides and Ghosh (2008). These papers use variants of the GMM estimation technique to estimate the LRR model using a rich set of return data. Furthermore, Ma (2007) finds that the LRR model is weakly identified leading

---

[1] Prior to Bansal and Yaron (2004), some other papers that assume the predictable component in consumption growth include Cecchetti, Lam, and Mark (1990, 1993), Constantinides (1990), Gabaix and Laibson (2001), Kandel and Stambaugh (1991), Mehra and Prescott (1985), and Parker (2001).

[2] See Bekaert, Engstrom, and Xing (2009), Chen (2006), Chen, Collin-Dufresne, and Goldstein (2009), Hansen, Heaton, and Li (2008), Kaltenbrunner and Lochstoer (2006), Kiku (2006), Lettau and Ludvigson (2005), Malloy, Moskowitz, and Vissing-Jorgensen (2009), Parker and Julliard (2005), Piazzesi and Schneider (2006), Yang (2006), and Zurek (2007) for other related papers. Bansal (2007) shows how the two channels of Bansal and Yaron (2004) could explain the asset pricing puzzles. Sargent (2007) provides an interesting discussion of the growing importance of the LRR model.

to the spurious inference that could explain the resolution of the asset pricing puzzles.

3. The model

In this section, we provide a brief summary of a consumption-based asset pricing model proposed by Bansal and Yaron (2004).

High relative risk aversion can reconcile, at least formally, consumption smoothness and relatively low correlation between consumption growth and returns with high equity premia. Unfortunately, high relative risk aversion may lead to high risk-free rates—the risk-free rate puzzle (Weil, 1989). Epstein and Zin (1989, 1991) and Weil (1989) show how to resolve the risk-free rate puzzle by a more general form of power utility. Under simple power utility, the elasticity of intertemporal substitution is the reciprocal of the coefficient of relative risk aversion. The more general form of utility does away with the tight link between the coefficients.

Bansal and Yaron (2004) stress that a potential explanation of the puzzles may be provided through the flexible Epstein–Zin–Weil utility. They start with the following joint dynamics for real consumption growth and real dividend growth under the assumption that both are driven by a persistent state variable x_t:[3]

$$\begin{aligned} g_{t+1} &= \mu + x_t + \sigma \eta_{t+1} \\ x_{t+1} &= \rho x_t + \varphi_e \sigma e_{t+1} \\ g_{m,t+1} &= \mu_m + \phi_m x_t + \varphi_m \sigma u_{m,t+1} \\ \eta_{t+1}, e_{t+1}, u_{m,t+1} &\sim i.i.d.Normal(0, 1), \end{aligned} \tag{1}$$

where $\mu, \sigma, \rho, \varphi_e, \mu_m, \phi_m$, and φ_m are constants; and g_{t+1} and $g_{m,t+1}$ are the log-growth rates of consumption and dividends on asset m over $[t, t+1]$. We provide detailed definitions of all model parameters in Table 1. All return variables are with continuous compounding.

To derive an equilibrium solution for equity premium, Bansal and Yaron (2004) start with the equilibrium Euler condition:

$$E_t\left[e^{m_{t+1}+r_{t+1}}\right] = 1. \tag{2}$$

As mentioned earlier, they assume the Epstein–Zin–Weil specification for stochastic discount factor, m_{t+1}, and Campbell and Shiller's (1988) log-linear approximation for asset returns:

$$r_{m,t+1} = k_{0,m} + k_{1,m} z_{m,t+1} - z_{m,t} + g_{m,t+1}, \tag{3}$$

[3] We consider homoskedastic version of their model which corresponds to their Case I. In the extended heteroskedastic version of their model, the conditional volatility of consumption growth, σ^2 (see Equation (1)), follows a mean-reverting Gaussian process.

Table 1

Definitions of model parameters

We report the definitions of original model parameters (Equations (1)–(4)), empirical implementation parameters (Equation (13)), and equilibrium parameters (Equations (6)–(12)).

	Original model parameters (Equations (1)–(4))
μ	Unconditional mean of g
σ	Conditional SD of g
ρ	First-order autocorrelation of conditional mean of g
φ_e	Ratio of conditional SDs of x and g
μ_n	Unconditional mean of g_m
ϕ_m	Sensitivity of g_m to x (leverage parameter)
φ_m	Ratio of conditional SDs of g_m and g
$k_{0,m} \approx 3.44$	A function of average price-dividend ratio for asset m
$k_{1,m} \approx 0.968$	A function of average price-dividend ratio for asset m
$A_{0,m}$	An intercept in the affine relation between the log price-dividend ratio and x
$A_{1,m}$	Sensitivity of the log price-dividend ratio to x
$k_0 \approx 3.653$	A function of average price-consumption ratio
$k_1 \approx 0.974$	A function of average price-consumption ratio
	Empirical implementation parameters (Equation (13))
α_m	Unconditional mean of asset m's return net of gm
σ_m	SD of measurement error for asset m's return
α_f	Unconditional mean of the risk-free rate of return
σ_f	SD of measurement error for the risk-free rate of return
ψ	Intertemporal elasticity of substitution (IES)
	Equilibrium parameters (Equations (6)–(12))
$\beta_{m,e}$	Return beta (conditional) with respect to conditional mean of g
$\lambda_{m,e}$	The price of expected consumption growth risk
γ	Relative risk aversion (RRA)
θ	Measure of relative magnitude of γ and ψ(IES)

where

$$k_{0,m} = \ln\left(1 + \frac{P}{D}\right)$$

$$k_{1,m} = \frac{P/D}{1 + P/D}$$

$$k_0 = \ln\left(1 + \frac{P}{C}\right)$$

$$k_1 = \frac{P/C}{1 + P/C}, \quad (4)$$

where $r_{m,t+1}$ is the return on an observed portfolio paying observed dividend d_t, and P/D and P/C are the average price-dividend ratio for asset m and average price-consumption ratio, respectively.

Bansal and Yaron (2004) solve the model (2)–(3) for the log price-dividend ratio $z_t \equiv p_t - d_t$. The solution is a deterministic affine relation between z_t and the latent conditional mean of consumption growth, x_{t+1}:

$$z_{m,t} = A_{0,m} + A_{1,m} x_t. \tag{5}$$

Equity premium adjusted for Jensen's inequality in this model is given by

$$E_t\left[r_{m,t+1} - r_{f,t+1}\right] + \frac{\mathrm{var}_t[r_{m,t+1}]}{2} = \beta_{m,e}\lambda_{m,e}\sigma^2 \tag{6}$$

$$\mathrm{var}_t[r_{m,t+1}] = \left(\varphi_m^2 + \beta_{m,e}^2\right)\sigma^2 \tag{7}$$

$$\beta_{m,e} = \kappa_{1,m} A_{1,m} \varphi_e \tag{8}$$

$$A_{1,m} = \frac{\phi_m - \frac{1}{\psi}}{1 - \kappa_{1,m}\rho} \tag{9}$$

$$\lambda_{m,e} = (1 - \theta)B \tag{10}$$

$$B = \kappa_1 \frac{\varphi_e\left(1 - \frac{1}{\psi}\right)}{1 - \kappa_1 \rho} \tag{11}$$

$$\theta = \frac{1 - \gamma}{1 - \frac{1}{\psi}}. \tag{12}$$

Potential benefits are evident in that coefficient ψ need not be small even if γ is high. High relative risk aversion does not have to imply high or highly volatile risk-free rate. Also, persistent consumption growth may contribute to the explanation of the equity premium. Return "beta" on the consumption growth factor $\beta_{m,e}$, as well as the price of expected consumption growth risk $\lambda_{m,e}$, as Bansal and Yaron (2004) point out, increase with ρ, the persistence parameter in the latent process that drives consumption and dividend growth rates. Thus, theoretically, it is plausible to produce any risk premia as well as return volatility.

4. Estimation

Despite the beneficial theoretical properties of the model, its empirical justification is relatively scarce. The success of the Epstein–Zin utility with a persistent latent factor relies mostly on the idea of high persistence of expected consumption growth. Yet, empirical evidence on the high persistence is mixed at best (Campbell, 2003). Recently, Bansal, Gallant, and Tauchen (2007) provide a version of simulated GMM evidence; Bansal, Kiku, and Yaron (2007) exploit asset pricing Euler conditions to estimate the model; and Constantinides and Ghosh (2008) back out unobserved state variables using the affine structure of the model and estimate it by GMM.

Generally, likelihood-based approaches are preferred due to their full-information nature and high statistical efficiency. In practice, however, they are rarely used as the closed form of the likelihood function is not available. In such situations incomplete information methods, such as GMM and quasi-maximum likelihood (QML) estimation, are the preferred choices. Fortunately, due to its affine nature, the model by Bansal and Yaron (2004) admits a closed form expression for the joint conditional density of data and a state variable. However, because the state variable is unobserved, we choose to estimate the model using Bayesian MCMC methods. These methods allow us to generate random samples from the joint distribution of the parameters and state variables given the data and obtain their empirical marginal densities, which contain all necessary information required for inference and testing. The Bayesian MCMC approach has distinct advantages over both likelihood and the GMM methods especially in the context of latent variable problems.

First, samples of parameters and latent variables in the MCMC method are based on their full conditional densities, that is, their distributions conditional on all the other variables. These densities in the Bansal and Yaron (2004) model are available in closed form and constitute the basis of our estimation approach. By contrast, traditional ML inference requires the knowledge of the marginal likelihood of the data given parameters with state variables integrated out. The marginal likelihood is not available in closed form even in the simplest version of the model. At the same time, Bayesian MCMC approach allows us to obtain the true parameter density. By contrast, ML estimates rely on normality assumptions in the limit of large samples, which may be questionable in the case of consumption data.

Second, both likelihood-based methods and GMM methods used in Bansal, Gallant, and Tauchen (2007) and Bansal, Kiku, and Yaron (2007) require nonlinear optimization procedures, which may substantially complicate estimation. Our MCMC estimation is a conditional simulation approach. As such, it requires no optimization.

Third, an extra complication is the estimation risk that is normally present in classical econometric approaches (e.g., maximum likelihood (ML) estimation with Kalman Filter) to state space models with latent processes. Kim and Nelson (1999) present an extensive discussion and literature review on ML estimation of state space models with unobserved variables.

We circumvent the estimation risk problem using the Bayesian MCMC approach to draw inferences on the model parameters and the latent process simultaneously. The information in the data about the latent process is summarized in the marginal posterior distribution:

$$p(X \mid Y) = \int p(X, \Omega \mid Y) d\Omega,$$

where Ω is the parameter vector, X is the latent process and Y is the data. In classical estimation procedures we compute a latent variable estimate conditioning on the parameters. This approach ignores sampling variation of the parameter estimates. In MCMC implementation, all the remaining parameters are marginalized out of the likelihood.

Equilibrium restriction (5) implies that, for any asset m, the return net of dividend growth $r_{m,t+1} - g_{m,t+1}$, is a deterministic affine function of the latent process x_{t+1}, given the parameters and the lag of the latent process. It follows that net returns (returns net of dividend growth) for any two assets in such an economy are perfectly correlated conditional on parameters and one lag of the latent process. Renault (1997) suggests that such restrictions introduce singularities into the likelihood function of the estimation problem. In our case, the singularity results from the absence of residual in the return Equation (3). This situation is reminiscent of econometric difficulties whenever the number of the observed asset prices is larger than the number of the latent state variables. Examples span all parametric latent factor models and range from derivative models, such as Heston (1993), with a cross-section of option prices and an unobserved volatility process to the estimation of multifactor term structure models, such as Lamoureux and Witte (2002). A remedy often proposed for such a situation is to introduce measurement errors. The presence of these errors is the recognition that either the data are measured with errors or the model is intrinsically misspecified. In econometric practice, the measurement errors help avoid singularities in the likelihood function. As an example of the approach, Jacquier and Jarrow (2000) explore different ways to introduce errors into the Black-Scholes option pricing model.

We adopt an empirical specification that avoids the singularities in the likelihood. Specifically, we assume that observed returns, asset m return and the risk-free rate of return, contain an i.i.d. mean zero measurement error:

$$\begin{aligned} r_{m,t+1} &= \alpha_m + g_{m,t+1} + \kappa_{1,m} A_{1,m} x_{t+1} - A_{1,m} x_t + \sigma_m \varepsilon_{m,t+1} \\ r_{f,t+1} &= \alpha_f + \frac{1}{\psi} x_t + \sigma_f e_{f,t+1}, \end{aligned} \tag{13}$$

where α_m, σ_m are asset specific parameters, α_f and σ_f are the unconditional mean and the conditional SD of the risk-free rate, and ψ is the IES. We obtain the first equation in (13) by substituting the definition of the price-dividend ratio in (5) into the expression for return (3) and adding the measurement error. Appendix A.1 provides the main steps in the derivation of the risk-free return specification in (13). The market return expression in (13) is not an identity with measurement error attached to it. It is a result of substitution of model implied price-dividend ratio from (5) into the return identity (3). Therefore, the market return specification in (13) correctly describes the dynamics of the market return only to the extent of the validity of model-implied solution for price-dividend ratio in (5).

The estimation proceeds in two steps. First, we estimate the parameters of the model based on the following set of time series (exogenous) and equilibrium (endogenous) conditions:

$$r_{m,t+1} = \alpha_m + g_{m,t+1} + \kappa_{1,m} A_{1,m} x_{t+1} - A_{1,m} x_t + \sigma_m \varepsilon_{m,t+1} \tag{14}$$

$$r_{f,t+1} = \alpha_f + \frac{1}{\psi} x_t + \sigma_f e_{f,t+1} \tag{15}$$

$$g_{t+1} = \mu + x_t + \sigma \eta_{t+1} \tag{16}$$

$$x_{t+1} = \rho x_t + \varphi_e \sigma e_{t+1} \tag{17}$$

$$g_{m,t+1} = \mu_m + \phi_m x_t + \varphi_m \sigma u_{m,t+1} \tag{18}$$

$$A_{1,m} = \frac{\phi_m - \frac{1}{\psi}}{1 - \kappa_{1,m} \rho} \tag{19}$$

$$\varepsilon_{m,t+1}, \eta_{t+1}, e_{t+1}, u_{m,t+1} \sim i.i.d.\ Normal(0, 1). \tag{20}$$

Second, we recover the remaining parameters θ and γ from the equilibrium conditions of the Bansal and Yaron (2004) model given in Equation (6) using definitions in (7)–(12). To this end, we substitute the equity return $r_{m,t+1}$, and the risk-free rate dynamics implied by the model given in (13) into equilibrium condition (6).

Because it is possible to sample directly from full conditional densities of the model, we use a Gibbs sampler method to draw the parameter as well as latent state values. Construction of the Gibbs sampler requires the specification of full conditional posteriors for all parameters and state variables. By Bayes' rule, a full conditional on a variable is proportional to the joint likelihood of the data given the parameters and the latent process times the prior density on that variable, with all the remaining variables being treated as constants. For example, given the joint data likelihood $p(r_m, g_m, g \mid x, \Omega)$, and the joint prior $p(x_0, \Omega)$, the joint posterior for the parameters $\Omega = \{\sigma, \rho, \varphi_e, \mu, a_m, \mu_m, \varphi_m, \phi_m, \sigma_m, A_{1,m}\}$, and the state process $x \equiv \{x_1, \ldots, x_T\}$ is given by

$$p(x, \Omega \mid r_m, g_m, g) \sim p(r_m, g_m, g \mid x, \Omega) p(x_0, \Omega).$$

Once the entire set of Gibbs samples of $A_{1,m}$, ϕ_m, and ρ is available, we can easily recover the marginal posterior density of $\frac{1}{\psi}$ from (19).

Given the posterior density, the Gibbs sampler is implemented as follows. Starting at an initial point $(x^{(0)}, \Omega^{(0)})$:

(1) Draw $\Omega_j^{(0)} \sim p(\Omega \mid x^{(0)}, \Omega_{-j}^{(0)}, r_m, g_m, g)$ for each component of parameter vector j.
(2) Draw $x_k^{(1)} \sim p(x \mid x_{-k}^{(0)}, \Omega^{(0)}, r_m, g_m, g)$ for each component of the x-vector.

Iterating this procedure, the Gibbs sampler produces a sequence of random variables, $\{x^{(m)}, \Omega^{(m)}\}_{m=1}^{M}$, which, under suitable regularity conditions, converge to draws from the joint posterior density, $p(x, \Omega \,|\, r_m, g_m, g)$. We use well-established tools to track the value of M at which convergence is achieved. Only post-burn draws beyond M are used for inference.

Data-related constants $\kappa_{1,m}, \kappa_1$, and κ_0 are computed from the annual price-dividend and price-consumption data according to Equation (4). The values obtained in this way are

(1) $\kappa_{1,m} = 0.968$;
(2) $\kappa_1 = 0.974$;
(3) $\kappa_0 = 3.653$.

To generate samples of parameters and state process and to draw inferences we use the WinBUGS 1.4 package. The description of the package and the priors as well as the details of full conditional densities used in the estimation procedure are described in Appendix C.

5. Empirical results

Below we present empirical results based on the estimation of Bansal and Yaron (2004) model. First, we discuss the data used in the estimation, the method used to estimate parameters of the model, and the robustness of the method and the estimates. Second, based on parameter estimates, we describe how well the Bansal and Yaron (2004) model fits equity volatility, equity premium, and a variety of consumption growth moments under different prior specifications.

5.1. Data and parameter estimates

The data used in the estimation procedure are stock market index returns, personal consumption data on nondurables and services, a dividend series, and returns of the three-month T bill (as a proxy for the risk-free rate). Although annual consumption data are available from 1930, secondary market T bill data are only available from 1934. This uneven data panel determines our choice of the sample period. We use annual data for the period 1934–2005 for estimation and inference. Furthermore, we estimate model parameters based on two periods, a full sample for the period 1934–2005 and a post-war subsample for 1948–2005, to monitor the stability of parameter estimates across different subsamples.[4] Although some sample autocorrelation and variance properties appear to be sample dependent, parameter estimates are fairly robust to the choice of the sample period. Wherever not explicitly stated, we report the estimates for the entire sample over 1934–2005 period. All nominal values are deflated using the Consumer Price Index-All Urban Consumers (CPI-U) taken from the Bureau of Labor Statistics.

[4] We thank an anonymous referee for suggesting this extra robustness check.

Table 2

Summary statistics

We report summary statistics for the value-weighted composite NYSE/Amex/Nasdaq index returns, consumption growth, and dividend growth series for the period 1934–2005. The statistics are based on annual observations. The index returns are taken from CRSP. We use the BEA for personal consumption for nondurables and services and extract dividends from the return series with and without dividends. All nominal observations are deflated using the CPI-U taken from the Bureau of Labor Statistics.

	Index returns	Consumption growth	Dividend growth
Mean	0.069	0.034	0.028
Median	0.109	0.035	0.021
Std. dev.	0.176	0.016	0.137
Skewness	−0.726	0.511	0.593
Kurtosis	0.48	2.24	0.73

For a stock market index, we use the value-weighted composite NYSE/Amex/Nasdaq index returns taken from the CRSP database. Personal consumption data on nondurables and services (in billion of dollars) comes from the Bureau of Economic Analysis (BEA). Real annual personal consumption growth is based on CPI-deflated personal consumption data. Dividends and the real dividend growth series are extracted from index returns with and without dividends. We obtain the returns on the three-month T bill from the Federal Reserve web site (*http://wwww.federalreserve.gov/releases/h15/data*).

Table 2 reports summary statistics for the index return, consumption, and dividend growth. All three variables have nonzero excess kurtosis, an evidence of fat tails. However, the magnitude is not large at the annual frequency. The skewness estimates are negative for returns and positive for consumption and dividend growth.

In Table 3, we present the parameter estimation results based on the model in (14)–(20). We monitor the Gelman–Rubin statistic to track the end of the burn-in phase. With several chains generated for the same parameter during a Monte Carlo experiment, Gelman–Rubin statistic measures the relative within-chain and between-chain variances. In the stationary limit, the ratio must approach 1. The ratio serves as a useful diagnostic test of convergence (Robert and Casella, 2004). In most runs we achieve convergence after only 100 to 200 iterations. We report parameter estimates based on discarding the first 10,000 iterations of the burn-in phase and using the remaining 10,000 for inference. However, our numerical experiments show that collecting samples beyond 5,000 iterations does not change the results in most cases. The fast convergence is due to a simple affine nature of the underlying model that admits closed-form full conditional densities for model parameters.

The results in Table Panel A are based on the entire sample of the data from 1934 to 2005. Columns 2 and 3 report posterior mean and SD. 2.5% and 97.5% posterior confidence bands are in columns 4 and 6, and the posterior median in column 5. Our estimate of parameter ρ is 0.37 with 2.5% (97.5%) posterior band of −0.01 (0.72). The bands imply that only 2.5% of Gibbs draws were below −0.01 and the same percent of draws were above 0.72. Alternatively, these bands imply that in the

Table 3

Posterior parameter estimates

We report parameter estimates using annual observations for consumption and dividend growth and returns on NYSE/Amex/Nasdaq. Panel A displays the parameter estimates based on the entire sample period, 1934–2005, while Panel B displays the estimates based on a post-war subsample period, 1948–2005. All units correspond to annual frequency. $A_{1,m}$ is the sensitivity of the log *price-dividend* ratio to latent factor innovations; α_f is the unconditional mean of the risk-free rate of return; α_m is the unconditional mean of asset *m*'s return net of the log-growth rate of dividends (g_m); φ_m is the ratio of conditional SDs of g_m and g (log-growth rate of consumption); φ_e is the ratio of conditional SDs of x (log-growth rate of latent factor process) and g; μ_m is the unconditional mean of g_m; μ is the unconditional mean of g; ϕ_m is the leverage ratio on expected consumption growth (sensitivity of g_m to x); ψ is the intertemporal elasticity of substitution (IES); ρ is the persistence parameter in the latent process; and σ is the conditional SD of g. Columns 2 and 3 display the posterior mean and SD, columns 4 and 6 report the 2.5% and 97.5% posterior confidence bands, while column 5 shows the posterior median. The table shows that key parameter ρ has a mean estimate of 0.37, with a relatively low 97.5% posterior band of 0.72 for the entire sample period. This estimate is robust to the estimation sample period.

Parameter	Mean	Std. dev.	2.5% conf. band	Median	97.5% conf. band
Panel A: Parameter estimates based on the entire sample, 1934–2005					
$A_{1,m}$	1.943	1.532	−0.533	1.745	5.436
α_f	0.000	0.004	−0.009	0.000	0.009
α_m	0.052	0.016	0.021	0.052	0.083
φ_m	10.35	1.58	7.62	10.21	13.76
φ_e	0.964	0.194	0.643	0.941	1.400
μ_m	0.018	0.016	−0.014	0.018	0.049
μ	0.033	0.003	0.028	0.033	0.039
ϕ_m	1.475	0.919	0.099	1.350	3.517
$1/\psi$	0.312	0.488	−0.629	0.306	1.324
ρ	0.372	0.186	−0.009	0.378	0.721
σ	0.013	0.002	0.010	0.013	0.016
Panel B: Parameter estimates based on a post-war subsample, 1948–2005					
$A_{1,m}$	0.868	1.663	−1.847	0.646	4.750
α_f	0.009	0.004	0.001	0.009	0.018
α_m	0.057	0.017	0.023	0.057	0.090
φ_m	11.060	1.785	7.978	10.930	14.960
φ_e	0.973	0.180	0.672	0.956	1.375
μ_m	0.019	0.017	−0.015	0.019	0.052
μ	0.032	0.003	0.026	0.032	0.038
ϕ_m	1.178	0.933	0.047	0.965	3.502
$1/\psi$	0.618	0.498	−0.382	0.615	1.598
ρ	0.357	0.221	−0.105	0.365	0.766
σ	0.011	0.001	0.009	0.011	0.015

entire sample there is only 2.5% chance that parameter ρ is below −0.01 or above 0.72. The posterior band is a Bayesian equivalent of the confidence limit in classical econometrics. For example, a 97.5% posterior band of 0.72 for parameter ρ implies that at the conventional confidence level of 2.5%, the parameter is below 0.72.

Parameter estimates vary slightly across different subsamples. For example, the mean estimate of ρ is 0.36 in the post-war sample 1948–2005 (Table 3, Panel B).

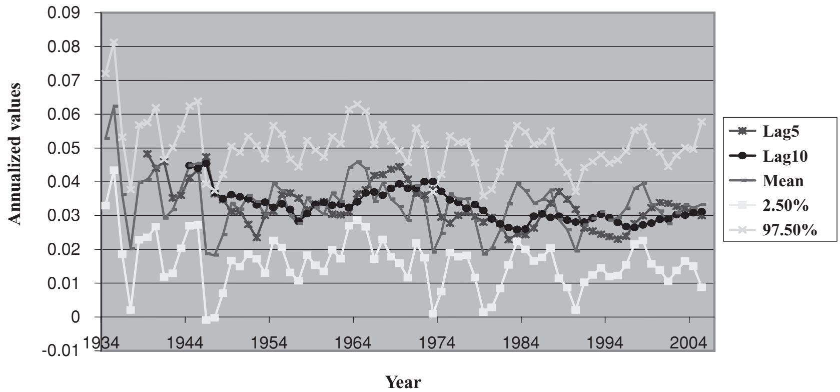

Figure 1

Time series plot of latent mean and corresponding moving average of consumption growth, 1934–2005

This graph depicts the annualized value of the latent mean and corresponding moving average of consumption growth for the period 1934–2005. Lag5 and Lag10 series correspond to moving average estimates for mean consumption growth based on five and ten lags, respectively. Series "mean" corresponds to the posterior mean estimate of the latent factor. We also report 2.5% and 97.5% posterior bands around the posterior mean estimate of the latent factor.

Parameter $A_{1,m}$, the sensitivity of equity returns to the latent process, is the least stable of all the parameters. However, the posterior SD is large, and the estimates of this parameter are not statistically different across the two samples. Also, as seen in Table 3, the precision of most estimates generally deteriorates in the post-war sample simply due to smaller sample size. In either case, however, the estimate of the persistence parameter ρ is very robust to the choice of the sample period.

The Bayesian approach also facilitates the estimation of the latent variable, the conditional mean of consumption growth. In Figures 1 and 2, we report our estimates of the latent variable for both 1934–2005 and 1948–2005 sample periods. In both cases, the mean is computed as the sum of the estimate of the latent state variable x_t for year t and the unconditional mean of consumption growth μ. We plot the estimated time series against the moving average of consumption growth. We report two moving average estimates based on five and ten annual lags of consumption growth. The figures also show the 2.5–97.5% posterior bands on our estimate of the conditional mean of consumption growth. Both the mean and the volatility of consumption growth have been slowly declining over the 1934–2005 period. A similar picture emerges upon inspection of both moving average estimates. Although our estimate of the latent state process has generally the same pattern as the moving average estimates, the latent process is substantially more volatile. This is not surprising because moving average estimates tend to smooth volatile time series. Indeed, we observe in both Figures 1 and 2 that the moving average based on ten lags of consumption growth is less volatile that that based on only five lags.

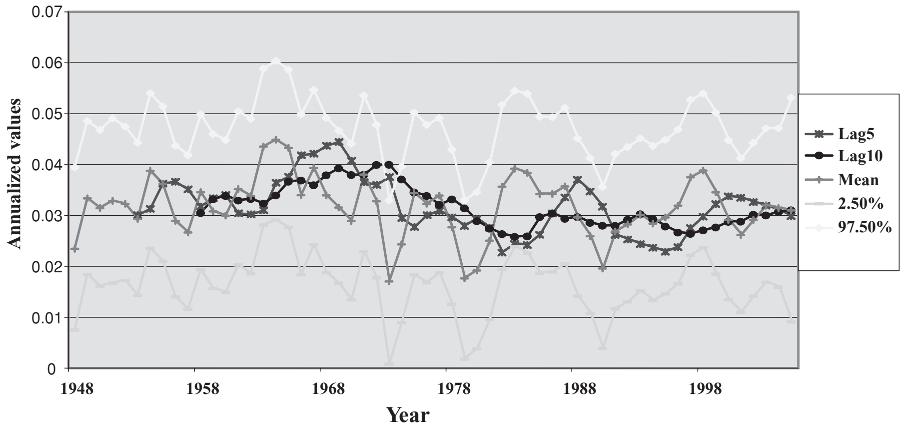

Figure 2

Time series plot of latent mean and corresponding moving average of consumption growth, 1948–2005

This graph depicts the annualized value of the latent mean and corresponding moving average of consumption growth for the postwar period 1948–2005. Lag5 and Lag10 series correspond to moving average estimates for mean consumption growth based on five and ten lags, respectively. Series "mean" corresponds to the posterior mean estimate of the latent factor. We also report 2.5% and 97.5% posterior bands around the posterior mean estimate of the latent factor.

Consumption parameter estimates can be potentially confounded by well-known issues associated with consumption data. The sensitivity of consumption variance and its covariance with other variables to aggregation are well understood (Grossman, Melino, and Shiller, 1987; Christiano, Eichenbaum, and Marshall, 1991; Heaton, 1993; Hansen, Heaton, and Yaron, 1996). For example, if consumption decisions happen with monthly frequency but consumption changes are observed only with annual frequency, the observed consumption growth will appear to be predictable even though it is not.[5]

To assess the effect of time aggregation and measurement error issues on our estimates within the present Epstein–Zin utility, we run a Gibbs sampler on monthly consumption growth and compare the parameter estimates to those obtained with annual frequency. We use the model-implied consumption growth specification at face value to generate parameter estimates:

$$g_{t+1} = \mu + x_t + \sigma \eta_{t+1} \quad (21)$$

[5] Christiano, Eichenbaum, and Marshall (1991) show that the predictability in quarterly seasonally adjusted consumption growth can be explained by temporal aggregation of consumption data. Heaton (1993) shows that nonseparabilities in the utility function is an important consideration for reconciling consumption behavior at different frequencies, with nonseparabilities being more important over shorter horizons.

Table 4

Comparison of parameter estimates for different frequencies of consumption growth data

We report parameter estimates of the consumption growth and latent factor processes using real consumption growth data for the period 1934–2005. The model we estimate is

$$g_{t+1} = \mu + \chi_t + \sigma\eta_{t+1}$$

$$\chi_{t+1} = \rho\chi_t + \varphi_e\sigma e_{t+1},$$

where μ, σ, ρ, and φ_e, are constants; and g_t and x_t are the log-growth rate of consumption and the latent factor process, respectively. We report only posterior means of different parameters for two different frequencies, monthly and annual. At monthly frequency we estimate the model twice. In column 2 we do not put any restriction on parameter ρ, whereas in column 3 we restrict this parameter to be nonnegative. The table shows that key parameter ρ has a lower mean estimate at monthly frequency (-0.30) compared to annual frequency, even if we restrict it to be nonnegative.

	Monthly frequency		
Parameter	$\rho \in (-1, 1)$	$\rho \in [0, 1)$	Annual frequency
φ_e	1.033	0.997	1.034
μ	0.00262	0.00262	0.033
ρ	−0.30	0.02	0.59
σ	0.004	0.004	0.014

$$x_{t+1} = \rho x_t + \varphi_e \sigma e_{t+1}. \tag{22}$$

Table 4 contains the comparison of the estimates.

At monthly frequency, the direct estimate of real consumption growth persistence of -0.30 is even lower (in fact, negative) than that at annual frequency ($\rho = 0.59$). Alternative restriction, such as enforcing a lower bound on parameter ρ at zero, makes little difference to the conclusions. The posterior mean and SD for ρ are 0.0194 and 0.0187, and 5% and 95% posterior bands are 0.0005 and 0.07, respectively (Table 4).

To learn about the effect of time aggregation in consumption data on the estimated first-order autocorrelation (FOA) coefficient we derive (see Appendix A.3) the consumption growth FOA coefficient when consumption is aggregated over T periods. We then compute the FOA based on the aggregate quantities. We can not rule out that the aggregation may lead to high perceived consumption growth autocorrelations in the aggregated data. However, the aggregation method (see Appendix A.3) implies certain restrictions on the magnitude of FOA. In the limit of continuous consumption decisions the FOA coefficient goes to 0.25. The FOA is a monotone function of T, declining if ρ is above 0.25 and increasing if ρ is below. Consequently, monthly ρ of -0.3 (Table 4) would imply an annual value of 0.24.

We obtain our estimates using a variety of criteria. The need for these extra restrictions results from nonuniqueness in the parameter values due to poor identification of the model which we discuss later, together with the criteria to mitigate the problem. Whatever criterion we use, our estimates of the persistence of the latent process are generally lower than those in other related studies. For example, our ρ-estimate of 0.37 based on annual data and the model in (14)–(20) (see Table 3) would correspond to 0.77 on a monthly basis upon de-aggregation (according to Appendix A.3). By contrast, Bansal, Kiku, and Yaron (2007) calibrate parameter ρ at 0.982 in monthly data. According to the aggregation effect, their estimate would correspond to an FOA of 0.89 at annual frequency ($T = 12$). Similarly, in Bansal, Gallant, and Tauchen (2007) the estimate of ρ is 0.987 at monthly frequency, which would correspond to 0.91 at annual frequency after the aggregation.

Parameter estimates are not unique, an issue related to the consumption model (16)–(17) being overparameterized.[6] We can understand the identification issue by looking at the expression for the unconditional variance of the consumption growth implied by the model:

$$\mathrm{var}(g_t) = \sigma^2 \left(1 + \frac{\varphi_e^2}{1 - \rho^2} \right). \tag{23}$$

It is not possible to break down the unconditional variance of consumption growth into conditional variance and the variance of the conditional mean unambiguously. In Table 5, we list the results of typical experiments with different priors on parameters σ and φ_e and find that there is a tradeoff between the two parameters. The relative magnitude of conditional variance of consumption growth and that of its latent mean x measured by $\varphi_e \equiv \sqrt{\frac{\tau_g}{\tau_x}}$ is known as signal-to-noise ratio. Here, τ_g and τ_x are the conditional precisions (i.e., inverse conditional variances) of consumption growth and of its latent mean, respectively. Low precision priors on conditional variance of x tend to favor lower estimates of φ_e. For example, for a high precision prior on x, $\varphi_e = 1.029$. Leaving the prior precision on τ_g unchanged and reducing the prior precision on τ_x reduces φ_e to 0.134. Parameter σ increases accordingly so that the model fits the unconditional variance of consumption growth (23). Unless we put further restriction on the prior about the signal-to-noise ratio, the data tell us little about the magnitude of the signal-to-noise ratio. One such restriction that we use is the predictive loss criterion (PLC). During the estimation stage (while running the Gibbs sampler), we generate artificial data from (16) to (17) for each draw in our MCMC algorithm given current parameter estimates. We compute PLC as the sum of squared differences between the actual and corresponding artificial data. Intuitively, PLC is a statistical measure of the distance between the data and the artificial sample. Our results suggest that lower precision priors on φ_e (hence, lower estimates of φ_e and lower signal-to-noise ratio) deliver lower PLC and, thus, better fit the data. However, at excessively low values of φ_e, the model implies too low consumption growth autocorrelations relative to sample estimates. In our experiments with the priors, the

[6] A related study Ma (2007) finds that weak identificaton of the Bansal and Yaron (2004) model leads to a spurious inference with inflated estimates of the persistence of the latent mean of consumption growth. We thank an anonymous referee for pointing out this study to us.

Table 5

Comparison of parameter estimates and model fit for consumption process for different precision priors on consumption growth and its latent mean

We report consumption parameter estimates for different priors on conditional variances of consumption growth and its latent mean. We use the entire sample of consumption available to us, from 1930 to 2005. We show results for different relative magnitudes of the signal-to-noise ratio, $\varphi_e = \sqrt{\frac{\tau_g}{\tau_x}}$, where τ_g is the conditional precision (i.e., inverse of conditional variance) of consumption growth and τ_x is the conditional precision of the latent process. PLC measures the distance of the actual data from the data based on the fitted model; γ_1 is the first-order autocorrelation in mean consumption growth; φ_e is the signal-to-noise ratio; μ is the unconditional mean of the log-growth rate of consumption (g); ρ is the persistence parameter in the latent process; and σ is the conditional SD of g. Mean and SD are posterior mean and SD, 2.5% and 97.5% are corresponding posterior bands for parameter estimates. The table shows that there is tradeoff between the priors on parameters σ and φ_e, so that the model fits the unconditional variance of consumption growth. A value of $\varphi_e \approx 1.034$ provides the best fit to the data and matches the first-order autocorrelation of consumption growth.

Parameter	Mean	SD	2.50%	97.5%	Mean	SD	2.50%	97.5%	Mean	SD	2.50%	97.5%
	$\varphi_e = 1.029$				$\varphi_e = 0.416$				$\varphi_e = 0.134$			
PLC	0.068	0.015	0.046	0.101	0.058	0.010	0.040	0.080	0.055	0.009	0.039	0.076
γ_1	0.328	0.104	0.146	0.554	0.144	0.085	0.039	0.361	0.022	0.025	0.004	0.089
φ_e	1.029	0.263	0.616	1.631	0.416	0.174	0.210	0.883	0.134	0.062	0.065	0.296
μ	0.032	0.004	0.024	0.042	0.033	0.003	0.027	0.040	0.033	0.002	0.029	0.038
ρ	0.550	0.118	0.302	0.769	0.632	0.088	0.438	0.790	0.631	0.073	0.474	0.760
σ	0.014	0.002	0.011	0.018	0.017	0.002	0.013	0.021	0.019	0.002	0.016	0.022
	$\varphi_e = 1.034$				$\varphi_e = 30.17$				$\varphi_e = 0.042$			
PLC	0.058	0.012	0.040	0.085	0.066	0.013	0.045	0.095	0.053	0.009	0.038	0.073
σ_1	0.260	0.145	0.050	0.552	0.433	0.101	0.237	0.632	0.002	0.003	0.000	0.009
φ_e	1.034	0.938	0.236	3.631	30.170	9.599	13.670	50.590	0.042	0.019	0.021	0.091
μ	0.033	0.004	0.026	0.040	0.031	0.004	0.022	0.037	0.033	0.002	0.029	0.038
ρ	0.589	0.112	0.341	0.776	0.434	0.101	0.237	0.632	0.622	0.080	0.431	0.753
σ	0.014	0.004	0.005	0.020	0.001	0.000	0.000	0.001	0.018	0.002	0.016	0.022

best fit to the data that also matches the FOA of consumption growth corresponds to φ_e around 1.034. Following the spirit of calibration exercises in Bansal and Yaron (2004) and Bansal, Kiku, and Yaron (2007), we use parameter estimates that match sample consumption growth autocorrelations in all further discussions.

The aggregation may not be the only issue as consumption data are also measured with error (for discussion of how consumption data are constructed see, e.g., Heaton, 1993). Thus, we also examine the sensitivity of the estimated FOA of −0.3 to different specifications of the consumption measurement error. In Appendix A.3 we derive the effect of a hypothetical AR(1) error in consumption growth on the FOA coefficient. It is a weighted average of the latent factor persistence and that of the measurement error, with the weight being equal to the contribution of the latent variable volatility to the total variance of the consumption growth, $w = \frac{\sigma_x^2}{\sigma_x^2 + \frac{\sigma^2}{1-\rho_u^2}}$, where ρ_u is the FOA of the measurement error. We re-estimated consumption growth process (21) and (22) at monthly frequency under different autocorrelations of the measurement error process. Consumption autocorrelation $\gamma_g = -0.3$ (an estimate of ρ in Table 4) is the estimate under white noise measurement error, that is, $\rho_u = 0$. We do not see much sensitivity of γ_g to different levels of ρ_u. For example, at $\rho_u = -0.99$, the posterior mean of the consumption growth FOA is -0.467 with 2.5–97.5% posterior band $[-0.6727, -0.2936]$. Similarly, at $\rho_u = 0.99$, the posterior mean of the consumption growth FOA is -0.267 with 2.5–97.5% posterior band $[-0.3851, -0.1413]$, not far away from the white noise results. For all other levels of measurement error autocorrelation, the posterior mean is always between -0.467 and -0.267.

High ρ would imply that most of the variation in the consumption growth is due to variation in the latent factor—conditional mean of the consumption growth (second term in (23)). However, our estimate of ρ is only moderate. Hence, both the conditional variance and the variance of the conditional mean have similar contributions to the unconditional variance of consumption growth. Fitting the consumption growth and the dividend growth simultaneously presents no problem despite the dividend growth being much more volatile than the consumption growth. The model can fit a much more volatile dividend growth due to an extra degree of freedom afforded by parameter φ_m. In Table 6, Panel B, we report the SD of the dividend growth in the data of 13% per year. Despite the small σ-estimate, the estimation procedure fits the dividend growth variance $\text{var}(g_{m,t}) = \sigma^2(\varphi_m^2 + \frac{(\phi_m\varphi_e)^2}{1-\rho^2})$ through a large φ_m-estimate, a measure of conditional SD of dividend growth. The resulting SD of the dividend growth implied by the model is 12.25% per year that compares to 13% in the data (Table 6, Panel B).

Next, we evaluate whether our estimates are consistent with the dependence of consumption growth persistence on the horizon over which the persistence is measured. To this end, we directly compare autoregressive integrated moving average (ARIMA) estimates of the consumption growth autocorrelations against those implied by the model, specifically, according to how fast they decay over time. In the model, the expression for autocovariance of the consumption growth at lag k follows from the first two equations in (1):

$$\text{cov}(g_{t+k}, g_t) = \frac{\rho^k \varphi_e^2 \sigma^2}{1 - \rho^2}. \tag{24}$$

Thus, model-implied autocovariances die out exponentially at the rate $\rho = 0.37$ (at the posterior mean). It follows from (23) and (24) that the autocorrelations also must decrease exponentially at the same speed:

$$AC(k) = \frac{\rho^k \varphi_e^2}{1 - \rho^2 + \varphi_e^2}. \tag{25}$$

Table 6 reports the SDs of the consumption and the dividend growth, their correlation and autocorrelations from the data (Panel B) and from the model (Panel A). Parameter estimates are consistent with autocorrelations in the data. Looking at the posterior means of autocorrelation estimates, there is some evidence that they decrease exponentially at the rate of ρ^k where ρ is the persistence parameter in the conditional mean of consumption growth and k is the lag. The exponential decay is especially evident for the 1930–2005 sample period, the entire history of consumption growth data that is available to us. In this period, consumption growth autocorrelations tend to decrease even at a faster rate than the estimated 0.59 for that period. However, they still remain within one SD of the model-implied values. In general, however, autocorrelation estimates are noisy. For various time periods (including the period 1948–2005, not reported here), we can not reject the statement that autocorrelations are not statistically different from zero at conventional levels. For example, all consumption growth autocorrelation estimates for the post-depression period 1934–2005 are within two SDs from zero (Table 6, Panel B).

Table 6

Moment estimates for consumption and dividend growth

We report moment estimates for consumption and dividend growth implied by the model, as well as directly from the data. Panel A reports model-implied consumption and dividend growth moment estimates using annual data on NYSE/Amex/Nasdaq index returns, consumption, and dividend growth for the period 1934–2005. We use MCMC parameter estimates to build posterior summaries of autocorrelations, γ, at different lags as well as unconditional SDs. Columns 2 and 3 display the posterior mean and SD, columns 4 and 6 report the 2.5% and 97.5% posterior confidence bands, while column 5 shows the posterior median. Panel B reports moment estimates obtained directly from the data using ARIMA procedure. We report mean, SD, autocorrelations at lags up to ten, as well as correlation between consumption and dividend growth estimates for two overlapping data periods, 1930–2005 and 1934–2005. The table displays evidence of autocorrelation in consumption growth (implied by the model as well as from the data). However, the model is not rich enough to reproduce negative autocorrelations observed in thp data.

Panel A: Posterior model estimates (MCMC output)

Autocorrelation (Lag)	Mean	Std. dev.	2.5% conf. band	Median	97.5% conf. band
	Consumption growth autocorrelations				
$\gamma(1)$	0.1703	0.1139	−0.0398	0.1636	0.4092
$\gamma(2)$	0.0755	0.0741	0.0004	0.0553	0.2670
$\gamma(5)$	0.0110	0.0269	0.0000	0.0020	0.0806
$\gamma(7)$	0.0041	0.0170	0.0000	0.0002	0.0360
$\gamma(10)$	0.0013	0.0108	0.0000	0.0000	0.0113
	Dividend growth autocorrelations				
$\gamma(1)$	0.0106	0.0192	−0.0014	0.0037	0.0643
$\gamma(2)$	0.0048	0.0104	0.0000	0.0011	0.0312
$\gamma(5)$	0.0007	0.0032	0.0000	0.0000	0.0061
$\gamma(7)$	0.0003	0.0020	0.0000	0.0000	0.0022
$\gamma(10)$	0.0001	0.0013	0.0000	0.0000	0.0005
	Unconditional correlation between consumption and dividend growth				
	0.0360	0.1244	−0.2056	0.0348	0.2823
	Unconditional SDs				
Consumption growth	0.0186	0.0017	0.0157	0.0184	0.0221
Dividend growth	0.1225	0.0108	0.1039	0.1218	0.1458

(*continued*)

Table 6 (*continued*)

Moment estimates for consumption and dividend growth

Panel B: Estimates from the data				
	Estimate	Std. error	Estimate	Std. error
Autocorrelation (Lag)	1934–2005		1930–2005	
	Consumption growth (g) autocorrelations			
$\gamma(1)$	0.224	0.118	0.416	0.115
$\gamma(2)$	−0.034	0.124	0.132	0.133
$\gamma(5)$	0.159	0.125	0.005	0.142
$\gamma(7)$	0.063	0.128	0.010	0.142
$\gamma(10)$	0.232	0.135	0.060	0.143
	Dividend growth (g_m) autocorrelations			
$\gamma(1)$	−0.177	0.118	−0.145	0.115
$\gamma(2)$	−0.148	0.121	−0.161	0.117
$\gamma(5)$	−0.025	0.124	−0.123	0.120
$\gamma(7)$	−0.119	0.125	−0.080	0.122
$\gamma(10)$	0.003	0.129	−0.005	0.126
	Estimate		Estimate	
Sample correlation of g and g_m	−0.005		0.120	
Sample SD of g	0.015		0.023	
Sample SD of g_m	0.130		0.123	

With respect to dividend growth autocorrelations, the model is not rich enough to reproduce negative autocorrelations observed in the data. As long as the estimate of ρ is positive with little probability mass on negative values, the model implies positive dividend growth autocovariances. Nevertheless, the fit is reasonable as the estimates from the data are not statistically different from zero at conventional significance levels. Similarly, the model-implied results are all very close to zero within two SDs margin. Covariance between consumption and dividend growth of −0.0046 in the data is well within one SD of the model posterior mean of 0.036.

We also perform a Monte Carlo assessment of how reliable our estimates are. We summarize the study in Appendix B. The results of the study indicate that our method provides reliable estimates of model parameters. Although an estimate of the leverage parameter (sensitivity of dividend growth to the latent mean of consumption growth) ϕ_m is somewhat biased, it is still within one SD from the true value. Of special importance is the fact that the estimates of the key parameter, the persistence of the latent consumption mean ρ, is very robust to alternative prior specifications.

5.2. Equity volatility

Breaking the singularity in the likelihood requires a measurement error in the return dynamics. The quality of the model can be gauged by the estimate of the volatility of the measurement error σ_m. The model is good, if most of the variation in returns is explained by the model and only a small portion of the total variation is due to the volatility of the measurement error. In such a case, the R^2 of regressing net return (return net of dividend growth) on the latent variable should be high. R^2 is however high only if the latent process has similar characteristics to those of the price-dividend ratio. Otherwise, a low R^2 reflects a conflict between the properties of the two processes. This conflict results in misrepresentation of equity volatility among other things. Below we build the R^2 measure to assess the magnitude of the equity volatility puzzle under the model.

One advantage of the MCMC procedure is that it emphasizes the most informative moments of the data. Looking at the return expression in (14), we find that parameter α_m plus the mean of dividend growth reflect the mean equity return. The last observation is due to the unconditional mean of the latent process x being zero in the model. The estimated value of parameter $\mathbf{A}_{1,m}$ is a compromise in the model's fit to the volatility of the index return in (14), the dividend growth and especially the consumption growth. The parameter is important because it measures the sensitivity of equity return to innovations in the latent variable (see (14)).

The following expressions based on (14) give both the conditional and unconditional volatility of index return explained by the model (i.e., under the assumption that $\sigma_m^2 = 0$) as well as the volatility of the risk-free rate based on (13):

$$\text{var}_t(r_{t+1}) = \sigma^2\left(\varphi_m^2 + \beta_{m,e}^2\right) \tag{26}$$

$$\text{var}(r_{t+1}) = \left(\frac{\sigma_x}{\psi}\right)^2 + \text{var}_t(r_{t+1}) = \left(\frac{\sigma_x}{\psi}\right)^2 + \sigma^2\left(\varphi_m^2 + \beta_{m,e}^2\right) \tag{27}$$

$$\text{var}(r_{f,t+1}) = \left(\frac{\sigma_x}{\psi}\right)^2 \tag{28}$$

$$\sigma_x^2 = \frac{\varphi_e^2\sigma^2}{1-\rho^2}, \tag{29}$$

where $\beta_{m,e}$ is given in (8). We report the volatilities in Table 7.

Campbell (1991) shows that most of the variation in returns comes from shocks to future returns rather than the dividend growth. Thus, according to the Campbell–Shiller (1988) identity (3), most of the return variation should come from price-dividend ratios. In the model, the price-dividend ratio is perfectly correlated with the latent variable. The model, however, prescribes two conflicting roles to the latent process x_t. On one hand, it inherits the properties of consumption growth including low volatility. On the other hand, as shown by expressions (27) and (28), it has to explain both the high equity return volatility and relatively low volatility of the risk-free rate.

We can think of unconditional equity volatility given by (27) as consisting of three parts. The first part is the unconditional volatility of the conditional mean, $(\frac{\sigma_x}{\psi})^2$,

Table 7

Conditional and unconditional equity volatilities

We report the conditional and unconditional volatility of the equity return that the model implies based on the parameter estimates reported in Table 3 for the period 1934–2005. Columns 2 and 3 display the posterior mean and SD, columns 4 and 6 report the 2.5% and 97.5% posterior confidence bands, while column 5 shows the posterior median. The first row reports the conditional equity volatility with measurement error included. The second row contains conditional equity volatility explained by the model (measurement error excluded). The third, fourth and fifth rows contain the SD of the measurement error, unconditional SD of equity return explained by the model, and the volatility of the latent state process, *x*, respectively. The last row is the portion of the equity variance explained by the model. The table shows that the model can explain only 50% of the total equity return variance.

Parameter	Mean	SD	2.5% posterior band	97.5% posterior band
$\sqrt{\text{var}_t(r_{t+1}) + \sigma_m^2}$	0.1906	0.0115	0.1699	0.2147
$\sqrt{\text{var}_t(r_{t+1})}$	0.1340	0.0116	0.1137	0.1592
σ_m	0.1349	0.0123	0.1128	0.1609
$\sqrt{\text{var}(r_{t+1})}$	0.1342	0.0116	0.1140	0.1594
σ_x	0.0135	0.0022	0.0101	0.0182
R^2	0.4977	0.0640	0.3763	0.6265

which is equal to the risk-free rate volatility under the model. Under the model, the conditional mean of equity return and that of the risk-free rate are perfectly correlated (through dependence on x) and have the same variance. This observation is closely related to zero predictability of excess equity returns by the price-dividend ratios under the model as we discuss in Appendix A.2. The second term in (27), which is a component of conditional equity volatility, is the conditional variance of dividend growth $\sigma^2\varphi_m^2$. Finally, the third term $\sigma^2\beta_{m,e}^2$ is the conditional variance of the price-dividend ratio. Formally, a low IES ψ can help match high return volatility. However, a low IES would also lead to a high risk-free rate volatility. The only way to resolve this conflict is through the conditional variance of the price-dividend ratio, the third term in (27). This term is roughly proportional to $\frac{1}{1-\rho}$ as a result of equilibrium considerations (see Bansal and Yaron, 2004, for details and also (8) and (19)). The conditional variance of the price-dividend ratio would be high if the persistence parameter ρ was very close to one. However, ρ is the autocorrelation of consumption growth rather than that of the price-dividend ratio, and is substantially below one in the data.

Numerically, the problem the model faces can be described as follows. Unconditional equity volatility consists of two components. One is the conditional equity volatility which is around 13.4% in annual data (Table 6). The other is the volatility of the latent process scaled down by ψ^2. The sample SD of equity return is around 17%. It implies that the scaled latent factor volatility must explain the remaining 3.7%. The estimated unconditional SD of the latent process x is 1.35%. Given these estimates, in order to match sample return volatility of 17% requires ψ of around 0.13, or, alternatively, $\frac{1}{\psi}$ of 7.7. However, the posterior mean of $\frac{1}{\psi}$ based on full data

(Table 3, Panel A) is only 0.31 with the 97.5% confidence limit of 1.32. The reason is that a low IES is incompatible with a low volatility of the risk-free rate. To sum up, the estimate of ψ is too large to generate reasonable equity volatility.

The estimates reported in Table 7 suggest that the model accounts only for $R^2 = \frac{\text{var}(r_{t+1})}{\text{var}(r_{t+1})+\sigma_m^2} \approx 50\%$ of the total return variance. The result demonstrates that Epstein–Zin utility with persistent consumption growth is not sufficient to reconcile the relatively smooth consumption growth and much more volatile return process.

5.3. Equity premium and Sharpe ratio

To estimate the equity premium and the Sharpe ratio, we substitute the two equations from (13) into the equilibrium condition (6) and using the definition in (10) we obtain an expression for parameters $\lambda_{m,e}$ and θ:

$$\lambda_{m,e} = (1-\theta)B = \frac{\alpha_m - \alpha_f + \mu_m + \frac{1}{2}\left(\text{var}_t(r_{t+1}) + \sigma_m^2\right)}{\sigma^2 \beta_{m,e}} \approx \frac{\alpha_m - \alpha_f + \mu_m}{\sigma^2 \beta_{m,e}} \quad (30)$$

$$1 - \theta = \frac{\lambda_{m,e}}{B} \approx \frac{\alpha_m - \alpha_f + \mu_m}{\sigma^2 \beta_{m,e} B}, \quad (31)$$

where $\beta_{m,e}$ and B are defined in (8) and (11), respectively. The conditional volatility of market return and the measurement error volatility have little influence on the equity premium and we omit them here.

The two variables that determine the equity risk premium are the equity beta on consumption growth $\beta_{m,e}$, and the price of consumption risk $\lambda_{m,e}$ (see Equation (6)). Our estimates in Table 8 suggest that the former is relatively moderate at 1.9. The last result is mainly due to the moderate persistence of consumption growth (see Equations (8) and (9) for the relation between the two quantities). However, the model can nevertheless fit the equity premium, even with the moderate equity beta on consumption growth. According to posterior means in Table 8, matching the premium requires an unusually high price of consumption growth risk $\lambda_{m,e}$, around 223. To match the unconditional equity premium of 7%, equilibrium restrictions (10) and (12) require high risk aversion of around 140. However, the precision of these estimates is low.

Alternatively, high equity premium within the current model can be explained by a high absolute value of parameter θ. It is clear from (31) that a major contribution to the magnitude of θ comes from the balance between low conditional volatility of consumption growth σ on one hand, and, on the other, relatively small parameters $\beta_{m,e}$ and B, both of which are changing with ρ at a rate of approximately $\frac{1}{1-\rho}$. To have a reasonably low risk aversion parameter estimate requires high persistence of the latent process that we do not observe in consumption data.

Even if we ignore the high risk aversion necessary to justify the equity premium, we still find it difficult to fit the unconditional Sharpe ratio observed in the data.

Table 8

Estimates of the unconditional equity premium and the Sharpe ratio

We report the Monte Carlo estimates of parameters determining equity return beta on consumption growth, $\beta_{m,e}$, and the price of consumption growth risk, $\lambda_{m,e}$. The equilibrium restrictions on parameters are given by the following system of two equations:

$$\textit{Unconditional Equity Premium} = \beta_{m,e}\lambda_{m,e}\sigma^2,\ \lambda_{m,e} = (1-\theta)B$$

where θ, σ, and B are constant parameters. ψ is the intertemporal elasticity of substitution (IES). Unconditional Sharpe ratio is obtained as the ratio of unconditional equity premium divided by the unconditional equity volatility with the measurement error included. Posterior statistics for the parameter estimates related to equity premium are based on annual data on NYSE/Amex/Nasdaq index returns, consumption, dividend growth, and T-bill rates for the period 1934–2005. All units correspond to annual frequency. This table shows that the low persistence of consumption growth causes a moderate value for the equity beta on consumption growth. Therefore, matching the unconditional equity premium of 7% requires unusually high price of consumption growth risk.

Parameter	Mean	SD	2.5% posterior band	97.5% posterior band
$1/\psi$	0.312	0.488	−0.629	1.324
B	1.107	0.925	−0.496	3.111
$\beta_{m,e}$	1.854	1.368	−0.445	4.798
σ	0.013	0.002	0.010	0.016
Unconditional equity premium	0.070	0.023	0.025	0.114
Unconditional Sharpe ratio	0.370	0.119	0.149	0.601

The volatility puzzle described above implies unrealistically high Sharpe ratios under the model. We can remedy this problem only by including the measurement error into equity volatility for the purpose of the Sharpe ratio calculation. We report the unconditional Sharpe ratio of 0.37 in Table 8.

5.4. Is latent process related to price-dividend ratios?

The model implies that price-dividend ratios are perfectly correlated with the latent process x_t. This property allows for an alternative look at this restriction. We run again an MCMC estimation procedure that uses price-dividend ratio as an alternative to the specification with the latent process x in (16) and (17). This procedure should be equivalent to the one with the latent factor as long as the model holds. To stress this point, we write the equation for the consumption growth in terms of log-price-dividend ratio z and substitute state z for state x solving for z in terms of x from (5):

$$\begin{aligned} g_{t+1} &= \mu - A_{0,m}A_{1,m}^{-1} + A_{1,m}^{-1}z_t + \sigma\eta_{t+1} \\ z_{t+1} &= (1-\rho)A_{0,m} + \rho z_t + A_{1,m}\varphi_e\sigma e_{t+1} \\ \eta_{t+1}, e_{t+1} &\sim \text{i.i.d. } \textit{Normal}(0,1). \end{aligned} \tag{32}$$

Table 9

Parameter estimates with conditional mean of consumption growth substituted out of the model

We report posterior statistics for parameter estimates based on a version of the model that captures the relation between consumption growth and price-dividend ratio using annual data for the period 1934–2005. In this version of the model, we substitute the latent variable out of the model by solving for it in terms of price-dividend ratio:

$$g_{t+1} = \mu - A_{0,m}A_{1,m}^{-1} + A_{1,m}^{-1}z_t + \sigma\eta_{t+1}$$

$$z_{t+1} = (1-\rho)A_{0,m} + \rho Z_t + A_{1,m}\varphi_e\sigma\ell_{t+1},$$

where μ, σ, ψ_e, $A_{1,m}$ and A_{0m} are constants; g_t and z_t are the log-growth rate of consumption and the log price-dividend ratio process, respectively. All units correspond to annual frequency. Columns 2 and 3 display the posterior mean and SD, columns 4 and 6 report the 2.5% and 97.5% posterior confidence bands, while column 5 shows the posterior median. The table shows that the estimate of the persistence parameter (0.948) is very different from the estimates that we got before using specifications with unobserved latent process (0.372 from Table 3, Panel A).

Parameter	Mean	St. dev.	2.5% posterior band	Median	97.5% posterior band
$A_{1,m}^{-1}$	−0.003	0.006	−0.016	−0.003	0.008
φ_e	−0.022	0.040	−0.101	−0.021	0.055
μ	0.033	0.032	0.015	0.035	0.045
ρ	0.948	0.030	0.884	0.950	0.997
σ	0.021	0.002	0.018	0.021	0.025

We report the estimates of this specification in Table 9. In the model, both the consumption and the price-dividend ratios have the same persistence properties. The rate of decay of autocorrelations of both variables should be the same according to the model. However, the estimate of ρ in Table 9 is different from those obtained from specifications with unobserved latent process x_t. The process does not share either high persistence or high volatility of the price-dividend ratio process. For example, the estimate for parameter ρ of 0.95 implies that the price-dividend ratio shocks have a half life of $\ln 0.5/\ln 0.948 \approx 13$ years, a well known fact about the high persistence of observed price-dividend ratios (Campbell, 2003). Despite a small predictable component in consumption growth, it is not volatile enough to proxy for the price-dividend ratios as the model would suggest.

Another dimension, along which we assess the model, is the test of return predictability directly from regressions of multiperiod returns on price dividend ratios. For example, Fama and French (1988) find that high dividend yields tend to be followed by high returns. The strength of the forecasts builds with horizon in that both the R^2 and the regression coefficients increase with horizon. In Appendix A.2, we show that the present model is not rich enough to reproduce these predictability patterns. In particular, one of the implications of the model is that continuously compounded equity returns in excess of the risk-free rate are orthogonal to log price-dividend ratios. The reason is that all returns in the model, including the risk-free rate, have a single-factor structure. Conditional means of all returns are affine functions of the latent process x_t with loading on this factor being the same for all returns and equal to the inverse of IES. Consequently, excess equity returns are no longer functions of the latent state variable. Because the only source of correlation between any two quantities in the model is the dependence on the common factor x, the model implies that excess returns are not correlated with (i.e., unforecastable by) any other variable. For example, dividend yields, dividend growth, consumption growth, interest rates, and price-dividend ratios have no predictive power with respect to multiperiod excess returns.

6. Conclusion

We provide an alternative estimation approach to the LRR model of Bansal and Yaron (2004). We emphasize an important role for the consumption growth persistence in explaining equity return volatilities and equity premia. Although empirical work tends to favor calibration to certain moments of consumption and dividend growth dynamics, we estimate the dynamics directly using the Bayesian MCMC method and use the estimates for inference and asset pricing implications. The method allows a joint estimation of the entire time series of the unobserved conditional mean of consumption growth and parameter values without estimation risk. Nor does it require any optimization techniques.

Using equity returns, consumption growth, dividend growth, and interest rate data we find a small persistent component in the mean consumption growth. According to the model's restrictions, the autocorrelations of the latent consumption growth mean must decline with horizon at the same rate as that of consumption growth. Our estimation results suggest that like consumption growth, the latent state variable is only moderately persistent with the corresponding half life of consumption growth shocks of around 1.3 years. At the same time, the model implies a deterministic affine relation between the price-dividend ratios and the latent process, thus imposing perfect correlation between the two processes (one-factor structure). One consequence of the affine restriction is that the price-dividend ratio autocorrelations must have the same decay rates as those of the mean of consumption growth. Since price-dividend ratios are highly persistent, the model's affine restriction is not supported by the data. When recast in terms of price-dividend ratios, the estimates of the model suggest a half life of price-dividend shocks of approximately 13 years.

We find that the moderate persistence of the conditional mean of consumption growth contributes to the equity volatility puzzle. Our estimation results suggest that the inverse of IES is not large; 0.3 in annual data. We find that the model fails to match equity volatility due to low IES combined with low volatility of the conditional mean of consumption growth. The model explains only 50% of the total return volatility.

In the present one-factor asset pricing model, consumption growth is the only pricing factor. Equity premium is determined by the equity beta on consumption growth and by the price of consumption growth risk. Our estimates show that the model can fit the equity premium only through the high price of consumption risk. The latter is directly related to the relative risk aversion. We find that the relative risk aversion consistent with the equity premium is 140.

We estimate a version of the Bansal and Yaron (2004) model with persistent conditional mean and homoskedastic innovations in consumption growth. An important avenue for future research is Bayesian estimation of the extended model based on heteroskedastic consumption growth, cross-section of asset returns, and built-in cointegration between consumption and dividend processes. Furthermore, although we use traditional construction of the consumption growth process using data on nondurables and services, we omit durables from consumption. An extension to our work along the lines of Yogo (2006) would be to incorporate durable consumption into the model.

Appendix A. Derivations

A.1. Derivation of risk-free return dynamics in (13)

In Bansal and Yaron (2004) model, the risk-free rate between t and $t+1$, $r_{f,t+1}$ satisfies the following condition:

$$r_{f,t+1} = -\theta \ln(\delta) + \frac{\theta}{\psi} E_t(g_{t+1}) - (\theta - 1) E_t(r_{a,t+1}) - \frac{1}{2} var_t(m_{t+1}), \tag{A1}$$

where $r_{a,t+1}$ is the return on an asset that pays consumption as dividend, $E_t(r_{a,t+1})$ and $\text{var}_t(m_{t+1})$ are the conditional mean of $r_{a,t+1}$ and conditional variance of the Epstein–Zin stochastic discount factor, respectively. Both $E_t(r_{a,t+1})$ and $var_t(m_{t+1})$ are constant (not functions of time) in the homoskedastic version of the Bansal and Yaron (2004) model. Using dynamics of consumption growth from (1), taking conditional expectation, simplifying, and adding a measurement error we rewrite (A1) as follows:

$$r_{f,t+1} = \alpha_f + \frac{1}{\psi} x_t + \sigma_f e_{f,t+1}. \tag{A2}$$

A.2. Multiperiod return beta on P/D ratio

First, we derive the regression beta of a k-period excess return on price-dividend ratio implied by the model, that is, beta in the following regression:

$$r_{t \to t+k} = a + b_k z_t + \varepsilon_{t+k}, \tag{A3}$$

where z_t is the $\ln(P_t/D_t)$ ratio.

The expression for the regression coefficient for a k-period return is

$$b_k = \frac{\text{cov}(r_{t \to t+k}, z_t)}{\text{var}(z_t)}. \tag{A4}$$

Also, let $r_{t\to t+k} = \sum_{j=1}^{k} r_j$, where $r_j \equiv r_{t+j-1\to t+j}$.

The dynamics of z, which is perfectly correlated with latent variable x according to (5), inherit the properties of the latent process:

$$z_{m,t+1} = (1-\rho)A_{0m} + \rho z_{m,t} + A_{1m}\varphi_e \sigma e_{t+1}. \quad \text{(A5)}$$

Taking expectations and assuming stationarity ($|\rho| < 1$), we obtain the unconditional variance of price-dividend ratio as

$$\text{var}(z_t) = A_{1m}^2 \text{var}(x_t). \quad \text{(A6)}$$

Below we measure all second moments in terms of $\text{var}(x_t)$. For example, in these units, $\text{var}(z_t) = A_{1m}^2$.

Substituting the dividend growth dynamics from (18) into (14) and subtracting the dynamics of risk-free rate given in (13) leads to the following process for excess return:

$$\begin{aligned} r_{t\to t+1} = {} & \alpha_m - \alpha_f + \mu_m + \kappa_{1,m}A_{1,m}x_{t+1} + \left(\phi_m - \frac{1}{\psi} - A_{1,m}\right)x_t \\ & + \varphi_m \sigma u_{t+1} + \sigma_m \varepsilon_{m,t+1} - \sigma_f e_{f,t+1}. \end{aligned} \quad \text{(A7)}$$

Recognizing that $\text{var}(x_{t+j}, x_t) = \rho^j$, we write down the unconditional covariance between r_j and the price-dividend ratio at time t as

$$\text{cov}(r_j, z_t) = \rho^{j-1}\left[\rho k_{1m}A_{1m}^2 + A_{1m}(\phi_m - \frac{1}{\psi} - A_{1m})\right]. \quad \text{(A8)}$$

Using the definition of A_{1m} from (19), we conclude that $\text{cov}(r_j, z_t) = 0$. The interpretation of this result is that the model predicts zero forecastability of excess log returns by log price-dividend ratio. The lack of predictability follows directly from the definition (A4) and identity $\text{cov}(r_{t\to t+k}, z_t) = \sum_{j=1}^{k} \text{cov}(r_j, z_t)$. The last result implies two things. First, excess returns in the model are all white noise with constant drift. Therefore, no variable can forecast them. Second, most of the difference in the volatilities of equity and risk-free rates must come from conditional volatilities. The volatilities of the conditional means of all returns in the model (including the risk-free rate) are identical and relatively small. The latter property is inherited from the small volatility of the latent state variable x_t.

However, the regression beta of a multiperiod return (as opposed to excess return) on price-dividend ratio is not zero. We obtain this result by simply removing the $\frac{1}{\psi}$ term from (A7) and (A8):

$$\text{cov}(r_j, z_t) = \rho^{j-1}\left[\rho k_{1m}A_{1m}^2 - A_{1m}(A_{1m} - \phi_m)\right] \quad \text{(A9)}$$

$$\text{cov}(r_{t\to t+k}, z_t) = \sum_{j=1}^{k} \text{cov}(r_j, z_t) = \frac{1-\rho^k}{1-\rho}\left[\rho k_{1m}A_{1m}^2 - A_{1m}(A_{1m} - \phi_m)\right] \quad \text{(A10)}$$

$$b_k = \frac{\text{cov}(r_{t\to t+k}, z_t)}{\text{var}(z_t)} = \frac{1-\rho^k}{1-\rho}\left[\rho k_{1m} - 1 + \frac{\phi_m}{A_{1m}}\right] = \frac{1-\rho^k}{1-\rho}\frac{1}{\psi A_{1m}}. \qquad \text{(A11)}$$

The regression coefficient in (A11) is increasing in the horizon k (as long as b_k is positive).

A.3. Effect of measurement error and time-aggregation on consumption growth autocorrelation

To assess the effect of autocorrelated measurement error on estimated consumption growth autocorrelation, we assume that the innovation of the measured consumption growth is an autocorrelated process:

$$g_{t+1} = \mu + x_t + \sigma u_{t+1} \qquad \text{(A12)}$$

$$x_{t+1} = \rho x_t + \varphi_e \sigma e_{t+1}, \qquad \text{(A13)}$$

where g_{t+1} is consumption growth from t to $t+1$, ρ is the true autocorrelation, μ, σ and φ_e are constant parameters, and u_{t+1} is the innovation that includes the autocorrelated measurement error. We assume that $u_{t+1} \sim N(\rho_u u_t, 1)$, that is, an AR(1) process. Given the assumption, the sample estimate of the consumption growth autocorrelation is

$$\gamma_g = \frac{\text{cov}(g_{t+1}, g_t)}{\text{var}(g_t)} = \frac{\rho\sigma_x^2 + \rho_u\sigma^2}{\sigma_x^2 + \frac{\sigma^2}{1-\rho_u^2}} = w\rho + (1-w)\rho_u\left(1-\rho_u^2\right), \qquad \text{(A14)}$$

where $\sigma_x^2 = \frac{(\varphi_e\sigma)^2}{1-\rho^2}$ from (A13) and $w = \frac{\sigma_x^2}{\sigma_x^2 + \frac{\sigma^2}{1-\rho_u^2}}$. The measured first-order consumption growth autocorrelation is the weighted average of the measurement error persistence and that of the latent factor.

To understand the effect of time aggregation, we assume that the consumption decisions are made every period with frequency $\frac{1}{\tau} = T$ and the data generating process is given by

$$g_{t+\tau} = \mu + x_t + \sigma e_{t+\tau}, \qquad \text{(A15)}$$

$$x_{t+\tau} = \rho x_t + \varphi_e \sigma e_{t+\tau}, \qquad \text{(A16)}$$

where $g_{t+\tau}$ and $x_{t+\tau}$ are the consumption growth and latent process from t to $t+\tau$, respectively. ρ is the true autocorrelation, σ is a constant parameter, and $e_{t+\tau} \sim N(0, 1)$. Consumption growth mean reverts at the speed of ρ at the time scale τ.

Assume that we aggregate consumption over T periods for estimation and inference purposes (e.g., using annual consumption when actual consumption decisions are made with monthly frequency). In this case, consumption growth over T periods based on the aggregate consumption measure is

$$G_{t+1} = C_{t+1} - C_t \equiv \sum_{s=1}^{T} c_{t+s\tau} - \sum_{s'=1-T}^{0} c_{t+s'\tau} = \sum_{s=1-T}^{T-1} (T - |s|) g_{t+(s+1)\tau}, \quad \text{(A17)}$$

where $c_{t+s\tau}$ is consumption over a period $t + (s-1)\tau \to t + s\tau$.

Use of aggregate consumption does affect its variance and correlation with other variables through a well-known smoothing effect. This effect, in fact, may be responsible for errors in equity premium calculations. If ρ is the autocorrelation on a scale τ as in (A16), upon aggregation we have the following expression for the consumption growth covariance[7]

$$\text{cov}(G_{t+1}, G_t) = \sum_{s=1-T}^{T-1} (T - |s|) \sum_{s'=1-T}^{T-1} (T - |s'|) \rho^{|T+s-s'|}, \quad \text{(A18)}$$

$$\text{var}(G_{t+1}) = \sum_{s=1-T}^{T-1} (T - |s|) \sum_{s'=1-T}^{T-1} (T - |s'|) \rho^{|s-s'|}. \quad \text{(A19)}$$

It is possible to show that the consumption growth autocorrelation in this case has the following limiting property as the autocorrelation scale τ goes to zero:[8]

$$\gamma_G \equiv \frac{\text{cov}(G_{t+1}, G_t)}{\text{var}(G_{t+1})} \xrightarrow{T \to \infty} \frac{1}{4}. \quad \text{(A20)}$$

Appendix B. Simulation study

We simulate 50 artificial data sets of 76 observations each for a given set of model parameters. We choose model parameters so that they closely resemble our estimates of the model based on actual data. The length of each artificial sample is typical of the length of available annual consumption data sample (e.g., 1930–2005). The data for consumption growth, dividend growth, equity returns, and the risk-free rates are generated from the model (14)–(18). We run 15,000 burn-in interations and collect the next 15,000 iterations for inference. Our choice of the burn-in sample length is motivated by monitored convergence using Gelman–Rubin diagnostic. Burn-in time is longer than in a single-sample inference because convergence does not happen simultaneously in all samples. As we mention in the main text, the model is overparameterized. In the simulation study we use the same two criteria (restrictions) as in the actual data estimation to choose the optimal signal-to-noise ratio, φ_e. In choosing the optimal prior on φ_e, we make sure that the estimates are consistent with

[7] We measure all variances and covariances in terms of the variance of g_t, which is given in (23).

[8] The continuous limit is well-known (Grossman, Melino, and Shiller, 1987; Heaton, 1993).

Table 10

Simulation study estimates

We report parameter averages for a set of 50 simulated artificial data sets. The simulations are performed for a fixed set of (true) parameter values for consumption growth, dividend growth, equity returns, and the risk-free rates based on our empirical model. Each data set contains 76 observations to match the length of available annual consumption data sample in our study (1930–2005). We run 15,000 burn-in iterations and collect the next 15,000 iterations for inference. Mean and SD are posterior mean and SD, 2.5% and 97.5% are corresponding posterior bands for parameter estimates, TRUE are the corresponding true parameter values used to generate the artificial data sets and RMSE is the root mean squared error of parameter estimates, μ is the unconditional mean of g (log-growth rate of consumption); σ is the conditional SD of g; μ_m is the unconditional mean of g_m (log-growth rate of dividends on asset m); ϕ_m is the leverage ratio on expected consumption growth; φ_e is the ratio of conditional SDs of x (log-growth rate of latent factor process); ρ is the persistence parameter in the latent process; φ_d is the ratio of conditional SDs of dividend and consumption growth; α_m is the unconditional mean of asset m's return net of g_m; α_f is the unconditional mean of the risk-free rate of return; $A_{1,m}$ is the sensitivity of the log *price-dividend* ratio to latent factor innovations; ψ is the intertemporal elasticity of substitution (IES); and γ_l is the first-order autocorrelation in mean consumption growth. This table shows that our parameter estimates can be reliably backed out from simulated data with a high degree of accuracy.

Parameter	Mean	SD	2.50%	97.50%	TRUE	RMSE
μ	0.031	0.001	0.028	0.033	0.030	0.002
σ	0.0133	0.0004	0.0126	0.0142	0.0130	0.001
μ_m	0.019	0.005	0.010	0.029	0.020	0.005
ϕ_m	0.305	0.238	−0.1601	0.766	0.500	0.307
φ_e	1.028	0.035	0.959	1.093	1.000	0.045
ρ	0.429	0.054	0.321	0.534	0.400	0.062
φ_d	9.778	0.391	9.041	10.570	10.000	0.450
α_m	0.054	0.004	0.046	0.063	0.050	0.006
α_f	0.000	0.002	−0.00351	0.003	0.000	0.002
$A_{1,m}$	2.343	0.403	1.560	3.151	2.000	0.529
$1/\psi$	−0.898	0.108	−1.111	−0.687	−1.000	0.148
γ_i	0.264	0.041	0.186	0.346	0.259	0.041

(1) the first-order serial correlation in consumption growth in our simulated samples and
(2) the smallest PLC for consumption and dividend growth data.

In Table 10, we report the results of the simulations for the optimal set of priors. Reported parameter estimates are averages across all artificial data sets. We also compute posterior means, 2.5–97.5% confidence intervals for simulated estimates, and the root mean squared error (RMSE). Our estimates for all model parameters, with the exception of parameter ϕ_m, are accurate. The estimate of this parameter, the sensitivity of dividend growth to the latent consumption growth mean, is somewhat biased. However, the estimate is still within one SD from the true value.

Appendix C. WinBUGS program details

BUGS, or Bayesian Inference using Gibbs Sampling, is a software developed in Spiegelhalter, Thomas, Best, and Gilks (1995) at MRC Biostatistics Unit in Cambridge, England. We use an interactive Windows version of the original program, WinBUGS 1.4. This program is designed for Bayesian analysis of complex statistical models using MCMC methods. The implementation stage requires declarations about the model, the data (includes initial conditions for parameters and latent variables), and the prior specifications.

The model specification involves representing the model in the form of conditionally independent blocks of parameters and state variables. The conditional independence of a block implies that the full joint distribution of all quantities can be represented as a product of conditional densities of independent blocks. For example, for the model in (14)–(20), using Bayes rule, we can write the joint log-posterior for parameters Ω and x_t as follows (T is the number of observations):

$$\begin{aligned} \ln p\big(\Omega, x_t \,|\, r_m^T, g_m^T, g^T\big) &= const + \sum_{t=0}^{T-1}[\ln p(r_{m,t+1} \,|\, g_{m,t+1}, x_{t+1}, x_t, \Omega) \\ &\quad + \ln p(g_{m,t+1} \,|\, x_t, \Omega)] + \sum_{t=0}^{T-1}[\ln p(g_{t+1} \,|\, x_t, \Omega) \\ &\quad + \ln p(x_{t+1} \,|\, x_t, \Omega)] + \ln \pi(x_0 \,|\, \Omega) + \ln \pi(\Omega). \quad \text{(C1)} \end{aligned}$$

The model stage in WinBUGS requires just the specification of the densities comprising the likelihood. In our case these densities directly follow from the dynamics in (14)–(20):

$$\begin{aligned} &p(r_{m,t+1} \,|\, g_{m,t+1}, x_{t+1}, x_t, \Omega) \sim N\big(\alpha_m + g_{m,t+1} + \kappa_{1,m} A_{1,m} x_{t+1} - A_{1,m} x_t, \sigma_m^2\big) \\ &p(g_{m,t+1} \,|\, x_t, \Omega)] \sim N(\mu_m + \phi_m x_t, (\varphi_m \sigma)^2) \qquad \text{(C2)} \end{aligned}$$

$$\begin{aligned} &p(g_{t+1} \,|\, x_t, \Omega) \sim N\big(\mu + x_t, \sigma^2\big) \\ &p(x_{t+1} \,|\, x_t, \Omega) \sim N\big(\rho x_t, (\varphi_e \sigma)^2\big). \qquad \text{(C3)} \end{aligned}$$

The next step in the model stage is specification of the priors. Current version of WinBUGS generally does not support the use of either nonconjugate or improper priors. Instead, we use "just" proper priors—low precision conjugate densities. For the present study we use the following set of priors: $x_0 \sim N(0,\ 1)$, $\alpha_m \sim N(0,\ 100)$, $\alpha_f \sim N(0,\ 10)$, $A_{1,m} \sim N(0,\ 100)$, $\frac{1}{\sigma_m^2} \sim G\alpha(1,\ 0.001)$, $\frac{1}{\sigma_f^2} \sim G\alpha(1,\ 0.01)$, $\mu \sim N(0,\ 1)$, $\frac{1}{\sigma^2} \sim G\alpha(1,\ 0.001)$, $\mu_m \sim N(0,\ 100)$, $\phi_m \sim N(0,\ 1000)\,\mathbb{I}_{\phi_m \geq 0}$, $\tau_d \sim G\alpha(1,\ 0.001)$, $\rho \sim U(-1,\ 1)$, $\tau_x \sim G\alpha(1,\ 0.001)$, where τ_d and τ_x are the precisions of the dividend growth, $g_{m,t+1}$, and the latent state process, x_{t+1}, respectively. The remaining parameters φ_e, φ_m, and $\frac{1}{\psi}$ are simple functions of the above parameters:

$$\varphi_e = \frac{1}{\sigma\sqrt{\tau_x}}$$

$$\varphi_m = \frac{1}{\sigma\sqrt{\tau_d}}$$

$$\frac{1}{\psi} = \phi_m - A_{1,m}(1 - k_{1,m}\rho). \tag{C4}$$

The package correctly infers the full conditional posteriors for both parameters and latent process. During the estimation based on full data (i.e., with risk-free rate included in the data set) $\frac{1}{\psi}$ appears twice: in the risk-free rate equation and as part of $A_{1,m}$ (see the two equations in (13)). It implies that $\frac{1}{\psi}$ and $A_{1,m}$ can no longer be considered as two independent variables. To enforce this link we explicitly impose restriction (C4) during the estimation.

References

Bansal, R., 2007. Long run risks and financial markets, *Federal Reserve Bank of St. Louis Review* 89(4), 1–17.

Bansal, R., R. Dittmar, and D. Kiku, 2009. Cointegration and consumption risks in equity returns, *Review of Financial Studies* 22(3), 1343–1375.

Bansal, R., R. Dittmar, and C. Lundblad, 2005. Consumption, dividends, and the cross-section of equity returns, *Journal of Finance* 60, 1639–1672.

Bansal, R., R. Gallant, and G. Tauchen, 2007. Rational pessimism, rational exuberance, and asset pricing models, *The Review of Economics and Statistics* 74, 1005–1023.

Bansal, R. and I. Shaliastovich, 2007. A long-run risks explanation of predictability puzzles in bond and currency markets. *Working paper*, Duke University.

Bansal, R. and A. Yaron, 2004. Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.

Bansal, R., D. Kiku, and A. Yaron, 2007. Risks for the long run: Estimation and inference. *Working paper*, Duke University.

Bekaert, G., E. Engstrom, and Y. Xing, 2009. Risk, uncertainty, and asset prices, *Journal of Financial Economics* 91(1), 59–82.

Breeden, D.T., 1979. An intertemporal asset pricing model with stochastic consumption and investment opportunities, *Journal of Financial Economics* 7, 265–296.

Breeden, D.T., M.R. Gibbons, and R.H. Litzenberger, 1989. Empirical test of the consumption-oriented CAPM, *Journal of Finance* 44, 231–262.

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

Campbell, J.Y., 2003. Consumption-Based Asset Pricing, in: George Constantinides, Milton Harris, and Rene Stulz, eds., *Handbook of the Economics and Finance,* ch. 13, vol. IB (North-Holland, Amsterdam).

Campbell, J.Y. and J.H. Cochrane, 1999. By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.

Campbell, J.Y. and N.G. Mankiw, 1989. Consumption, income, and interest rates: Reinterpreting the time series evidence, in: O.J. Blanchard and S. Fischer, eds., *NBER Macroeconomics Annual* (The MIT Press, Cambridge, Massachusetts).

Campbell, J.Y. and R.J. Shiller, 1988. The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 4, 195–228.

Cecchetti, S.G., P.S. Lam, and N.C. Mark, 1990. Mean reversion in equilibrium asset prices, *American Economic Review* 80, 398–418.

Cecchetti, S.G., P.S. Lam, and N.C. Mark, 1993. The equity premium and the risk-free rate: Matching the moments, *Journal of Monetary Economics* 31, 21–45.

Chen, H., 2006. Macroeconomic conditions and the puzzles of credit spreads and capital structure. *Working paper*, University of Chicago.

Chen, L., P. Collin-Dufresne, and R.S. Goldstein, 2009. On the relation between the credit spread puzzle and the equity premium puzzle, *Review of Financial Studies* 22(9), 3367–3409.

Christiano, L., M. Eichenbaum, and D. Marshall, 1991. The permanent income hypothesis revisited, *Econometrica* 59, 371–396.

Cochrane, J.H., 1994. Permanent and transitory components of GDP and stock prices, *Quarterly Journal of Economics* 109, 241–265.

Cochrane, J.H., 2006. Financial markets and the real economy, in: J.H. Cochrane, ed., *International Library of Critical Writings in Financial Economics*, vol. 18 (Edward Elgar, London).

Colacito, R. and M. Croce, 2006. Risks for the long-run and the real exchange rate. *Working paper*, University of North Carolina.

Constantinides, G., 1990. Habit formation: A resolution of the equity premium puzzle, *Journal of Political Economy* 98, 519–543.

Constantinides, G.M. and A. Ghosh, 2008. Asset pricing tests with long run risks in consumption growth. *Working paper*, University of Chicago.

Epstein, L. and S. Zin, 1989. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–968.

Epstein, L. and S. Zin, 1991. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical investigation, *Journal of Political Economy* 99, 263–286.

Eraker, B., 2008. Affine general equilibrium models, *Management Science* 54(12), 2068–2080.

Fama, E. and K. French, 1988. Dividend yields and expected stock returns, *Journal of Financial Economics* 22, 3–27.

Gabaix, X. and D. Laibson, 2001. The 6D bias and the equity premium puzzle, in: B. Bernanke and K. Rogoff, eds., *NBER Macroeconomics Annual* (The MIT Press, Cambridge, Massachusetts).

Grossman, S., A. Melino, and R.J. Shiller, 1987. Estimating the continuous time consumption based asset pricing model, *Journal of Business and Economic Statistics* 5, 315–327.

Hall, R.E., 1988. Intertemporal substitution in consumption, *Journal of Political Economy* 96, 221–273.

Hansen, L.P., J.C. Heaton, and N. Li, 2008. Consumption strikes back? Measuring long-run risk, *Journal of Political Economy* 116(2), 260–302.

Hansen, L.P., J.C. Heaton, and A. Yaron, 1996. Intertemporal substitution in consumption, *Journal of Business and Economic Statistics* 14, 262–280.

Hasseltoft, H., 2007. The long-run risk model: Dynamics and cyclicality of interest rates. *Working paper*, Stockholm School of Economics.

Heaton, J., 1993. The interaction between time-nonseparable preferences and time aggregation, *Econometrica* 61, 353–385.

Heaton, J., 1995. An empirical investigation of asset pricing with temporally dependent preference specifications, *Econometrica* 63, 681–717.

Heston, S., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options, *Review of Financial Studies* 6, 327–343.

Jacquier, E. and R. Jarrow, 2000. Bayesian analysis of contingent claim model error, *Journal of Econometrics* 94, 145–180.

Johannes, M. and N. Polson, 2009. MCMC methods for continuous-time financial econometrics, in: Y. Ait Sahalia and L. Hansen, eds., *Handbook of Financial Econometrics*, ch. 13, vol. 2 (Elsevier Science, Kidlington, Oxford).

Kaltenbrunner, G. and L.A. Lochstoer, 2006. Long-run risk through consumption smoothing. *Working paper*, London Business School.

Kandel, S. and R.F. Stambaugh, 1991. Asset returns and intertemporal preferences, *Journal of Monetary Economics* 27, 39–71.

Kiku, D., 2006. Is the value premium a puzzle? *Working paper*, University of Pennslyvania.

Kim, C.J. and C.R. Nelson, 1999. *State-Space Models with Regime Switching* (The MIT Press, Cambridge, Massachusetts).

Lamoureux, C. and D. Witte, 2002. Empirical analysis of the yield curve: The information in the data viewed through the window of Cox, Ingersoll, and Ross, *Journal of Finance* 57, 1479–1520.

Lettau, M. and S. Ludvigson, 2001. Consumption, aggregate wealth and expected stock returns, *Journal of Finance* 56, 815–849.

Lettau, M. and S. Ludvigson, 2005. Expected returns and expected dividend growth, *Journal of Financial Economics* 76, 583–626.

Lucas, R.E., 1978. Asset prices in an exchange economy, *Econometrica* 46, 1429–1446.

Ma, J., 2007. The long-run risk in consumption and equity premium puzzle: New evidence based on improved inference. *Working paper*, University of Alabama.

Malloy, C.J., T.J. Moskowitz, and A. Vissing-Jorgensen, 2009. Long-run stockholder consumption risk and asset returns, *Journal of Finance*.

Mankiw, N.G. and M.D. Shapiro, 1986. Risk and return: Consumption beta versus market beta, *The Review of Economics and Statistics* 68, 452–459.

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

Parker, J. and C. Julliard, 2005. Consumption risk and the cross-section of asset returns, *Journal of Political Economy* 113, 185–222.

Parker, J.A., 2001. The consumption risk of the stock market, *Brookings Papers on Economic Activity* 2, 279–348.

Piazzesi, M. and M. Schneider, 2006. Equilibrium yield curves, in: D. Acemoglu, K. Rogoff, and M. Woodford, eds., NBER Macroeconomics Annual (The MIT Press, Cambridge, Massachusetts).

Renault, E., 1997. *Econometric Models of Option Pricing Errors, in Advances in Economics and Econometrics: Theory and Applications,* vol. 3 (Cambridge University Press, Cambridge).

Robert, C.P. and G. Casella, 2004. *Monte Carlo Statistical Methods*, 2nd ed. (Springer, New York).

Rubinstein, M., 1976. The valuation of uncertain income streams and the pricing of options, *Bell Journal of Economics* 7, 407–425.

Sargent, T.J., 2007. Commentary, *Federal Reserve Bank of St. Louis Review* 4, 301–303.

Spiegelhalter, D.J., A. Thomas, N. Best, and W.R. Gilks, 1995. *BUGS: Bayesian Inference Using Gibbs Sampling, Version 0.50, Technical Report* (Medical Research Council Biostatistics Unit, Institute of Public Health, Cambridge).

Sundaresan, S., 1989. Intertemporally dependent preferences and the volatility of consumption and wealth, *Review of Financial Studies* 3, 73–89.

Weil, P., 1989. The equity premium puzzle and the risk-free rate puzzle, *Journal of Monetary Economics* 24, 401–421.

Yang, W., 2006. Time-varying exposure to long-run consumption risk. *Working paper*, University of Rochester.

Yogo, M., 2006. A consumption-based explanation of expected stock returns, *Journal of Finance* 61, 539–580.

Zurek, P., 2007. Momentum and long-run risks. *Working paper*, University of Pennsylvania.