Long-Run Risk and Hidden Growth Persistence

Pakos, Michal

Michal Pakos

17 April 2013

Long-Run Risk and Hidden Growth Persistence

Michal Pakoš[1]

CERGE-EI

FEBRUARY 2013

[1]CERGE-EI, a joint workplace of Charles University in Prague and the Economics Institute of the Academy of Sciences of the Czech Republic, Politických Vězňů 7, 111 21 Prague 1, Czech Republic. Email: michal.pakos@cerge-ei.cz. URL: home.cerge-ei.cz/pakos. Financial support of the Czech Science Foundation (grant no. P403/11/2288) is gratefully acknowledged. Helpful comments from Byeongju Jeong, John Campbell, John Cochrane, George Constantinides, Max Gillman, Rick Green, Lars Hansen, John Heaton, Jan Hanousek, Burton Hollifield, Štepán Jurajda, Michal Kejak, Jan Kmenta, Christopher Otrok, Luboš Pástor, Bryan Routledge, Sergey Slobodyan, Chris Telmer, Pietro Veronesi, Evangelia Vourvachaki, Amir Yaron, Stan Zin, and Petr Zemčík, the participants of the Netherlands Econometric Study Group (NESG) Annual Meeting (Leuven, June 2010), the discussant (Ivan Shaliastovich) at the Western Finance Association Meeting 2011 (Santa Fe, New Mexico), and the participants of the 2012 Meetings of the Society for Computational Economics 2012 (Prague, Czech Republic) as well as the 2012 European Meeting of the Econometric Society (Málaga, Spain). The paper previously circulated under the name "Consumption, Asset Prices and Persistent Economic Uncertainty".

Abstract

An extensive literature has analyzed the implications of hidden shifts in the dividend growth *rate*. However, corresponding research on learning about growth *persistence* is completely lacking. Hidden persistence is a novel way to introduce long-run risk into standard business-cycle models of asset prices because it tightly intertwines the cyclical and long-run frequencies. Hidden persistence magnifies endogenous changes in the forecast variance of the long-run dividend growth rate despite homoscedastic consumption innovations. Not only does changing forecast variance make discrimination between protracted spells of anemic growth and brief business recessions difficult, it also endogenously induces additional variation in asset price discounts due to the preference for early uncertainty resolution.

Keywords: Asset Pricing, Learning, Hidden Persistence, Forecast Variance, Economic Uncertainty, Business Cycles, Long-Run Risk, Peso Problem, Timing Premium

JEL: E13, E21, E27, E32, E37, E44, G12, G14

1 Introduction

Asset prices contain both cyclical as well as long-run components. The cyclical component has been documented empirically by Fama and French (1988a,b, 1989, 1990), Ferson and Harvey (1991), Campbell and Diebold (2009), Backus, Routledge, and Zin (2010), Lustig and Verdelhan (2013) and many others. The long-run-risk component has been shown to arises from either consumption smoothing as in Lochstoer and Kaltenbrunner (2010) or learning about the mean growth rate as in **?**. In response, two separate strands within the finance literature have developed. The first one accentuates solely the cyclical component in the expected dividends and discount rates as the key driver of asset prices. Important theoretical contributions that model the cyclical component as a Markov chain include Abel (1994), Brandt, Zeng, and Zhang (2004), Cagetti et al. (2002), Cecchetti, Lam, and Mark (1990, 1993, 2000a), Ju and Miao (2012) , Kandel and Stambaugh (1990), Lettau, Ludvigson, and Wachter (2008) and Veronesi (2004). Such important literature however is only partially successful in accounting for the salient features of asset prices.

The second strand of the asset pricing literature accentuates the long-run component. Bansal and Yaron (2004); Bansal, Gallant, and Tauchen (2007); Bansal and Shaliastovich (2010), Hansen, Heaton, and Li (2008) and Parker and Julliard (2005) are the key contributions. The long-run risk literature is significantly more successful in accounting for a variety of observed asset-pricing phenomena. However, the exclusive reliance on the long-run risk component suggests that the variation in expected business conditions is irrelevant. In fact, Bansal, Kiku, and Yaron (2010) argue that the cyclical risk premiums are negligible and as a result changing forecasts of business conditions do not matter for asset pricing. [1] It seems however that such

[1] In fact, the cyclical risk premiums are significant only for low values of the inter temporal elasticity of substitution which leads to counter-factually high and volatile risk-free rates.

conclusion directly contradicts the extensive empirical literature that accentuates the cyclical component in asset prices.

The main thesis of our paper is that expected business conditions matter for asset pricing even in long-run risk models. We build upon the important research agenda of Bansal, Kiku, and Yaron (2010) and model consumption and dividend growth rate dynamics at two frequencies jointly: the cyclical one and the long-run one. We do not posit complete information as Bansal et al., however. In fact, we assume that perfect discrimination among the two components is not feasible. We model such incomplete information in the manner that has become standard in the modern business-cycle literature: subject the consumption and dividends to hidden shifts across various phases of the business cycle. This is consistent with Burns and Mitchell (1946) who emphasize the different dynamics of economic aggregates across the various phases of the business cycle.

Subjecting growth persistence to hidden Markov shifts is novel but could not be more relevant. First, vast literature has analyzed the implications of hidden Markov shifts only in the rate of the growth. Considering persistence is a natural extension. For example, AR(1) process has two key parameters: the long-run mean to which the process tends to revert but also the persistence, the speed of the adjustment toward the mean.

Second, stochastic persistence introduces long-run components into asset prices in a very natural manner. Moreover, the speed of adjustment in the AR(1) process for the expected consumption growth in Bansal and Yaron (2004) suddenly has a structural interpretation in terms of the average duration of high and low growth periods.

Third, hidden growth persistence nests both the business-cycle and long-run risk models of asset prices within a single framework, thereby allowing to evaluate relative contributions of each component to risk premiums.

Fourth, the fact that the persistence is actually hidden intertwines the cyclical and the long-run components, effectively bringing the long-run asset-price dynamics to the business-cycle frequency. The reason is that the inference problem with respect to the hidden duration of the recession or lost decade gives rise to "Peso" problem, a situation where even a small probability that the growth might be protracted affects asset prices dramatically. [2]

Fifth, hidden persistence dramatically magnifies the fluctuations in the forecast variance of the long-run dividend growth. Learning about persistence is thus an alternative way, in comparison to the exogenously fluctuating consumption volatility in Bansal and Yaron, to induce persistent variation in economic uncertainty and thus to generate endogenous changes in the conditional distribution of asset prices.

Sixth, the relevance of modeling growth persistence is also empirical and stems from a simple back-of-the-envelope calculation showing that decade-long spells of anemic growth in commonly calibrated two-state models with constant transition probability matrices occur about every 22,000th recession which is far too rare. [3]

The virtue of our modeling approach is that the hidden persistence is naturally nested within the Hamilton (1989) hidden Markov chain setting when a subset of hidden states have exactly the same growth rate and just differ in their average

[2]The Peso problem refers to a situation when the possibility of some infrequent event, such as significantly higher persistence of low growth, has an effect on asset prices. The event may or may not have even happened in the sample but it must be difficult to accurately predict using economic history. Note that learning about the hidden persistence of economic growth fits this description well. See Evans (1996) for a review of the Peso literature.

[3]Assume that the recession has mean duration of about a year. That is consistent with the point estimate λ_{22} presented in Table 2. Then, the probability of randomly drawing an economic slowdown of length 10 years or more equals

$$P\left(T > 10 | \,\text{recession}\right) = \int_{10}^{\infty} e^{-t} \mathrm{d}t \approx 0.00005.$$

Thus, it takes about $1/0.00005 \approx 22,000$ full recessions for the lost decade to arise. That however is far too rare for it to matter for asset pricing. In our calibration, lost decades occur once a century on average.

duration. Such stochastic duration, or persistence, may be economically thought of as the outcome of several, in our case two, distinct types of economic slowdowns, for example. The first type includes business-cycle slowdowns: the fairly common but brief fluctuations in economic activity that tend to last on average about 3 to 5 quarters. The second type may be thought of as "lost decades", that is, economic slowdowns that are less likely to occur but tend to be very protracted (Japan's lost decade is a prime example).

Although our underlying Markov chain technically features three rather than two states, the growth rate shifts between high-growth and low-growth only. That happens because the third state is needed solely to model the slowdown persistence. In fact, as in Tauchen (1986a,b) and Tauchen and Hussey (1991), our Markov chain may be thought of as a discretized AR(1) process where the persistence parameter is in addition hidden. Statistically speaking, the probability distribution of the recession duration has heavier tails than the usual exponential distribution that arises in continuous-time Markov chains with constant transition probability intensities. The fact that the second and third states feature exactly the *same* growth rate but just one tends to last so much longer is substantially different from the Markov-chain setting in Rietz (1988), Barro (2006) and Johannes, Lochstoer, and Mou (2012) where the extension of the two-state chain of Mehra and Prescott (1985b) to three states is dictated by the need to model quite a rare disastrous growth rate, one in which the economy shrinks at a large negative rate well below the rate observed during recessions. Moreover, except for Johannes et al., these studies do not focus on learning as their driving Markov model is in fact observable.

We follow the long-run risk literature and endow the representative investor with the preferences of Kreps and Porteus (1978) and Epstein and Zin (1989, 1991) as extended to the continuous-time setting by Duffie and Epstein (1992). Thus, our investor cares deeply about the timing profile of the resolution of economic uncertainty. Relaxing the independence axiom of the expected utility framework is relevant as Dreeze and Modigliani (1972) observe that the utility in dynamic models is derived not only from the level of state-contingent consumption over time but also from the way in which the uncertainty about future resolves over time. Such timing of the resolution of uncertainty entails no utility gain or loss in a framework built around expected utility theory of von Neuman and Morgenstern. In the words of Chew and Epstein (1989) and Ma (1998), investors do not demand extra risk premiums or price discounts for entering actuarially fair gambles the uncertainty of which resolves far in the future rather than now. Thus, the so-called "timing premiums in expected returns" are absent. In our setting, however, the significant and persistent fluctuations in the investor's beliefs regarding future rate as well as persistence of economic growth do induce corresponding variation and persistence of the forecast variance of the long-run dividend growth and thus generate additional "timing discounts" in asset prices. The setting of Markov shifts dramatically improves the match of the salient features of the conditional distributions of asset prices when the shifts are hidden rather than observable as in Abel (1994).

There are in addition a few important related papers. Cecchetti, Lam, and Mark (2000b) study near-rational investors who tend to underestimate the average duration of economic slowdowns. Shorter slowdowns however tend to make the risk premiums for cyclical fluctuations in Bansal, Kiku, and Yaron (2010) lower. Next, Johannes, Lochstoer, and Mou (2012) explore Bayesian model averaging of the two-state (à la Mehra and Prescott (1985a)) and three-state (à la Rietz (1988) and Barro (2006)) hidden Markov chain for consumption growth. Their model does not feature a long-run dividend risk. Finally, Ai (2010) and Edgea, Laubacha, and Williams (2007) study learning about the long-run mean of AR(1) process. However, Kalman filter does not allow to learn about the persistence of AR(1) process.

The paper is organized as follows. We present the formal model and derive the theoretical asset pricing implications of hidden growth persistence in Section 2. We describe data in Section 3 and the results of statistical inference in Section 4. We describe the empirical asset pricing implications of hidden growth persistence in Section 5. We conclude in Section 6. The detailed mathematical proofs are relegated to online Appendix.

2 Model

We start the analysis by describing the investor's preferences. After that, we specify the dynamics of the dividends from all the assets. We then solve the recursive Bayesian inference problem. Next, we decompose the total variance of the long-run dividend growth rate into the variance of the conditional mean plus the mean of the conditional variance, show that the former changes endogenously whereas the latter is a constant and interpret economic uncertainty in terms of the variability of the mean long-run forecasts. Finally, we set up consumption-portfolio problem and using the first-order conditions we find the equilibrium asset prices.

2.1 Preferences

Consider an endowment economy with a single long-lived asset which yields completely perishable consumption good. Furthermore, endow the representative investor with the recursive utility

$$v_t = E_t\left[\int_t^\infty u\left(c_\tau, v_\tau\right) \mathrm{d}\tau\right],$$

defined over the perishable consumption rate stream c_t and the continuation utility v_t. [4] The normalized inter-temporal utility aggregator u over the consumption rate and the continuation utility is

$$u(c,v) = \frac{\delta}{1-\frac{1}{\psi}} \frac{c^{1-\frac{1}{\psi}} - ((1-\gamma)\,v)^{\frac{1-\frac{1}{\psi}}{1-\gamma}}}{((1-\gamma)\,v)^{\frac{1-\frac{1}{\psi}}{1-\gamma}-1}}. \tag{2.1}$$

The parameter δ is the subjective discount rate, the parameter γ is the coefficient of the relative risk aversion in terms of atemporal and actuarially-fair gambles over the wealth, and the parameter ψ is the elasticity of inter-temporal substitution. The preference for early resolution of uncertainty is reflected in the convexity of the inter temporal aggregator $u\,(c,v)$ with respect to the second argument; for example, agents prefer early resolution of uncertainty for $\gamma > \frac{1}{\psi}$. The expected utility of von Neuman and Morgenstern is nested as a special case of $\gamma = \frac{1}{\psi}$.

2.2 Asset Markets and Dividends

We consider two asset classes: equities and bonds. For equities, we in addition distinguish between unlevered and levered equity. Unlevered equity, sometimes also called consumption claim, corresponds to the standard Lucas tree with which the representative agent is endowed. It is the only asset in positive net supply and we normalize it to one. The levered equity corresponds to the aggregate equity market and it may be thought of as a levered consumption claim. For bonds, we consider only purely discount real bonds that pay zero coupons, $D_t^b = 0$. We denote the universe of assets $\mathcal{A} = \{u, l, b\}$ where u is the unlevered equity, l is levered equity and b is a purely discount real bond.

[4] We agree that modeling consumer durable goods and housing is worthwhile but unfortunately it augments the state space model in the dynamic programming exercise. See in particular Aït-Sahalia, Parker, and Yogo (2004), Pakoš (2006, 2011), Piazzesi, Schneider, and Tuzel (2007) and Yogo (2006).

We specify the dynamics of the dividend streams D_t^u and D_t^l from the equities as a bivariate hidden Markov model in logs:

$$\mathrm{d}\log D_t^u = \mu_{S_t}^u \mathrm{d}t + \sigma^u \mathrm{d}Z_t^u, \tag{2.2}$$

$$\mathrm{d}\log D_t^l = \mu_{S_t}^l \mathrm{d}t + \sigma^l \mathrm{d}Z_t^l, \tag{2.3}$$

where Z_t^u and Z_t^l are two uncorrelated Brownian Motions. The model is an extension of Hamilton (1989) and Cecchetti, Lam, and Mark (1993). The instantaneous dividend volatilities σ^u and σ^l are constant whereas the predictable components $\mu_{S_t}^e$ for $e \in \mathcal{E} = \{u, l\}$ are driven by the common Markov chain S_t with the state space

$$\mathcal{S} = \{1 = \text{expansion}, 2 = \text{recession}, 3 = \text{lost decade}\}.$$

The unobservability of the underlying state induces endogenously time-varying uncertainty due to inference problems. All dividend parameters are estimated by maximum likelihood from the postwar U.S. consumption and dividend data.

2.2.1 Expected Endowment Rate

Our specification of the hidden Markov chain differs from the outstanding literature. Standard models of business-cycle fluctuations in the economic activity are naturally modeled as a *two-state* Markov chain with the state space $\{1 = \text{high growth}, 2 = \text{low growth}\}$. The high-growth state corresponds to expansions and the low-growth state to recessions. [5] Furthermore, the transition probability matrix is typically assumed constant. In our setting of very brief decision interval h, such assumption

[5] Lettau, Ludvigson, and Wachter (2008) in their explorations of the role of the Great Moderation in understanding the equilibrium asset prices also consider independent regime switches in the consumption volatility. As we shall see hereafter, the volatility of the $time-aggregated$ consumption growth rate does change in our setting endogenously due to the hidden regime shifts and we do not need to model the variation in consumption volatility as an additional Markov chain; see for example Gillman, Kejak, and Pakoš (2012).

amounts to specifying the transition probability matrix

$$P(h) = \begin{pmatrix} 1 - \lambda_{12}\, h & \lambda_{12}\, h \\ \lambda_{21}\, h & 1 - \lambda_{21}\, h \end{pmatrix}$$

in terms of positive transition probability intensities λ_{ij} for $i \neq j$. [6] Note that the duration of high-growth and low-growth periods is exponentially distributed as follows

$$\begin{aligned} \text{duration of high growth}|S_t = 1 &\sim \text{Exp}(\lambda_{12}), \\ \text{duration of low growth}|S_t = 2 &\sim \text{Exp}(\lambda_{21}), \end{aligned}$$

therefore the hazard rate of a transition is constant and the distribution displays memoryless property.

Our setting nests this canonical model of the business-cycle fluctuations in dividend growth rates. Our natural extension to a three-state Markov chain with the state space $\mathcal{S}$ defined above features one high-growth state $\mu^e_{\text{expansion}}$ for $e \in \mathcal{E}$ and two low-growth state of *equal* growth $\mu^e_{\text{recession}} = \mu^e_{\text{lost decade}}$ for $e \in \mathcal{E}$. We let the state 1 correspond to the business cycle expansion and let the states 2 and 3 correspond to two different types of slowdowns in economic activity. The first type is fairly common but also fairly brief; it has all the characteristics of the business-cycle recession. The second type is much rarer and occurs on average once a century. Moreover, it is also very persistent as in our calibration it lasts on average ten years. Following the experience of Japan in the 1990s, we refer to this state as the "lost decade". A statistical way to think of our setting is to picture a nature that tosses a biased coin the outcome of which decides whether the business-cycle recession is

[6] The intensity represents the instantaneous risk of moving from state i to state j, $\lambda_{ij} = \lim_{h \downarrow 0} \text{Prob}\left(S_{t+h} = j | \, S_t = i\right)/h$. The intensities form a matrix whose rows sum to zero so that the diagonal entries are defined by $\lambda_{ii} = -\sum_{j \neq i} \lambda_{ij}$. See for example Karlin and Taylor (1975) for a thorough introduction to the theory of continuous-time Markov chain.

going to be the short type or the long type.

The setting of our three-state model is observationally equivalent to the well-accepted two-state models for the dividend growth rates. The key difference is that the hazard rate of the low-growth state is stochastic rather than constant. In other words, the nature tosses a biased coin and its outcome decides the average duration of the slowdown in economic activity.

With these ideas in mind, we specify the transition probability matrix as

$$P(h) = \begin{pmatrix} 1-(\lambda_{12}+\lambda_{13})\, h & \lambda_{12}\, h & \lambda_{13}\, h \\ \lambda_{21}\, h & 1-\lambda_{21}\, h & 0 \\ \lambda_{31}\, h & 0 & 1-\lambda_{31}\, h \end{pmatrix}.$$

The transition intensities $\lambda_{12}, \lambda_{13}, \lambda_{21}, \lambda_{31}$ are all strictly positive. In addition, although we rule out "instantaneous" transitions between the two types of slowdowns in economic activity, these states nonetheless do "communicate" because $\mathrm{Prob}\left(S_{t+T} = j \,|\, S_t = i\right)$ is strictly positive for any positive interval T and any $i, j \in \mathcal{S}$.

We believe that the three-state Markov chain is by far the most convenient mathematical device to introduce a tractable model of the business fluctuations in which the economy from time to time experiences low-growth periods of dramatically longer duration than the typical NBER downturn. This result cannot be obtained with two-state model calibrated to NBER recessions that features constant transition probability matrix. The reason is that the duration of the states in the Markov chain are exponentially distributed but with a relatively small mean. And although particular realizations indeed differ, the differences are insufficient to generate a long recession. For example, when the average duration of the downturn is a year, lost decades tend to occur about every 22,000th recession which is very rare. In contrast, our setting with hidden persistence is different because the distribution of the downturn duration is a time-varying mixture of exponentials and that implies heavier tails of the duration distribution.

2.3 Signal-Extraction Problem

We assume that the Markov state of the economy S_t is hidden from the investor. However, its value $S_t = i$ determines the mean dividend growth rate $E\left(\mathrm{d}\ln D_t^e \middle| S_t = i\right) = \mu_i^e \mathrm{d}t$ for each $e \in \mathcal{E}$ over the interval $(t, t+\mathrm{d}t)$. The sum of the mean growth rate and the idiosyncratic shocks $\mathrm{d}Z_t^e$ that continuously hit the economy determines the realized dividend growth rate $\mathrm{d}\ln D_t^e$. The idiosyncratic shocks mask the mean growth rate and thus the true hidden state. As a result, realized dividends are but a noisy signal of the underlying state of the economy. The investor's inference problem is to optimally extract the current hidden state from the history of the dividend rate signals $\mathcal{I}_t = \left\{\left(D_\tau^u, D_\tau^l\right) \text{ for } \tau \leq t\right\}$. The outcome of the inference is a discrete posterior distribution $\pi_t = (\pi_{1,t}, \pi_{2,t}, \pi_{3,t})'$ with

$$\pi_{i,t} = \mathrm{Prob}\left(S_t = i \middle| \mathcal{I}_t\right). \tag{2.4}$$

Furthermore, the dividend growth rates in equations (2.2) and (2.3) may be decomposed into the predictable part $E\left(\mathrm{d}\log D_t^e \middle| \mathcal{I}_t\right)$ and the unpredictable part $\mathrm{d}\log D_t^e - E\left(\mathrm{d}\log D_t^e \middle| \mathcal{I}_t\right)$. The predictable part is related to the posterior distribution as

$$E\left(\mathrm{d}\log D_t^e \middle| \mathcal{I}_t\right) = E\left(\mu_{S_t}^e \middle| \mathcal{I}_t\right)\mathrm{d}t = \left(\sum_{i\in\mathcal{S}} \mu_i^e \pi_{i,t}\right)\mathrm{d}t = \widetilde{\mu}^e\left(\pi_t\right)\mathrm{d}t.$$

The unpredictable part is the instantaneous forecast error. We normalize it to have variance one and denote it by tildes:

$$\mathrm{d}\widetilde{Z}_t^e = \frac{1}{\sigma^e}\left(\mathrm{d}\log D_t^e - E\left(\mathrm{d}\log D_t^e \middle| \mathcal{I}_t\right)\right).$$

The decomposition into predictable and unpredictable parts implies that the realized dividend growth rate may be written as

$$\mathrm{d}\log D_t^e \quad = \quad \widetilde{\mu}^e(\pi_t)\,\mathrm{d}t + \sigma^e \mathrm{d}\widetilde{Z}_t^e \text{ for } e \in \mathcal{E}. \tag{2.5}$$

In the language of stochastic calculus, the process $\widetilde{Z}_t = \left(\widetilde{Z}_t^u, \widetilde{Z}_t^l\right)$ is known as the innovation process. Liptser and Shiryaev (1977) show that the innovation process is a standard bivariate Brownian motion with respect to the investors filtration $\mathcal{I} = \{\mathcal{I}_t : t > 0\}$. In the language of economics, the unpredictable part $\mathrm{d}\widetilde{Z}_t^e$ is a source of dividend news regarding the hidden state S_t because

$$\mathrm{d}\widetilde{Z}_t^e \quad = \quad \left(\frac{\mu_{S_t}^e - \widetilde{\mu}^e(\pi_t)}{\sigma^e}\right)\mathrm{d}t + \mathrm{d}Z_t^e.$$

Nonzero realizations of $\mathrm{d}\widetilde{Z}_t^e$ may come from the I.I.D. dividend shock $\mathrm{d}Z_t^e$ but also from the hidden transitions of the Markov state S_t. The term $\left[\mu_{S_t}^e - \widetilde{\mu}^e(\pi_t)\right]/\sigma^e$ measures the regime change in units of the volatility of the noise and may be thought of as the signal-to-noise ratio. Large signal-to-noise ratio tends to be particularly informative about the regime change.

When we recursively apply the Bayes rule to eq. (2.4), we obtain the filtering equations of Wonham (1964):

$$\mathrm{d}\pi_t \quad = \quad \eta(\pi_t)\,\mathrm{d}t + \sum_{e\in\mathcal{E}} \nu^e(\pi_t)\,\mathrm{d}\widetilde{Z}_t^e, \tag{2.6}$$

where the exact functional form of the drift $\eta = \eta(\pi)$ and volatility $\nu^e = \nu^e(\pi)$ is presented in Appendix A in Lemma 1.

2.4 Precision of Consumption and Dividend Forecasts and Economic Uncertainty

Economic uncertainty is measured by the precision of the long-run consumption and dividend forecasts; for example, times of heightened uncertainty are characterized by increased variance of the forecast errors. Such economic uncertainty is priced using the Epstein-Zin preferences that are configured so that an early resolution of uncertainty is preferred. From an asset pricing perspective, persistent variation in the forecast precision of the long-run consumption and dividend growth rates induces a corresponding variation in the uncertainty (or timing) discounts in asset prices.

We decompose the forecast error variance of the long-run dividend growth rates into the sum of two terms. The first term is the forecast error variance of the mean long-run dividend growth rates. The second term is the forecast error variance of the long-run idiosyncratic shocks. This decomposition is relevant because it identifies two sources of variation in the forecast precision and thus allows us to compare hidden Markov model for consumption and dividends to the canonical long-run risk model of Bansal and Yaron (2004) and Bansal, Kiku, and Yaron (2010). We find that when consumption and dividends are subject to hidden Markov shifts, the forecast error variance of the mean long-run growth growth fluctuates in response to the perceived economic uncertainty. In contrast, the forecast error variance of the mean long-run growth is constant in the Bansal-Yaron economies which means that unless the variance of the long-run idiosyncratic shocks changes the forecasts of the long-run growth display constant precision. Second, our model of the consumption and dividend growth rate implies that variance of the long-run idiosyncratic shocks is constant whereas it follows an exogenous mean-reverting process in the Bansal-Yaron economies.

In order to calculate the variance decomposition of the long-run forecast errors, we first calculate the long-run forecasts and the forecast error variance conditional on the hidden Markov state. We then condition down to the investor's coarser information sets. Finally, we compare the hidden Markov model of expected consumption and dividends to the corresponding AR(1) model from Bansal-Yaron economies.

2.4.1 Forecasts Conditional on the Hidden Markov State

The instantaneous forecasts for the growth rates are $E\left(\mathrm{d}\log D_t^e \middle| S_t = i\right) = \mu_i^e \mathrm{d}t$ for each $e \in \mathcal{E}$ and $i \in \mathcal{S}$ and the corresponding forecast errors are $\mathrm{d}\log D_t^e - E\left(\mathrm{d}\log D_t^e \middle| \mathcal{I}_t\right)$. The instantaneous homoscedasticity of the forecast errors $\mathrm{var}_t\left[\mathrm{d}\log D_t^e - E\left(\mathrm{d}\log D_t^e \middle| \mathcal{I}_t\right)\right] = \left(\sigma^e\right)^2 \mathrm{d}t$ is a consequence of the Girsanov Theorem. In contrast, the long-run forecasts of the mean dividend growth rates $g_{T|i}^e = E\left(\int_t^{t+T} \mathrm{d}\log D_\tau^e \middle| S_t = i\right) = E\left(\int_t^{t+T} \mu_\tau^e \mathrm{d}\tau \middle| S_t = i\right)$ can be calculated in semi-closed form as described in Appendix D.4. Note that because the chain may transition a number of times during the long forecast period $(t, t+T)$, the approximation for $g_{T|i}^e$ as $\mu_i^e \times T$ is imprecise. In addition, the instantaneous forecasts μ_i^e as well as the long-run ones $g_{T|i}^e$ are both conditional on S_t and therefore evolve as dependent Markov chains as well.

2.4.2 Forecasts Conditional on the Investor's Information

The investor's information sets $\mathcal{I}_t$ do not contain the hidden state S_t. Thus, the optimal forecasts with respect to $\mathcal{I}_t$ must be calculated as $g_{T|\pi}^e = \sum_{i\in\mathcal{S}} g_{T|i}^e \pi_i$. Furthermore, the forecast error variance $v_{T|\pi}^e = \mathrm{var}\left(\int_t^{t+T} \mathrm{d}\log D_\tau^e \middle| \mathcal{I}_t\right)$ may be decomposed into a sum of the forecast error variance of the mean long-run growth rate,

$$\mathrm{var}\left(\int_t^{t+T} \mu_\tau^e \mathrm{d}\tau \middle| \mathcal{I}_t\right) = \sum_{i\in\mathcal{S}} \left(g_{T|i}^e\right)^2 \pi_i - \left(\sum_{i\in\mathcal{S}} g_{T|i}^e \pi_i\right)^2,$$

and the variance of the long-run idiosyncratic shocks,

$$\mathrm{var}\left(\int_t^{t+T}\sigma^e \mathrm{d}Z_\tau \middle| \mathcal{I}_t\right) = (\sigma^e)^2 T.$$

Both the mean long-run growth forecast as well as its variance necessarily depend on the posterior distribution $\pi = (\pi_1, \pi_2, \pi_3)'$ and thus change in response to the degree of the perceived economic uncertainty. The variance of the long-run idiosyncratic shocks equals $(\sigma^e)^2 T$ and is thus constant for a given forecast horizon.

2.4.3 Comparison to Bansal-Yaron Economies

Our cash-flow model with hidden Markov shifts differs from the canonical long-run risk models of Bansal and Yaron (2004) and Bansal, Kiku, and Yaron (2010) in several important dimensions. First, the variance of the mean long-run growth $\mathrm{var}\left(\int_t^{t+T}\mu_\tau^e \mathrm{d}\tau \middle| \mathcal{I}_t\right)$ is constant in the Bansal-Yaron economies because the shocks to the expected consumption growth are by assumption homoscedastic. In contrast, when expected consumption growth follows hidden Markov chain, shocks to the instantaneous expected consumption growth $\widetilde{\mu}^e(\pi)$ are necessarily heteroscedastic as proved by Veronesi (1999, Proposition 6). Second, the variance of the long-run idiosyncratic consumption shocks in Bansal-Yaron economies is specified exogenously but is constant in our setting.

The key implication is that the setting with hidden regime shifts does not need the assumption of exogenous consumption volatility changes in order to generate fluctuating precision in the long-run growth forecasts. In fact, even with constant volatility of the idiosyncratic consumption innovations $\mathrm{var}\left(\int_t^{t+T}\sigma^e \mathrm{d}Z_\tau \middle| \mathcal{I}_t\right)$, the forecast error variance of the long-run growth changes endogenously in response to the variation in economic uncertainty as measured by the dispersion of the posterior distribution π.

2.5 Consumption-Portfolio Problem

The recursive structure of the investor's consumption-portfolio problem leads to the Hamilton-Jacobi-Bellman equation

$$0 = \max_{\{c,\omega_a \text{ for } a\in\mathcal{A}\}} \left\{ u(c,v) + \frac{1}{\mathrm{d}t} E_t(\mathrm{d}v) \right\} \tag{2.7}$$

subject to the dynamic budget constraint

$$\mathrm{d}w_t = w_t \mathrm{d}r_t^w - c_t \mathrm{d}t, \tag{2.8}$$

where w_t is the investor's financial wealth and the return on the wealth portfolio $\mathrm{d}r_t^w$ is defined by

$$\mathrm{d}r_t^w = \sum_{a\in\mathcal{A}} \omega_{a,t} \mathrm{d}r_t^a + \left(1 - \sum_{a\in\mathcal{A}} \omega_{a,t}\right) r_t^f \mathrm{d}t,$$

where r_t^f is the short-term riskless rate and $\omega_{a,t}$ is the portfolio share of the asset a.

The first-order condition w.r.t. to the consumption rate c yields the standard condition that the marginal utility of consumption u_c equals the marginal utility of wealth v_w. The first-order conditions w.r.t. portfolio weights ω_a for each $a \in \mathcal{A}$ state that the risk premiums equal the covariance of the realized returns with the negative of the marginal utility growth; see equation (2.15) in the next section.

In order to make analytical progress, we conjecture that the value function is separable across the beliefs (π_1, π_2) and the financial wealth w,

$$v(w,\pi_1,\pi_2) = h(\pi_1,\pi_2) \frac{w^{1-\gamma}}{1-\gamma}, \tag{2.9}$$

where π_3 is given implicitly as $1 - \pi_1 - \pi_2$. The first-order condition with respect to the consumption rate implies that the function $h = h(\pi_1, \pi_2)$ depends on the equilibrium price-dividend ratio of the Lucas tree,

$$h(\pi_1, \pi_2) = \delta^{\frac{1-\gamma}{1-\frac{1}{\psi}}} \left(\frac{P^u}{D^u}\right)^{\frac{1-\gamma}{\psi-1}}.$$

Equilibrium in the goods market stipulates that the aggregate demand for the consumption good equals the aggregate supply, $c_t = D_t^u$. Equilibrium in the capital market stipulates that the aggregate demands for the unlevered and levered equity and real bonds equal their fixed supply, $\omega_t^u = 1$, $\omega_t^l = 0$ and $\omega_t^b = 0$ for any given $t > 0$.

2.6 Marginal Utility and Asset Prices

When the innovations in the dividends are Brownian motions, the equilibrium asset prices necessarily follow diffusion as shown in Huang (1987). In our case, the return dynamics for each $a \in \mathcal{A}$ has the following structure:

$$\mathrm{d}r_t^a = \eta_t^a \mathrm{d}t + \sum_{e \in \mathcal{E}} \sigma_{a,t}^e \mathrm{d}\widetilde{Z}_t^e, \tag{2.10}$$

where

$$\mathrm{d}r_t^a = \frac{\mathrm{d}P_t^a + D_t^a \mathrm{d}t}{P_t^a}.$$

The expected return η_t^a and the return volatilities $\left(\sigma_{a,t}^u, \sigma_{a,t}^l\right)$ are determined by market clearing in general equilibrium. In fact, the first-order conditions with respect to the portfolio weights ω_a state

$$E_t(\mathrm{d}r_t^a) - r^f \mathrm{d}t = -\mathrm{cov}_t\left(\mathrm{d}r_t^a, \frac{\mathrm{d}M_t}{M_t}\right) \tag{2.11}$$

where $M_t = \exp\left(\int_0^t u_v\left(c_\tau, v_\tau\right) \mathrm{d}\tau\right) u_c\left(c_t, v_t\right)$ is the marginal utility of the representative investor. Itô lemma says that the growth rate of the marginal utility M_t evolves as

$$\frac{\mathrm{d}M_t}{M_t} = -r^f\left(\pi_t\right)\mathrm{d}t - \sum_{e\in\mathcal{E}} \Lambda^e\left(\pi_t\right)\mathrm{d}\widetilde{Z}_t^e. \tag{2.12}$$

As usual, the negative of the expected marginal utility growth rate equals the short riskless rate r^f. Furthermore, the volatility $\Lambda^e\left(\pi_t\right)$ is the so-called risk-price function and it measures the contribution of a marginal exposure to the shock $\mathrm{d}\widetilde{Z}_t^e$ for $e \in \mathcal{E}$ to the total expected return on any asset $a \in \mathcal{A}$; see eq. (2.15). The functional forms for Λ^u and Λ^l

$$\Lambda^u\left(\pi_1, \pi_2\right) = \gamma\sigma^u + \left(1 - \frac{1-\gamma}{1-\frac{1}{\psi}}\right)\sum_{i=1}^{2}\nu_i^u\left(\frac{\Phi_i^u}{\Phi^u}\right), \tag{2.13}$$

$$\Lambda^l\left(\pi_1, \pi_2\right) = \left(1 - \frac{1-\gamma}{1-\frac{1}{\psi}}\right)\sum_{i=1}^{2}\nu_i^l\left(\frac{\Phi_i^u}{\Phi^u}\right) \tag{2.14}$$

reveal that the risk-price functions depend on the price-dividend ratio of the Lucas tree $\Phi^u = P^u/D^u$ as well as the respective partial derivatives $\Phi_i^u = \partial\Phi^u/\partial\pi_i$. The interpretation of the constant term $\gamma\sigma^u$ in Λ^u is as the Lucas-Breeden component. In fact, eq. (2.15) implies that part of the risk premiums on assets $a \in \mathcal{A}$ derives from the covariance with the consumption growth times the risk aversion coefficient. We refer to it as the "short-run risk" and it plays a marginal role in accounting for asset pricing phenomena in our model. The remaining component(s) in Λ^e for $e \in \mathcal{E}$ reflect the equilibrium compensation for the late resolution of the dividend uncertainty associated with investing into long-lived assets. Such "timing" premium is an equilibrium outcome of the investor's hedging demand for the asset in response to the fluctuations in their own uncertainty; see Merton (1973) and Veronesi (1999).

Expected utility of von Neuman and Morgenstern is nested as a special case $\gamma = 1/\psi$ with zero timing risk premiums as the independence axiom implies that agents exhibit neither preference nor dislike for the timing of resolution of the dividends uncertainty and therefore $\Lambda^u = \gamma\sigma^u$ and $\Lambda^l = 0$.

Using the definitions of the risk prices (2.13) and (2.14) allows us to rewrite the first-order conditions with respect to the portfolio weights ω^a for $a \in \mathcal{A}$ as

$$\eta^a - r^f \quad = \quad \sum_{e\in\mathcal{E}} \Lambda^e \sigma_a^e \tag{2.15}$$

by direct substitution from equations (2.10) and (2.12). Finally, although asset prices P_t^a for $a \in \mathcal{A} \setminus \{b\}$ have unit roots and tend to drift upwards, the respective price-dividend ratios

$$\Phi^a\left(\pi_{1t}, \pi_{2t}\right) \quad \equiv \quad \frac{P_t^a}{D_t^a} \tag{2.16}$$

are only functions of the posterior distribution that evolves as a stationary stochastic process. Proposition 1 in Appendix B then allows us to express these as well as the zero-coupon real bond price $P_t^b \equiv P^b\left(\pi_{1t}, \pi_{2t}, t; T\right)$ as solutions of certain differential Fichera boundary value problems.

3 Data Description

We measure nominal consumption as the sum of nominal personal consumption expenditures on nondurable goods (PCND) plus services (PCESV). The data are quarterly from 1952:I to 2011:IV and are available from the Federal Reserve Economic Data at the Federal Reserve Bank of St. Louis. We convert the nominal consumption series to real per-capita basis by dividing by the population (POP) and deflating by the end-of-quarter consumer price index (CPIAUCSL). Furthermore, we update Bansal, Gallant, and Tauchen (2007) by constructing the asset return and dividend series from 1952:I to 2011:IV. We measure the return on the market portfolio as the value-weighted return, and the dividend as the sum of total dividends, on all the NYSE and the AMEX common stocks. The data are monthly from 1952:1 to 2011:12 and are available from the Center for Research in Security Prices (CRSP) at the University of Chicago. We start by constructing the monthly stock market valuation series as the month-end value of the nominal capitalizations of the NYSE and the AMEX. We compute the monthly dividend series on the NYSE and AMEX as the difference between the nominal value weighted total return and the value-weighted capital gain. We apply this dividend yield to the preceding month's market capitalization to obtain an implied monthly nominal dividend series. We remove the large seasonal component in monthly nominal dividend series by using the X-12-ARIMA Seasonal Adjustment Program, available from the U.S. Census Bureau. We compute the quarterly real geometric return and dividend growth rate series as the monthly nominal continuously-compounded return and dividend on the NYSE and AMEX, cumulated over the quarter to form nominal quarterly series and deflated by the consumer price index (CPIAUCSL). The real dividend series is in addition converted to real per-capita basis by dividing by the population (POP). The real interest rate is the one-month nominal Treasury bill log yield that is aggregated to quarterly values and deflated by the consumer price index (CPIAUCSL).

Table 1 presents the summary statistics. The annualized quarterly mean log consumption growth rate is about 1.87 percent with the standard error of about 0.10 whereas the annualized mean log dividend growth rate is about 2.06 percent but higher standard error of about 0.36. We cannot reject the null hypothesis for the equality of the means as the bootstrap confidence interval at 5 percent significance level for the series difference comes out (-0.37, 0.28). Furthermore, the annualized consumption volatility of 1.26 percent is almost an order of magnitude smaller than the dividend volatility of about 10.38. In addition, there are economically negligible first-order autocorrelations in consumption and dividend growth rates. As regards the asset market data, the mean short riskless interest rate is about 1.01 percent per year with volatility of about 1.34 percent whereas the annual equity premium is about 5.51 percent. We easily reject the null hypothesis of zero equity premium. Furthermore, the annual return volatility comes out about 16.55 percent and the annualized equity premium autocorrelation is insignificant. In addition, the average annual Sharpe ratio in our sample is about 0.33. Finally, the mean dividend yield is about 3.19 percent per year with a large annual volatility of 34 percent and annual persistence of 0.82 with the standard error of 0.09.

4 Inference

We estimate the unknown parameters for unlevered and levered dividends in the bivariate continuous-time Hidden Markov model in equations (2.2)–(2.3) by maximum likelihood from the realized quarterly consumption and dividend series. Data used as well as the parameter estimates that we report in Table 2 are expressed per quarter. In addition, Figure 1 plots the realized consumption and dividend log growth rates, the realized dividend to consumption ratio and the estimated beliefs for the sample period 1952:I – 2011:IV.

In our approach we take into account the fact that the outcomes are discretely observed by calculating the transition probabilities as the elements of the matrix exponential of the transition intensity matrix

$$\lambda = \begin{pmatrix} -(\lambda_{12}+\lambda_{13}) & \lambda_{12} & \lambda_{13} \\ \lambda_{21} & -\lambda_{21} & 0 \\ \lambda_{31} & 0 & -\lambda_{31} \end{pmatrix}.$$

Lost decades are arguably relatively rare events and we cannot hope to identify all these parameters based on $60 \times 4 = 240$ quarters of data that we have. We thus impose the arguably plausible restriction that lost decades last on average a decade and thus the inverse transition intensity $\left|\lambda_{33}^{-1}\right|$ equals 40 quarters. Furthermore, we posit that lost decades, such as the Great Depression, tend to occur only once a century and thus restrict the stationary probability of the chain as $\overline{\pi}_{\text{lost decade}} = 0.1$.

Next, the fact that the data are sampled at quarterly frequency makes it arguably difficult to detect within-quarter transitions. We thus allow only for end-of-quarter transitions and calculate the mean quarterly growth rate $E\left(\int_t^{t+1} \mathrm{d}\ln D_\tau^e \middle| S_t = i\right)$ as $E\left(\int_t^{t+1} \mu_\tau^e \mathrm{d}\tau \middle| S_t = i\right) = \mu_i^e \times \int_t^{t+1} \mathrm{d}\tau = \mu_i^e$ where we use the fact that Z_t^e is a zero-mean martingale that is statistically independent of the Markov chain S_t. In addition, we are able to calculate the conditional quarterly variance as $\operatorname{var}\left(\int_t^{t+1} \mathrm{d}\ln D_\tau^e \middle| S_t = i\right)$ as $\operatorname{var}\left(\int_t^{t+1} \sigma^e \mathrm{d}Z_\tau^e\right) = (\sigma^e)^2 \times \int_t^{t+1} \mathrm{d}\left[Z^e\right]_\tau = (\sigma^e)^2$ for all $e \in \mathcal{E}$ where we again invoke the independence between Z_t^e and S_t. The restriction of only end-of-quarter transitions moreover allows us to easily construct the outcome probability $\operatorname{Prob}\left(D_t^u, D_t^l \middle| S_t = i\right)$ as the bivariate Gaussian distribution with mean $\left(\mu_i^u, \mu_i^l\right)'$ and variance-covariance matrix $\operatorname{diag}\left\{(\sigma^u)^2, \left(\sigma^l\right)^2\right\}$. Note that we restrict the growth rates μ_2^e and μ_3^e in the two slowdowns in economic activity to be exactly equal which is our way of modeling hidden persistence. [7] In addition, we restrict the dividend volatilities σ^u and σ^l to be constant across the hidden regimes which is exactly consistent with our parsimonious specification in equations (2.2)–(2.3) when

[7] The reason for that was discussed before; equal growth rates in the downturn states effectively introduce stochastic average downturn duration into the well-accepted two-state Markov model of business-cycles.

regime shifts may occur only at the end of the quarter.

Overall, the vector of unknown parameters is $p = \left(\lambda_{11}, \lambda_{12}, \lambda_{13}, \lambda_{22}, \mu_1^u, \mu_2^u, \sigma^u, \mu_1^l, \mu_2^l, \sigma^l\right)'$. The construction of the likelihood function $L\left(p|\mathcal{I}_T\right)$ follows Hamilton (1989). Next, we maximize the likelihood by initiating the Markov chain in its stationary distribution. We then invoke several global optimization algorithms over 10×1 parameter vector p where we must impose upper and lower bounds on all the estimated parameters. We subsequently polish the optimum so found to greater accuracy using local derivative-free algorithm. The achieved log-likelihood is $\log L = 1256.13$. Table 2 reports the point estimates as well as the standard errors for the transition intensities, and consumption and dividend growth rates and volatilities.

4.1 Transition Probability Intensities

The average duration of the high-growth expansion period comes out about $\left|\hat{\lambda}_{11}^{-1}\right| \approx 23.8$ quarters or about 6 years whereas the average duration of the brief low-growth recession is about $\left|\hat{\lambda}_{22}^{-1}\right| \approx 4.2$ quarter, or about a year. Furthermore, the probability $q_{12} = \frac{\lambda_{12}}{\lambda_{12}+\lambda_{13}}$ of a transition from expansion to short-lived recession, conditional on the transition occurring, comes out $\hat{q}_{12} \approx 0.92$ and thus about every 12th–13th recession tends to be a protracted lost decade. Moreover, the fact that the stationary distribution of the chain is estimated to be $\overline{\pi} = (0.73, 0.17, 0.10)$ implies that a particular century tends to experience about 73 years of good-times that are interrupted by about 17 years of the brief business-cycle recessions, each of the average duration of about a year, where the remaining $100 - 73 - 17 \times 1 = 10$ years are spent in the period of the long anemic economic growth (i.e., lost decade).

4.2 Instantaneous Forecasts

The instantaneous rates of consumption growth in the high-growth state and low-growth state are estimated to be quite conservative; consumption tends to grow on average at the annualized rates of about $0.592\% \times 4 = 2.37\%$ and $-0.267\% \times 4 = -1.07\%$, respectively. Conservative estimates for the growth rates $\hat{\mu}_1^u$ and $\hat{\mu}_2^u = \hat{\mu}_3^u$ decrease the consumption risk in the economy which tends to worsen the asset-pricing implications. Furthermore, we estimate the consumption volatility to be $\hat{\sigma}^u = 0.543\% \times 2 = 1.09\%$ p.a. The estimates for dividends are conservative as well. The high-growth rate is estimated to be $\hat{\mu}_1^l = 0.764\% \times 4 = 3.06\%$ whereas the growth rate in the recession period is $\hat{\mu}_2^l = -7.55\%$. Conservative estimates for the growth rates $\hat{\mu}_1^l$ and $\hat{\mu}_2^l = \hat{\mu}_3^l$ decrease the dividend risk in the economy which tends to worsen the asset-pricing implications for the aggregate stock market. Finally, we estimate the dividend volatility to be $\hat{\sigma}^l = 5.08 \times 2 = 10.1\%$ p.a.

Compared to the latest asset-pricing literature, the calibration of the two-state HMM model in Ju and Miao (2012, Table II) specifies the consumption growth rate in the low-growth period to be -6.79% per year. Our consumption growth rate of -1.07% per year is arguably more conservative. In addition, their calibration of the dividend growth rate in the low-growth period is $-6.79\% \times 3 = -20.37\%$ per year, where the number 3 reflects the dividend leverage, whereas our dividend growth rate of -7.55% per year is again arguably more conservative.

4.3 Long-Run Forecasts

The annual consumption growth rate forecasts are $g^u_{T=4|i=1} = 2.16\%$ in expansions whereas $g^u_{T=4|i=2} = 0.09\%$ in downturns and $g^u_{T=4|i=3} = -0.9\%$ in lost decades where time is measured in quarters. The reason why the consumption growth is forecast to grow at positive rate of 0.09% during downturn is simple; the average duration of a downturn is one year and hence a transition to the expansion state is very likely.

In addition, the decade-long forecasts are also quite relevant. In particular, consumption growth is forecast to grow $g^u_{T=40|i=1} = 1.81\%$ expansions, $g^u_{T=40|i=2} = 1.51\%$ in downturns and $g^u_{T=40|i=3} = 0.01\%$ in lost decades. In fact, the lost decade $s = 3$ may be thought of as a protracted, decade-long, period of anemic growth during which the consumption level is forecast to remain the same. Note that the estimation procedure does not impose any constraints on the consumption growth rates themselves; it restricts only the average duration $\left|\lambda_{33}^{-1}\right| = 10$ years and the relative frequency in century-long series (i.e., $\overline{\pi}_{\text{lost decade}} = 0.1$). To that extent, the fact that the decade-long consumption-growth forecast comes out zero is dictated by the realized consumption and dividends series but nonetheless it is consistent with our interpretation of the long recession as the lost decade.

As regards dividends, the decade-long forecasts come out $g^l_{T=40|i=1} = 1.32\%$ in expansions, $g^l_{T=40|i=2} = 0.42\%$ in downturns and $g^l_{T=40|i=3} = -4.23\%$ in lost decades. For example, the cumulative drop in dividends over the whole 10-year duration of the rare lost decade is forecast to be $-4.23\% \times 10 = -42.30\%$ which is quite a plausible number. For example, the calibration of the two-state HMM model in Ju and Miao (2012, Table II) implies that the cumulative drop in the dividends during their down-state is about $-6.79\% \times 3 \times 2 = -40.74\%$, where the number 3 reflects the dividend leverage and the number 2 is the average duration of the down state. However, our model predicts dividends to fall by about -42% once a century whereas Ju and Miao predict the dividends to fall by about -40% during each business-cycle downturn. In fact, our model predicts that the mean cumulative decline in realized dividends over the business-cycle recessions is about $-3.98\% \times 1 = -3.98\%$. This quantitative comparison shows that much less dividend risk, measured in terms of the difference in the hidden growth rates $\mu^u_1 - \mu^u_2$ and $\mu^l_1 - \mu^l_2$, is needed when the growth persistence itself is subject to change, as opposed to the common two-state models of asset prices that feature constant persistence.

5 Asset-Pricing Implications

We first assess the performance of the model in terms of the unconditional moments for unlevered and levered equity as well as the short- and long-term yields on real zero-coupon bonds. Because we do not have closed-form solutions for the moments, we use Monte Carlo methods to integrate conditional moments with respect to the stationary density for the state variables – the posterior distribution. In addition, we use Hamilton (1989) filter to estimate the beliefs $\widehat{\pi}_{i,t} = \widehat{\text{Prob}}\,(S_t = i | \, \mathcal{I}_t)$ and then assign asset prices to such sequence of estimated beliefs. We plot the so-calculated asset prices in Figure 3. Second, we discuss the properties of the price-consumption and price-dividend ratios and explain their convexity as the outcome of the investor's hedging demands against his own uncertainty. Third, we calculate the long-horizon equity premium and equity volatility by proper time aggregation which leads to a Fichera boundary value problem to be solved numerically. We explain how the variation in economic uncertainty gives rise to endogenous changes in risk premiums and volatility. Fourth, we evaluate the serial correlation in the dividend growth rates and excess equity returns. We calculate variance ratios and find that in contrast to the dividend growth rates excess returns display pronounced mean reversion with variance ratios well below one. The results of the predictive regressions furthermore indicate that long-run dividend growth rates are close to unpredictable whereas the results for the excess returns display increasing slope coefficients and R^2s, consistent with the empirical finance literature. Fifth, we calculate the average excess returns, volatility and Sharpe ratios across different phases of the business cycle. Our model is the first long-run risk model to account for the decreasing pattern of excess returns, volatility and Sharpe ratios over the phases of the expansion and the increasing pattern of the excess returns, volatility and Sharpe ratio over the recession. Our Monte Carlo results produce average risk premiums and Sharpe ratio over the expansion and recessions that are surprisingly close to the point estimates reported in Lustig and Verdelhan (2013, Table 2).

5.1 Unconditional Moments

Table 4 in Panel A reports the pricing moments from the Monte Carlo simulation to be compared to Table 1. First, the riskless interest rate measured as the yield-to-maturity on a real zero-coupon bond maturing in 3 months comes out about 1.20 percent per year with annual volatility of about 0.59 percent whereas the equity premium on the levered long-lived asset is about 5.89 percent per year with a large volatility $\sigma\left(E_t\left(R_{t+1}-R_t^f\right)\right)$ of about 3.61 percent. The average conditional volatility of returns $E\left(\sigma_t\left(R_{t+1}-R_t^f\right)\right)$ comes out about 15.88 percent and varies a lot across the cycle; its standard deviation $\sigma\left(\sigma_t\left(R_{t+1}-R_t^f\right)\right)$ is about 2.32 percent. Note that the cautionary observation in Abel (1999) that one may be accidentally accounting for the large equity premium with a large term premium does not apply as the yield-to-maturity on a long-term zero-coupon bond $-\left(\frac{1}{30}\right)E\left(\log P_{t,30}^b\right)$ comes out about 0.08 percent per year. The average dividend-yield $E\left(D_t^l/P_t^l\right)$ is large, about 4.26 percent with the volatility $\sigma\left(\log D_t^l-\log P_t^l\right)$ of 11.96 percent, about a half in comparison to the sample estimates in the literature, which is large especially if we take into account the fact that the economy grows at either high or low rate. The dividend yield is also quite persistent, its first-order auto-correlation is about 0.80 when annualized. When the estimated beliefs $\widehat{\pi}_t$ are inserted into the model, the performance reported in Table 4 (Panel B) partially weakens but many important moments remain almost intact. The equity premium comes about more than 5.00

percent per year, for example. We conclude that the model does a surprisingly good job of accounting for the average returns on financial assets.

5.2 Price-Consumption and Price-Dividend Ratios

Figure 2 (top left) plots the price-consumption ratio Φ^u as a function of the posterior probabilities (π_1, π_2) and Appendix describes the numerical approach used. The graph lies in the plausible interval 96 and 106 and has the additional two properties. First, the price-consumption ratio is a strictly increasing function on the whole simplex domain. This result is directly implied by our ordering of the states as $\mu_1^e > \mu_2^e = \mu_3^e$ for $e \in \mathcal{E}$. Second, the price-consumption ratio is also a convex function on the whole domain. This result, first analyzed by Veronesi (1999) with time-additive exponential utility, has the following economic intuition. The high-growth state tends to last long enough for a high confidence about the regime to develop. As the economy is constantly hit by Gaussian shocks, consumption news tend to positive as well as negative with equal probability. The positive news are however fairly uninformative; there is no higher growth state to shift to and therefore the news only reinforce the already high confidence about the up state. Negative news, though, are confusing. They may originate not only from a hidden regime shift but also from a disappointing idiosyncratic innovation in consumption growth without any change of the regime at all. It is difficult to disentangle the correct source of the news but the Bayes rule in fact suggest to partially adjust the beliefs π_1 as well as $\pi_2 + \pi_3$ toward $1/2$ in response to negative consumption growth innovation. The end result is that the expected consumption growth $\widetilde{\mu}^u = \widetilde{\mu}^u(\pi)$ falls and so does the value of the Lucas tree. What is more, such rise in the economic uncertainty tends to lengthen the time until the uncertainty partially resolves in terms of the posterior odds. That is particularly disliked by the Epstein-Zin households and leads to further discounts in the price of the tree. [8] The intuition is similar when good news comes during bad times. The expected consumption growth rises and this effect pushes the Lucas tree price upward. It is the simultaneous rise in the economic uncertainty, though, that works in the opposite direction due to a rise in average time to the partial resolution of uncertainty which again induces the timing discounts in the prices of all long-lived assets, including the Lucas tree. Figure 2 (top left) plots the price-dividend ratio Φ^l. We find that the properties are similar to the price-consumption ratio but tend to be more pronounced.

5.3 Time-Varying Equity Premium and Volatility

Figure 2 (middle left) plots the annual equity risk premium $E\left(r^l - r^f \,\middle|\, \pi_1, \pi_2\right)$ and equity volatility $\sigma\left(r^l - r^f \,\middle|\, \pi_1, \pi_2\right)$ as functions of the beliefs (π_1, π_2). [9] Both functions display pronounced variation. The equity premium varies from less than 1 percent per year to almost 25 percent per year whereas the equity volatility changes from less than 11 percent to almost 25 percent. Furthermore, both moments display pronounced convexity and each tends to attain its maximum when the economic uncertainty is largest, which happens when investor's uncertainty is about expansions versus lost decades (i.e., $(\pi_1, \pi_2, \pi_3) \approx \left(\frac{1}{2}, 0, \frac{1}{2}\right)$). However, what is interesting is that even in case of the uncertainty being solely between the two types of recessions (i.e. $(\pi_1, \pi_2, \pi_3) \approx \left(0, \frac{1}{2}, \frac{1}{2}\right)$), the moments still display economically significant variation between about 1–9 and 11–17 percents, respectively. What is more, this variation arises despite the fact that the transition intensities λ_{23} and λ_{32} of reaching lost decade from the regular downturn, and vice versa, are both exactly zero. The intuition is that the probability of transitioning between the two recessions is positive

[8] As in Bansal and Yaron (2004), our parameter configuration implies a preference for the early resolution of uncertainty.

[9] Note that we carefully time aggregate all reported quantities.

over finite intervals such as a year although over infinitesimal intervals it exactly zero. In turn, this fact implies that the fluctuating uncertainty as to the recession type induces a corresponding variation in the asset prices at the business-cycle frequency. In other words, the brief downturn state is not a redundant state of the underlying Markov chain as might at first have appeared at a glance. The reason is that without its presence, there would no variation in asset prices in response to changing forecasts of business conditions because the expansion lasts about 6 years and the lost decade 10 years. And these are long enough periods subject to rare transitions during which investors are able to develop large and persistent confidence about the underlying hidden state relatively easily.

5.3.1 Serial Correlation and Predictability of Dividends and Excess Returns

The conclusion of the extensive but careful empirical literature is that the variance ratios of the long-horizon excess returns are smaller than one and tend to decline with the horizon whereas those for the realized dividend growth are around one. These results suggest a negative serial correlation in realized excess returns at long-horizons. Furthermore, the predictive regressions for excess stock returns suggest that the slopes as well as the R^2s are relatively large and tend to rise with the forecast horizon. These results pose a challenge for a broad family of the equilibrium consumption-based asset pricing moments. In particular, the recent important study of Ju and Miao (2012) explores the asset-pricing implications of Bayesian learning. Their setting features solely learning about the mean growth rate. They consider several nested preference specification, including the Epstein-Zin and the recursive ambiguity models. Their conclusion is that the setting with two hidden growth rates and Epstein-Zin preferences has a significant difficulty in explaining the slopes and

R^2s from the predictive regressions.

In contrast, our findings obtained in a two-state model with hidden growth persistence as reported in Table 5 provide strong evidence in favor of large predictable variation in expected returns but not in dividend growth rates. In fact, not only do the variance ratios tend to hover around one for the dividend growth rate but they decline steeply below one for the equity excess returns. Furthermore, when we run the predictive regressions, we find that slopes and R^2s are dramatically higher for the return series in comparison to Ju and Miao (2012, Table V), and in fact match the increasing pattern found in the empirical finance literature. This novel result underscores only further the importance of modeling the Markov shifts in growth persistence in addition to shifts in the growth rate. It is the unobservability of the growth persistence in fact that endogenously induces large counter cyclical variation in economic uncertainty as measured by the forecast variance of the long-run dividend growth.

5.3.2 Business-Cycle Variation in the Risk-Return Trade-off

The recent empirical study of Lustig and Verdelhan (2013) explores the variation in the risk-return trade off across the phases of the business cycle. For each successive 3-month period after the peak and trough, they construct the annual holding-period equity returns; the first two sample moments are reported in their Table 2. Their main finding is that the returns are not only higher on average during recessions than expansions but tend to be also more volatile. They furthermore provide evidence that Sharpe ratios are significantly higher during recessions.

Our learning model with hidden growth persistence is able to account for these tantalizing results as follows. First, upon transitioning to the expansion state, the economic uncertainty in terms of the forecast variance of the long-run dividend growth tends to be secularly declining and the timing premiums in expected returns that compensate for late resolution of the dividend uncertainty is falling. Speaking quantitatively, the mean annual excess return, volatility and Sharpe ratio on the levered equity declines from about 10.0 percent, 18.4 percent and 0.52 to about 6.0 percent, 16.2 percent and 0.35, respectively. The mean equity premium across the whole expansion period comes out about 5.24 percent whereas the mean Sharpe ratio is about 0.32. These numbers compare surprisingly well to Lustig and Verdelhan (2013, Table 2) for the postwar U.S. sample 1945:1 – 2009:12 whose point estimates of the mean annual equity premium for the expansion period are about 5.28 percent and about 0.38 for the mean Sharpe ratio.

In contrast, a transition to the recession state is accompanied by a rise in the economic uncertainty. The forecast variance of the long-run dividend growth becomes increasingly larger because long duration of the recession in economic activity significantly raise the risk of a protracted lost decade period due to learning about the growth persistence. And again, investors that strictly prefer early resolution of uncertainty rationally demand extra premiums (timing premiums) in expected returns for bearing the risk of an unfavorable realization of future dividends. Speaking quantitatively, the mean annual excess return, volatility and Sharpe ratio on the levered equity rise from about 6.4 percent, 16.4 percent and 0.38 to about 10.5 percent, 18.5 percent and 0.56, respectively. The mean equity premium across the whole recession period comes out about 9.27 percent whereas the mean Sharpe ratio is about 0.50. These numbers again compare surprisingly well to Lustig and Verdelhan (2013, Table 2) for the postwar U.S. sample 1945:1 – 2009:12 who estimate the mean annual equity premium of about 9.09 percent and the mean Sharpe ratio about 0.45 during recessions.

6 Conclusion

We generalize Bansal, Kiku, and Yaron (2010) by introducing incomplete information about the cyclical and long-run components in dividend growth rates and show that cyclical components carry significant risk premiums in asset prices even when the elasticity of intertemporal substitution is greater than one. Our long-run risk model with incomplete information successfully explains a wide variety of dynamic asset pricing phenomena. The model is able to produce the level and variability of the riskless interest rate as well as the equity risk premium with low but positive yields on real zero-coupon bonds. Furthermore, the model is able to generate large variation in the conditional moments of asset prices over the business cycle. Last but not least, our model is the first long-run risk model that is able to explain the variation in the risk return tradeoff over the various phases of the business cycle that was recently uncovered by Lustig and Verdelhan (2013).

The success of the model is attributable to the following three novel results. First, we show how to map hidden shifts in persistence into the well-known Markov chain setting of Hamilton (1989). An extensive literature has analyzed the implications of hidden shifts in the dividend growth rate but corresponding research on learning about growth persistence is still missing. However, hidden persistence is a novel way to introduce long-run risk into standard business-cycle models of asset prices because it tightly intertwines the cyclical and long-run frequencies.

Second, we show that the forecast variance of the long-run dividend growth fluctuates endogenously in response to the variation in economic uncertainty as measured by the dispersion of the posterior distribution. Our variance decomposition identifies two sources of such variation in the forecast variance. The first source is related to the precision of the mean long-run growth. The mean long-run growth features constant precision in the models of Bansal and Yaron (2004) and Bansal, Kiku, and Yaron (2010) but varies endogenously in our setting. The second source is the variation in the consumption volatility which is constant in our setting but varies exogenously in the Bansal-Yaron economies. In fact, Bansal and Yaron are implicitly arguing that times of larger economic uncertainty are characterized by larger variance of *idiosyncratic* shocks to the consumption growth rather than by *fluctuating precision* of the forecast errors of the mean long-run growth as the Bansal-Yaron forecast errors display the same precision regardless of the degree of the economic uncertainty.

Finally, we show that the interaction between the preference for early resolution of uncertainty and the variability in the forecast precision of the long-run dividend growth is critical. Learning about persistence induces significant fluctuations in the forecast variance of the long-run dividend growth rates despite homoscedastic consumption innovations. The outcome is an economic environment in which rational investors tend to confuse the relatively rare but protracted spells of anemic growth with more common but brief recessions. Such uncertainty takes long and variable lags until it is only partially, but never in fact fully, resolved in terms of the posterior odds.

References

Abel, Andrew B. 1994. “Exact Solutions for Expected Rates of Return under Markov Regime Switching: Implications for the Equity Premium.” *Journal of Money, Credit and Banking* 26 (3):345–361.

———. 1999. “Risk Premia and Term Premia in General Equilibrium.” *Journal of Monetary Economics* 43:3–33.

Ai, Hengjie. 2010. "Information Quality and Long-Run Risk: Asset Pricing Implications." *Journal of Finance* 65 (4):1333–1367.

Aït-Sahalia, Yacine, Jonathan A. Parker, and Motohiro Yogo. 2004. "Luxury Goods and the Equity Premium." *Journal of Finance* 59 (6):2959–3004.

Backus, D. K., B. R. Routledge, and S. E. Zin. 2010. "The Cyclical Component of US Asset Returns." *Working Paper, Carnegie Mellon University* .

Bansal, Ravi, A. Ronald Gallant, and George Tauchen. 2007. "Rational Pessimism, Rational Exuberance, and Asset Pricing Models." *Review of Economic Studies* 74 (4):1005–1033.

Bansal, Ravi, Dana Kiku, and Amir Yaron. 2010. "Long-Run Risks, the Macroeconomy, and Asset Prices." *American Economic Review Papers and Proceedings* 100 (2):542–546.

Bansal, Ravi and Ivan Shaliastovich. 2010. "Confidence Risks and Asset Prices." *American Economic Review, Papers and Proceedings* 100 (2):537–41.

Bansal, Ravi and Amir Yaron. 2004. "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles." *Journal of Finance* 59:1481–1509.

Barro, Robert J. 2006. "Rare Disasters and Asset Markets in the Twentieth Century." *Quarterly Journal of Economics* 121 (3):823–866.

Brandt, Michael W., Qi Zeng, and Lu Zhang. 2004. "Equilibrium Stock Return Dynamics Under Alternative Rules of Learning About Hidden States." *Journal of Economic Dynamics and Control* 28:1925–1954.

Burns, A.F. and W.C. Mitchell. 1946. *Measuring Business Cycles.* New York: National Bureau of Economic Research.

Cagetti, M., Lars Peter Hansen, Thomas Sargent, and Noah Williams. 2002. "Robustness and Pricing With Uncertain Growth." *Review of Financial Studies* 15:363–404.

Campbell, Sean and Francis X. Diebold. 2009. "Stock Returns and Expected Business Conditions: Half a Century of Direct Evidence." *Journal of Business and Economic Statistics* 27 (2):266–278.

Cecchetti, Stephen G., Pok-Sang Lam, and Nelson C. Mark. 1990. "Mean Reversion in Equilibrium Asset Prices." *American Economic Review* 80 (3):398–418.

———. 1993. "The equity premium and the risk-free rate: Matching the moments." *Journal of Monetary Economics* 31 (1):21–45.

———. 2000a. "Asset Pricing with Distorted Beliefs: Are Equity Returns Too Good to Be True?" *American Economic Review* 90 (4):787–805.

———. 2000b. "Asset Pricing with Distorted Beliefs: Are Equity Returns Too Good to Be True?" *American Economic Review* 90 (4):787–805.

Chew, S. H. and Larry Epstein. 1989. "The Structure of Preferences and Attitudes Towards the Timing of Resolution of Uncertainty." *International Economic Review* 30:103–117.

Dreeze, J. and M. Modigliani. 1972. "Consumption Decisions Under Uncertainty." *Journal of Economic Theory* 5:165–177.

Duffie, Darrell and Larry G. Epstein. 1992. "Stochastic Differential Utility." *Econometrica* 60 (2):353–394.

Edgea, Rochelle M., Thomas Laubacha, and John C. Williams. 2007. "Learning and shifts in long-run productivity growth." *Journal of Monetary Economics* 54 (8):2421–2438.

Epstein, Larry G. and Stanley E. Zin. 1989. "Substitution, Risk-Aversion and the Temporal Behavior of Consumption and Asset returns: A Theoretical Framework." *Econometrica* 57 (4):937–969.

———. 1991. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis." *Journal of Political Economy* 99 (2):263–86.

Evans, Martin D. 1996. "Peso Problems: Their Theoretical and Empirical Implications." In *Handbook of Statistics*, vol. 14, edited by G. S. Maddala and C. R. Rao. Elsevier Science B.V., 613–646.

Fama, Eugene and Kenneth French. 1990. "Stock Returns, Expected Returns, and Real Activity." *Journal of Finance* 45:1089–1108.

Fama, Eugene F. and Kenneth R. French. 1988a. "Dividend Yields and Expected Stock Returns." *Journal of Financial Economics* 22:3–26.

———. 1988b. "Permanent and Temporary Components of Stock Prices." *Journal of Political Economy* 96 (2):246–273.

———. 1989. "Business Conditions and Expected Returns on Stocks and Bonds." *Journal of Financial Economics* 25:23–49.

Ferson, Wayne E. and C. R. Harvey. 1991. "The Variation of Economic Risk Premiums." *Journal of Political Economy* 99:385–415.

Gillman, Max, Michal Kejak, and Michal Pakoš. 2012. "Endogenous Consumption Volatility and the Long-Run Risk." *Manuscript, New York University* .

Hamilton, James. 1989. "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle." *Econometrica* 57 (2):357–384.

Hansen, Lars P., John C. Heaton, and Nan Li. 2008. "Consumption Strikes Back? Measuring Long-Run Risk." *Journal of Political Economy* 116 (2):260–302.

Huang, Chi-Fu. 1987. "An Intertemporal General Equilibrium Asset Pricing Model: The Case of Diffusion Information." *Econometrica* 55 (1):117–142.

Johannes, Michael, Lars Lochstoer, and Yiqun Mou. 2012. "Learning About Consumption Dynamics." *Manuscript, Columbia Business School* .

Ju, Nengjiu and Jianjun Miao. 2012. "Ambiguity, Learning and Asset Returns." *Econometrica* 80 (2):559–592.

Kandel, S. and R. Stambaugh. 1990. "Expectations and Volatility of Consumption and Asset Returns." *Review of Financial Studies* 3 (2):207–232.

Karlin, Samuel and Howard M. Taylor. 1975. *A First Course in Stochastic Processes*. Academic Press, 2nd ed.

Kreps, David M. and Evan L. Porteus. 1978. "Temporal Resolution of Uncertainty and Dynamic Choice Theory." *Econometrica* 46 (1):185–200.

Lettau, Martin, Sydney C. Ludvigson, and Jessica A. Wachter. 2008. "The Declining Equity Premium: What Role Does Macroeconomic Risk Play?" *Review of Financial Studies* 21 (4):1653–1687.

Liptser, M. and A. Shiryaev. 1977. *Statistics of Random Processes: General Theory*. Springer-Verlag.

Lochstoer, Lars and Georg Kaltenbrunner. 2010. "Long-Run Risk Through Consumption Smoothing." *Review of Financial Studies* 23:3141–3189.

Lustig, Hanno N. and Adrien Verdelhan. 2013. "Business Cycle Variation in the Risk-Return Trade-Off." *Journal of Monetary Economics, Supplement Issue: October 15–16 2010 Research Conference on 'Directions for Macroeconomics: What Did We Learn form the Economic Crises', Sponsored by the Swiss National Bank* 59 (S).

Ma, Chenghu. 1998. "Attitudes toward the timing of resolution of uncertainty and the existence of recursive utility." *Journal of Economic Dynamics and Control* 23 (1):97–112.

Mehra, R. and E. Prescott. 1985a. "The Equity Premium: A Puzzle." *Journal of Monetary Economics* 15:145–161.

Mehra, Rajnish and Edward C. Prescott. 1985b. "The Equity Premium: A Puzzle." *Journal of Monetary Economics* 15:145–161.

Merton, Robert C. 1973. "An Intertemporal Capital Asset Pricing Model." *Econometrica* 45 (5):867–887.

Pakoš, Michal. 2006. "Asset Pricing With Durable Goods and Nonhomothetic Preferences." Carnegie Mellon University Working Paper Series.

———. 2011. "Estimating Intertemporal and Intratemporal Substitutions When Both Income and Substitution Effects Are Present: The Role of Durable Goods." *Journal of Business and Economic Statistics* 29 (3):439–454.

Parker, Jonathan A. and Christian Julliard. 2005. "Consumption Risk and the Cross Section of Expected Returns." *Journal of Political Economy* 115 (1):185–222.

Piazzesi, Monika, Martin Schneider, and Selale Tuzel. 2007. "Housing, Consumption, and Asset Pricing." *Journal of Financial Economics* 83:531–569.

Rietz, Thomas A. 1988. "The Equity Risk Premium: A Solution." *Journal of Monetary Economics* 22 (1):117–131.

Tauchen, George. 1986a. "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions." *Economics Letters* 20:177–181.

———. 1986b. "Statistical Properties of Generalized Method-of-Moments Estimators of Structural Parameters Obtained from Financial Market Data." *Journal of Business and Economic Statistics* 4 (4):397–416.

Tauchen, George and Robert Hussey. 1991. "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models." *Econometrica* 59 (2):371–96.

Veronesi, Pietro. 1999. "Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model." *Review of Financial Studies* 12 (5):975–1007.

———. 2004. "The Peso Problem Hypothesis and Stock Market Returns." *Journal of Economic Dynamics and Control* 28:707–725.

Wonham, W. M. 1964. "Some Applications of Stochastic Differential Equations to Optimal Nonlinear Filtering." *SIAM Journal on Control* 2 (3):347–369.

Yogo, Motohiro. 2006. "A Consumption-Based Explanation of Expected Stock Returns." *Journal of Finance* 61 (2):539–580.

Table 1. Summary Statistics [a] [b]

Time Series	Mean Est.	Mean (S.E)	Volatility Est.	Volatility (S.E)	Autocorrelation Est.	Autocorrelation (S.E)
Dividend Growth Rates						
Unlevered	1.87	(0.10)	1.26	(0.11)	0.00	(0.00)
Levered	2.06	(0.36)	10.38	(0.74)	0.01	(0.01)
Asset Prices						
Three-Month T-Bill Rate	1.01	(0.22)	1.34	(0.17)	0.08	(0.04)
Equity Premium	5.51	(0.78)	16.55	(0.95)	0.00	(0.00)
Sharpe Ratio	0.33	(0.01)	na	(na)	na	(na)
Price-Dividend Ratio	3.19	(0.10)	0.34	(0.05)	0.82	(0.09)

[a] All variables are in logs. Standard errors obtained by performing a block bootstrap with each block having geometric distribution with length 32 quarters; 50,000 experiments performed. Sample period for consumption is 1952.I – 2012.III whereas for dividends and asset prices it is 1952.I – 2011.IV due to data availability restriction.

[b] Magnitudes are annualized by multiplying by 400, 200 and taking to power 4, where applicable.

Table 2. Maximum Likelihood Estimates of the Hidden Markov Model for Consumption and Dividends [a]

Parameter	Estimate	(S.E.)	Parameter	Estimate	(S.E.)
Transition Probability Intensities			Instantaneous Consumption Growth Parameters		
λ_{11}	-0.042	(0.019)	μ_1^u	0.592	(0.046)
$\lambda_{12}/(\lambda_{12}+\lambda_{13})$	0.923	(0.034)	μ_2^u	-0.267 [d]	(0.179)
λ_{22}	-0.236	(0.157)	μ_3^u	-0.267 [d]	(0.179)
λ_{33}	-0.025 [b]		σ^u	0.543 [c]	(0.029)
Stationary Distribution of the Markov Chain			Instantaneous Dividend Growth Parameters		
$\overline{\pi}_1$	0.773		μ_1^l	0.764	(0.358)
$\overline{\pi}_2$	0.127		μ_2^l	-1.888 [d]	(1.040)
$\overline{\pi}_3$	0.100 [b]		μ_3^l	-1.888 [d]	(1.040)
			σ^l	5.081 [c]	(0.235)

[a] We estimate the parameters of a three-state continuous-time hidden Markov model that is quarterly observed from the bivariate time series of consumption and dividends. Data as well as the reported parameter estimates are quarterly. The achieved log-likelihood is $\log(L) = 1256.13$. The initial probability is assumed to be drawn from the stationary distribution of the hidden Markov chain. The symbol $\overline{\pi}_s$ denotes the stationary probability of the state $s \in \mathcal{S}$. Sample period is 1952:I–2011:IV.

[b] Held fixed in the estimation.

[c] We impose the restriction that the instantaneous volatility σ^e is constant across the hidden regimes for $e \in \mathcal{E}$.

[d] Hidden growth persistence introduced as the restriction that $\mu_2^e = \mu_3^e$ for $e \in \mathcal{E}$.

Table 3. Long-Horizon Forecasts of the Average Dividend Growth Rates [a]

	$E\left(\frac{1}{T}\int_t^{t+T} \mathrm{d}\log D^u_\tau \middle\vert S_t=i\right)$			$E\left(\frac{1}{T}\int_t^{t+T} \mathrm{d}\log D^l_\tau \middle\vert S_t=i\right)$		
Forecast Horizon T	Expansion	Recession	Lost Decade	Expansion	Recession	Lost Decade
Instantaneous						
	2.37	-1.07	-1.07	3.06	-7.55	-7.55
Time-Aggregated						
Quarter Ahead	2.30	-0.70	-1.03	2.85	-6.41	-7.42
Year Ahead	2.16	0.09	-0.91	2.41	-3.98	-7.06
Decade Ahead	1.81	1.51	0.01	1.32	0.42	-4.23

[a] Results are expressed in percentages per year.

Figure 1. Asset Prices Implied by the Estimated Beliefs

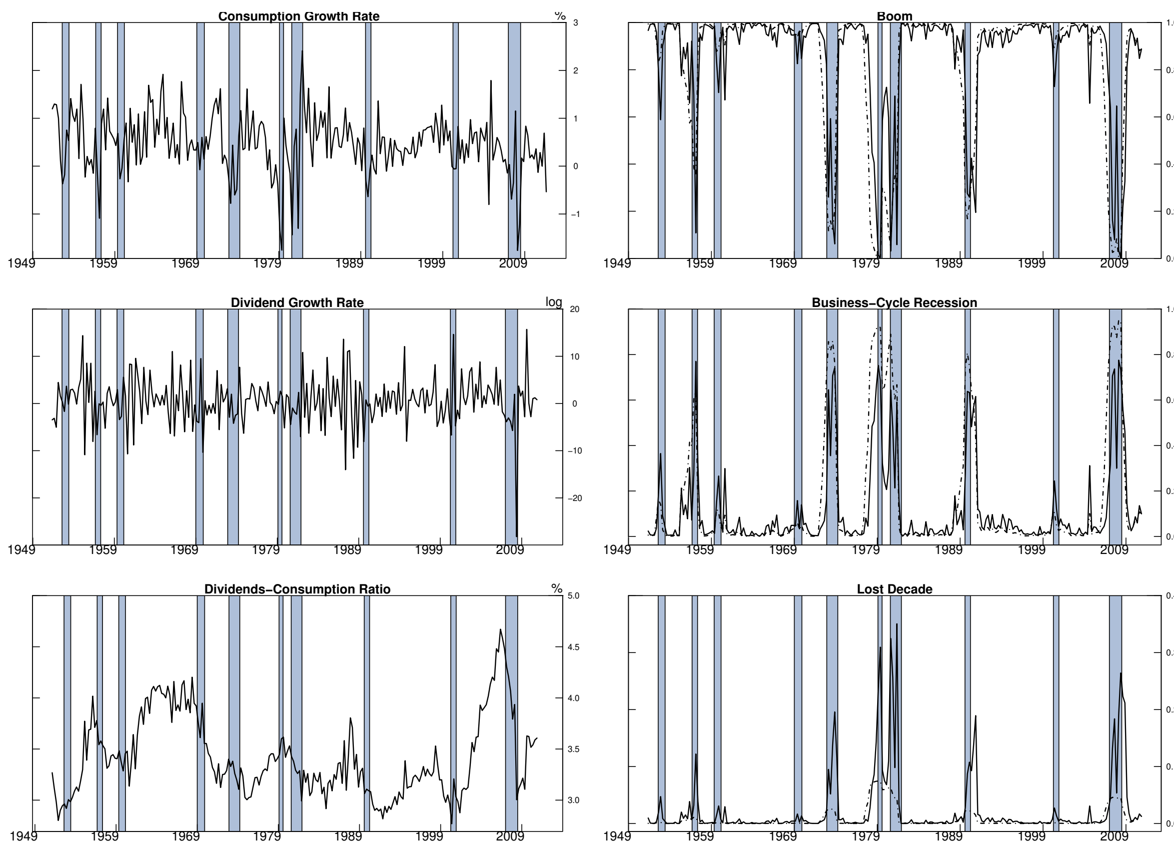

Notes. In the right panel, posterior probabilities $\text{Prob}\,(S_t = s|\,\mathcal{I}_t)$ plotted as solid lines, smoothed probabilities $\text{Prob}\,(S_t = s|\,\mathcal{I}_T)$ as dot-dashed lines. Sample period is 1952.I – 2011.IV.

Table 4. Asset Pricing Moments [a]

	Panel A. Beliefs from Monte Carlo Simulation						
Pricing Moment	Mean	S.D.	AC1	Pricing Moment	Mean	S.D.	AC1
3-Month Discount Bond				Levered Equity			
Yield-To-Maturity	1.20	0.59	0.54	Risk Premium	5.89	3.61	0.44
30-Year Discount Bond				Volatility	15.88	2.32	0.41
Yield-To-Maturity	0.08	0.09	0.65	Sharpe Ratio	0.35	0.15	0.47
Unlevered Equity				Dividend Yield	4.26	11.96	0.74
Dividend Yield	0.95	2.68	0.82				

	Panel B. Beliefs Implied by Consumption and Dividends						
Pricing Moment	Mean	S.D.	AC1	Pricing Moment	Mean	S.D.	AC1
3-Month Discount Bond				Levered Equity			
Yield-To-Maturity	1.57	0.48	0.55	Risk Premium	5.11	3.08	0.54
30-Year Discount Bond				Volatility	15.35	1.97	0.50
Yield-To-Maturity	0.10	0.08	0.57	Sharpe Ratio	0.32	0.13	0.56
Unlevered Equity				Dividend Yield	4.07	7.81	0.55
Dividend Yield	0.95	1.36	0.52				

[a] The reported asset-pricing moments in Panel A are annual and are based on Monte Carlo simulation of the length 4,000,000 quarters with the dividend growth rate calibration based on Table 2. The results in Panel B are based on the filtered probabilities Prob $(S_t = i | \mathcal{I}_t)$ obtained as in Hamilton (1989) for the sample period 1952:I–2011:IV. In both cases, data are properly time-aggregated from instantaneous to annual quantities. The preference parameter values are $\gamma = 10.0$, $\psi = 1.50$ and $\delta = 0.01$.

Figure 2. Asset Prices

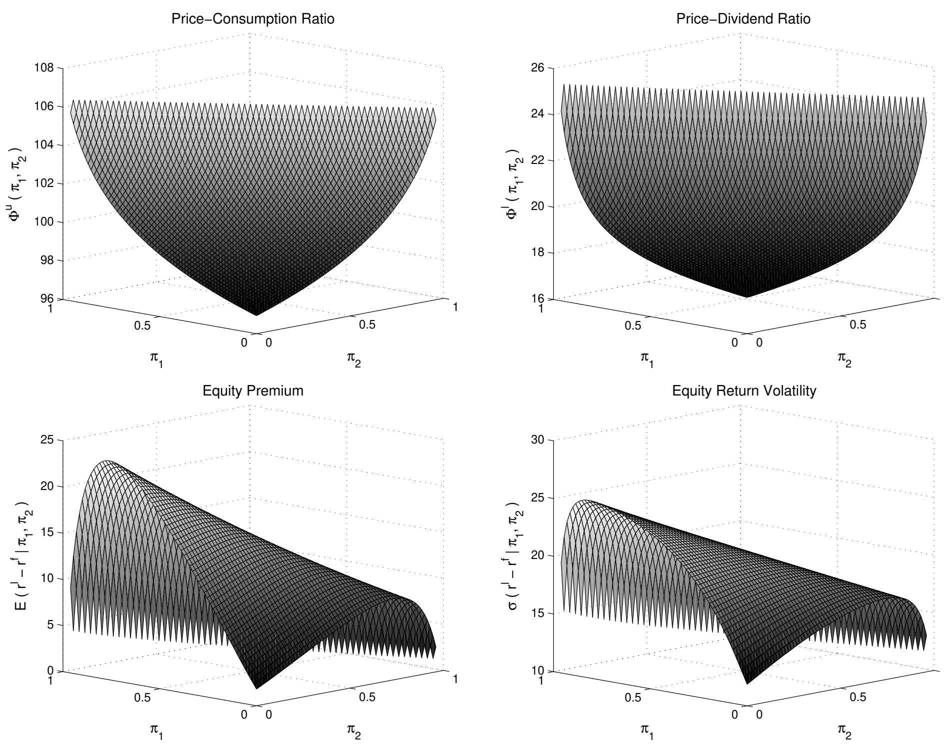

Notes. Mean equity premium and equity premium volatility obtained by appropriately time-aggregating the first-two conditional moments from instantaneous to annual (levered dividend claim).

Figure 3. Asset Prices Implied by the Estimated Beliefs

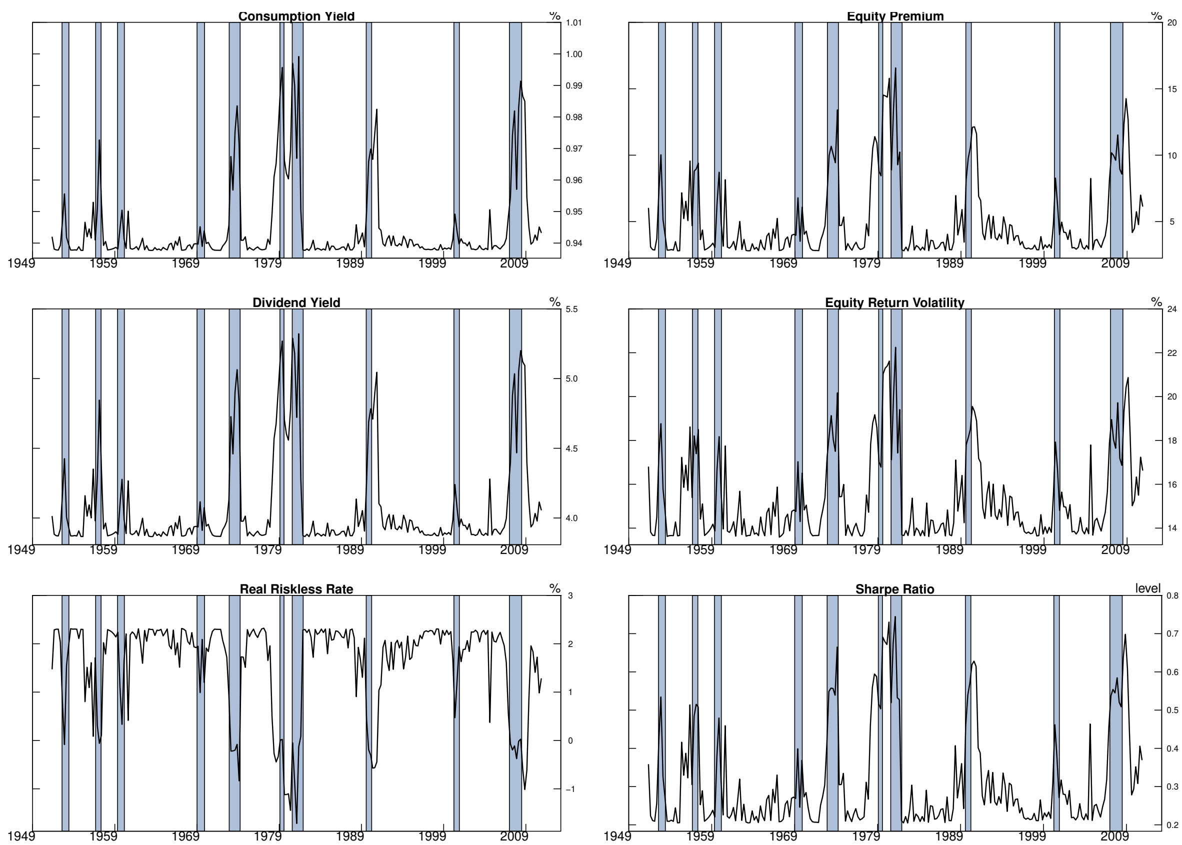

Notes. Sample period 1952.I – 2011.IV.

Table 5. Predictability and Persistence of Dividends and Excess Returns [a] [b]

Quarters Ahead	Levered Dividend Growth Rate			Levered Excess Return		
	Slope	R^2	VR	Slope	R^2	VR
4	-0.134	0.013	1.043	0.563	0.092	0.907
8	-0.156	0.013	1.059	0.875	0.125	0.790
12	-0.143	0.014	1.052	1.012	0.128	0.700
20	-0.100	0.014	1.011	1.106	0.111	0.578
32	-0.036	0.016	0.923	1.150	0.093	0.475

[a] The reported entries are median values of 100,000 Monte Carlo simulations, each consisting of 240 quarters of excess returns, dividend growth rates and dividend yields. For an experiment to qualify, it must have zero occurence of lost decade(s) so that it corresponds to the postwar U.S. experience as well as possible. Variance ratios computed in the same way as Cecchetti, Lam, and Mark (1990, 2000b). The model corresponds to our 3-state HMM with hidden growth persistence where the calibration of the dividend growth rates is based on Table 2, and where the preference parameter values are $\gamma = 10.0$, $\psi = 1.5$ and $\delta = 0.01$.

[b] The entries are based on the linear projection

$$\sum_{i=1}^{n} x_{t+i} = \text{Intercept} + \text{Slope} \times (d_t - p_t) + \varepsilon_{t+i}$$

for the log excess returns $x_{t+1} = r^l_{t+1} - r^f_t$ as well as the log dividend growth rate. The superscript l stands for the levered equity. Time is measured in quarters and small letters are in logs.

Table 6. Business-Cycle Variation in the Risk-Return Tradeoff [a]

	Recession: Buy in nth 3-month period after peak and sell 1 year later					Average	Expansion: Buy in nth 3-month period after trough and sell 1 year later					Average
	1st	2nd	3rd	4th	5th		1st	2nd	3rd	4th	5th	
Equity Premium												
Mean	6.43	8.34	9.34	10.00	10.53	9.27	10.04	9.70	8.34	7.07	6.14	5.24
S.D.	(3.12)	(2.65)	(2.24)	(2.07)	(2.05)	(3.26)	(4.19)	(5.17)	(5.28)	(4.85)	(4.24)	(3.14)
Equity Volatility												
Mean	16.44	17.58	18.03	18.29	18.50	17.89	18.42	18.39	17.61	16.81	16.21	15.58
S.D.	(1.98)	(1.56)	(1.32)	(1.26)	(1.28)	(1.90)	(2.46)	(3.03)	(3.19)	(3.02)	(2.70)	(2.05)
Equity Sharpe Ratio												
Mean	0.38	0.46	0.51	0.54	0.56	0.50	0.52	0.50	0.44	0.39	0.35	0.32
S.D.	(0.13)	(0.11)	(0.09)	(0.08)	(0.08)	(0.13)	(0.16)	(0.19)	(0.20)	(0.18)	(0.16)	(0.13)

[a] The reported asset-pricing moments are annual (not annualized) and are based on Monte Carlo simulation of the length 4,000,000 quarters with the dividend calibration based on Table 2. The hidden Markov chain state space is $\mathcal{S} = \{1, 2, 3\}$. Expansion corresponds to the state $s = 1$ and the recession to the state $s = 2$.