Learning about Rare Disasters: Implications For Consumption and Asset Prices*

MAX GILLMAN[1], MICHAL KEJAK[2] and MICHAL PAKOŠ[2]

[1]*University of Missouri–St. Louis and* [2]*CERGE-EI, Prague*

Abstract. Rietz (1988) and Barro (2006) subject consumption and dividends to rare disasters in the *growth rate*. We extend their framework and subject consumption and dividends to rare disasters in the *growth persistence*. We model growth persistence by means of two hidden types of economic slowdowns: recessions and lost decades. We estimate the model based on the postwar US data using maximum likelihood and find that it can simultaneously match a wide array of dynamic pricing phenomena in the equity and bond markets. The key intuition for our results stems from the inability to discriminate between the short and the long recessions *ex ante*.

JEL Classification: E13, E21, E32, E43, E44, G12

1. Introduction

Rietz (1988) proposes to model rare disasters as sudden cataclysms: short but deep declines in the standards of living. Using the large economic declines in the USA associated with World War I, the Great Depression, and World War II, Barro (2006) calibrates the probability of the disasters and argues that it is possible to account for the level of the equity premium.[1] Rietz and Barro consider a constant probability of disaster. Farhi and Gabaix (2011), Gabaix (2008, 2012), Gourio (2008, 2012, 2013), Gourio, Siemer, and Verdelhan (2013), Seo and Wachter (2013), Tsai and Wachter

---

[1] In addition, Brown, Goetzmann, and Ross (1995) study the long-term survival of financial markets while Barro and Jin (2011) and Barro and Ursua (2012) analyze the large economic declines in international macroeconomic data.

(2013), and Wachter (2013) extend their work by making the probability variable while Martin (2013) exploits cumulant-generating functions. In accordance with Timmermann (1993), Gourio (2012) suggests learning as a fruitful way to endogenize the disaster probability. Nevertheless, learning in this Barro–Rietz framework inevitably plays marginal role because the deep declines in consumption are learned almost instantaneously.

This article proposes an alternative model of rare disasters as protracted stagnation in the standards of living, so-called "lost decades" in the macroeconomics literature on depressions (Hayashi and Prescott, 2002; Kydland and Zarazaga; 2002, Bergoeing, Kehoe, and Soto, 2002). Interpreting disasters as protracted stagnation makes learning rather slow and generates a sizable increase in the magnitude as well as in the variation of economic uncertainty, thus dramatically enhancing the match of a broad range of macroeconomic and finance phenomena.

In the language of macroeconomics, the uncertainty shocks of Bloom (2009) arise endogenously as the consumption volatility fluctuates due to learning, contrary to the exogenous specification in the long-run risk models of Bansal and Yaron (2004) and Bansal, Gallant, and Tauchen (2007), Bansal, Kiku, and Yaron (2010), Bansal *et al.* (2012). In a related paper, Orlik and Veldkamp (2013) propose a different way to endogenize these uncertainty shocks.

In the language of finance, learning induces a procyclical variation in consumption and dividend forecasts and a countercyclical variation in the Epstein and Zin (1989, 1991) discount rates in response to changes in the average time to the (partial) resolution of uncertainty, and thus our model can simultaneously match a wide array of dynamic pricing phenomena in the equity and bond markets.

We follow Mehra and Prescott (1985) and consider a version of Lucas (1978) representative-agent model of asset pricing with exogenous, stochastic, and perishable dividends, as extended to a continuous-time incomplete-information setting by Veronesi (2004) and David and Veronesi (2013). Similarly to Pakoš (2013), we extend the regime-switching models of Cecchetti, Lam, and Mark (1990), David (1997), David and Veronesi (2013), Hamilton (1989) and Veronesi (1999, 2000, 2004) by subjecting consumption and dividends to hidden shifts in the growth rate and growth persistence as well. The variability in the growth persistence is modeled by considering two types of recessions with identical repressed growth rates but different mean duration: the former corresponds to a regular business-cycle recession while the latter is a rare lost decade, which happens on average once a century.

From the perspective of Mehra and Prescott (1985) and Weil (1989), our underlying hidden chain is not Markov with exponentially distributed sojourn times but rather *semi*-Markov[2] with the sojourn times following any distribution, in our case a time-varying mixture of two exponential distributions, one for each recession type. Modeling multiple recessions with different mean duration inculcates a tail uncertainty about the sojourn times as in Weitzman (2013), interpreted as long-run risk in Pakoš (2013). In related studies, Branger, Rodrigues, and Schlag (2012), and quite recently Jin (2014), emphasize the interplay between rare events and long-run risk.

In comparison to the model of Rietz (1988), semi-Markov chains can be reformulated as Markov ones by augmenting their state space. Such reformulation in our setting leads to a Markov chain with three states: expansion, short recession, and long recession, subject to the restriction that the recessions share exactly the same growth rate. We think of "a low-probability, depression-like third state" of Rietz (1988) as a decade-long stagnation in consumption with the disaster probability (the subjective belief about the third state) fluctuating in response to changing economic conditions.

Our model of hidden growth persistence is closely related to Cogley and Sargent (2008) who study learning about the mean duration of recessions.[3] In their setting, the duration distribution of expansions as well as recessions is governed by fixed but unknown parameters, while their representative agent is endowed with pessimistic priors based on the negative experience of the Great Depression. Such a calibrated model matches well to many pricing puzzles in the equity market. In a related study, Collin-Dufresne, Johannes, and Lochstoer (2013) point out that the best unbiased estimate of a fixed but hidden parameter is a martingale that induces permanent shocks. They extend Cogley and Sargent (2008) by using the recursive preferences of Epstein and Zin (1989, 1991) so as to inculcate long-run risk into asset price dynamics.

Our analysis differs from Cogley and Sargent (2008) and Collin-Dufresne, Johannes, and Lochstoer (2013) in the following ways. First, rather than using pessimistic priors from the Great Depression, we instead estimate the consumption and dividend parameters by maximum likelihood from the postwar US data from 1952 to 2011. Second, the persistence in our model follows a hidden two-state Markov chain rather than being a fixed parameter, which has the advantage that the risk premiums are stationary. Third, each slowdown in economic activity confronts the investor with a

---

[2] See Howard (1971, Chapter 10).
[3] Weitzman (2007) and Johannes, Lochstoer, and Mou (2012) also emphasize the importance of Bayesian updating about unknown structural parameters.

Peso-type problem about the mean duration of the recession, generating a tail uncertainty as in Weitzman (2013).

The additional related literature includes Backus, Chernov, and Martin (2011), Bates (2000), Branger, Rodrigues, and Schlag (2012), Santa-Clara and Yan (2010), and quite recently Schreindorfer (2014). These studies suggest to measure the frequency and size of such disasters using the option price data and other derivatives on US equity indexes. Furthermore, while working on our article, we have come across a study of Lu and Siemer (2013) who study learning about rare events in a framework without a tail uncertainty about the recession duration.

The article is organized as follows. In Section 2, we present the formal model and derive the theoretical implications of variable growth persistence for consumption and asset prices. In Section 3 we present the results of estimation. Section 4 describes the quantitative implications of learning for consumption and asset prices while in Section 5 we present sensitivity analysis and discuss our preliminary results for option pricing. We conclude in Section 6. Detailed mathematical proofs are found in the Supplementary Appendix.

2. Model

We start by briefly describing the representative investor's preferences and specifying the hidden semi-Markov model of the cash-flow growth rates. We then go on to solve the investor's optimal consumption-portfolio problem. In order to do this, we first solve the inference problem by introducing the posterior distribution over the discrete number of hidden states and derive a recurrent relationship for its law of motion by applying the Bayes rule. We then discuss how the variation in the posterior distribution generates time-varying endogenous economic uncertainty in terms of changing forecasts as well as changing forecast error variances of the T-period cash-flow growth rates. Second, we take the posterior distribution and make it a part of the state vector in the dynamic programming problem. This is relevant as it makes the optimization problem Markovian, leading to the standard Hamilton–Jacobi–Bellman (HJB) equation. Using the derived first-order conditions and the guess-and-verify method, we then derive the pricing equations for unlevered and levered equity, real zero-coupon bonds as well as European options. The section additionally relates the real yield curves and the bond risk premiums to the term structure of the T-period forecasts as well as the T-period forecast error variances of the consumption growth rate.

2.1 PREFERENCE SPECIFICATION

The representative agent maximizes the recursive utility function of Epstein and Zin (1989, 1991) over his consumption stream c_t and the continuation utility J_t defined by the recursion

$$J_t = E_t \left\{ \int_t^\infty U(c_\tau, J_\tau) \mathrm{d}\tau \right\}, \tag{1}$$

with the CES utility aggregator

$$U(c, J) = \frac{\delta}{1 - \frac{1}{\psi}} \frac{c^{1-\frac{1}{\psi}} - ((1-\gamma)J)^{\frac{1}{\theta}}}{((1-\gamma)J)^{\frac{1}{\theta}-1}}. \tag{2}$$

In these expressions, E_t denotes the conditional expectation operator, $\delta \geq 0$ is the rate of time preference, $\gamma \geq 0$ is the coefficient of the relative risk aversion, $\psi \geq 0$ is the magnitude of the elasticity of the intertemporal substitution, and $\theta = (1-\gamma)/(1-\psi^{-1})$ is a measure of the nonindifference to the timing of the resolution of uncertainty as we relax the independence axiom of the expected utility.

The investor prefers early resolution of the uncertainty when the current marginal utility $\frac{\partial U}{\partial c}$ falls as relatively more of the consumption occurs in the future, measured by higher continuation utility J. In this case, the cross-derivative $\frac{\partial}{\partial J}\left(\frac{\partial U}{\partial c}\right)$ is negative which happens for $\gamma > \psi^{-1}$. The expected utility is nested for $\frac{\partial}{\partial J}\left(\frac{\partial U}{\partial c}\right) = 0$ which happens for $\gamma = \psi^{-1}$.

2.2 ASSET MARKETS

We endow the investor with a single Lucas tree called unlevered equity (or consumption claim), and denote it with the superscript u. The asset yields a continuous flow of dividends at the rate D_t^u. In addition, we distinguish between the total wealth, which is unobservable, and the aggregate equity market. We thus introduce levered equity denoted with the superscript l. The levered equity yields a continuous flow of dividends at the rate D_t^l, which is different from D_t^u. We refer to the unlevered and levered dividends jointly as the cash flows and distinguish them using the superscript $e \in E = \{u, l\}$. We furthermore introduce real zero-coupon bonds with maturities of up to 30 years (superscript b) and European call options (superscript c). For the simplicity of notation, we denote the class of the securities $A = E \cup \{b, c\}$. We assume that all assets in A except the unlevered equity are in zero net supply.

Our bivariate time series model for the cash-flow growth rates generalizes the standard Markov-trend model in logs introduced by Hamilton (1989) to

a semi-Markov setting by subjecting the instantaneous cash-flow growth rates $dg_t^e = d \log D_t^e$ for $e \in E$ to hidden semi-Markov shifts:

$$dg_t^e = \mu_{s_t}^e dt + \sigma^e dz_t^e. \tag{3}$$

The predictable component $\mu_{s_t}^e \in \{\underline{\mu}^e, \bar{\mu}^e\}$ with $\underline{\mu}^e < \bar{\mu}^e$ is driven by a two-state semi-Markov chain s_t which is hidden in the standard Brownian noise $z_t = (z_t^u, z_t^l)'$. We assume for tractability that z_t and s_t are statistically independent processes.

Our cash flow model in Equation (3) implies that the forecast error variance of the cash-flow growth rate over the next instant

$$(\sigma^e)^2 dt = \text{var}_t \{dg_t^e - E_t \{dg_t^e\}\} \tag{4}$$

is constant. Nonetheless, our learning model with hidden shifts generates a predictable variation in the T-period forecast error variances when the instantaneous cash-flow growth rates dg_t^e are time-aggregated from the infinitesimal decision intervals to their T-period counterparts $\int_t^{t+T} dg_\tau^e$ as shown in detail in Section 2.4. Our setting thus differs significantly from the extensive long-run risk literature where the predictable variation in the forecast error variance $(\sigma_t^e)^2$ is exogenous rather than endogenous due to learning.[4]

2.2.a. *Semi-Markov Chain*

Current literature models business-cycle fluctuations in terms of two-state hidden Markov chains with the state space

$$S = \{s_1 = \text{expansion}, s_2 = \text{recession}\}. \tag{5}$$

In a continuous-time setting, it is natural to express the transition probabilities in terms of the hazard rates of transition

$$\lambda_i = \sum_{s_j \in S \setminus \{s_i\}} \lambda_{ij}, \tag{6}$$

where λ_{ij} denotes the nonnegative transition intensity for any $s_i, s_j \in S$ and $i \neq j$. If the hazard rates are constant, the density of the sojourn time τ_i for $i = 1, 2$ is given by the exponential distribution

$$f_{\tau_i}(t) = \lambda_i \exp(-\lambda_i t) \tag{7}$$

---

[4] See in particular Bansal and Yaron (2004) and Bansal, Gallant, and Tauchen (2007), Bansal, Kiku, and Yaron (2010), Bansal *et al.* (2012).

for nonnegative t. Exponential distribution tends to be a common choice for modeling sojourn times due to the mathematical tractability allowed by the memoryless property

$$P\{\tau_i > x + y | \tau_i > x\} = P\{\tau_i > y\}. \tag{8}$$

However, exponential distributions have the drawback that they feature light right tails if the hazard rate is inferred from the macroeconomic data, which in other words means that long recessions are extremely rare. In fact, the following back-of-the-envelope calculation suggests that the probability of observing an economic recession with a duration of more than 10 years $\{\tau_2 \geq 10\}$ equals

$$P\{\tau_2 > 10 | s = s_2\} = \int_{10}^{\infty} f_{\tau_2}(\tau)d\tau = \exp(-10\lambda_2) = \exp(-10) = 0.00005, \tag{9}$$

when the mean duration of the recession state λ_2^{-1} is four quarters. In that case, the number of slowdowns until the first appearance of a lost decade follows the geometric distribution with the mean $(0.00005)^{-1} = 22,000$. In other words, it takes on average about 22,000 transitions in order to draw at least a decade-long recession, which is arguably implausible due to the extreme rareness of the event.[5,6]

In order to model the long-lasting recessions in a more plausible way, we propose to generalize the standard two-state Markov chain setting with the state space S to a two-state semi-Markov chain setting where the probability law governing the recession sojourn time is a mixture of two exponential distributions. Although semi-Markov chains can be arguably less tractable, there are special instances when they can be easily represented in terms of restricted Markov chains by augmenting their state space. As shown in Murphy (2012, Section 17.6), the two-state semi-Markov chain can be expressed in terms of a three-state augmented Markov chain, in our case with two substates for the downturn which differ in the mean duration. The first substate corresponds to the common business-cycle recession and has the mean duration λ_2^{-1}. The second substate corresponds to the rare but

---

[5] As a piece of anecdotal evidence, consider the lost decade experienced by Japan at the end of the 20th century.

[6] Diebold, Lee, and Weinbach (1994) suggest modeling transition intensities as logistic functions of certain exogenous variables. The drawback, however, is that in a general equilibrium setting one needs to specify the dynamics of those exogenous variables in fine detail.

protracted recession where we set the mean duration λ_3^{-1} equal to forty quarters. The augmented state space is

$$\tilde{S} = \{\tilde{s}_1, \tilde{s}_2, \tilde{s}_3\}, \tag{10}$$

where $\tilde{s}_1 = (s_1, \lambda_1)$, $\tilde{s}_2 = (s_2, \lambda_2)$, and $\tilde{s}_3 = (s_2, \lambda_3)$. The semi-Markov property implies equality of the growth rates across the recession types, that is, $\mu_1^e = \bar{\mu}^e$ and $\mu_2^e = \underline{\mu}^e = \mu_3^e$ for each $e \in E$. As a result of two recession types, the sojourn times of a low-growth epoch follow a mixture of two exponential densities with different means.[7]

Furthermore, we assume that upon leaving the expansion state nature tosses a biased coin according to the Bernoulli probability distribution $(q, 1-q)$ for some $q \in (0,1)$ where the outcome of the toss decides the type of the downturn. As a result, the transition probability matrix between times t and $t+T$ equals the matrix exponential $\exp\{\lambda T\}$, where the transition intensity matrix[8] λ is given by

$$\lambda = \begin{pmatrix} -\lambda_1 & q\lambda_1 & (1-q)\lambda_1 \\ \lambda_2 & -\lambda_2 & 0 \\ \lambda_3 & 0 & -\lambda_3 \end{pmatrix}. \tag{11}$$

We note that the two-state model without long recessions is nested for $q=1$ because then the hazard rate of entering the long recession $(1-q)\lambda_1$ is zero.

The invariant distribution $\bar{\pi} = (\bar{\pi}_1, \bar{\pi}_2, \bar{\pi}_3)'$ is given as the left eigenvector that corresponds to the zero eigenvalue of the transition intensity matrix, subject to the restriction that $\sum_{i=1}^{3} \bar{\pi}_i = 1$. In particular, the invariant probability

$$\bar{\pi}_3 = \frac{(1-q)\lambda_3^{-1}}{\lambda_1^{-1} + q\lambda_2^{-1} + (1-q)\lambda_3^{-1}} \tag{12}$$

equals the average time spent in the long recession, $(1-q)\lambda_3^{-1}$, divided by the average length of one whole cycle, $\lambda_1^{-1} + q\lambda_2^{-1} + (1-q)\lambda_3^{-1}$. This result is important in the empirical section where we propose to model the rare long recession as a lost decade that occurs on average once a century, thus setting $\lambda_3^{-1} = 10$ years and $\bar{\pi}_3 = 0.1$, exactly in line with Equation (12).

---

[7] Such a mixture density is called hyperexponential distribution. Hyperexponential density is thus the probability distribution that governs the sojourn time spent in recession when the type is hidden.

[8] Identification requires that we rule out instantaneous transitions between short and long recessions by setting $\lambda_{23} = \lambda_{32} = 0$. This assumption, however, is not particularly restrictive because the transition probabilities $P\{s_{t+T} = \tilde{s}_j | s_t = \tilde{s}_i\}$ are positive for any finite interval $T > 0$ and $i, j = 1, 2, 3$.

2.3 INFERENCE PROBLEM

The investor's inference problem is to extract the current but hidden state s_t from the history of the cash-flow signals $\mathcal{F}_t = \{(g_\tau^u, g_\tau^l) \text{ for } \tau \leq t\}$. For that purpose, we define the belief

$$\pi_{i,t} = P\{s_t = \tilde{s}_i | \mathcal{F}_t\} \text{ for } i = 1, 2, 3, \tag{13}$$

and introduce the so-called "innovation process" $\tilde{z}_t^e$ the increment of which is the normalized forecast error of the cash-flow growth rates over the next instant,

$$d\tilde{z}_t^e = \frac{1}{\sigma^e} \left(dg_t^e - E_t\{dg_t^e\} \right). \tag{14}$$

First, Liptser and Shiryaev (1977) show that $\tilde{z}_t^e$ is a Brownian motion in the investor's filtration which makes the investor's intertemporal optimization problem Markovian by allowing us to treat the beliefs $\pi_t = (\pi_{1,t}, \pi_{2,t}, \pi_{3,t})'$ as part of the state vector that reflects the variation in the investment opportunity set perceived by the investor. Second, it is straightforward to see that the innovation process is correlated with the hidden semi-Markov chain s_t. Third, the innovation process enables us to express the cash-flow dynamics in Equation (3) as the sum of the predictable part $m_t^e dt = E_t\{dg_t^e\}$ and the cash-flow news $dg_t^e - E_t\{dg_t^e\}$ by using Equation (14),

$$dg_t^e = m_t^e dt + \sigma^e d\tilde{z}_t^e \text{ for } e \in E. \tag{15}$$

We can then apply the Bayes rule and obtain the following law of motion for the beliefs π_t:

$$d\pi_{i,t} = \eta_{i,t} dt + \sum_{e \in E} v_{i,t}^e d\tilde{z}_t^e. \tag{16}$$

2.3.a. *Intuitive explanation*

Although the reference to the formal proof is provided in the Supplementary Appendix B, the outline of the intuition behind Equation (16) is relatively straightforward. First, the predictable part given by the drift $\eta_{i,t} = \sum_{j=1}^{3} \pi_{j,t} \lambda_{ji}$ reflects the dynamics of the perfectly observable semi-Markov chain when augmented to the three-state Markov chain. Second, the volatility

$$v_1^e = \pi_{1,t}(1 - \pi_{1,t}) \left(\frac{\bar{\mu}^e - \underline{\mu}^e}{\sigma^e} \right) \tag{17}$$

and

$$v_{i,t}^e = -\pi_{1,t}\pi_{i,t}\left(\frac{\bar{\mu}^e - \underline{\mu}^e}{\sigma^e}\right) \text{ for } i = 2, 3 \qquad (18)$$

measures the weight that the investor puts in his sequential updating on the normalized news $\mathrm{d}\tilde{z}_t^e$ in Equation (16). Indeed, the weight $v_{1,t}^e$ is proportional to the economic uncertainty about the underlying instantaneous growth rate measured by the prior variance

$$\mathrm{var}_t\left\{\mu_{s_t}^e\right\} = \pi_{1,t}(1 - \pi_{1,t})(\bar{\mu}^e - \underline{\mu}^e) \qquad (19)$$

of the Bernoulli distribution over the growth rates $\bar{\mu}^e$ and $\underline{\mu}^e$ at the beginning of the instant $(t, t + \mathrm{d}t)$.

Speaking more formally, the Bayes rule says that the posterior odds equal the prior odds times the likelihood ratio. Thus, the increment in the log of the odds

$$O_{1,23} = \frac{\pi_{1,t}}{1 - \pi_{1,t}} \qquad (20)$$

in favor of the expansion equals the log-likelihood ratio.[9] The log-likelihood ratio in favor of the hypothesis $H_0 : \mu_{s_t}^e = \bar{\mu}^e$ against the alternative $H_1 : \mu_{s_t}^e = \underline{\mu}^e$, conditional on the new data $\mathrm{d}g_t^e$ and no regime shifts in the interval $(t, t + \mathrm{d}t)$, equals

$$-\frac{1}{2}\left\{\frac{(\mathrm{d}g_t^e - \bar{\mu}^e\,\mathrm{d}t)^2}{(\sigma^e)^2\,\mathrm{d}t} - \frac{(\mathrm{d}g_t^e - \underline{\mu}^e\,\mathrm{d}t)^2}{(\sigma^e)^2\,\mathrm{d}t}\right\}. \qquad (21)$$

As a result of Equation (15), the increment in the log odds due to the arrival of the new information is given by

$$\mathrm{d}\log O_{1,23} = \mathcal{O}(\mathrm{d}t) + \left(\frac{\bar{\mu}^e - \underline{\mu}^e}{\sigma^e}\right)\mathrm{d}\tilde{z}_t^e, \qquad (22)$$

and we recover the diffusion term $v_{1,t}^e$ in Equation (16) by applying Itô lemma to Equation (20). As can be seen in Equation (22), good cash-flow news $\mathrm{d}\tilde{z}_t^e$ always raises the posterior odds $O_{1,23}$ in favor of the expansion.

---

[9] In the sequential Bayesian updating, the posterior for the previous instant $(t - \mathrm{d}t, t)$ becomes the prior for the next instant $(t, t + \mathrm{d}t)$.

Furthermore, the total weight $v_{2,t}^e + v_{3,t}^e = -v_{1,t}^e < 0$ is split[10] across the beliefs $\pi_{2,t}$ and $\pi_{3,t}$ according to the prior odds at the beginning of the instant $(t, t + dt)$ in favor of the short recession

$$O_{23} = \frac{\pi_{2,t}}{\pi_{3,t}} \qquad (23)$$

as $v_{2,t}^e / v_{3,t}^e$. As a result, good cash-flow news $d\tilde{z}_t^e$ during the short recession lowers not only the beliefs $\pi_{2,t}$ and $\pi_{3,t}$ but also brings down the relatively high odds in favor of the shorter recession. Indeed, suppose we know that $s_t \neq \tilde{s}_1$ and we try to discriminate between short and long recessions $O_{1,23}$.[11]

Speaking more formally, let us denote T the random time spent in the low-growth state $s_t \in \{\tilde{s}_2, \tilde{s}_3\}$ and recall that it follows an exponential distribution with mean λ_2^{-1} in the short recession and λ_3^{-1} in the long recession. The Bayes rule implies that the increment in the log of the odds O_{23} again equals the log-likelihood ratio

$$d \log O_{23} = (\lambda_2 - \lambda_3)dt \qquad (24)$$

which is basically a special case of Equation (16) for $\pi_{1,t} = 0$. As a result, the posterior odds in favor of the short recession O_{23} have tendency to decrease with the amount of time spent in the low-growth state and the learning about the growth persistence is time-consuming in proportion to the difference between the hazard rates for the short recession and for the long recession, $\lambda_2 - \lambda_3$.

2.3.b *Endogenous disaster probability*

Rare consumption disasters in our semi-Markov model manifest themselves as unfavorable draws of the recession duration. When we represent the two-state semi-Markov chain in terms of a restricted three-state Markov chain, subject to the equality constraint $\mu_2^e = \underline{\mu}^e = \mu_3^e$ for each $e \in E = \{u, l\}$, we identify the rare consumption disasters as the long recessions $s_t = \tilde{s}_3$ and the disaster probability as the belief

$$\pi_{3,t} = P\{s_t = \tilde{s}_3 | \mathcal{F}_t\}. \qquad (25)$$

The endogenous variation in the disaster probability $\pi_{3,t}$ comes from the fluctuations in the posterior odds in favor of the expansion $O_{1,23}$ in Equation (20) and the posterior odds about the type of the recession O_{23}

---

[10] Note that the restriction that the beliefs sum to one $\sum_{i=1}^3 \pi_{i,t} = 1$ implies that the increment $d(\sum_{i=1}^3 \pi_{i,t}) = d(1)$ is zero and hence the drifts as well as volatility must sum to zero as well, $\sum_{i=1}^3 \eta_{i,t} = 0 = \sum_{i=1}^3 v_{i,t}^e$.

[11] We can equivalently think of the analysis as being conditional on $\{s_t \neq \tilde{s}_1\}$.

in Equation (23). In fact, each recession $s_t = \tilde{s}_2$ carries with it the subjective risk that it may correspond to the lost decade regime $\tilde{s}_3$ due to the unobservability of the recession type. Such a novel model of consumption disasters is an example of a Peso problem, which refers to a situation in which the possibility of some infrequent event (such as a long recession) has an effect on asset prices.[12]

2.4 FLUCTUATING ECONOMIC UNCERTAINTY

It is well-known in the literature that the variation in economic uncertainty is the key to successfully explaining the variation in asset prices.[13] Economic uncertainty can be measured by the degree of difficulty in making precise forecasts of future cash-flow growth rates measured by the term structure of the forecast error variances of the T-period-ahead cash-flow growth rates.

We show that the introduction of hidden regime shifts generates endogenous variation in the forecast error variance of the cash-flow growth rates and thus in economic uncertainty. The key to showing this consists in time-aggregating the growth rates from the infinitesimal decision intervals to their T-period intervals.

In order to time-aggregate the instantaneous cash-flow growth rates dg_t^e for each $e \in E$, let us denote the T-period growth rate

$$g_{t,T}^e = \int_t^{t+T} dg_\tau^e, \qquad (26)$$

the mean T-period growth rate

$$\mu_{t,T}^e = \int_t^{t+T} \mu_{s_\tau}^e d\tau, \qquad (27)$$

and the T-period innovation

$$z_{t,T}^e = \int_t^{t+T} \sigma^e dz_\tau^e. \qquad (28)$$

As a result of the time aggregation, the cash-flow model leads to

$$g_{t,T}^e = \mu_{t,T}^e + z_{t,T}^e \qquad (29)$$

---

[12] See Evans (1996) for a review of the Peso literature.
[13] See in particular Bansal and Yaron (2004); Bansal, Gallant, and Tauchen (2007); Bansal, Kiku, and Yaron (2010); Bansal *et al.* (2012).

with the following variance decomposition of the T-period cash-flow growth rate,

$$\text{var}_t\left\{g^e_{t,T}\right\} = \text{var}_t\left\{\mu^e_{t,T}\right\} + \text{var}_t\left\{z^e_{t,T}\right\}. \tag{30}$$

Expressed in words, the variance of the T-period cash-flow growth rate $\text{var}_t\left\{g^e_{t,T}\right\}$ is given by the sum of the forecast error variance of the mean T-period growth rate (first term) and the forecast error variance of the T-period innovations (second term). In particular, the volatility of the annual consumption growth rate corresponds to $\sigma^u_{t,1}$. The commonly used autoregressive processes imply that the forecast error variance of the mean T-period growth rate (first term), $\text{var}_t\left\{\mu^e_{t,T}\right\}$, is constant.[14]

We show in the following two sections that learning about hidden regime shifts can generate countercyclical fluctuations in the forecast error variance of the T-period cash-flow growth rates with constant σ^e and hence constant

$$\text{var}_t\left\{z^e_{t,T}\right\} = (\sigma^e)^2 T. \tag{31}$$

2.4.a. *Time-varying forecast error variance*

Let us denote the T-period forecast conditional on the hidden state

$$m^e_{t,T|i} = E\left\{g^e_{t,T}\middle|\mathcal{F}_t, s_t = \tilde{s}_i\right\} \tag{32}$$

and the T-period forecast error variances conditional on the hidden state

$$\left(\sigma^e_{t,T|i}\right)^2 = \text{var}\left\{g^e_{t,T}\middle|\mathcal{F}_t, s_t = \tilde{s}_i\right\} \tag{33}$$

and

$$\left(\upsilon^e_{t,T|i}\right)^2 = \text{var}\left\{\mu^e_{t,T}\middle|\mathcal{F}_t, s_t = \tilde{s}_i\right\}. \tag{34}$$

The conditional moments given the hidden state s_t vary due to the possibility of a regime change with the more persistent state of lower hazard rate of transitioning λ_i displaying lower volatility $\upsilon^e_{t,T|i}$. Furthermore, variance

---

[14] For example, Bansal and Yaron (2004) in their Model I specify the expected consumption growth rate as an AR(1) process subject to homoscedastic innovations. Therefore, their model generates constant forecast error variance of the T-period consumption growth rate which they relax in their Model II by introducing exogenous variation in the consumption variance $\left(\sigma^e_t\right)^2$ as an AR(1) process.

decomposition conditional on the hidden state analogous to Equation (30) can be expressed as

$$\left(\sigma_{t,\,T|i}^{e}\right)^{2}=\left(\upsilon_{t,\,T|i}^{e}\right)^{2}+(\sigma^{e})^{2}T. \tag{35}$$

We then condition down to the investor's information set $\mathcal{F}_{t}$ which does not contain the hidden state s_{t}. The mean T-period cash-flow growth rate $m_{t,\,T}^{e}=E_{t}\left\{g_{t,\,T}^{e}\right\}$ is given by

$$m_{t,\,T}^{e}=\sum_{i=1}^{3}m_{t,\,T|i}^{e}\pi_{i,\,t}. \tag{36}$$

Furthermore, using the decomposition that the variance equals the variance of the conditional mean plus the mean of the conditional variance[15]

$$\mathrm{var}_{t}\{x\}=\mathrm{var}_{t}\{E\{x|\mathcal{F}_{t},s_{t}\}\}+E_{t}\{\mathrm{var}\{x|\mathcal{F}_{t},s_{t}\}\} \tag{37}$$

yields the following decomposition of the corresponding T-period cash-flow variance $\mathrm{var}_{t}\left\{g_{t,\,T}^{e}\right\}$ in Equation (30)

$$\left(\sigma_{t,\,T}^{e}\right)^{2}=\underbrace{\sum_{i=1}^{3}\left(\underbrace{\left(\upsilon_{t,\,T|i}^{e}\right)^{2}+(\sigma^{e})^{2}T}_{\text{Variance under Complete Info}}\right)\pi_{i,\,t}}_{\text{Mean Variance under Incomplete Info}}+\underbrace{\sum_{i=1}^{3}\left(m_{t,\,T|i}^{e}\right)^{2}\pi_{i,\,t}-\left(\sum_{i=1}^{3}m_{t,\,T|i}^{e}\pi_{i,\,t}\right)^{2}}_{\text{Variance of Mean Growth under Incomplete Info}}.$$

$$\tag{38}$$

As we can see from Equation (36) and Equation (38), the variation in both the mean T-period growth rate forecast $m_{t,\,T}^{e}$ as well as the volatility $\sigma_{t,\,T}^{e}$ depends on the evolution of the beliefs $\left(\pi_{1,\,t},\pi_{2,\,t},\pi_{3,\,t}\right)'$.[16] In addition, the T-period forecast $m_{t,\,T}^{e}$ attains its maximum when the confidence in favor of the expansion state is the highest and its minimum when the confidence in favor of the lost decade is the highest, whereas the corresponding forecast

---

[15] Note that the conditional moments E_{t} and var_{t} are conditional only on $\mathcal{F}_{t}$ which does not include the hidden state s_{t}.

[16] This is consistent with Veronesi (1999) who shows in Proposition 6 that if expected consumption growth rate follows a hidden Markov chain then shocks to the instantaneous expected dividend growth rate are necessarily heteroscedastic. Although he does not time aggregate the cash-flow growth rates from infinitesimal decision intervals to the finite ones, his two-state Markov chain setting is able to generate time-varying consumption volatility after the aggregation. However, such variation would be quantitatively smaller in comparison to our two-state semi-Markov setting.

error variance $\left(\sigma_{t,T}^e\right)^2$ attains values based on the magnitude of the economic uncertainty measured by the dispersion of the beliefs.

2.4.b. *Countercyclical consumption and dividend volatility*

A two-state continuous-time Markov chain can be expressed as a linear combination of two independent compensated Poisson processes leading to a continuous-time AR(1) process with innovations that are non-Gaussian and heteroscedastic, having the instantaneous variance proportional to the persistence of the state, $\left(\bar{\mu}^e - \underline{\mu}^e\right)^2 \lambda_{s_t} dt$ for each $s_t \in S$.[17] It can be shown that the analogous result carries over to our two-state semi-Markov setting.

Our empirical estimates in Section 3.2.1 confirm that the transition hazard rates satisfy $\lambda_2 > \lambda_1$ and we thus obtain the ordering on the forecast error volatility of the mean T-period cash-flow growth rate conditional on the hidden state s_t as $\upsilon_{t,T|2}^e > \upsilon_{t,T|1}^e$. Note that we omit the discussion related to the rare state by assuming $\pi_{3,t} \approx 0$. In view of Equation (35), the volatility of the T-period cash-flow growth rate under complete information is thus countercyclical because the mean duration of the recessions $\lambda_{s_t}^{-1}$ is empirically shorter in comparison to the expansions.

Furthermore, in case of incomplete information about the underlying state the variance of the T-period cash-flow growth rate decomposed in Equation (38) has two terms, the mean of the conditional variance under complete information plus the variance of the conditional mean hidden due to incomplete information. The first term $\upsilon_{t,T|1}^e \pi_{1,t} + \upsilon_{t,T|2}^e \left(1 - \pi_{1,t}\right)$ is a decreasing function of the belief $\pi_{1,t}$ due to the ordering of the conditional volatility $\upsilon_{t,T|i}^e$ displayed above and it is thus countercyclical. The second term

$$\left(m_{t,T|1}^e\right)^2 \pi_{1,t} + \left(m_{t,T|2}^e\right)^2 \left(1 - \pi_{1,t}\right) - \left(m_{t,T|1}^e \pi_{1,t} + m_{t,T|2}^e \left(1 - \pi_{1,t}\right)\right)^2 \quad (39)$$

is a quadratic function of the belief $\pi_{1,t}$ and it is an increasing function in the belief $\pi_{1,t}$ for $\pi_{1,t} < \frac{1}{2}$ but decreasing for $\pi_{1,t} > \frac{1}{2}$ due to the ordering of the T-period-ahead forecasts $m_{t,T|1}^e > m_{t,T|2}^e$ implied by $\mu_1^e > \mu_2^e$. As we see later in our parametrization using maximum likelihood estimates, the second term is usually dominated by the first one and thus the total cash-flow volatility $\sigma_{t,T}^e$ remains countercyclical even after accounting for incomplete information.

---

[17] See Hamilton (1989) for discrete time treatment.

2.4.c. *Cash-flow dynamics over the phases of the business cycle*

The forecast as well as the forecast error variance of the T-period cash-flow growth rate vary monotonically over the separate phases of the business cycle. First, transitioning to the high-growth state $s_t = \tilde{s}_1$ is associated with a gradual improvement in the T-period forecast $m_{t,T}^e = E_t(g_{t,T}^e)$ as well as the T-period forecast error variance $(\sigma_{t,T}^e)^2 = \text{var}_t\{g_{t,T}^e\}$. These gradual changes are driven by the rise in the posterior odds $O_{1,23}$ in Equation (20) as the high-growth state is being recognized. Second, transitioning to the low-growth state $s_t = \tilde{s}_2$ is associated with a gradual deterioration in the forecast and the forecast error variance of the T-period cash-flow growth rate. Again, these gradual changes are driven not only by falling posterior odds $O_{1,23}$ as the low-growth regime is being recognized, but also rising posterior odds in favor of the long recession O_{23}^{-1} in Equation (23) as the likelihood of a protracted slowdown is increasing.[18]

2.5 INVESTOR'S PROBLEM

The investor's financial wealth W_t comprises the unlevered and the levered equity as well as the real zero-coupon bond with a given maturity T and the riskless cash account offering the continuously compounded rate of return r_t. We denote the share of each asset $a \in A$ in the wealth portfolio W_t as ω_t^a and let the investor decide continuously how much to consume and how much to save out of his current wealth W_t. The dynamic budget constraint takes the standard form as in Merton (1971),

$$\mathrm{d}W_t = \left(\sum_{a \in A} \omega_t^a(\mathrm{d}R_t^a - r_t\mathrm{d}t) + r_t\mathrm{d}t\right)W_t - c_t\mathrm{d}t, \qquad (40)$$

where we still need to specify the law of motion for the asset return $\mathrm{d}R_t^a$.

According to Equation (14), the increment in the innovation process $\mathrm{d}\tilde{z}_t^e$ is the normalized instantaneous forecast error of the cash-flow growth rate $\mathrm{d}g_t^e$ for each $e \in E = \{u, l\}$ and it is to be thought of as the news about the current hidden state $s_t \in \tilde{S}$. In informationally efficient asset markets, news arrival leads to an instant revision in the price of each asset $a \in A$ generating a

---

[18] Of course, the economic uncertainty in $s \in \{\tilde{s}_3\}$ will eventually decline once the unobservable state is recognized but it takes a long time in comparison to the mean duration of the short recession $\tilde{s}_2$.

surprise return (also called news, innovation or forecast error) in proportion to the asset volatility $\vartheta_t^{a,e}$,

$$\mathrm{d}R_t^a - E_t\{\mathrm{d}R_t^a\} = \sum_{e \in E} \vartheta_t^{a,e} \mathrm{d}\tilde{z}_t^e, \tag{41}$$

where the net return is defined as usual

$$\mathrm{d}R_t^a = \frac{\mathrm{d}P_t^a + D_t^a \mathrm{d}t}{P_t^a}. \tag{42}$$

The realized return $\mathrm{d}R_t^a$ is composed of the predictable part given by the expected return $E_t(\mathrm{d}R_t^a) = \zeta_t^a \mathrm{d}t$ and the unpredictable part given by the surprise return in Equation (41),

$$\mathrm{d}R_t^a = \zeta_t^a \mathrm{d}t + \sum_{e \in E} \vartheta_t^{a,e} \mathrm{d}\tilde{z}_t^e. \tag{43}$$

In our model, the expected return ζ_t^a and the return volatility $\vartheta_t^{a,e}$ for each $e \in E$ and each asset $a \in A$ are determined jointly by market clearing in general equilibrium.

The investor's consumption-portfolio problem is to maximize his lifetime utility defined recursively in Equation (1) subject to the dynamic budget constraint in Equation (40) leading to the standard HJB equation[19]

$$0 = \max_{\{c, \omega^u, \omega^l, \omega^b\}} \left\{ U(c, J)\mathrm{d}t + E_t\{\mathrm{d}J(W, \pi_1, \pi_2)\} \right\}, \tag{44}$$

where the posterior distribution becomes a part of the state vector in addition to the wealth W.[20] Itô lemma applied to the continuation utility $J = J(W, \pi_1, \pi_2)$, along with the budget constraint in Equation (40) and the dynamics of the return $\mathrm{d}R_t^a$ in Equation (43), then leads to a nonlinear partial differential equation of the second order for J.

2.5.a First-order conditions and equilibrium

The first-order condition for the consumption rate c states that the marginal utility of consumption equals the marginal utility of wealth $\frac{\partial U}{\partial c} = \frac{\partial J}{\partial W}$. The first-order condition for the portfolio weight ω^a for the asset $a \in A$ states that the total demand for asset a equals the myopic demand plus the

---

[19] See Duffie and Epstein (1992a, 1992b).
[20] Note that the belief π_3 is given implicitly due to the restriction that the probabilities sum to one.

intertemporal hedging demand that arises from the fluctuations in the investor's own uncertainty about the state of the macroeconomy (Merton, 1973; Veronesi, 1999).

In equilibrium, the conditions $c_t = D_t^u$, $\omega_t^u = 1$, $\omega_t^l = 0$, and $\omega_t^b = 0$ must hold for the asset and the goods markets to clear.

2.5.b. *Value function and wealth-consumption ratio*

The first-order condition for the consumption rate implicitly defines the optimal policy function $c = c(W, \pi_1, \pi_2)$. Invoking the homotheticity of the recursive preferences $\frac{\partial \log c}{\partial \log W} = 1$ implies that the policy function $c(W, \pi_1, \pi_2)$ is separable across the financial wealth W and the beliefs (π_1, π_2). The separability of the policy function in turns implies, through the first-order condition, that the value function is also separable across W and (π_1, π_2),

$$J(W, \pi_1, \pi_2) = \delta^\theta [\Phi^u(\pi_1, \pi_2)]^{\frac{\theta}{\psi}} \frac{W^{1-\gamma}}{1 - \gamma}, \qquad (45)$$

where we choose to parametrize it in terms of the equilibrium wealth-consumption ratio

$$\Phi^u(\pi_{1,t}, \pi_{2,t}) = W_t/c_t. \qquad (46)$$

The conjecture in Equation (45) reduces the nonlinear PDE, coming from the HJB equation for the continuation utility J in Section 2.5, to the nonlinear degenerate-elliptic partial differential equation of the second order for Φ^u, presented in Proposition D.1 in the Supplementary Appendix D.

When the investor prefers early resolution of uncertainty (i.e., $\theta < 0$), the cross-derivative of the marginal utility $\frac{\partial}{\partial \pi_i}\left(\frac{\partial J}{\partial W}\right)$ is negative. The intuition for these results is simple. A positive short-run news $d\tilde{z}_t^e$ always raises the posterior odds in favor of the expansion $O_{1,23}$, and thus the beliefs π_1 and π_2 go up. This in turn leads to an improvement in the T-period forecasts of future consumption growth rate $m_{t,T}^u$, raising the duration of the consumption stream, and so delaying the mean time to the (partial) resolution of uncertainty about the consumption stream. This is disliked by the investor with a preference for early uncertainty resolution, and the marginal utility falls. The fall in the marginal utility is larger when the increase in the duration is bigger, allowing us to order the cross-derivatives as $\frac{\partial}{\partial \pi_1}\left(\frac{\partial J}{\partial W}\right) < \frac{\partial}{\partial \pi_2}\left(\frac{\partial J}{\partial W}\right) < 0$.

As a corollary, the wealth-consumption ratio is procyclical, being an increasing function of the beliefs $\frac{\partial \Phi^u}{\partial \pi_1} > \frac{\partial \Phi^u}{\partial \pi_2} > 0$.

2.6 STATE-PRICE DENSITY

The absence of arbitrage implies the existence of a positive state-price density process M_t which in case of the Epstein–Zin preferences in Equation (1) is given by the formula[21]

$$M_t = \exp\left(\int_0^t \frac{\partial U}{\partial J}\mathrm{d}\tau\right)\frac{\partial U}{\partial c}.\tag{47}$$

The following proposition presents the law of motion for the state-price density and decomposes the corresponding risk prices into the Lucas-Breeden component reflecting the covariance with the consumption growth and the variable timing component reflecting the changing forecasts of the time to the (partial) resolution of the consumption uncertainty in terms of the posterior odds in Equation (20) and in Equation (23). In our parametrization of the preferences, late resolution of the uncertainty is disliked by the investor and the cross-derivative $\frac{\partial}{\partial J}\left(\frac{\partial U}{\partial c}\right)$ is negative, as argued in Sections 2.1 and 2.5.2. In case of the expected utility, we recover the standard consumption-based capital asset pricing model with zero timing components because the independence axiom implies that the marginal utility of consumption does not depend on the continuation utility and hence $\frac{\partial}{\partial J}\left(\frac{\partial U}{\partial c}\right)$ is zero which happens for $\gamma = \psi^{-1}$.

Proposition 2.1
Let the equilibrium state-price density M_t be given by Equation (47). Then,

- *M_t satisfies*

$$\log M_t = -\theta\delta t - (1-\theta)\int_0^t \left(\Phi^u(\pi_{1,\tau},\pi_{2,\tau})\right)^{-1}\mathrm{d}\tau - x_t,\tag{48}$$

with

$$x_t = \gamma \log D_t^u + (1-\theta)\log \Phi^u(\pi_{1,t},\pi_{2,t}).\tag{49}$$

- *M_t evolves according to the stochastic differential equation*

$$\frac{\mathrm{d}M_t}{M_t} = -r_t\mathrm{d}t - \sum_{e\in E}\Lambda_t^e\mathrm{d}\bar{z}_t^e,\tag{50}$$

---

[21] See Duffie and Epstein (1992a, 1992b) and Schroder and Skiadas (1999).

where

- *the instantaneous riskless interest rate* $r_t = r(\pi_{1,t}, \pi_{2,t})$ *is given by Equation (86)*
- *the risk price functions* $\Lambda_t^e = \Lambda^e(\pi_{1,t}, \pi_{2,t})$ *for each* $e \in E$ *are given by*

$$\Lambda_t^e = \underbrace{\gamma \sigma^u \delta_{e,u}}_{\text{Lucas-Breeden Component}} + \underbrace{(1-\theta) \sum_{i=1}^{2} \nu_{i,t}^e \frac{1}{\Phi_t^u} \frac{\partial \Phi^u}{\partial \pi_i} (\pi_{1,t}, \pi_{2,t})}_{\text{Time-Varying Timing Component}}, \qquad (51)$$

where the symbol $\delta_{e,u}$ is the Kronecker delta.[22]

Proof

See the Supplementary Appendix C. □

Proposition 2.1, together with the equilibrium conditions, allows us to express the first-order condition for the portfolio weights ω_t^a in Problem (44) for each asset $a \in A$ as the restriction that the risk premium equals the negative of the covariance with the state-price density growth rate, $E_t(dR_t^a - r_t dt) = -\text{cov}_t\{\frac{dM_t}{M_t}, dR_t^a - r_t dt\}$, that is,

$$\zeta_t^a - r_t = \sum_{e \in E} \Lambda_t^e \vartheta_t^{a,e}. \qquad (52)$$

As can be seen, the risk prices Λ_t^e measure the increase in the asset risk premiums in response to the marginal increase in the exposure to the Brownian shock $d\tilde{z}_t^e$ for each $e \in E$.

2.6.a *Risk prices*

The risk prices Λ_t^e for $e \in E$ in Equation (51) measure the sensitivity of the growth rate of the marginal utility of wealth to the news carried by the Brownian shocks[23] $d\tilde{z}_t^e$,

$$\underbrace{-M_t^{-1}(dM_t - E_t\{dM_t\})}_{\text{Marginal Utility Growth Rate Surprise}} = \underbrace{\Lambda_t^u d\tilde{z}_t^u}_{\text{Response to Consumption Surprise}} + \underbrace{\Lambda_t^l d\tilde{z}_t^l}_{\text{Response to Dividend Surprise}}$$

$$(53)$$

---

[22] Recall that the Kronecker delta satisfies $\delta_{e,e} = 1$ and $\delta_{e,u} = 0$ for $u \neq e$.
[23] The levered dividend shock $d\tilde{z}_t^l$ enters because it is correlated with the hidden state s and is thus also a source of news as shown in (16).

The innovations in the cash-flow growth rate $\mathrm{d}\tilde{z}_t^e$ are i.i.d. and correspond to the short-run cash-flow news. Good short-run cash-flow news always increases the posterior odds $O_{1,23}$ in Equation (20) in favor of the expansion state and generates good long-run cash-flow news in terms of the improved forecasts of the T-period cash-flow growth rates $m_{t,T}^e$.[24] The effect of short-run cash-flow news on the long-run growth prospects is called the "cash-flow effect" in the literature. Furthermore, the good short-run news also changes the duration of the consumption stream as well as the forecast error variance and tends to lengthen the mean time to the (partial) resolution of the consumption uncertainty leading to a rise in the discount rates. The effect of the short-run news on the discount rates is called the "discount-rate effect" in the literature. According to the analysis in Section 2.5.2, the cash-flow effect dominates the discount rate effect and the wealth-consumption ratio $\Phi^u(\pi_{1,t}, \pi_{2,t})$ is pro-cyclical rising unexpectedly on good short-run news $\mathrm{d}\tilde{z}_t^e$. We call the innovation in the wealth-consumption ratio coming from the long-run cash-flow as well as discount rate news simply the long-run news.

A positive piece of news in both the short-run and the long-run generates a positive surprise in the return on the wealth portfolio,

$$\mathrm{d}R_t^a - E_t\{\mathrm{d}R_t^a\} = \underbrace{(\mathrm{d}g_t^a - E_t\{\mathrm{d}g_t^a\})}_{\text{short-run news}} + \underbrace{(\Phi_t^a)^{-1}(\mathrm{d}\Phi_t^a - E_t\{\mathrm{d}\Phi_t^a\})}_{\text{long-run news}}, \text{ for } a = u,$$

(54)

and a negative surprise in the marginal utility of wealth. When we invoke Equations (45) and (54), and apply Itô lemma to $\Phi_t^u = \Phi^u(\pi_{1,t}, \pi_{2,t})$, we easily recover the formulas for the risk prices Λ_t^e in Proposition 2.1.

2.7 LEVERED EQUITY PRICES

The absence of arbitrage implies that unlevered and levered equity prices equal the expected discounted value of the future dividend stream,

$$M_t P_t^a = E_t\left\{\int_t^\infty M_\tau D_\tau^a \mathrm{d}\tau\right\} \text{ for } a \in \{u, l\}.$$

---

[24] Recall that the shocks to the beliefs (π_1, π_2) are persistent in proportion to the mean duration of the states λ_j^{-1} for $j = 1, 2, 3$ as well as the conditional probability of transitioning to the short recession out of the expansion q.

The equity price P_t^a trends upward, making it more tractable to solve for the equilibrium price function in terms of the corresponding price–dividend ratio[25]

$$\Phi_t^a = \frac{P_t^a}{D_t^a}. \tag{55}$$

Proposition 2 in the Supplementary Appendix D exploits the martingale property of the gain process $M_t \Phi_t^a D_t^a + \int_0^t M_\tau D_\tau^a d\tau$, and derives the Fichera boundary value problems to be solved numerically as described in the Supplementary Appendix D for the ratio

$$\Phi_t^a = \Phi^a(\pi_{1,t}, \pi_{2,t}). \tag{56}$$

2.7.a *Procyclical price–dividend ratio*

A preference for early resolution of uncertainty (i.e., $\theta < 0$) implies that equity prices rise on good news. This happens because a positive innovation in the cash-flow growth rate $d\tilde{z}_t^e$ (short-run news) raises the posterior odds in favor of the expansion state in Equation (20) improving the T-period forecasts in the cash-flow growth rate (cash-flow effect). Although the increase in the posterior odds tends to lower future mean discount factors $M_{t+\tau}$ (discount rate effect), the cash-flow effect dominates the discount rate effect in our parametrization. The dominance of the cash-flow effect then implies that increasing the belief $\pi_{i,t}$ for each $i = 1, 2$ necessarily lowers the belief $\pi_{3,t} = 1 - \pi_{1,t} - \pi_{2,t}$, *ceteris paribus*, which improves the growth prospects and leads to an increase in the price–dividend ratio Φ_t^a. Hence, the derivative $\frac{\partial \Phi^a}{\partial \pi_i}$ is positive. In fact, the belief $\pi_{1,t}$ corresponds to the expansion state and its increase improves growth prospects more than the corresponding increase in $\pi_{2,t}$, allowing us to order the derivatives $\frac{\partial \Phi^a}{\partial \pi_1} > \frac{\partial \Phi^a}{\partial \pi_2} > 0$. However, short recessions in our parametrization last on average about 1 year which indicates that the long-run improvement in the growth prospects is about the same, hence, the derivatives $\frac{\partial \Phi^a}{\partial \pi_1}$ and $\frac{\partial \Phi^a}{\partial \pi_2}$ are of comparable magnitudes.

2.7.b. *Conditional return moments*

The procyclical variation in equity prices in turn generates a corresponding countercyclical variation in the conditional moments of the equity returns.

---

[25] Observe that the price–dividend ratio for the unlevered equity in equilibrium must equal the wealth-consumption ratio, $\frac{P_t^u}{D_t^u} = \Phi_t^u = \frac{W_t}{c_t}$, and its pricing equation was partially analyzed already in Section 2.5.1.

To see this, let us look first at the conditional equity return volatility $\sigma_t^{a,e}$ for $e \in E$, which measures the sensitivity of the equity return to the cash-flow news $d\tilde{z}_t^e$ as shown in Equation (41). According to the analog of Equation (54) for $a = u$, we can also decompose the news for $a = l$ into short-run news and long-run news. The long-run news can be further decomposed using Itô lemma as $d\Phi_t^a - E_t\{d\Phi_t^a\} = \sum_{i=1}^2 \frac{\partial \Phi^a}{\partial \pi_i}(d\pi_{i,t} - E_t\{d\pi_{i,t}\})$, and thus the total sensitivity to news $\vartheta_t^{a,e} = \vartheta^{a,e}(\pi_1, \pi_2)$ equals the sum of the sensitivity to the short-run news, which is constant by assumption, and the sensitivity to the long-run news, which depends on the beliefs,

$$\underbrace{\vartheta^{a,e}}_{\text{Total News Sensitivity}} = \underbrace{\sigma^e \delta_{a,e}}_{\text{Short−Run News Sensitivity}} + \underbrace{\sum_{i=1}^2 v_i^e \frac{1}{\Phi^a} \frac{\partial \Phi^a}{\partial \pi_i}}_{\text{Long−Run News Sensitivity}} , \quad (57)$$

where the symbol $\delta_{a,e}$ is the Kronecker delta.

The volatility $\vartheta_t^{a,e}$ in Equation (57) is positive and countercyclical. Positive short-run news $d\tilde{z}_t^e$ leads to a rise in the odds in favor of the expansion state in Equation (20), increasing $\pi_{1,t}$ but decreasing the sum $\pi_{2,t} + \pi_{3,t}$ by exactly the same amount. But the belief $\pi_{3,t}$ is relatively small due to the inherent rareness of long recessions and thus the news $d\pi_{i,t} - E_t\{d\pi_{i,t}\}$ for $i = 1, 2$ are of comparable magnitude but opposite sign. The inequality $\frac{\partial \Phi^a}{\partial \pi_1} > \frac{\partial \Phi^a}{\partial \pi_2}$ from the previous section then implies that the long-run news is always positively correlated with the short-run news. In addition, the sensitivity of beliefs to short-run news $v_{1,t}^e$ in Equation (16) is proportional to the prior variance $\text{var}_t\{\mu_{s_t}^e\}$ in Equation (19), which tends to be large during times of heightened economic uncertainty measured by the dispersion of the beliefs and leads to countercyclical variation in the magnitude of the long-run news, and hence, in the equilibrium equity volatility $\vartheta_t^{a,e}$ for each $e \in E$.

Second, the first-order condition in Equation (52) says that the equity risk premium can be decomposed into the Lucas-Breeden component (short-run risk premium) plus the timing premium due to the nonindifference to the timing of the resolution of uncertainty about future consumption growth inherent in the Epstein–Zin preferences (the long-run risk premium),

$$\underbrace{E_t(dR_t^a - r_t dt)}_{\text{Equity Risk Premium}} = \underbrace{\gamma \, \text{cov}_t(dR_t^a - E_t\{dR_t^a\}, dg_t^u - E_t\{dg_t^u\})}_{\text{Short−Run Equity Risk Premium}}$$

$$+ \underbrace{(1 - \theta)\text{cov}_t\left(dR_t^a - E_t\{dR_t^a\}, \frac{1}{\Phi^u}(d\Phi_t^u - E_t\{d\Phi_t^u\})\right)}_{\text{Long−Run Equity Risk Premium}}.$$

$$(58)$$

The equity risk premium inherits the property of countercyclical variation from the volatility $\vartheta_t^{a,e}$ as well as the risk prices Λ_t^e. We note again that although positive short-run news does generate positive long-run news, the magnitude of the long-run news is countercyclical as explained above, which tends to lower both covariances through $\nu_{i,t}^e$ in Equation (16) in times of high confidence when $\pi_{1,t} \approx 0$ or $\pi_{1,t} \approx 1$.

The return variance $\text{var}_t\{dR_t^a\} = \left(\sum_{e \in E} (\vartheta_t^{a,e})^2\right)dt$ as well as the expected return $E_t\{dR_t^a\} = \zeta_t^a dt$ are instantaneous moments corresponding to infinitesimal decision intervals and thus must be time-aggregated to finite intervals as described in Proposition E.1 in the Supplementary Appendix E in order to be comparable to the data.

2.8 REAL ZERO-COUPON BOND PRICES

Absence of arbitrage implies that the price of the real zero-coupon bond P_t^b with a given maturity T equals the expected discounted value of the principal payment

$$M_t P_t^b = E_t\{M_T P_T^b\}, \tag{59}$$

where we normalize the principal $P_T^b \equiv 1$. Proposition D.2 in the Supplementary Appendix D exploits the martingale property of the deflated price $M_t P_t^b$ and derives the partial differential equation for the bond price P^b as a function of the beliefs (π_1, π_2) and time t,

$$P_t^b = P^b(\pi_{1,t}, \pi_{2,t}, t; T). \tag{60}$$

As discussed in Section 2.6.1, and applied to price–dividend ratios later in Section 2.7.1, the effect of uncertainty on asset prices can be decomposed into the procyclical cash-flow effect and the countercyclical discount-rate effect. As the cash-flow effect is not present in case of nondefaultable zero-coupon bonds, their prices are driven solely by the countercyclical variation in the discount rates. The countercyclical variation in the bond prices generates a corresponding procyclical variation in the bond risk premium but countercyclical variation in the bond return volatility and the bond yields. Such countercyclical fluctuations in the bond prices imply negative surprise in the bond returns during the good times and positive surprise in the bond return during the bad times. The real bonds thus carry negative risk premiums exactly because they help to smooth consumption.[26]

---

[26] As before, the expected bond return $E_t\{dR_t^b\} = \zeta_t^b dt$ as well as the bond return variance $\text{var}_t\{dR_t^b\} = \left(\sum_{e \in E} (\vartheta_t^{a,b})^2\right)dt$ correspond to infinitesimal decision intervals. In order to be

2.8.a. *Real yield curves*

The intertemporal price between consumption today and consumption in T periods ahead equals the gross yield-to-maturity $1 + Y_t^{(T)} = \exp\left(y_t^{(T)}\right)$ on a real zero-coupon bond that matures in T periods. The functional dependence of the annualized yield on the beliefs $\left(\pi_{1,t}, \pi_{2,t}, \pi_{3,t}\right)'$ presented in the previous proposition is driven by three distinct effects. First, the subjective discount rate δ measures the investor's desire for immediate consumption with more impatient investors demanding higher yields in order to willingly accept lower consumption today relative to the one in T periods ahead. The second effect reflects the desire for a smooth consumption growth profile. The increased desire to borrow against improved economic prospects measured by $m_{t,T}^u$ shifts the demand for consumption to the right which, however, cannot be met in an endowment economy without a corresponding change in the equilibrium yield $Y_t^{(T)}$. The strength of such an effect, moreover, is measured by the elasticity of intertemporal substitution ψ. The consumption smoothing effect explains why yields are high in times of good economic prospects and low in times of bad economic prospects. The third and last motive is related to the desire to save. Such precautionary saving is inherently related to the degree of economic uncertainty measured by the forecast error volatility of the T-period consumption growth rate $\sigma_{t,T}^u$. As discussed in Section 2.4, the model with hidden regime shifts endogenously generates the variation in the forecast error variance in response to fluctuations in the posterior distribution $\left(\pi_{1,t}, \pi_{2,t}, \pi_{3,t}\right)'$.

The following proposition links the real yield curve to the optimal forecasts and the forecast error variances of the T-period consumption growth rate.

Proposition 2.2
Denote $y_t^{(T)}$ the continuously-compounded yield-to-maturity on a T-period real zero-coupon bond and $r_{t+1,T}^b$ the corresponding continuously-compounded holding period return. Then, the annualized yield-to-maturity on the T-period bond is given by

$$y_t^{(T)} \approx \theta\delta + (1-\theta)(\Phi^u(\bar{\pi}))^{-1} + \left(\frac{1}{T}\right)\gamma\, m_{t,T}^u - \left(\frac{1}{2T}\right)\gamma^2\left(\sigma_{t,T}^u\right)^2, \qquad (61)$$

where $\bar{\pi} = (\bar{\pi}_1, \bar{\pi}_2, \bar{\pi}_3)$ denotes the invariant distribution of the Markov chain.

---

comparable to their discrete-time counterparts, they must be time aggregated to $E_t\{R_{t+1,T}^b\}$ and $\mathrm{var}_t\{R_{t+1,T}^b\}$ as explained in Proposition E.1 in the Supplementary Appendix E.

Proof

See the Supplementary Appendix F.   □

Our quantitative results discussed later show that the real yield curve is driven predominantly by the intertemporal substitution effect in response to changing forecasts of future mean consumption growth rates $m_{t,T}^u / T$. The term structures of the T-period mean forecasts $m_{t,T}^u / T$, and hence of the T-period real log yields $y_t^{(T)}$, slope down during the expansions but up during the recessions because of the mean-reverting nature of the instantaneous consumption growth rate $\mu_{s_t}^u$. The slope is moreover steeper during the long recessions because of the dramatically more inferior short- and medium-term forecasts of the consumption growth rate averaged over the T periods.

In addition, our model has implications for the volatility of the real yield curve. The volatility curve of the real yields $\sigma\{y_t^{(T)}\}$ for $T > 0$ depends on the variability of the mean T-period consumption growth rate forecasts $\sigma\{m_{t,T}^u / T\}$. Such forecast variability necessarily declines with the forecast horizon T and hence the model generates a downward-sloping volatility curve for the real yield curve.

2.8.b. *Bond risk premiums*

The annual bond risk premium depends crucially on the time series properties of consumption as the following proposition shows.

Proposition 2.3

The annual geometric risk premium on the T-period real zero-coupon bond is given by

$$E_t\left\{r_{t+1,T}^b - y_t^{(1)}\right\} \approx -\frac{\gamma^2}{2}\left(\text{var}_t\left\{(E_{t+1} - E_t)g_{t,T}^u\right\} - \text{var}_t\left\{(E_{t+1} - E_t)g_{t,1}^u\right\}\right). \tag{62}$$

Proof

See the Supplementary Appendix F.   □

The above proposition is consistent with the findings in Campbell (1986). First, bond prices carry a zero risk premium when the consumption growth rate $g_{t,1}^u$ is I.I.D. because then $(E_{t+1} - E_t)g_{t,T}^u = (E_{t+1} - E_t)g_{t,1}^u$. Second, bond prices carry a negative risk premium when the expected consumption growth rate is positively autocorrelated,

$$\text{var}_t\left\{(E_{t+1} - E_t)g_{t,T}^u\right\} > \text{var}_t\left\{(E_{t+1} - E_t)g_{t,1}^u\right\}. \tag{63}$$

Moreover, the magnitude of the bond risk premium is an increasing function of the consumption growth rate persistence.

2.9 EUROPEAN OPTIONS

Absence of arbitrage implies that the price of the European call option P_t^c with the given maturity time T and the strike price $\bar{P}^l$ equals the expected discounted value of the option payoff at the maturity

$$M_t P_t^c = E_t\{M_T P_T^c\}, \tag{64}$$

where the option price at the maturity equals the final payoff, $P_T^c = \max(P_T^l - \bar{P}^l, 0)$.

Proposition 4 in the Supplementary Appendix D exploits the martingale property of the discounted option price $M_t P_t^c$ and derives the partial differential equation for the no-arbitrage price

$$P_t^c = P^c(\pi_{1,t}, \pi_{2,t}, x_t, t; T, \bar{P}^l) \tag{65}$$

as a function of the beliefs $(\pi_{1,t}, \pi_{2,t})$, the log of the levered equity price to the strike price $x_t = \log(P_t^l/\bar{P}^l)$ and time t.

3. Empirical Section

We follow the empirical strategy used by Cecchetti, Lam, and Mark (2000) and estimate our two-state semi-Markov model for consumption and dividends hidden in I.I.D. Gaussian shocks in Equation (3) by the maximum likelihood method. We discuss the point estimates and their standard errors. We then demonstrate the plausibility of these parameter estimates by calculating the variance ratios as well as the long-run forecasts of the consumption and dividend growth rates. The fact that the decade-long optimal forecast of the consumption growth rate during the long recession comes out close to zero motivates our interpretation of the long recession as the lost decade. Finally, we show that learning about hidden growth persistence endogenizes the variation in the probability of the consumption disaster specified exogenously in Gourio (2012, 2013), Seo and Wachter (2013) and Wachter (2013).

3.1 SUMMARY STATISTICS

Our data construction of US time series is similar to Bansal, Gallant, and Tauchen (2007) and it is described in full in the Supplementary Appendix A.1.

Table I. Summary statistics: US data[a]

[a]The sample means and the sample standard deviations are in percentages. The cash-flow growth rates are geometric averages while returns are reported with simple compounding. AC1 denotes first-order autocorrelation. Standard errors obtained by performing a block bootstrap with each block having geometric distribution. The sample period is quarterly 1952:I–2011:IV.

Time series	Mean Estimate (SE)	Standard deviation Estimate (SE)	AC1 Estimate (SE)
Cash-flow growth rates			
Consumption	1.87 (0.10)	1.26 (0.11)	0.00 (0.00)
Levered dividends	2.06 (0.36)	10.38 (0.74)	0.01 (0.01)
Securities			
Risk-free rate	1.06 (0.45)	2.15 (0.34)	0.65 (0.04)
Equity risk premium	7.26 (1.59)	17.29 (1.10)	−0.13 (0.00)
Price–dividend ratio	33.24 (3.45)	33.59 (4.75)	0.82 (0.09)

Table I presents the summary statistics for consumption and aggregate equity market dividends in the USA. The geometric growth rates of the series hover most of the time around their unconditional means of 1.87% and 2.06% but the dividend series are more volatile with the annualized standard deviation of 10.38% compared to 1.26% for the consumption series. The first-order annualized autocorrelations in both series are negligible. The skewness coefficient is negative −0.44 and −0.45 due to the marked tendency to experience declines during economic downturns while the excess kurtosis of about 1.41 and 3.78 along with the quantile-to-quantile plots and the Kolmogorov–Smirnov tests against the null of Gaussian distribution (not reported) favor a leptokurtic distribution such as a Gaussian mixture density. Table I also presents the summary statistics for the aggregate equity market. The average equity risk premium (simple compounding) is about 7.26% with a volatility of about 17.29% per year. The null hypothesis of zero equity risk premium can be rejected at 5% significance level. The average price–dividend ratio is about 31.24 per year with annual volatility of about 33.59% and the first-order annualized autocorrelation coefficient of about 0.82.

The long-term annual data for real per-capita consumer expenditures for forty-two countries is from Barro and Ursua (2012) and it is described in the Supplementary Appendix A.2. The confidence interval for the relative frequencies of the periods with negative T-year growth rate are plotted in Figure 1 for $T = 1, \ldots, 30$. We classify periods of negative 10-year growth rate as lost decades and plot in Figure 2 the histogram of their relative frequencies, considering each country separately. The relative frequency of
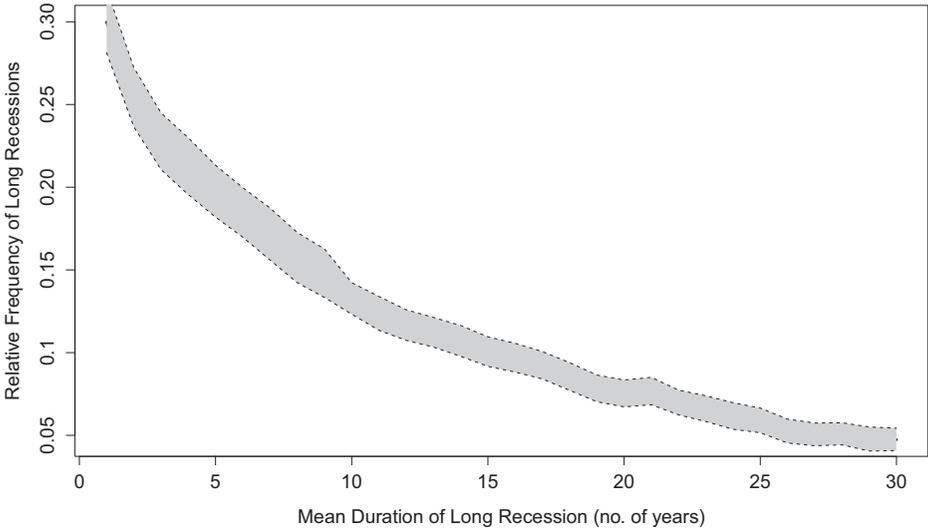
Figure 1. Relative frequency of long recessions: international evidence. *Notes*: The figure uses the international consumption data from Barro and Ursua (2012) as described in more detail in the Supplementary Appendix A.2. The confidence bounds correspond to percentile intervals from bootstrap with a random block length having geometric distribution.

lost decades displays a large cross-sectional variation: negligible in Indonesia and Philippines but more than 30% in Venezuela. The summary statistics in Table II reveals that the lost decades in the USA occur in the 1790–2009 sample with a relative frequency of about 12%, with a standard error of 5%, and average a mean growth rate of about −0.68%, with a standard error of 0.09%. This is significantly below a mean growth rate in the full Barro–Ursua cross-section of about −1.41%, with a standard error of 0.15%, but nonetheless it is surprisingly close to our maximum likelihood estimate $\underline{\mu}^e = -0.79\%$ in the USA 1952:I–2011:IV sample.

3.2 MAXIMUM LIKELIHOOD

The cash-flow model in Equation (3) depends on the parameter vector

$$\theta = \left(\bar{\mu}^u, \underline{\mu}^u, \bar{\mu}^l, \underline{\mu}^l, \sigma^u, \sigma^l, \lambda_1, \lambda_2, \lambda_3, q \right)', \tag{66}$$

subject to the restriction that the long-run geometric means of the consumption and the dividend growth rates $E\left\{ \mu_{s_t}^e \right\}$ for $e \in E$ are equal.

Rare long recessions are not observed in the US postwar data and thus we cannot estimate their mean duration λ_3^{-1} as well as their relative
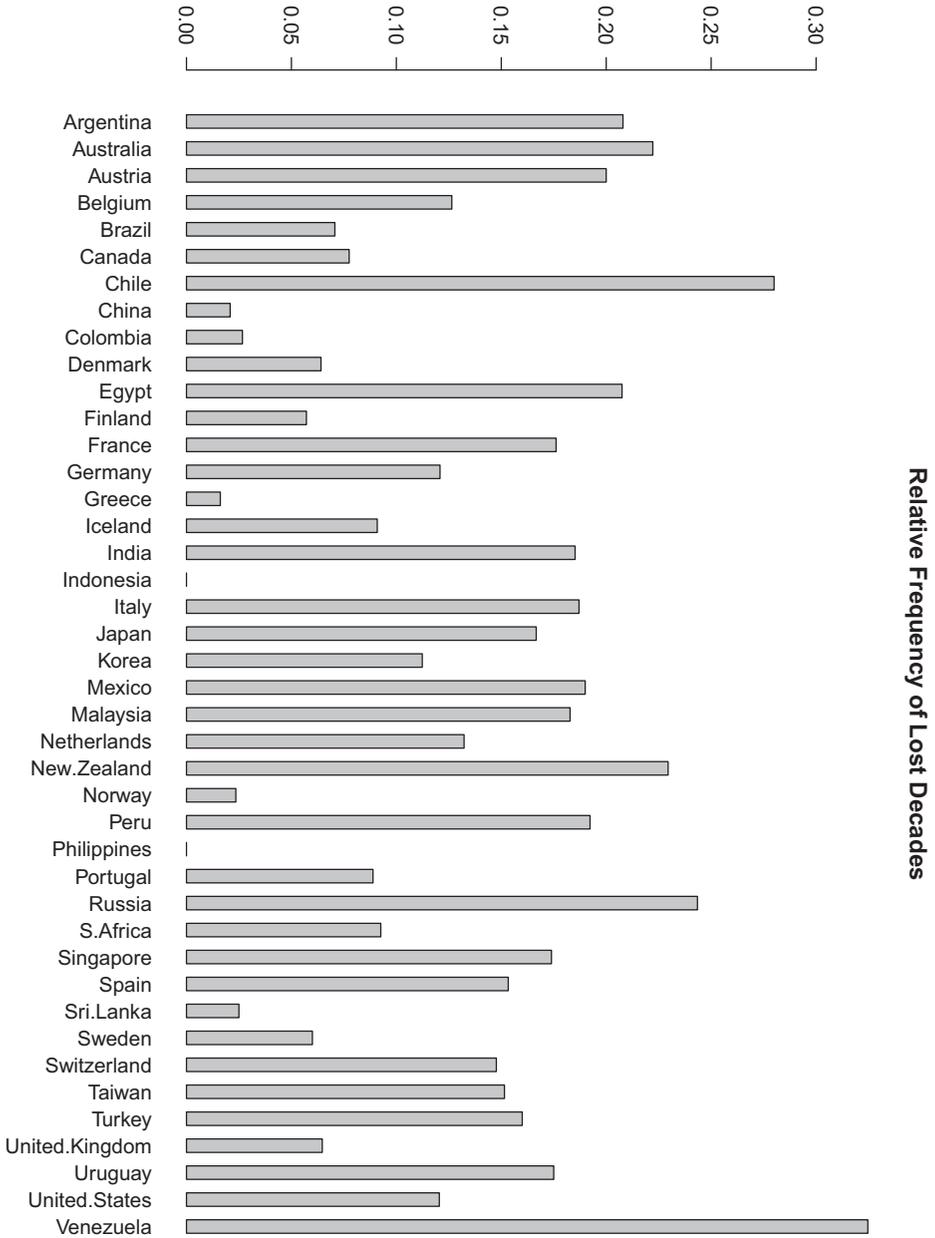
Figure 2. Relative frequency of lost decades: country-level evidence. *Notes:* The international data are from Barro and Ursua (2012) as described in more detail in the Supplementary Appendix A.2.

Table II. Summary statistics: international data[a]

[a]The international data are from Barro and Ursua (2012) as described in more detail in the Supplementary Appendix A.2. The growth rates are decade-long geometric means. Standard errors are obtained by performing a block bootstrap with each block having geometric distribution.

Time series	Mean Estimate (SE)	Standard deviation Estimate (SE)	Relative frequency Estimate (SE)
International sample			
Full	1.80 (0.10)	2.11 (0.11)	
Only declines	−1.41 (0.15)	1.75 (0.21)	0.13 (0.01)
US sample, 1790–2009			
Full	1.58 (0.25)	1.19 (0.13)	
Only declines	−0.68 (0.09)	0.44 (0.05)	0.12 (0.05)

frequency $\bar{\pi}_3$. For this reason, we calibrate these parameters based on the Barro–Ursua sample for the USA (1790–2009). Table II reveals that long recessions with mean duration of 10 years occur in this sample with the relative frequency of about 12% with a standard error of 5%. We set $\lambda_3^{-1} = 10$ years and $\bar{\pi}_3 = 10\%$ which implies exactly one lost decade per century on average. Our choice is also consistent with the broader evidence in the cross-section of forty-two countries in the Barro–Ursua sample where the relative frequency of lost decades is 13%.

We demonstrate in Section 3.2.1 that our choice of λ_3 and $\bar{\pi}_3$, together with the maximum likelihood estimates of the remaining parameters $\left(\bar{\mu}^u, \underline{\mu}^u, \bar{\mu}^l, \sigma^u, \sigma^l, \lambda_1, q\right)'$, implies empirically plausible magnitudes of the T-period-ahead forecasts of the dividend growth rate $m_{t,T}^e$ for each $e \in E$.

3.2.a. *Parameter estimates*

Table III reports the parameter values and the asymptotic standard errors of the maximum likelihood estimates for the transition intensities λ_i, consumption and the dividend growth rates μ_i^e as well as the volatility σ^e for each $e \in E$ and $i = 1, 2, 3$.

First, consumption is estimated to grow instantaneously at the annualized rate of about $\bar{\mu}^u = 2.65\%$ in expansions, and about $\underline{\mu}^u = -0.79\%$ in recessions, which is consistent with the sample mean growth rate of about −0.68%, with a standard error of 0.09%, experienced by the USA during its lost decades between 1790 and 2009 but not used in the estimation (Table II).[27]

---

[27] Table II shows that the average is about −1.41% per year with a standard error 0.15% in the Barro-Ursua cross-section of forty-two countries.

Table III. Cash-flow model: maximum likelihood estimates[a]

[a]We estimate the parameters of a two-state continuous-time hidden semi-Markov model that is discretely observed from the bivariate time series of consumption and dividend growth rates. We map the semi-Markov chain into a restricted three-state Markov chain by imposing the restriction $\mu_2^e = \underline{\mu}^e = \mu_3^e$ for each $e \in E$. The instantaneous volatility σ^e is constant across the hidden regimes for each $e \in E$. The maximum likelihood estimation allows for transitions of the continuous-time chain only at the quarter ends. The initial probability π_0 is set equal to the invariant distribution of the chain denoted by the symbol $\bar{\pi} = (\bar{\pi}_1, \bar{\pi}_2, \bar{\pi}_3)$. The reported parameter estimates are annualized. The sample period is quarterly 1952:I–2011:IV.

[b]The parameter λ_2 is restricted by Equation (12).

[c]We set the mean duration of the long recession equal to 10 years and so the hazard rate $\lambda_3^{-1} = 0.100$.

[d]We posit that the long recessions occur on average once a century and so the invariant probability $\bar{\pi}_3 = 0.100$.

Parameter	Estimate (SE)	Parameter	Estimate (SE)	Parameter	Estimate (SE)
Transition hazard rates		Instantaneous Consumption growth rate		Instantaneous Levered dividend growth rate	
λ_1	0.168 (0.038)	$\bar{\mu}^u$	2.650 (0.092)	$\bar{\mu}^l$	4.278 (0.716)
λ_2	0.944[b] (0.314)	$\underline{\mu}^u$	−0.786 (0.358)	$\underline{\mu}^l$	−6.330 (2.088)
λ_3	0.100[c]	σ^u	1.086 (0.041)	σ^l	10.162 (0.332)
Invariant distribution of the chain		Probability of the short recession			
$\bar{\pi}_1$	0.773 (0.088)	q	0.923 (0.034)		
$\bar{\pi}_2$	0.127 (0.088)				
$\bar{\pi}_3$	0.100[d]				

Next, the aggregate dividend is estimated to grow instantaneously at the annualized rate of about $\bar{\mu}^l = 4.28\%$ in expansions and about $\underline{\mu}^l = -6.33\%$ in recessions while the annualized estimate of the consumption volatility comes out around $\sigma^u = 1.09\%$, whereas it is about $\sigma^l = 10.16\%$ for the aggregate dividends.[28]

In addition, Table IV reports the long-run forecasts of consumption and dividend growth rates, conditional on all the hidden states. The annual consumption growth rate forecasts are $g_{T=1|1}^u = 2.44\%$ in expansions, whereas $g_{T=1|2}^u = 0.37\%$ in downturns and $g_{T=1|3}^u = -0.63\%$ in lost decades. In addition, the annualized decade-long forecasts are $g_{T=10|1}^u = 2.09\%$

---

[28] For comparison, David and Veronesi (2013) calibrate consumption volatility six times higher at 6.34% per year.

Table IV. Cash-flow model: implied T-period forecasts[a]

[a]The optimal T-period-ahead forecast from Equation (36) is conditional on the hidden state $m_{t,T|i}^{u} = E\{g_{t,T}^{u} \mid s_t = i\}$.

Forecast	T-period forecast $m_{t,T|i}$					
Horizon	Consumption growth rate			Dividend growth rate		
T	Expansion	Recession	Lost Decade	Expansion	Recession	Lost Decade
Quarter	2.58	−0.42	−0.74	4.07	−5.19	−6.20
Year	2.44	0.37	−0.63	3.63	−2.76	−5.83
Decade	2.09	1.79	0.29	2.54	1.64	−3.00

expansions, $g_{T=10|2}^{u} = 1.79\%$ in downturns, and $g_{T=10|3}^{u} = 0.29\%$ in lost decades. The lost decade $\tilde{s}_t = 3$ may in fact be thought of as a protracted, decade-long, period of anemic growth during which the consumption level is forecast to stagnate.[29]

As regards dividends, the decade-long forecasts come out $g_{T=40|1}^{l} = 2.54\%$ in expansions, $g_{T=40|2}^{l} = 1.64\%$ in downturns, and $g_{T=40|3}^{l} = -3.00\%$ in lost decades. For example, the cumulative drop in dividends over the whole 10-year duration of the lost decade, which happens once a century, is $-3.00\% \times 10 = -30.0\%$. This quantitative exercise demonstrates that much less consumption and dividend risk, measured in terms of the difference in the hidden growth rates $\bar{\mu}^{u} - \underline{\mu}^{u}$ and $\bar{\mu}^{l} - \underline{\mu}^{l}$ per unit of the volatility σ^{u} and σ^{l}, 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.

Second, the mean duration of the expansion λ_1^{-1} comes out almost 6 years and according to Equation (12) the mean duration of the short recession λ_2^{-1} comes out slightly above 1 year. The transition probability to the short recession, conditional on leaving the expansion state, $q = q\lambda_1/(q\lambda_1 + (1-q)\lambda_1)$, is estimated around 0.92. These estimates imply that the invariant distribution is $\bar{\pi} = (0.773, 0.127, 0.100)$, with each century experiencing on average 77 years of good times interrupted by about thirteen

---

[29] We note that the estimation procedure restricts only the mean duration of the lost decades λ_3^{-1} and their relative frequency in a century-long series (i.e., $\bar{\pi}_3 = 0.1$). The fact that the decade-long consumption-growth forecast comes out close to zero is dictated by the realized consumption and dividends series.

brief business-cycle recessions and about one lost decade. The Great
Recession of 2008 is exactly the 13th recession after the Great Depression.

3.2.b. *Hidden state estimates*

We use the maximum likelihood estimates from Table III and follow
Hamilton (1989) in order to obtain the time series of the filtered beliefs

$$\hat{\pi}_{i,t} = P\Big\{s_t = \tilde{s}_i|\mathcal{F}_t, \hat{\theta}\Big\} \tag{67}$$

for each $i = 1, 2, 3$ as well as the smoothed beliefs $P\{s_t = \tilde{s}_i|\mathcal{F}_T, \hat{\theta}\}$, which are
conditional on the whole sample period. The filtered belief $\hat{\pi}_{3,t}$ corresponds
to the probability of the consumption disaster.

Figure 3 shows that the filtered beliefs nicely track the NBER recessions
when $\hat{\pi}_{1,t}$ falls and $\hat{\pi}_{2,t}$ and $\hat{\pi}_{3,t}$ both rise due to the Peso problem discussed
in Section 2.3. Speaking quantitatively, the estimated disaster probability
$\hat{\pi}_{3,t}$ reaches magnitudes of almost 30% in the recessions while the corres-
ponding smoothed probability is estimated almost always below 10%. This
fact suggests that the US economy did not experience significantly pro-
tracted recessions *ex post* in the postwar sample but investors nonetheless
worried about such a possibility on an ongoing basis.

4. Implications for Consumption and Asset Prices

We assess the performance of the model by comparing the model-implied
unconditional and conditional asset-pricing moments to their sample coun-
terparts. The estimates of the dividend growth rate model in Equation (3) are
from Table III while the utility aggregator in Equation (2) is configured with
the relative risk aversion $\gamma = 10.0$, the elasticity of intertemporal substitution
$\psi = 150$, and the subjective discount rate $\delta = 0.015$. Our assumed level of the
relative risk aversion is 10, a value considered plausible by Mehra and
Prescott (1985) and used also by Bansal and Yaron (2004).

In order to obtain the conditional asset-pricing moments, we need to solve
the partial differential equations (87) and (88) for the unlevered and levered
price–dividend ratios $\Phi^u(\pi_1, \pi_2)$ and $\Phi^l(\pi_1, \pi_2)$, and further (98) for the
whole term structure of the zero-coupon real bond prices $P^b(\pi_1, \pi_2, t; T)$,
and finally (123) and (124) for the first two moments of the time-aggregated
annual levered equity return $M_i^l(\pi_1, \pi_2, t; T)$ for $i = 1, 2$. We then calculate
the (gross) bond yields $Y^{(T)}(\pi_1, \pi_2, t)$ as

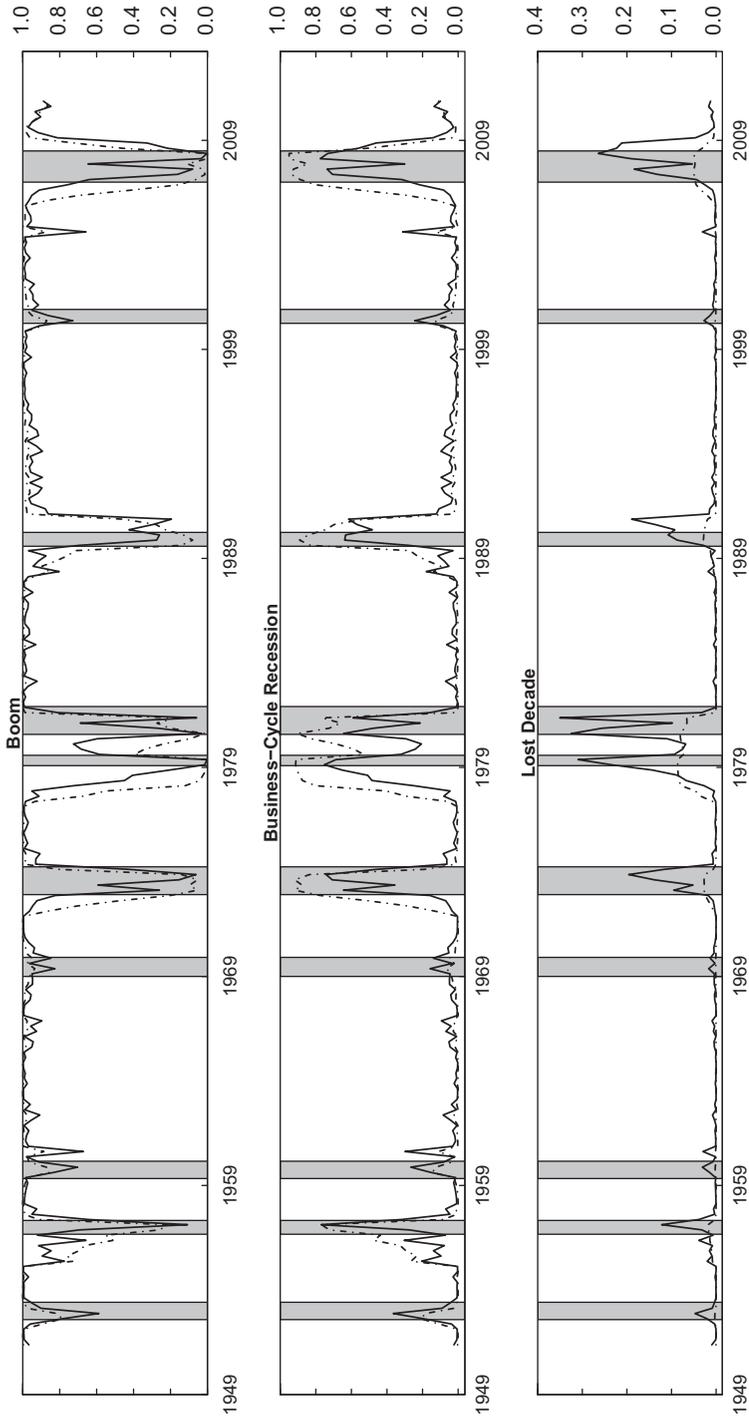$$Y^{(T)}(\pi_1, \pi_2, t) = \big(P^b(\pi_1, \pi_2, t; T)\big)^{-\frac{1}{T}}, \tag{68}$$

Figure 3. Time series of beliefs. *Notes*: Posterior probabilities $P\{S_t = \tilde{s}_i | \mathcal{F}_t, \hat{\theta}_{ML}\}$ plotted as solid lines and smoothed probabilities $P\{S_t = \tilde{s}_i | \mathcal{F}_T, \hat{\theta}_{ML}\}$ as dot-dashed lines. The sample period is quarterly from 1952:I to 2011:IV.

and the annual holding-period return $R^{b,T}(\pi_1, \pi_2, t)$ on the T-period zero-coupon real bond as

$$R^{b,T}(\pi_1, \pi_2, t) = \frac{P^b(\pi_1, \pi_2, t; T-1)}{P^b(\pi_1, \pi_2, t; T)}. \tag{69}$$

We additionally calculate the levered equity risk premium $E^l(\pi_1, \pi_2, t)$ as the expected levered equity return in excess of the 1-year yield,

$$M^l(\pi_1, \pi_2, t) = M_1^l(\pi_1, \pi_2, t; 1) - Y^{(1)}(\pi_1, \pi_2, t), \tag{70}$$

the levered equity volatility $V^l(\pi_1, \pi_2, t)$ as the second moment minus the first moment squared

$$V^l(\pi_1, \pi_2, t) = \left(M_2^l(\pi_1, \pi_2, t; 1) - \left(M_1^l(\pi_1, \pi_2, t; 1) \right)^2 \right)^{\frac{1}{2}}. \tag{71}$$

The unconditional moments are obtained by Monte Carlo integration with respect to the invariant distribution of the beliefs. First, we use the Euler–Maruyama scheme in order to solve the stochastic differential equation in Equation (3) in the finer filtration $\mathcal{G}_t = \sigma(\mathcal{F}_t \cup \{s_\tau : 0 \leq \tau \leq t\})$ for $t > 0$ and obtain the time series of the cash flows D_t^e for $e \in E$, the shocks z_t^e, and the hidden states s_t. We then invoke Equation (14) in order to construct the time series of the instantaneous cash-flow forecast errors $d\tilde{z}_t^e$ as well as the beliefs $\pi_t = (\pi_{1t}, \pi_{2t}, \pi_{3t})'$. The time series obtained $\left(D_t^u, D_t^l, s_t, \pi_{1,t}, \pi_{2,t}, \pi_{3,t} \right)'$ contains both the hidden state as well as the beliefs about that hidden state, given the coarse information set $\mathcal{F}_t$ available to the investor.

Second, we construct the time series of the T-period expected consumption growth rate $m_{t,T}^e$ and the T-period consumption growth rate volatility $\sigma_{t,T}^e$ by means of Equations (36) and (38). Furthermore, we use Equations (68) and (69) in order to construct the real bond yields $Y_t^{(T)}$ and the annual holding period returns $R_t^{b,T}$ for maturities up to 30 years. In addition, we construct the equity risk premium M_t^l, equity return volatility V_t^l, equity Sharpe ratio M_t^l/V_t^l, and the equity price–dividend ratios Φ_t^e for $e \in E$ by means of Equations (55), (70), and (71). Finally, we calculate the unconditional moments, such as the mean, the standard deviation or the first-order auto-correlation, as the corresponding sample statistics.

In addition, we evaluate the performance of the model based on the beliefs estimated from the actual postwar US consumption and dividend data. Such a stringent test is often absent in the literature because generating plausible historical posterior beliefs based on the actual consumption data is challenging; see the recent paper of Ju and Miao (2012, Figure 3) for a notable exception. As we document below, our model generates plausible

and comparable dynamics for the levered equity prices using both the historical and the simulated posterior beliefs.

4.1 CONSUMPTION

In Table V we report the means, the standard deviations, and the first-order autocorrelations of the realized and the simulated data as well as the corresponding annual variance ratios for the horizons up to 5 years. As can be seen, the simulated cash-flow model in Equation (3) nicely matches the salient features of the consumption and dividend data.

Furthermore, the analysis in Section 2.4.3 predicts that the annual consumption growth rate forecast $m_{t,1}^u$ is procyclical and the consumption growth rate volatility $\sigma_{t,1}^u$ is countercyclical. Table VII confirms these predictions: the mean forecast $m_{t,1}^u$ is about 1.04% in recessions and 2.08% during expansions, and the consumption volatility $\sigma_{t,1}^u$ is about 1.74% in recessions but only 1.36% in expansions. In addition, Section 2.4.3 predicts a rising pattern for the annual forecast, and a falling one for the annual volatility, over the expansion, and *vice versa* for the recession. This prediction is confirmed in Table VII as well.

Table V. Consumption and dividend moments[a]

[a]The reported entries for the mean, standard deviation, and the first-order autocorrelation are quarterly moments annualized by multiplying by 4, 2 and taking to power of 4, respectively. The population moments are obtained from a Monte Carlo simulation of the length 4 million quarters and then time-aggregated from the infinitesimal quantities to their quarterly counterparts using the parameter estimates from Table III. Annual variance ratios are computed in the same way as in Cecchetti, Lam, and Mark (1990, 2000).

	Annualized Moments			Variance ratios				
	Mean	SD	AC1	One Year	Two Years	Three Years	Four Years	Five Years
Sample								
	Panel A: 1952:I–2011:IV							
Consumption growth	1.89	1.26	0.01	1.00	1.42	1.68	1.06	1.64
	(0.10)	(0.12)	(0.00)		(0.17)	(0.28)	(0.40)	(0.44)
Dividend growth	2.06	10.38	0.01	1.00	1.33	0.94	1.10	1.25
	(0.36)	(0.72)	(0.01)		(0.13)	(0.31)	(0.34)	(0.38)
Population								
	Panel B: With lost decades							
Consumption growth	1.87	1.29	0.00	1.00	1.40	1.72	1.98	2.21
Dividend growth	1.95	10.38	0.00	1.00	1.09	1.16	1.22	1.27

Our model of consumption given by Equation (3) is thus able to endogenously generate countercylical uncertainty shocks in Bloom (2009), Bloom *et al.* (2012), and Baker and Bloom (2012).

4.2 ASSET PRICES

Table VI in Panel A presents the model-implied levered-equity and real-bond pricing moments and compares them to their sample counterparts. Table VI in Panel B repeats the same analysis using the inferred beliefs from the actual postwar US consumption and dividend data. Table VII presents the variation in the moments of the conditional distributions of the levered-equity and real-bond prices over the various phases of the business cycle, and compares them to their empirical evidence in Lustig and Verdelhan (2013).

4.2.a. *Levered equity*

In this Section, we refer to Tables VI and VII jointly.

First, the unconditional risk premium of about $E\{M_t^l\} = 6.29\%$ per year from Panel A in Table VI compares well to the sample estimate of 7.26% with a standard error of 1.59% in Table I. The conditional risk premium varies significantly over time, having a standard deviation of about $\sigma\{M_t^l\} = 4.11\%$. In addition, Table VII reveals that such variation in the risk premium occurs at the business cycle frequency: the mean equity risk premium in short recessions comes out about 8.88%, significantly above the mean risk premium of 5.52% during expansions. These numbers compare surprisingly well to the point estimates of 11.31% with the standard error 2.20% and 5.28% with standard error 1.87% obtained by Lustig and Verdelhan (2013).

Second, the mean return volatility from Panel A in Table VI is about $E\{V_t^l\} = 15.78\%$ per year and also displays a large variation over time, having the standard deviation of about $\sigma\{V_t^l\} = 3.59\%$. The total volatility of the realized excess return is given by the square root of the mean of the conditional variance $E\{(V_t^l)^2\}$ plus the variance of the conditional mean $\sigma^2\{M_t^l\}$. It comes out about 16.31% per year, close to the point estimate of 17.29% with a standard error 1.10% in Table I. As before, Table VII reveals the strong business-cycle variation: the mean volatility in short recessions comes out about 19.43%, which is significantly higher than the mean volatility of 14.71% during expansions.

Third, the mean Sharpe ratio from Panel A in Table VI is about $E\{M_t^l/V_t^l\} = 0.37$ with a standard deviation of $\sigma\{M_t^l/V_t^l\} = 0.15$. As before, Table VII reveals that the large variation in the Sharpe ratio M_t^l/V_t^l is tightly linked to the business cycle: the mean Sharpe ratio during short recessions is about 0.43 and during expansions about 0.35. These

Table VI. Asset pricing moments: two-state Markov chain versus two-state semi-Markov chain[a]

[a]The reported entries are the asset-pricing moments obtained as follows. First, we solve Equation (87) for the price-consumption ratio $\Phi^u(\pi_1,\pi_2)$. Second, we solve Equation (88) for the price–dividend ratio $\Phi^l(\pi_1,\pi_2)$. Third, we solve for the whole term structure of the zero-coupon real bond prices $P^b(\pi_1,\pi_2,0;T)$ by solving Equation (98) for T up to 30 years. Fourth, we solve Equation (123) and Equation (124) for the first two moments of the time-aggregated annual levered equity return $M^l_i(\pi_1,\pi_2,0;T)$. We then calculate the (gross) bond yields as $Y^{(T)}(\pi_1,\pi_2) = (1/P^b(\pi_1,\pi_2,0;T))^{1/T}$, the stock risk premium as $M^l_1(\pi_1,\pi_2,0;1) - Y^{(1)}(\pi_1,\pi_2,0;1)$, the stock volatility as $M^l_2(\pi_1,\pi_2,0;1) - (M^l_1(\pi_1,\pi_2,0;1))^2$, the short-term yield as $Y^{(1)}(\pi_1,\pi_2)$ and the long-term yield as $Y^{(30)}(\pi_1,\pi_2)$, and the annual holding period return on T-period zero-coupon real bond as $\mathrm{HPR}^{(T)}(\pi_1,\pi_2) = (P^b(\pi_1,\pi_2,0;T-1)/P^b(\pi_1,\pi_2,0;T))$. These conditional moments being dependent on the beliefs (π_1,π_2) are then conditioned down to their unconditional counterparts with Monte Carlo integration by simulating a sample path of the beliefs from Equation (72) of the length 1 million years. The cash-flow parameter estimates are from Table III and the preference parameter values are set to $\gamma = 10.0$, $\psi = 1.50$ and $\delta = 0.015$.

| | Panel A. simulated beliefs (4 million quarters of artificial data) | | | | | | Panel B. historical beliefs (239 quarters from 1952:II to 2011:IV) | | | | | |
| | Two-state Markov model | | | Two-state semi-Markov model | | | Two-state Markov model | | | Two-state semi-Markov model | | |
	Mean	SD	AC1	Mean	SD	AC1	Mean	SD	AC1	Mean	SD	AC1
Levered equity												
Risk premium	0.95	0.44	0.27	6.29	4.11	0.35	0.90	0.44	0.26	5.58	3.52	0.23
Volatility	11.65	0.57	0.29	15.78	3.59	0.47	11.59	0.58	0.28	14.99	2.99	0.40
Sharpe ratio	0.08	0.03	0.27	0.37	0.15	0.29	0.08	0.03	0.26	0.35	0.14	0.11
Price-dividend ratio	111.81	0.02	0.33	22.48	12.20	0.75	111.88	0.02	0.35	23.26	8.08	0.43
Indexed bond prices												
Short-term yield	2.51	0.40	0.33	1.78	0.90	0.54	2.52	0.44	0.35	2.03	0.79	0.45
Long-term yield	2.39	0.02	0.33	0.80	0.10	0.58	2.39	0.02	0.35	0.83	0.08	0.45
30-Year term premium	-0.13	0.63	-0.07	-1.08	2.34	0.01	-0.13	0.72	-0.03	-1.30	2.46	-0.09

Table VII. Business-cycle variation in consumption and levered equity prices[a]

[a]The asset-pricing moments are constructed analogously to Table VI. We let n denote the number of quarters since the beginning of each epoch (i.e., the high-growth one $\tilde{s} \in \{\tilde{s}_1\}$ and the low-growth one $\tilde{s} \in \{\tilde{s}_2, \tilde{s}_3\}$). The reported entries in the table are the means of the asset-pricing moments conditional jointly on the epoch type $\{\{\tilde{s}_1\}, \{\tilde{s}_2, \tilde{s}_3\}\}$ and $n \in \{1, 2, 3, 4, 5\}$.

[b]*Source:* Panel II in Table II in Lustig and Verdelhan (2013).

| Conditional moment | Business-cycle recession | | | | | | Business-cycle expansion | | | | | |
| | Starting in n-th quarter after the regime shift and ending 1 year later | | | | | Mean | Starting in n-th quarter after regime shift and ending 1 year later | | | | | Mean |
	$n=1$	$n=2$	$n=3$	$n=4$	$n=5$		$n=1$	$n=2$	$n=3$	$n=4$	$n=5$	
Consumption growth												
Mean	1.94	1.43	0.93	0.65	0.51	1.04	1.34	1.84	2.01	2.06	2.07	2.08
Volatility	1.47	1.72	1.78	1.79	1.79	1.74	1.72	1.54	1.42	1.38	1.36	1.36
Levered equity												
Data[b]												
Mean return	7.53	14.13	11.45	12.96	10.49	11.31	7.45	2.77	1.89	5.59	8.67	5.28
	(5.27)	(5.20)	(5.38)	(4.66)	(4.70)	(2.20)	(5.05)	(5.30)	(4.92)	(4.36)	(4.59)	(1.87)
Sharpe ratio	0.43	0.78	0.62	0.82	0.67	0.66	0.51	0.19	0.14	0.42	0.67	0.38
	(0.31)	(0.31)	(0.31)	(0.31)	(0.31)	(0.14)	(0.30)	(0.32)	(0.29)	(0.29)	(0.30)	(0.14)
Two-state semi-Markov model												
Mean risk premium	6.86	9.01	10.15	10.94	11.54	8.88	10.96	10.55	8.98	7.53	6.50	5.52
Conditional volatility	15.87	18.11	19.54	20.52	21.24	19.43	19.89	18.45	16.92	15.84	15.16	14.71
Sharpe ratio	0.41	0.49	0.51	0.53	0.54	0.43	0.53	0.53	0.49	0.43	0.39	0.35
Price–dividend ratio	22.88	21.60	20.69	20.06	19.61	19.11	20.42	21.51	22.35	22.92	23.27	23.47

numbers again compare surprisingly well to the point estimates of 0.66 with standard error 0.14 and 0.38 with standard error 0.14 in Lustig and Verdelhan (2013).

Fourth, the mean price–dividend ratio comes out as $E\{\Phi_t^l\} = 22.48$ and displays volatility of nearly $\sigma\{\log \Phi_t^l\} = 12.20\%$ per year. The price–dividend ratio volatility is below the sample counterpart of 33.59% with a standard error of 4.75% reported in Table I. It is arguably difficult to match the volatility of prices in a model where the mean consumption growth rate switches between high and low values only. Nonetheless, the model can match the persistence of the price–dividend ratio measured by the first-order autocorrelation AC1, which comes out about 0.75 in annual data in Panel A of Table VI and compares favorably with the point estimate of 0.82 with a standard error of 0.09 in Table I.

Fifth, Panel B in Table VI reveals that the model generates comparable values for unconditional moments for levered equity prices using the historical beliefs as well.

4.2.b. *Real bonds*

Speaking quantitatively, the average short-term yield 1.78% is above the average long-term yield of about 0.80%. The short-term yield is also more volatile, with the standard deviation of about 0.90%, than the long-term yield with the standard deviation of about 0.10%. As a result, the average yield curve is mildly downward sloping, which is consistent with the empirical findings in Ang, Bekaert, and Wei (2008), Campbell, Shiller, and Viceira (2009), and Piazzesi and Schneider (2007).

Furthermore, our model with learning features significantly lower persistence of the consumption growth rate, which, according to Proposition 2.3, generates a negligible bond risk premium on a 30-year zero-coupon bond of about -1.08%. The negative sign for the bond risk premium comes from the countercyclical variation in the real bond prices, which fall in good times and rise in bad times, as explained in Section 2.8. In contrast, the related long-run risk literature can generate a sizable equity risk premium only if the expected consumption growth rate is highly persistent, in which case the risk premium on real zero-coupon bonds is highly negative as discussed in Beeler and Campbell (2012).

4.3 CONSUMPTION AND ASSET PRICES OVER THE PHASES OF THE BUSINESS CYCLE

As explained in Section 2.4.3, a regime shift to the high-growth state $s = \tilde{s}_1$ is associated with rising T-period forecasts of the cash-flow growth rate

$m_{t,T}^u = E_t(\mu_{t,T}^u)$ as well as falling economic uncertainty measured by the corresponding T-period cash-flow volatility $\sigma_{t,T}^u = \sigma_t\{g_{t,T}^u\}$. These predictable changes occur gradually while the regime shift is being recognized in terms of the posterior odds in favor of the high-growth state $O_{1,23}$ in Equation (20). From the perspective of the equity pricing, the shift is associated with falling levered-equity risk premiums, return volatility, and Sharpe ratios as well as a rising levered-equity price–dividend ratio, as shown in Table VII. From the perspective of the real-bond pricing, the shift is associated with rising bond yields and holding-period excess returns in Table VIII. From the perspective of macroeconomics, the shift is associated with rising annual forecasts of the consumption growth rate $m_{t,1}^u$ and falling annual consumption growth rate volatility $\sigma_{t,1}^u$, as also shown in Table VII.

In contrast, a regime shift to the low-growth state $s \in \{\tilde{s}_2, \tilde{s}_3\}$ is associated with falling T-period consumption growth-rate forecasts $m_{t,T}^u$ as well as rising economic uncertainty measured by the corresponding T-period consumption volatility $\sigma_{t,T}^u$. These predictable changes also occur gradually over time as the regime shift is being recognized in terms of the posterior odds $O_{1,23}$ and O_{23}. The posterior odds $O_{1,23}$ are informative about the growth rate (i.e., high growth rate versus low growth rate), whereas the posterior odds O_{23} in Equation (23) are informative about the growth rate persistence (i.e., high persistence versus low persistence). In fact, the uncertainty about the persistence does not arise in a standard two-state Markov chain setting. From the perspective of the equity pricing, the shift is associated with rising levered-equity risk premiums, return volatility, and Sharpe ratios as well as a falling levered-equity price–dividend ratio, as shown in Table VII. From the perspective of the real-bond pricing, the shift is associated with falling bond yields and holding-period excess returns in Table VIII. From the perspective of macroeconomics, the shift is associated with falling annual forecasts of the consumption growth rate $m_{t,1}^u$ and rising annual consumption growth-rate volatility $\sigma_{t,1}^u$, as also shown in Table VII.

The predictable variation in the cash-flow forecasts and the discount rates depends crucially on the fact that the economic uncertainty is declining over time after the regime shift to the high-growth state, but rising after the shift to the low-growth state. This single learning mechanism has the power to generate (i) the procyclical variation in the price–dividend ratio; (ii) the countercyclical variation in mean risk premium, return volatility, and Sharpe ratio; (iii) the rising pattern of the risk premiums, the return volatility, and the Sharpe ratios during recessions and falling pattern during the expansions; (iv) the leverage effect; (v) the mean reversion of excess returns; and (vi) the predictability of consumption volatility from the price–dividend

Table VIII. Business-cycle variation in real bond prices[a]

[a]The asset-pricing moments are constructed analogously to Table VI. We let n denote the number of quarters since the beginning of each epoch (i.e., the high-growth one $\tilde{s} \in \{\tilde{s}_1\}$ and the low-growth one $\tilde{s} \in \{\tilde{s}_2, \tilde{s}_3\}$). The reported entries in the table are the means of the asset-pricing moments conditional jointly on the epoch type $\{\{\tilde{s}_1\}, \{\tilde{s}_2, \tilde{s}_3\}\}$ and $n \in \{1, 2, 3, 4, 5\}$.

Asset	Business-cycle recession						Business-cycle expansion					
	Starting in n-th quarter after the regime shift and ending a year later					Mean	Starting in n-th quarter after regime shift and ending a year later					Mean
	$n=1$	$n=2$	$n=3$	$n=4$	$n=5$		$n=1$	$n=2$	$n=3$	$n=4$	$n=5$	
Yield-to-maturity												
1-year	1.78	1.20	0.86	0.64	0.50	0.71	0.81	1.13	1.48	1.76	1.94	2.09
3-year	1.49	1.08	0.81	0.64	0.52	0.71	0.82	1.09	1.33	1.51	1.63	1.71
5-year	1.30	0.98	0.78	0.65	0.57	0.71	0.80	1.01	1.19	1.32	1.40	1.46
10-year	1.06	0.88	0.76	0.68	0.64	0.71	0.77	0.89	1.00	1.07	1.12	1.15
30-year	0.88	0.78	0.72	0.68	0.66	0.69	0.73	0.79	0.85	0.89	0.91	0.93
Risk premium												
3-year	0.62	0.17	−0.11	−0.26	−0.33	0.00	−1.63	−1.33	−1.07	−0.90	−0.79	−0.71
5-year	0.94	0.26	−0.18	−0.40	−0.51	−0.02	−2.41	−1.95	−1.56	−1.30	−1.14	−1.04
10-year	1.16	0.31	−0.23	−0.51	−0.65	−0.05	−3.01	−2.42	−1.94	−1.61	−1.40	−1.27
30-year	1.23	0.32	−0.25	−0.56	−0.71	−0.07	−3.27	−2.63	−2.10	−1.74	−1.51	−1.36

ratio. In particular, Table VII reveals that the variation in the conditional moments of the levered-equity prices over the various phases of the business cycle compares quantitatively to the empirical evidence in Lustig and Verdelhan (2013).

5. Robustness of Results

In this section we perform the following robustness analysis. First, we assess the performance of the two-state Markov chain model, which is nested in our semi-Markov setting for $q = 1$. We then examine the sensitivity of our results to the choice of the two parameters $(\lambda_3, \bar{\pi}_3)$, which cannot be estimated from the short sample. We end by briefly assessing the implications for the European options.

5.1 TWO-STATE MARKOV CHAIN AS A NESTED MODEL

The standard two-state Markov chain is nested in our framework for $q = 1$. Table VI reveals that the model without lost decades performs marginally better than Mehra and Prescott (1985) because the states s_i for $i = 1, 2$ are hidden, which generates a small uncertainty premium due to the preference for early resolution of uncertainty coming from the Epstein–Zin preferences. Despite that, the semi-Markov model with lost decades dramatically outperforms the Markov model in every dimension reported.

5.2 SENSITIVITY TO NONESTIMATED PARAMETERS

The cash-flow model parameters $(\lambda_3, \bar{\pi}_3)$ are difficult to estimate given our short sample. In our empirical approach, which is discussed in Section 3.2, we choose thoughtfully the mean duration of the long recession $\lambda_3^{-1} = 10$ years and the invariant distribution of the lost decade $\bar{\pi}_3 = 0.1$. We then invert the constraint in Equation (12) for the mean duration of the short recession λ_2^{-1} as a function of q, λ_1^{-1}, λ_3^{-1} and $\bar{\pi}_3$. Such choice of the parameters $(\lambda_3, \bar{\pi}_3)$ implies that each century features on average about one lost decade.

In order to examine the sensitivity of our results to the above choice of λ_3 and $\bar{\pi}_3$, it is more natural to consider the pair (λ_2, λ_3), and then obtain $\bar{\pi}_3$ from the constraint in Equation (12). We consider the hazard rate of the short recession $\lambda_2 \in \{0.5, 1.0, 1.5\}$ and the hazard rate of the long recession $\lambda_3 \in \{0.08, 0.10, 0.12\}$, obtaining $3 \times 3 = 9$ candidate cases to consider. Table IX presents the asset-pricing implications for the levered equity and

Table IX. Asset-pricing moments: sensitivity analysis[a]

[a]The reported entries are constructed by following the steps as in Table VI. In each case, we modify the calibration in Table III by varying the durations of the short recession $\lambda_2 \in \{0.5, 1.0, 1.5\}$ and the long recession $\lambda_3 \in \{0.08, 0.10, 0.12\}$. The stationary probability $\bar{\pi}$ is specific to each case.

		Transition intensity of the lost decade, duration of the lost decade in years								
		$\lambda_3 = 0.12$, 8.33 years			$\lambda_3 = 0.1$, 10 years			$\lambda_3 = 0.08$, 12.5 years		
	Asset-pricing Moments	Mean	SD	AC1	Mean	SD	AC1	Mean	SD	AC1
		$\bar{\pi} = (0.83, 0.08, 0.09)$			$\bar{\pi} = (0.81, 0.08, 0.10)$			$\bar{\pi} = (0.79, 0.08, 0.13)$		
$\lambda_2 = 1.5$, 0.67 years	Equity premium	6.92	5.18	0.29	7.12	5.18	0.28	7.22	5.19	0.28
	Equity volatility	15.84	4.01	0.41	15.76	3.78	0.38	15.62	3.58	0.35
	Equity Sharpe ratio	0.40	0.19	0.21	0.42	0.20	0.22	0.43	0.21	0.26
	Price–dividend ratio	24.09	12.51	0.72	22.82	12.79	0.74	21.74	13.12	0.77
	One-year bond yield	1.84	0.94	0.44	1.73	0.96	0.42	1.61	0.98	0.40
	30-Year bond yield	0.82	0.11	0.53	0.79	0.09	0.50	0.78	0.08	0.46
	30-Year term premium	−1.16	2.73	0.00	−1.04	2.46	0.01	−0.90	2.20	0.02
		$\bar{\pi} = (0.80, 0.12, 0.08)$			$\bar{\pi} = (0.78, 0.12, 0.10)$			$\bar{\pi} = (0.76, 0.12, 0.12)$		
$\lambda_2 = 1.0$, 1 year	Equity premium	6.53	4.46	0.35	6.75	4.47	0.32	6.90	4.51	0.31
	Equity volatility	16.14	3.98	0.48	16.06	3.77	0.45	15.94	3.57	0.42
	Equity Sharpe RAtio	0.37	0.16	0.25	0.39	0.17	0.24	0.40	0.18	0.26
	Price–dividend ratio	23.97	12.47	0.72	22.69	12.62	0.74	21.62	12.77	0.75
	One-year bond yield	1.84	0.90	0.51	1.74	0.92	0.50	1.64	0.93	0.48
	30-Year bond yield	0.84	0.12	0.58	0.81	0.10	0.55	0.79	0.09	0.51
	30-Year term premium	−1.15	2.84	0.00	−1.04	2.57	0.00	−0.91	2.31	0.02
		$\bar{\pi} = (0.71, 0.22, 0.08)$			$\bar{\pi} = (0.69, 0.22, 0.09)$			$\bar{\pi} = (0.68, 0.21, 0.11)$		
$\lambda_2 = 0.5$, 2.0 years	Equity premium	5.38	2.87	0.30	5.70	3.02	0.25	5.93	3.17	0.22
	Equity volatility	16.03	3.17	0.54	16.10	3.10	0.49	16.07	3.00	0.45
	Equity Sharpe ratio	0.32	0.11	0.17	0.34	0.12	0.15	0.35	0.13	0.16
	Price–dividend ratio	23.72	11.53	0.73	22.31	11.91	0.75	21.17	12.22	0.76
	One-year bond yield	1.81	0.82	0.59	1.71	0.84	0.59	1.61	0.85	0.58
	30-Year bond yield	0.94	0.13	0.69	0.87	0.12	0.67	0.83	0.10	0.64
	30-Year term premium	−1.03	2.82	−0.03	−0.97	2.62	−0.02	−0.87	2.38	−0.01

Transition intensity of the recession, duration of the recession in years

the real bonds in the form of a 2D matrix.[30] We additionally report the implied invariant distribution $\bar{\pi} = (\bar{\pi}_1, \bar{\pi}_2, \bar{\pi}_3)'$. As can be observed, the asset-pricing moments for the candidate calibrations are of comparable magnitudes to the benchmark results in Table VI. Interestingly, higher mean duration of the short recessions λ_2^{-1} lowers the equity premium because it weakens the Peso problem by slowing down the learning about the recession type.

---

[30] The calibration in the middle of the matrix $\lambda_2^{-1} = 1$ and $\lambda_3^{-1} = 10$ is closest to the choice in Section 3.2 with $\lambda_2^{-1} = 0.944^{-1}$ and $\lambda_3^{-1} = 10$.

5.3 DISCUSSION OF EUROPEAN OPTIONS

Modeling consumption disasters as large negative declines in the realized consumption growth rate results in option prices less in line with the data. In the words of Backus, Chernov, and Martin (2011) on p. 1994:

> The consumption-based calibration has a steeper smile, greater curvature, and lower at-the-money volatility. This follows, in part, from its greater risk-neutral skewness and excess kurtosis.... They suggest higher risk-neutral probabilities of large disasters and lower probabilities of less extreme outcomes.

Our results suggest, in contrast, that modeling consumption disasters as protracted periods of anemic consumption growth rate results in option prices consistent with the data. Because the gist of our article is not option pricing, our analysis is necessarily brief.

We first solve numerically the partial differential equation in Equation (106) in Proposition D.3 for the equilibrium price $P^c(\pi_1, \pi_2, t)$ of the European call option with 3-months to maturity, as described in the Supplementary Appendix D. We then calculate the implied volatility function $I(\pi_1, \pi_2, \frac{\bar{P}^l}{P^l})$ by equating the equilibrium option price P^c to the Black–Scholes formula. The average implied volatility curve as a function of the moneyness $\log(\bar{P}^l/P^l)$ is then constructed as the sample mean $\frac{1}{T}\sum_{t=1}^{T} I(\hat{\pi}_{1,t}, \hat{\pi}_{2,t}, \frac{\bar{P}^l}{P^l})$, where the beliefs $\hat{\pi}_{1,t}$ and $\hat{\pi}_{2,t}$ are estimated from the consumption and dividend data. We consider nine candidate calibrations for the two nonestimated parameters $\lambda_2 \in \{0.5, 1.0, 1.5\}$ and $\lambda_3 \in \{0.08, 0.10, 0.12\}$.

The results are plotted in Figure 4, a counterpart to Figure 5 in Backus, Chernov, and Martin (2011). As can be observed, our model seems to be able to resolve the discrepancy stated in the quote by Backus *et al*.: the implied volatility curves in our model are mildly downward sloping and display negligible curvature.

6. Conclusion

Our model is a minimal extension of the Mehra–Prescott–Rietz asset-pricing framework to an incomplete information setting that can explain a broad range of dynamic phenomena in macroeconomics and finance. We show that the success of the model is attributable to the interplay of two key factors. First, we extend the standard two-state hidden Markov model to a two-state hidden semi-Markov setting which introduces variable growth persistence. Learning about growth persistence dramatically magnifies the level as well as
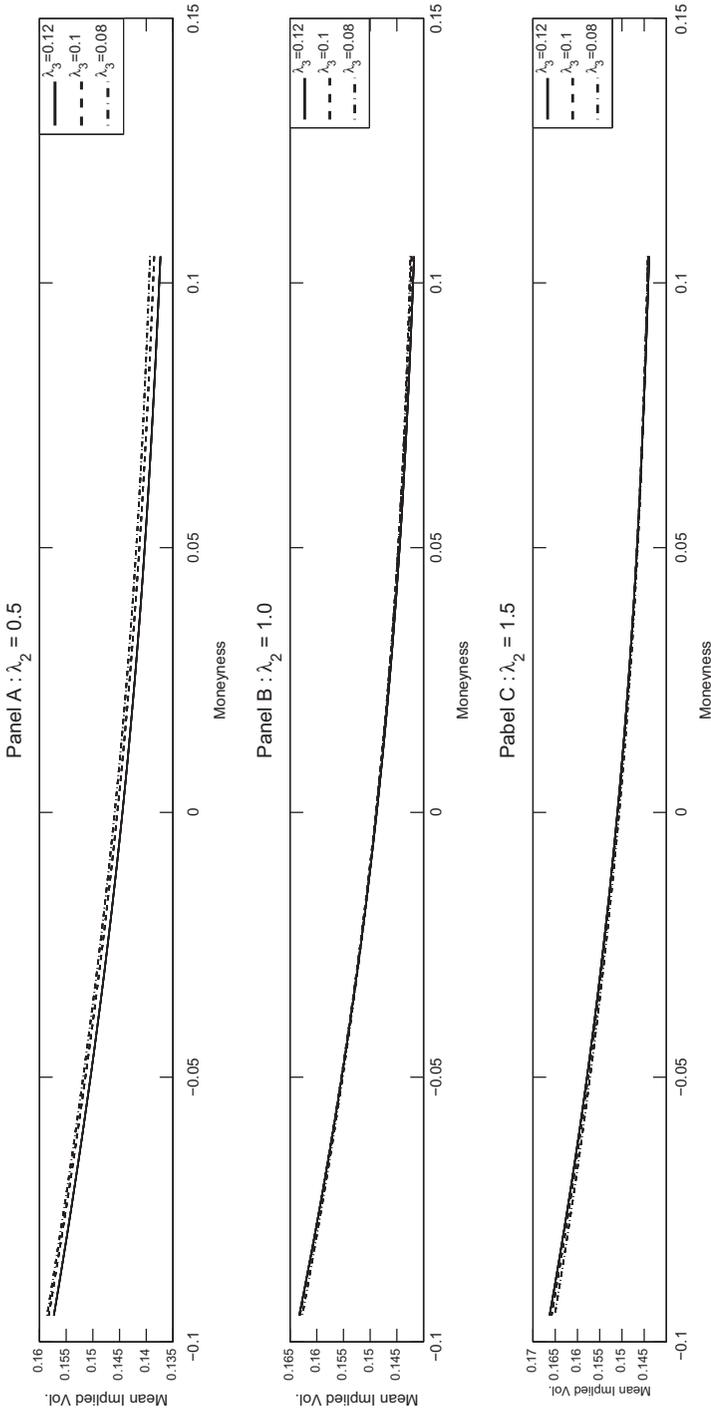
Figure 4. Implied volatility curves on three-month levered-equity call option. *Notes*: The moneyness is measured in terms of $\log(\bar{P}^l/P^l)$ where $\bar{P}^l$ is the strike price and P^l is the price of the levered equity (the underlying asset).

the variation in economic uncertainty. Second, we relax the independence axiom of the expected utility by using the recursive Epstein–Zin preferences configured so that early resolution of uncertainty is preferred. This makes assets with uncertain future payoffs comparably less valuable, increasing their equilibrium risk premiums.

It is the interplay of the fluctuations in the economic uncertainty due to learning, and higher risk premiums due to the preference for the early resolution of uncertainty, which makes the asset desirability not only lower, but also fluctuating over time in response to the variation in the average time needed to resolve the payoff uncertainty.

We estimate the model using maximum likelihood on the US postwar sample of consumption and dividend series. The model can generate endogenously the following array of consumption and asset-pricing phenomena:

- consumption growth rate:
  - procyclical variation in the T-period forecasts, including
    - their rising pattern during expansions,
    - their falling pattern during recessions,
  - countercyclical variation in the T-period forecast error volatility, including
    - their falling pattern during expansions,
    - their rising pattern during recessions,

for any forecast horizon T;

- equity prices:
  - procyclical variation in the price–dividend ratios, including
    - their rising pattern during expansions,
    - their falling pattern during recessions,
  - countercyclical variation in risk premiums, return volatility, and Sharpe ratios, including
    - their rising pattern during recessions,
    - their falling pattern during expansions.

These effects naturally induce the leverage effect, the mean reversion of excess returns, as well as the predictability of consumption volatility from price–dividend ratio;

- real bond prices:
  - average real yield curve,
  - level, variability, and persistence of real yields,
  - mean bond risk premiums.

Additionally, our preliminary results for option prices suggest that our model could be useful for understanding the behavior of implied volatility of S&P 500 index options observed in the data.

The recent study of David and Veronesi (2013) uses expected-utility preferences and shows that learning is important for understanding the comovement of stocks and nominal bonds. We show that additionally relaxing the independence axiom by employing the recursive utility is key for understanding the comovement of consumption and asset prices in general. Such comovement depends crucially on the fact that the economic uncertainty in our semi-Markov model is declining over time after the shift to the high-growth state, but rising over time after the shift to the low-growth state.

To sum up, modeling rare consumption disasters in terms of protracted but mild recessions rather than deep short declines in the realized consumption growth rate helps to understand consumption and asset prices through the lens of rare disasters. The fact that the probability of the rare event is endogenous, being a result of Bayesian updating rather than an exogenously specified stochastic process, makes our results more credible.

Supplementary Material

Supplementary appendices A, B, C, D, E, and F are available at *Review of Finance* online.

References

Ang, A., Bekaert, G., and Wei, M. (2008) The term structure of real rates and expected inflation, *Journal of Finance* **63**, 797–849.

Backus, D., Chernov, M., and Martin, I. (2011) Disasters implied by equity index options, *Journal of Finance* **66**, 1969–2012.

Baker, S. and Bloom, N. (2012) Does uncertainty drive business cycles? Using disasters as a natural experiment. Working Paper, Stanford University.

Bansal, R., Gallant, A. R., and Tauchen, G. (2007) Rational pessimism, rational exuberance, and asset pricing models, *Review of Economic Studies* **74**, 1005–1033.

Bansal, R., Kiku, D., Shaliastovich, I., and Yaron, A. (2012) Volatility, macroeconomy and asset prices. NBER Working Paper No. 18104.

Bansal, R., Kiku, D., and Yaron, A. (2010) Long-run risks, the macroeconomy, and asset prices, *American Economic Review Papers and Proceedings* **100**, 542–546.

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

Barro, R. J. (2006) Rare disasters and asset markets in the twentieth century, *Quarterly Journal of Economics* **121**, 823–866.

Barro, R. J. and Jin, T. (2011, 09) On the size distribution of macroeconomic disasters, *Econometrica* **79**, 1567–1589.

Barro, R. J. and Ursua, J. F. (2012) Rare macroeconomic disasters, *Annual Review of Economics* **4**, 83–109.

Bates, D. S. (2000) Post-'87 crash fears in the S&P 500 futures option market, *Journal of Econometrics* **94**, 181–238.

Beeler, J. and Campbell, J. Y. (2012) The long-run risks model and aggregate asset prices: an empirical assessment, *Critical Finance Review* **1**, 141–182.

Bergoeing, R., Kehoe, T. J., and Soto, R. (2002) A decade lost and found: Mexico and Chile in the 1980s, *Review of Economic Dynamics* **5**, 166–205.

Bloom, N. (2009) The impact of uncertainty shocks, *Econometrica* **77**, 623–685.

Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., and Terry, S. J. (2012) Really uncertain business cycles. NBER Working paper No. 18245.

Branger, N., Rodrigues, P., and Schlag, C. (2012) The role of volatility shocks and rare events in long-run risk models. Manuscript, Goethe University Frankfurt.

Brown, S. J., Goetzmann, W. N., and Ross, S. A. (1995) Survival, *Journal of Finance* **50**, 853–873.

Campbell, J. Y. (1986) Bond and stock returns in a simple exchange model, *Quarterly Journal of Economics* **101**, 185–804.

Campbell, J. Y., Shiller, R. J., and Viceira, L. M. (2009) Understanding inflation-indexed bond markets, *Brookings Papers on Economic Activity* 79–120.

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

Cecchetti, S. G., Lam, P.-S., and Mark, N.C. (2000) Asset pricing with distorted beliefs: are equity returns too good to be true?, *American Economic Review* **90**, 787–805.

Cogley, T. and Sargent, T. J. (2008) The market price of risk and the equity premium: a legacy of the great depression?, *Journal of Monetary Economics* **55**, 454–476.

Collin-Dufresne, P., Johannes, M., and Lochstoer, L. A. (2013) Parameter learning in general equilibrium: the asset pricing implications. Working paper, Columbia Business School.

David, A. (1997) Fluctuating confidence in stock markets: Implications for returns and volatilities, *Journal of Financial and Quantitative Analysis* **32**, 427–462.

David, A. and Veronesi, P. (2013) What ties return volatilities to price valuations and fundamentals?, *Journal of Political Economy* **121**, 682–746.

Diebold, F. X., Lee, J.-H., and Weinbach, G. C. (1994) *Nonstationary Time-Series Analysis and Cointegration (Advanced Texts in Econometrics)*. Oxford University Press, Oxford.

Duffie, D. and Epstein, L. G. (1992a) Asset pricing with stochastic differential utility, *Review of Financial Studies* **5**, 411–436.

Duffie, D. and Epstein, L. G. (1992b) Stochastic differential utility, *Econometrica* **60**, 353–394.

Duffie, D. and Lions, P.-L. (1992) PDE solutions of stochastic differential utility, *Journal of Mathematical Economics* **21**, 577–606.

Epstein, L. G. and Zin, S. E. (1989) Substitution, risk-aversion and the temporal behavior of consumption and asset returns: a theoretical framework, *Econometrica* **57**, 937–969.

Epstein, L. G. and Zin, S. E. (1991) Substitution, risk aversion, and the temporal behavior of consumption and asset returns: an empirical analysis, *Journal of Political Economy* **99**, 263–286.

Evans, M. D. (1996) Peso problems: their theoretical and empirical implications, in: G. S. Maddala and C. R. Rao (eds), *Handbook of Statistics*, Vol. 14, Elsevier Science B.V., New York, pp. 613–646.

Farhi, E. and Gabaix, X. (2011) Rare disasters and exchange rates. Working paper, Harvard University.

Friedman, A. (2006a) *Stochastic Differential Equations and Applications*, Vol. I, Dover, Mineola, New York.

Friedman, A. (2006b) *Stochastic Differential Equations and Applications*, Vol. II, Dover, Mineola, New York.

Gabaix, X. (2008) Variable rare disasters: a tractable theory of ten puzzles in macro-finance. *American Economic Review: Papers and Proceedings of the One Hundred Twentieth Annual Meeting of the American Economic Association* **98**, 64–67.

Gabaix, X. (2012) Variable rare disasters: an exactly solved framework for ten puzzles in macro-finance, *The Quarterly Journal of Economics* **127**, 645–700.

Gourio, F. (2008) Time-series predictability in the disaster model, *Finance Research Letters* **5**, 191–203.

Gourio, F. (2012) Disaster risk and business cycles, *The American Economic Review* **102**, 2734–2766.

Gourio, F. (2013) Credit risk and disaster risk, *AEJ: Macroeconomics* **5**, 1–34.

Gourio, F., Siemer, M., and Verdelhan, A. (2013) International risk cycles, *Journal of International Economics* **89**, 471–484.

Hamilton, J. (1989) A new approach to the economic analysis of nonstationary time series and the business cycle, *Econometrica* **57**, 357–384.

Hayashi, F. and Prescott, E. C. (2002) The 1990s in Japan: a lost decade, *Review of Economic Dynamics* **5**, 206–235.

Howard, R. A. (1971) *Dynamic Probabilistic Systems, Vol. II: Semi-Markov and Decision Processes*, John Wiley and Sons, Inc., New York.

Jin, T. (2014) Rare events and long-run risks for asset prices: empirical identification and evaluation, Manuscript, Harvard University.

Johannes, M., Lochstoer, L., and Mou, Y. (2012) Learning about consumption dynamics, Manuscript, Columbia Business School.

Ju, N. and Miao, J. (2012) Ambiguity, learning and asset returns, *Econometrica* **80**, 559–592.

Kydland, F. E. and Zarazaga, C. E. J. M. (2002) Argentina's lost decade, *Review of Economic Dynamics* **5**, 152–165.

Liptser, M. and Shiryaev, A. (1977) *Statistics of Random Processes: General Theory*, Springer, New York.

Lu, Y. K. and Siemer, M. (2013) Learning, rare disasters, and asset prices. Working Paper, Federal Reserve Board No. 2013-85.

Lucas, R. E. (1978) Asset prices in an exchange economy, *Econometrica* **46**, 1429–1445.

Lustig, H. N. and Verdelhan, A. (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 from the Economic Crises', Sponsored by the Swiss National Bank* **59**.

Ma, T. and Yu, Y. Q. (1989) The Keldyš-Fichera boundary value problem for degenerate quasilinear elliptic equations of second order, *Differential and Integral Equations* **2**, 379–388.

Martin, I. (2013) Consumption-based asset pricing with higher cumulants, *Review of Economic Studies* **80**, 745–773.

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

Merton, R. C. (1971) Optimum consumption and portfolio rules in a continuous-time model, *Journal of Economic Theory* **3**, 373–413.

Merton, R. C. (1973) An intertemporal capital asset pricing model, *Econometrica* **41**, 867–887.

Murphy, K. P. (2012) *Machine Learning: A Probabilistic Perspective*, MIT Press, Cambridge, Massachusetts.

Orlik, A. and Veldkamp, L. (2013) Understanding uncertainty shocks and the role of the black swan, Working paper, New York University, Stern School of Business.

Pakoš, M. (2013) Long-run risk and hidden growth persistence. *Journal of Economic Dynamics and Control* **37**, 1911–1928.

Piazzesi, M. and Schneider, M. (2007) *Equilibrium Yield Curves*, NBER Macroeconomics Annual 2006, MIT Press, 389–442.

Rietz, T. A. (1988) The equity risk premium: a solution, *Journal of Monetary Economics* **22**, 117–131.

Santa-Clara, P. and Yan, S. (2010) Crashes, volatility, and the equity premium: lessons from S&P 500 options, *Review of Economics and Statistics* **92**, 435–451.

Schreindorfer, D. (2014) Tails, fears and equilibrium option prices. Job Market Paper, Tepper School of Business, Carnegie Mellon University, pp. 1–57.

Schroder, M. and Skiadas, C. (1999) Optimal consumption and portfolio selection with stochastic differential utility, *Journal of Economic Theory* **89**, 68–126.

Seo, S. B. and Wachter, J. A. (2013) Option prices in a model with stochastic disaster risk. Working paper, University of Pennsylvania (Wharton).

Timmermann, A. G. (1993) How learning in financial markets generates excess volatility and predictability of stock returns, *Quarterly Journal of Economics* **108**, 1135–1145.

Tsai, J. and Wachter, J. A. (2013) Rare booms and disasters in a multi-sector endowment economy. Working paper, University of Pennsylvania (Wharton).

Veronesi, P. (1999) Stock market overreaction to bad news in good times: a rational expectations equilibrium model, *Review of Financial Studies* **12**, 975–1007.

Veronesi, P. (2000) How does information quality affect stock returns, *Journal of Finance* **55**, 807–837.

Veronesi, P. (2004) The Peso problem hypothesis and stock market returns, *Journal of Economic Dynamics and Control* **28**, 707–725.

Wachter, J. A. (2013) Can time-varying risk of rare disasters explain aggregate stock market volatility?, *Journal of Finance* **68**, 987–1035.

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

Weitzman, M. (2013) A precautionary tale of uncertain tail fattening, *Environmental and Resource Economics* **55**, 159–173.

Weitzman, M. L. (2007) Subjective expectations and asset-return puzzles, *American Economic Review* **97**, 1102–1130.

Wonham, W. M. (1964) Some applications of stochastic differential equations to optimal nonlinear filtering, *SIAM Journal on Control* **2**, 347–369.