ECONOMETRICA

AMBIGUITY, LEARNING, AND ASSET RETURNS

NENGJIU JU
*Hong Kong University of Science and Technology, Clear Water Bay, Kowloon,
Hong Kong and SAIF, Shanghai Jiaotong University, China*

JIANJUN MIAO
*Boston University, Boston, MA 02215, U.S.A. and CEMA, Central University of
Finance and Economics and AFR, Zhejiang University, China*

AMBIGUITY, LEARNING, AND ASSET RETURNS

By Nengjiu Ju and Jianjun Miao[1]

We propose a novel generalized recursive smooth ambiguity model which permits a three-way separation among risk aversion, ambiguity aversion, and intertemporal substitution. We apply this utility model to a consumption-based asset-pricing model in which consumption and dividends follow hidden Markov regime-switching processes. Our calibrated model can match the mean equity premium, the mean risk-free rate, and the volatility of the equity premium observed in the data. In addition, our model can generate a variety of dynamic asset-pricing phenomena, including the procyclical variation of price–dividend ratios, the countercyclical variation of equity premia and equity volatility, the leverage effect, and the mean reversion of excess returns. The key intuition is that an ambiguity-averse agent behaves pessimistically by attaching more weight to the pricing kernel in bad times when his continuation values are low.

Keywords: Ambiguity aversion, learning, asset-pricing puzzles, model uncertainty, robustness, pessimism, regime switching.

1. INTRODUCTION

Under the rational expectations hypothesis, there exists an objective probability law governing the state process, and economic agents know this law which coincides with their subjective beliefs. This rational expectations hypothesis has become the workhorse in macroeconomics and finance. However, it faces serious difficulties when confronted with asset markets data (see Section 4.1 for a detailed discussion on the empirical facts). Most prominently, Mehra and Prescott (1985) showed that for a standard rational, representative–agent model to explain the high equity premium observed in the data, an implausible high degree of risk aversion is needed, resulting in the equity premium puzzle. Weil (1989) showed that this high degree of risk aversion generates an implausibly high risk-free rate, resulting in the risk-free rate puzzle. Shiller (1981) found that equity volatility is too high to be justified by

changes in the fundamental. In addition, a number of empirical studies document puzzling links between aggregate asset markets and macroeconomics: Price–dividend ratios move procyclically (Fama and French (1989)) and conditional expected equity premiums move countercyclically (Campbell and Shiller (1988a, 1988b) and Fama and French (1989)). Excess returns are serially correlated and mean reverting (Fama and French (1988b) and Poterba and Summers (1988)). Excess returns are forecastable; in particular, the log dividend yield predicts long-horizon realized excess returns (Campbell and Shiller (1988a, 1988b) and Fama and French (1988a)). Conditional volatility of stock returns is persistent and moves countercyclically (Bollerslev, Chou, and Kroner (1992)).

In this paper, we develop a representative–agent consumption-based asset-pricing model that helps explain the preceding puzzles simultaneously by departing from the rational expectations hypothesis. Our model has two main ingredients. First, we assume that consumption and dividends follow a hidden Markov regime-switching model. The agent learns about the hidden state based on past data. The posterior state beliefs capture fluctuating economic uncertainty and drive asset return dynamics. Second, and more importantly, we assume that the agent is ambiguous about the hidden state and his preferences are represented by a generalized recursive smooth ambiguity model that allows for a three-way separation among risk aversion, ambiguity aversion, and intertemporal substitution. Our recursive model nests some popular utility models with hidden states as special cases, including Epstein–Zin preferences (Epstein and Zin (1989)), smooth ambiguity preferences (Klibanoff, Marinacci, and Mukerji (2009)), multiplier preferences (Hansen (2007), Hansen and Sargent (2007, 2010)), and risk-sensitive preferences (Hansen and Sargent (2001), Tallarini (2000)). Ambiguity aversion is manifested through a pessimistic distortion of the pricing kernel in the sense that the agent attaches more weight on low continuation values in recessions. It is this pessimistic behavior that allows our model to explain the asset-pricing puzzles.

We motivate our adoption of the recursive ambiguity model in two ways. First, the Ellsberg paradox (Ellsberg (1961)) and related experimental evidence demonstrate that the distinction between risk and ambiguity is behaviorally meaningful. Roughly speaking, risk refers to the situation where there is a probability measure to guide choice, while ambiguity refers to the situation where the decision maker is uncertain about this probability measure due to cognitive or informational constraints. Knight (1921) and Keynes (1936) emphasized that ambiguity may be important for economic decision-making. We assume that the agent in our model is ambiguous about hidden states in consumption and dividend growth. Our adopted utility model captures ambiguity-sensitive behavior. Our second motivation is related to the robustness theory developed by Hansen (2007) and Hansen and Sargent (2001, 2007, 2008). Specifically, the agent in our model may fear model misspecification. He is concerned about model uncertainty and, thus, seeks robust decision-making. We may interpret our model of ambiguity as a model of robustness in the presence of model uncertainty.

Our modeling of learning echoes with Hansen's (2007) suggestion that one should put econometricians and economic agents on comparable footings in terms of statistical knowledge. When estimating the regime-switching consumption process, econometricians typically apply Hamilton's (1989) maximum likelihood method and assume that they do not observe the hidden state. However, the rational expectations hypothesis often assumes economic agents are endowed with more precise information than econometricians. A typical assumption is that agents know all parameter values underlying the consumption process (e.g., Cecchetti, Lam, and Mark (1990, 2000)). In this paper, we show that when agents are concerned about statistical ambiguity, there are important quantitative implications from removing the information gap between them and econometricians, while standard Bayesian learning has small quantitative effects.[2]

Learning is naturally embedded in the recursive smooth ambiguity model. In this model, the posterior of the hidden state and the conditional distribution of the consumption process given a state cannot be reduced to a compound predictive distribution, unlike in the standard Bayesian analysis. It is this irreducibility that captures sensitivity to ambiguity or model uncertainty. An important advantage of the smooth ambiguity model over other models of ambiguity such as the maxmin expected utility (or multiple-priors) model of Gilboa and Schmeidler (1989) is that it achieves a separation between ambiguity (beliefs) and ambiguity attitude (tastes). This feature allows us to do comparative statics with respect to the ambiguity aversion parameter holding ambiguity fixed, and to calibrate it for quantitative analysis. Another advantage is that under standard regularity conditions, we can apply the usual differential analysis to the smooth ambiguity model. Deriving its pricing kernel becomes tractable. By contrast, the widely applied maxmin expected utility model lacks this smoothness property.

Our paper is related to a growing body of literature that studies the implications of ambiguity and robustness for finance and macroeconomics.[3] We contribute to this literature by (i) proposing a novel generalized recursive smooth

---

[2]There is a large literature on learning in asset pricing using the standard Bayesian framework. Notable works include Brandt, Zeng, and Zhang (2004), Brennan and Xia (2001), David (1997), Detemple (1986), Dothan and Feldman (1986), Veronesi (1999, 2000), and Weitzman (2007), among others.

[3]See Cao, Wang, and Zhang (2005), Chen and Epstein (2002), Epstein and Miao (2003), Epstein and Wang (1994), Garlappi, Uppal, and Wang (2007), Miao (2009), and Routledge and Zin (2009) for asset-pricing applications of the multiple-priors utility model. See Anderson, Hansen, and Sargent (2003), Cagetti, Hansen, Sargent, and Williams (2002), Hansen and Sargent (2001), Hansen, Sargent, and Tallarini (1999), Liu, Pan, and Wang (2005), Maenhout (2004), and Uppal and Wang (2003) for models of robustness and applications. Maccheroni, Marinacci, and Rustichini (2006) provided an axiomatic foundation for one of Hansen and Sargent's robustness formulations—the multiplier preferences. See Backus, Routledge, and Zin (2004) and Hansen and Sargent (2008) for a survey.

ambiguity model and tractable homothetic specifications, and (ii) putting this utility model in quantitative work to address a variety of asset-pricing puzzles.

We now discuss closely related papers. In the maxmin framework, Epstein and Schneider (2007) modeled learning under ambiguity using a set of priors and a set of likelihoods. Both sets are updated by Bayes' rule in a suitable way. Applying this learning model, Epstein and Schneider (2008) analyzed asset-pricing implications. Leippold, Trojani, and Vanini (2008) embedded this model in the continuous-time framework of Chen and Epstein (2002). In contrast to our paper, there is no distinction between risk aversion and intertemporal substitution, and no separation between ambiguity and ambiguity attitudes in the preceding three papers. Hansen and Sargent (2007) formulated a learning model that allows for two forms of model misspecification: (i) misspecification in the underlying Markov law for the hidden states and (ii) misspecification of the probabilities assigned to the hidden states. Hansen and Sargent (2010) applied this learning model to study time-varying model uncertainty premia. Hansen (2007) surveyed models of learning and robustness. He analyzed a continuous-time model similar to our log-exponential specification, but he did not consider the general homothetic form and did not conduct a thorough quantitative analysis as in our paper. Our paper is also related to Abel (2002), Brandt, Zeng, and Zhang (2004), and Cecchetti, Lam, and Mark (2000), who modeled the agent's pessimism and doubt in specific ways and showed that their modeling helps to explain many asset-pricing puzzles. Our adopted smooth ambiguity model captures pessimism and doubt with a decision theoretic foundation.

The remainder of the paper proceeds as follows. Section 2 presents our generalized recursive smooth ambiguity model. Section 3 analyzes its asset pricing implications in a Lucas-style model. Section 4 calibrates the model and studies its quantitative implications. Section 5 concludes. The Appendix contains proofs.

2. GENERALIZED RECURSIVE SMOOTH AMBIGUITY PREFERENCES

In this section, we introduce the generalized recursive smooth ambiguity model adopted in our paper. In a static setting, this utility model delivers essentially the same functional form that has appeared in some other papers, for example, Ergin and Gul (2009), Klibanoff, Marinacci, and Mukerji (2005), Nau (2006), and Seo (2009). These papers provide various different axiomatic foundations and interpretations. Our adopted dynamic model was axiomatized by Hayashi and Miao (2011) and is closely related to Klibanoff, Marinacci, and Mukerji (2009). Here we focus on the utility representation and refer the reader to the preceding papers for axiomatic foundations.

We start with a static setting in which a decision maker's ambiguity preferences over consumption are represented by the utility function

$$(1) \qquad v^{-1}\left(\int_\Pi v\left(u^{-1}\left(\int_S u(C)\,d\pi\right)\right)d\mu(\pi)\right), \quad \forall C\!:\!S \to \mathbb{R}_+,$$

where u and v are increasing functions and μ is a subjective prior over the set Π of probability measures on S that the decision maker thinks possible. We have defined utility in (1) in terms of two certainty equivalents. When we define $\phi \equiv v \circ u^{-1}$, it is ordinally equivalent to the smooth ambiguity model of Klibanoff, Marinacci, and Mukerji (2005):

$$(2) \qquad \int_\Pi \phi\left(\int_S u(C)\,d\pi\right)d\mu(\pi) \equiv \mathbb{E}_\mu \phi(\mathbb{E}_\pi u(C)).$$

A key feature of this model is that it achieves a separation between ambiguity, identified as a characteristic of the decision maker's subjective beliefs, and ambiguity attitude, identified as a characteristic of the decision maker's tastes.[4] Specifically, ambiguity is characterized by properties of the subjective set of measures Π. Attitudes toward ambiguity are characterized by the shape of ϕ or v, while attitudes toward pure risk are characterized by the shape of u. In particular, the decision maker displays risk aversion if and only if u is concave, while he displays ambiguity aversion if and only if ϕ is concave or, equivalently, if and only if v is a concave transformation of u. Intuitively, an ambiguity-averse decision maker prefers consumption that is more robust to the possible variation in probabilities. That is, he is averse to mean-preserving spreads in the distribution μ_C induced by the prior μ and the consumption act C. This distribution represents the uncertainty about the ex ante utility evaluation of C, $\mathbb{E}_\pi u(C)$ for all $\pi \in \Pi$. Note that there is no reduction between μ and π, in general: It is possible when ϕ is linear. In this case, the decision maker is ambiguity neutral and the smooth ambiguity model reduces to the standard expected utility model. It is the *irreducibility of compound distributions* that allows for ambiguity-sensitive behavior, as pointed out by Hansen (2007), Klibanoff, Marinacci, and Mukerji (2005), Segal (1987), and Seo (2009). This modeling is supported by the experimental evidence reported by Halevy (2007).

Klibanoff, Marinacci, and Mukerji (2005) showed that the multiple-priors model of Gilboa and Schmeidler (1989), $\min_{\pi \in \Pi} \mathbb{E}_\pi u(C)$, is a limiting case of the smooth ambiguity model with infinite ambiguity aversion. An important advantage of the smooth ambiguity model over other models of ambiguity, such

---

[4]The behavioral foundation of ambiguity and ambiguity attitude is based on the theory developed by Ghirardato and Marinacci (2002) and Klibanoff, Marinacci, and Mukerji (2005). Epstein (1999) provided a different foundation. The main difference is that the benchmark ambiguity neutral preference is the expected utility preference according to Ghirardato and Marinacci (2002), while Epstein's (1999) benchmark is the probabilistic sophisticated preferences.

as the multiple-priors utility model, is that it is tractable and admits a clear-cut comparative statics analysis. Tractability is revealed by the fact that the well developed machinery for dealing with risk attitudes can be applied to ambiguity attitudes. In addition, the indifference curve implied by (2) is smooth under regularity conditions, rather than kinked as in the case of the multiple-priors utility model. More importantly, comparative statics of ambiguity attitudes can be easily analyzed using the function ϕ or v only, holding ambiguity fixed. Such a comparative statics analysis is not evident for the multiple-priors utility model since the set of priors Π in that model may characterize ambiguity as well as ambiguity attitudes.

We may alternatively interpret the utility model defined in (1) as a model of robustness in which the decision maker is concerned about model misspecification and thus seeks robust decision making. Specifically, each distribution $\pi \in \Pi$ describes an economic model. The decision maker is ambiguous about which is the right model specification. He has a subjective prior μ over alternative models. He is averse to model uncertainty and, thus, evaluates different models using a concave function ϕ.

We now embed the static model (1) in a dynamic setting. Time is denoted by $t = 0, 1, 2, \dots$. The state space in each period is denoted by S. At time t, the decision maker's information consists of history $s^t = \{s_0, s_1, s_2, \dots, s_t\}$ with $s_0 \in S$ given and $s_t \in S$. The decision maker ranks adapted consumption plans $C = (C_t)_{t \geq 0}$. That is, C_t is a measurable function of s^t. The decision maker is ambiguous about the probability distribution on the full state space S^∞. This uncertainty is described by an unobservable random state z in the space Z. The hidden state z can be interpreted in several different ways. It could be an unknown model parameter, a discrete indicator of alternative models, or a hidden Markov state that evolves over time in a regime-switching process (Hamilton (1989)).

The decision maker has a prior μ_0 over the hidden state z. Each value of z gives a probability distribution π_z over the full state space. The posterior μ_t and the conditional likelihood $\pi_{z,t}$ can be obtained by Bayes' rule. Inspired by Kreps and Porteus (1978) and Epstein and Zin (1989), we propose the generalized recursive ambiguity utility model

$$(3) \qquad V_t(C) = W\big(C_t, \mathcal{R}_t(V_{t+1}(C))\big),$$

$$\mathcal{R}_t(\xi) = v^{-1}\big(\mathbb{E}_{\mu_t}\{v \circ u^{-1}\mathbb{E}_{\pi_{z,t}}[u(\xi)]\}\big),$$

where $V_t(C)$ is the continuation value at date t, $W : \mathbb{R}^2 \to \mathbb{R}$ is a time aggregator, $\mathcal{R}_t$ is an uncertainty aggregator that maps an s^{t+1}-measurable random variable ξ to an s^t-measurable random variable, and u and v admit the same interpretation as in the static setting. When $v \circ u^{-1}$ is linear, (3) reduces to the recursive utility model of Epstein and Zin (1989) with hidden states. In this model, the posterior μ_t and the likelihood $\pi_{z,t}$ can be reduced to a *predictive distribution*, which is the key idea underlying the Bayesian analysis. When

$v \circ u^{-1}$ is nonlinear, μ_t and $\pi_{z,t}$ cannot be reduced to a single distribution for decision making in (3), leading to ambiguity-sensitive behavior as in the static model (1).

The irreducibility of compound distributions also allows for sensitivity to temporal resolution of uncertainty (Kreps and Porteus (1978), Hansen (2007), and Strzalecki (2009)). Strzalecki (2009) showed that, given independent and identically distributed (IID) ambiguity and the standard discounted aggregator $W(c, y) = h(c) + \beta y$, where h is some von Neumann–Morgenstem index, the recursive multiple-priors model with $\mathcal{R}_t(\xi) = \min_{\pi \in \Pi} \mathbb{E}_\pi[\xi]$ is the only one among a family of dynamic ambiguity preferences that satisfies the indifference to the timing of the resolution of uncertainty. Thus, to allow for non-indifference, one has to consider a time aggregator W other than the discounted aggregator (e.g., Epstein and Zin (1989)) or an uncertainty aggregator $\mathcal{R}_t$ other than the Gilboa–Schmeidler form.

Our generalized recursive smooth ambiguity model in (3) permits a three-way separation among risk aversion, ambiguity aversion, and intertemporal substitution. However, attitudes toward timing are not independently separated from intertemporal substitution, risk aversion, or ambiguity aversion. In application, it proves tractable to consider the homothetic specification

$$(4) \qquad W(c, y) = [(1 - \beta)c^{1-\rho} + \beta y^{1-\rho}]^{1/(1-\rho)}, \quad \rho > 0, \neq 1,$$

and u and v are given by

$$(5) \qquad u(x) = \frac{x^{1-\gamma}}{1 - \gamma}, \quad \gamma > 0, \neq 1,$$

$$(6) \qquad v(x) = \frac{x^{1-\eta}}{1 - \eta}, \quad \eta > 0, \neq 1,$$

where $\beta \in (0, 1)$ is the subjective discount factor, $1/\rho$ represents the elasticity of intertemporal substitution (EIS), γ is the risk aversion parameter, and η is the ambiguity aversion parameter. We then have

$$(7) \qquad V_t(C) = \left[(1 - \beta)C_t^{1-\rho} + \beta\{\mathcal{R}_t(V_{t+1}(C))\}^{1-\rho}\right]^{1/(1-\rho)},$$

$$(8) \qquad \mathcal{R}_t(V_{t+1}(C)) = \left\{\mathbb{E}_{\mu_t}\left(\mathbb{E}_{\pi_{z,t}}[V_{t+1}^{1-\gamma}(C)]\right)^{(1-\eta)/(1-\gamma)}\right\}^{1/(1-\eta)}.$$

If $\eta = \gamma$, the decision maker is ambiguity neutral and (7) reduces to the recursive utility model of Epstein and Zin (1989) and Weil (1989). The decision maker displays ambiguity aversion if and only if $\eta > \gamma$. By the property of certainty equivalent, a more ambiguity-averse agent with a higher value of η has a lower utility level.

In the limiting case with $\rho = 1$, (7) reduces to

$$(9) \qquad U_t = (1 - \beta) \ln C_t$$
$$+ \frac{\beta}{1 - \eta} \ln \left\{ \mathbb{E}_{\mu_t} \exp \left(\frac{1 - \eta}{1 - \gamma} \ln \left(\mathbb{E}_{\pi_{z,t}} \exp((1 - \gamma) U_{t+1}) \right) \right) \right\},$$

where $U_t = \ln V_t$. This specification reduces to the multiplier model with hidden states studied by Hansen (2007) and Hansen and Sargent (2010). In particular, there are two risk-sensitivity adjustments in (9). The first risk-sensitivity adjustment for the distribution $\pi_{z,t}$ reflects the agent's concerns about the misspecification in the underlying Markov law given a hidden state z. The second risk-sensitivity adjustment for the distribution μ_t reflects the agent's concerns about the misspecification of the probabilities assigned to the hidden states. More generally, our model in (3) nests a version of the recursive multiplier model with hidden states in Hansen and Sargent (2007) as a special case when we set $W(c, y) = h(c) + \beta y$, $u(x) = -\exp(-x/\theta_1)$, and $v(x) = -\exp(-x/\theta_2)$ for $\theta_1, \theta_2 > 0$.[5]

If we further take the limit in (9) when $\gamma \to 1$, equation (9) becomes

$$(10) \qquad U_t = (1 - \beta) \ln C_t + \frac{\beta}{1 - \eta} \ln \left\{ \mathbb{E}_{\mu_t} \exp \left((1 - \eta) \mathbb{E}_{\pi_{z,t}} [U_{t+1}] \right) \right\}.$$

This is the log-exponential specification studied by Ju and Miao (2007). In this case, there is only one risk-sensitive adjustment for the state beliefs μ_t. Following Klibanoff, Marinacci, and Mukerji (2005), we can show that when $\eta \to \infty$, (10) becomes

$$(11) \qquad U_t = (1 - \beta) \ln C_t + \beta \min_z \mathbb{E}_{\pi_{z,t}} [U_{t+1}].$$

This utility function belongs to the class of the recursive multiple-priors model of Epstein and Wang (1994) and Epstein and Schneider (2003, 2007). The agent is extremely ambiguity averse by choosing the worst continuation value each period.

Our model in (3) nests the following Klibanoff, Marinacci, and Mukerji (2009) model as a special case:

$$(12) \qquad V_t(C) = u(C_t) + \beta \phi^{-1} \left(\mathbb{E}_{\mu_t} \phi \left(\mathbb{E}_{\pi_{z,t}} [V_{t+1}(C)] \right) \right).$$

In this model, risk aversion and intertemporal substitution are confounded. In addition, Ju and Miao (2007) found that when u is defined as in (5) and $\phi(x) =$

---

[5]This model intersects the Kreps–Porteus–Epstein–Zin recursive expected utility model only when $\theta_1 = \theta_2$. By contrast, the multiplier preference model with full information (e.g., Anderson, Hansen, and Sargent (2003), Barillas, Hansen, and Sargent (2009), and Hansen and Sargent (2001)) is a special case of the Kreps–Porteus–Epstein–Zin model.

$x^{1-\alpha}/(1-\alpha)$ for $x > 0$ and $1 \neq \alpha > 0$, the model (12) is not well defined for $\gamma > 1$. Thus, they considered (7) and (8) with $\gamma = \rho$ and $\alpha \equiv 1 - (1-\eta)/(1-\gamma)$, which is ordinally equivalent to (12) when $\gamma \in (0,1)$. The utility function in (12) is always well defined for the specification $\phi(x) = -\exp(-x/\theta)$, $\theta > 0$. The nice feature of this specification is that it has a connection with risk-sensitive control and robustness, as studied by Hansen (2007) and Hansen and Sargent (2007). A disadvantage of this specification is that the utility function generally does not have a homogeneity property. Thus, the curse of dimensionality makes numerical analysis of the decision maker's dynamic programming problem complicated, except for the special case where $u(c) = \ln(c)$ as in (10) (see Ju and Miao (2007) and Collard et al. (2009)). As a result, we focus on the homothetic specification (7) and (8) in our analysis below.

3. ASSET PRICING IMPLICATIONS

3.1. *The Economy*

We consider a representative–agent pure-exchange economy. There is only one consumption good with aggregate consumption given by C_t in period t. The agent trades multiple assets. Among these assets, we focus on the risk-free bond with zero net supply and equity that pays aggregate dividends D_t in period $t = 0, 1, 2, \ldots$. Let $R_{f,t+1}$ and $R_{e,t+1}$ denote their gross returns from period t to period $t+1$, respectively. We specify aggregate consumption by a regime-switching process as in Cecchetti, Lam, and Mark (1990, 1993, 2000) and Kandel and Stambaugh (1991):

$$(13) \qquad \ln\left(\frac{C_{t+1}}{C_t}\right) = \kappa_{z_{t+1}} + \sigma \varepsilon_{t+1},$$

where ε_t is an IID standard normal random variable, and z_{t+1} follows a Markov chain which takes values 1 or 2 with transition matrix (λ_{ij}), where $\sum_j \lambda_{ij} = 1$, $i, j = 1, 2$. We may identify state 1 as the boom state and state 2 as the recession state in that $\kappa_1 > \kappa_2$.

In a standard Lucas-style model (Lucas (1978)), dividends and consumption are identical in equilibrium. This assumption is clearly violated in reality. There are several ways to model dividends and consumption separately in the literature (e.g., Cecchetti, Lam, and Mark (1993)). Here, we follow Bansal and Yaron (2004) and assume

$$(14) \qquad \ln\left(\frac{D_{t+1}}{D_t}\right) = \zeta \ln\left(\frac{C_{t+1}}{C_t}\right) + g_d + \sigma_d e_{t+1},$$

where e_{t+1} is an IID standard normal random variable and is independent of all other random variables. The parameter $\zeta > 0$ can be interpreted as the

leverage ratio on expected consumption growth as in Abel (1999). This parameter and the parameter σ_d allow us to calibrate volatility of dividends (which is significantly larger than consumption volatility) and their correlation with consumption. The parameter g_d helps match the expected growth rate of dividends. Our modeling of the dividend process is convenient because it does not introduce any new state variable in our model.

The model of consumption and dividends in (13) and (14) is a nonlinear counterpart of the long-run risk processes discussed in Campbell (1999) and Bansal and Yaron (2004) in that both consumption and dividends contain a common persistent component of Markov chain. Garcia, Meddahi, and Tedongap (2008) showed that the processes in (13) and (14) can be obtained by discretizing the long-run risks model Case I in Bansal and Yaron (2004). Unlike Case II in Bansal and Yaron (2004), we assume that volatility σ is constant and independent of regimes. In the Bansal–Yaron model, fluctuating volatility of consumption growth is needed to generate time-varying expected equity premia. Our assumption of constant σ intends to generate this feature through endogenous learning rather than exogenous fluctuations in consumption volatility.

Unlike the long-run risks model of Bansal and Yaron (2004), the regime-switching model can be easily estimated by the maximum likelihood method. Following Veronesi (2000) and Hansen (2007), we put economic agents and econometricians on equal footings by assuming that economic regimes are not observable. What is observable in period t is the history of consumption and dividends $s^t = \{C_0, D_0, C_1, D_1, \ldots, C_t, D_t\}$. In addition, the representative agent knows the parameters of the model (e.g., ζ, g_d, σ, and σ_d). He has ambiguous beliefs about the hidden states. His preferences are represented by generalized recursive smooth ambiguity utility defined in (7). To apply this utility model, we need to derive the evolution of the posterior state beliefs. Let $\mu_t = \Pr(z_{t+1} = 1|s^t)$.[6] The prior belief μ_0 is given. By Bayes' rule, we can derive

$$(15) \qquad \mu_{t+1} = \frac{\lambda_{11} f(\ln(C_{t+1}/C_t), 1)\mu_t + \lambda_{21} f(\ln(C_{t+1}/C_t), 2)(1 - \mu_t)}{f(\ln(C_{t+1}/C_t), 1)\mu_t + f(\ln(C_{t+1}/C_t), 2)(1 - \mu_t)},$$

where $f(y, i) = \frac{1}{\sqrt{2\pi}\sigma} \exp[-(y - \kappa_i)^2/(2\sigma^2)]$ is the density function of the normal distribution with mean κ_i and variance σ^2. By our modeling of dividends in (14), dividends do not provide any new information for belief updating or for the estimation of the hidden states.

---

[6]We abuse notation here since we have used μ_t to denote the posterior distribution over the parameter space in Section 2.

3.2. *Asset Pricing*

As is standard in the literature, we derive the pricing kernel or the stochastic discount factor to understand asset prices. Following Duffie and Skiadas (1994) or Hansen, Heaton, and Li (2008), we use the homogeneity property of generalized recursive smooth ambiguity utility (7) to show that its pricing kernel is given by

$$(16) \qquad M_{z_{t+1},t+1} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\rho} \left(\frac{V_{t+1}}{\mathcal{R}_t(V_{t+1})} \right)^{\rho-\gamma} \left(\frac{(\mathbb{E}_{z_{t+1},t}[V_{t+1}^{1-\gamma}])^{1/(1-\gamma)}}{\mathcal{R}_t(V_{t+1})} \right)^{-(\eta-\gamma)},$$

where $z_{t+1} = 1, 2$ and $\mathbb{E}_{z_{t+1},t}$ denotes the expectation operator for the distribution of the consumption process conditioned on the history s^t and the period-$t+1$ state z_{t+1}. Given this pricing kernel, the return $R_{k,t+1}$ on any traded asset k satisfies the Euler equation

$$(17) \qquad \mathbb{E}_t \big[M_{z_{t+1},t+1} R_{k,t+1} \big] = 1,$$

where $\mathbb{E}_t$ is the expectation operator for the predictive distribution conditioned on history s^t. We distinguish between the unobservable price of aggregate consumption claims and the observable price of aggregate dividend claims. The return on the consumption claims is also the return on the wealth portfolio, which is unobservable, but can be solved using equation (17).

A challenge in estimating our model empirically is that the continuation value V_{t+1} in (16) is not observable. One possible approach is to use the relation between continuation value and wealth (proved in the Appendix)

$$(18) \qquad \frac{W_t}{C_t} = \frac{1}{1-\beta} \left(\frac{V_t}{C_t} \right)^{1-\rho},$$

where W_t is the wealth level at time t. We can then represent the pricing kernel (16) in terms of consumption growth and the return on the wealth portfolio, as in Epstein and Zin (1989). However, the return on the wealth portfolio is also unobservable, which makes empirical estimation of our model difficult.

We now turn to the interpretation of our pricing kernel in (16). The last multiplicative factor in (16) reflects the effect of ambiguity aversion. In the case of ambiguity neutrality (i.e., $\eta = \gamma$), this term vanishes and the pricing kernel reduces to that for the recursive utility model of Epstein and Zin (1989) and Weil (1989). When the agent is ambiguity averse with $\eta > \gamma$, a recession is associated with a high value of the pricing kernel. Intuitively, the agent has a lower continuation value V_{t+1} in a recession state, causing the adjustment factor in (16) to take a higher value in a recession than in a boom.

To explain asset-pricing puzzles, a number of studies propose to adjust the standard pricing kernel. As Campbell and Cochrane (1999) argued, they have to answer the basic question, Why do people fear stocks so much? In the Campbell and Cochrane habit formation model, people fear stocks because stocks do poorly in recessions, times when consumption falls low relative to habits. Our model's answer is that people fear stocks because they are pessimistic and have low continuation values in recessions. This pessimistic behavior will reduce the stock price and raise the stock return. In addition, it will reduce the risk-free rate because the agent wants to save more for the future. More formally, using equation (17), we can derive

$$(19) \qquad \mathbb{E}_t[R_{e,t+1} - R_{f,t+1}] = \frac{-\operatorname{Cov}_t(M_{z_{t+1},t+1}, R_{e,t+1})}{\mathbb{E}_t[M_{z_{t+1},t+1}]}.$$

Because stocks do poorly in recessions when ambiguity-averse people put more weight on the pricing kernel, ambiguity aversion helps generate high negative correlation between the pricing kernel and stock returns, and, thus, raises equity premium as shown in equation (19).[7]

To better understand an agent's pessimistic behavior, we consider the special case of the unitary EIS ($\rho = 1$). In this case, our model reduces to the Hansen–Sargent robust control model (9) and the pricing kernel becomes

$$(20) \qquad M_{z_{t+1},t+1} = \beta \frac{C_t}{C_{t+1}} \frac{V_{t+1}^{1-\gamma}(\mathbb{E}_{z_{t+1},t}[V_{t+1}^{1-\gamma}])^{-(\eta-\gamma)/(1-\gamma)}}{[\mathcal{R}_t(V_{t+1})]^{1-\eta}}.$$

The expression $\beta C_t/C_{t+1}$ is the pricing kernel for the standard log utility. It is straightforward to show that the adjustment factor in (20) is a density with respect to the predictive distribution because we can use the law of iterated expectations to show that the conditional expectation of the adjustment factor is equal to 1. As a result, we can write the Euler equation (17) as $\hat{\mathbb{E}}_t[\beta C_t/C_{t+1}R_{k,t+1}] = 1$, where $\hat{\mathbb{E}}_t$ is the conditional expectation operator for the slanted predictive distribution. In this case, the model is observationally equivalent to an expected utility model with distorted beliefs. The distorted beliefs are endogenous and attach more weight to the recession state. A similar observation equivalence result also appears in the multiple-priors model. (See Epstein and Miao (2003) for a discussion.) An undesirable feature of the unitary EIS case is that the consumption–wealth ratio is constant in that $C_t = (1 - \beta)W_t$ by (18).

---

[7]Using a static smooth ambiguity model, Gollier (2011) analyzed the effect of ambiguity aversion on the pricing kernel and portfolio choice. He showed that ambiguity aversion may not generally reinforce risk aversion.

Ju and Miao (2007) considered further the special case (10) with $\rho = \gamma = 1$. In this log-exponential case, the pricing kernel becomes

$$M_{z_{t+1},t+1} = \beta \frac{C_t}{C_{t+1}} \exp\big((1 - \eta)\mathbb{E}_{z_{t+1},t}[\ln V_{t+1}]\big)$$
$$/\big(\mu_t \exp((1 - \eta)\mathbb{E}_{1,t}[\ln V_{t+1}])$$
$$+ (1 - \mu_t)\exp((1 - \eta)\mathbb{E}_{2,t}[\ln V_{t+1}])\big).$$

The agent slants his state beliefs toward the state with the lower continuation value or the recession state. Ju and Miao (2007) also showed that the return on equity satisfies $R_{e,t+1} = \frac{1}{\beta} C_{t+1}/C_t$ if dividends are equal to aggregate consumption, $C_t = D_t$. Consequently, this case cannot generate interesting stock returns dynamics.

We now turn to the general homothetic specification with $\rho \neq 1$. In this case, the effect of ambiguity aversion is not distorting beliefs because the multiplicative adjustment factor in (16) is not a probability density. Thus, unlike in the case of $\rho = 1$, our model with $\rho \neq 1$ is not observationally equivalent to an expected utility model because one cannot find a change in beliefs of an expected utility maximizer that can account for the ambiguity aversion behavior in our model. However, our interpretation of the ambiguity aversion behavior as attaching more weight (the preceding adjustment factor) to the recession state with a worse continuation value is still valid, but the weight may not be mixture linear in state beliefs.

Ju and Miao (2007) studied the power–power case with $\rho = \gamma \neq 1$, in which risk aversion and intertemporal substitution are confounded. They required $\rho = \gamma < 1$ to explain the procyclical variation of price–dividend ratios and other asset-pricing puzzles. In a continuous-time multiple-priors model, Leippold, Trojani, and Vanini (2008) also calibrated $\rho = \gamma < 1$. Unlike the present paper, they assumed that (i) dividends are equal to consumption, (ii) dividend growth takes finitely many values without regime shifts, and (iii) the agent receives an additional signal about dividends.

Let $P_{e,t}$ denote the date t price of dividend claims. Using equations (16) and (17) and the homogeneity property of V_t, we can show that the price–dividend ratio $P_{e,t}/D_t$ is a function of the state beliefs, denoted by $\varphi(\mu_t)$. By definition, we can write the equity return as

$$R_{e,t+1} = \frac{P_{e,t+1} + D_{t+1}}{P_{e,t}} = \frac{D_{t+1}}{D_t} \frac{1 + \varphi(\mu_{t+1})}{\varphi(\mu_t)}.$$

This equation implies that the state beliefs drive changes in the price–dividend ratio, and hence the dynamics of equity returns. In the next section, we show that ambiguity aversion and learning under ambiguity help amplify consump-

tion growth uncertainty, while Bayesian learning has a modest quantitative effect.

4. QUANTITATIVE RESULTS

We first describe stylized facts and calibrate our model. We then study properties of unconditional and conditional moments of returns generated by our model. Our model does not admit an explicit analytical solution. We thus solve the model numerically using the projection method (Judd (1998)) and run Monte Carlo simulations to compute model moments.[8] For comparison, we also solve two benchmark models. Benchmark model I is the fully rational model with Epstein–Zin preferences under complete information similar to that studied by Bansal and Yaron (2004). Benchmark model II incorporates learning and is otherwise the same as benchmark model I. This model is a special case of our general model with $\eta = \gamma$. A special case of benchmark model II with time-additive expected utility ($\eta = \gamma = \rho$) is similar to the continuous-time model of Veronesi (1999, 2000).

4.1. *Stylized Facts and Calibration*

We start by summarizing some asset-pricing puzzles documented in the empirical literature. We use a century-long annual data set of real returns on Standard and Poors 500 stocks and 6-month commercial paper (1871–1993), and U.S. per capita real consumption growth (1889–1994) from Campbell (1999) and Cecchetti, Lam, and Mark (2000). Panel A of Table I reports that the mean values of equity premium and risk-free rate are given by 5.75 and 2.66 percent, respectively.[9] In addition, the volatility of equity premium is 19.02 percent. These values are hard to match in a standard asset-pricing model under reasonable calibration. This fact is often referred to as the equity premium, risk-free rate, and equity volatility puzzles (see Campbell (1999) for a survey). Panel B of Table I reports that the log dividend yield predicts long-horizon realized excess returns. It also shows that the regression slopes and R^2s increase with the return horizon. This return predictability phenomenon was first documented by Campbell and Shiller (1988a, 1988b) and Fama and French (1988a). Panel B of Table I also reports variance ratio statistics for the equity premium. These ratios are generally less than 1 and fall with the horizon. This evidence suggests that excess returns are negatively serially correlated or asset prices are mean reverting (Fama and French (1988b) and Poterba and Summers (1988)).

---

[8]The Fortran codes and a technical appendix detailing our numerical method are available as Supplemental Material (Ju and Miao (2012)).

[9]Campbell (1999) and Campbell and Cochrane (1999) found similar estimates using log returns. We follow Cecchetti, Lam, and Mark (2000) and report arithmetic average returns in both data and model solutions. Mehra and Prescott (1985) also reported arithmetic averages.

TABLE I

STYLIZED FACTS OF EQUITY AND SHORT-TERM BOND RETURNS USING THE CENTURY-LONG
ANNUAL DATA SET[a]

A. First and second moments as a percentage

Mean equity premium	μ_{eq}	5.75
Mean risk-free rate	r_f	2.66
Standard deviation		
Equity premium	$\sigma(\mu_{eq})$	19.02
Risk-free rate	$\sigma(r_f)$	5.13

B. Predictability and persistence of excess returns

Horizon	Regression Slope	R^2	Variance Ratio
1	0.148	0.043	1.000
2	0.295	0.081	1.038
3	0.370	0.096	0.921
5	0.662	0.191	0.879
8	0.945	0.278	0.766

C. Autocorrelations and cross-correlations

$p_t - d_t, \Delta c_t$	$r_{t+1}^{ex}, \Delta c_t$	$|r_t^{ex}|, \Delta c_t$	$|r_{t+1}^{ex}|, |r_t^{ex}|$	$|r_{t+1}^{ex}|, p_t - d_t$
0.25	−0.16	−0.28	0.13	−0.11

[a]The regression slope and R^2 are for regressions of the k-year ($k = 1, 2, 3, 5, 8$) ahead equity premium on the current log dividend–price ratio. The variance ratio is the variance of the k-year equity premium divided by k times the variance of the 1-year equity premium. Panels A and B reproduce Table 1 in Cecchetti, Lam, and Mark (2000). In panel C, $p_t - d_t = \ln(P_{e,t}/D_t)$, $\Delta c_t = \ln(C_t/C_{t-1})$, and $r_t^{ex} = \ln(R_{e,t}/R_{f,t})$.

In addition to the preceding puzzles, we use our model to explain three other stylized facts: (i) procyclical variation in price–dividend ratios (Fama and French (1989)), (ii) countercyclical variation in conditional expected equity premium (Campbell and Shiller (1988a, 1988b) and Fama and French (1989)), and (iii) persistent and countercyclical variation in conditional volatility of equity premium (Bollerslev, Chou, and Kroner (1992)). Panel C of Table I illustrates these facts using simple autocorrelations and cross-correlations. In this panel, we use absolute returns as a proxy for volatility. Because aggregate consumption growth and output growth are positively correlated, we use consumption growth as a procyclical economic indicator. The last column of this panel also shows that the correlation between the log price–dividend ratio and the next period absolute log excess returns is negative. This result indicates that price declines increase volatility, a "leverage effect" found by Black (1976) and many others.

Now we calibrate our model at the annual frequency. We first calibrate parameters in the consumption process (13). Cecchetti, Lam, and Mark (2000) applied Hamilton's maximum likelihood method to estimate these parameters using the annual per capita U.S. consumption data covering the period 1890–1994. Table II reproduces their estimates. This table reveals that the high-

TABLE II

MAXIMUM LIKELIHOOD ESTIMATES OF THE CONSUMPTION PROCESS[a]

	λ_{11}	λ_{22}	κ_1	κ_2	σ
Estimate	0.978	0.516	2.251	−6.785	3.127
(t ratio)	(50.94)	(1.95)	(6.87)	(−3.60)	(13.00)

[a]The numbers in the last three columns are expressed in percentage. This table reproduces Table 2 in Cecchetti, Lam, and Mark (2000).

growth state is highly persistent, with consumption growth in this state being 2.251 percent. The economy spends most of the time in this state with the unconditional probability of being in this state given by $(1 - \lambda_{22})/(2 - \lambda_{11} - \lambda_{22}) = 0.96$. The low-growth state is moderately persistent, but very bad, with consumption growth in this state being −6.785 percent. The long-run average rate of consumption growth is 1.86 percent.

We next calibrate parameters in the dividend process (14). We follow Abel (1999) and set the leverage parameter $\zeta = 2.74$. We then follow Bansal and Yaron (2004) and choose $g_d = -0.0323$ so that the average rate of dividend growth is equal to that of consumption growth. We choose σ_d to match the volatility of dividend growth in the data. Using different century-long annual samples, this volatility is equal to 0.136 and 0.142, according to the estimates given by Cecchetti, Lam, and Mark (1990) and Campbell (1999), respectively. Here, we take 0.13 and find $\sigma_d = 0.084$. Our calibrated values of σ_d and ζ imply that the correlation between consumption growth and dividend growth is about 0.76. This value may seem high. However, Campbell and Cochrane (1999) argued that the correlation is difficult to measure and it may approach 1.0 in the very long run since dividends and consumption should share the same long-run trends.

Turn to the preference parameters. We follow Bansal and Yaron (2004) and set EIS to 1.5, implying $\rho = 1/1.5$. An EIS greater than 1 is critical to generate procyclical variation of the price–consumption ratio. Researchers in macroeconomics and finance generally believe that the risk aversion parameter is around 2. We thus set $\gamma = 2$, so as to demonstrate that the main force of our model comes from ambiguity aversion, but not risk aversion. We select the subjective discount factor β and the ambiguity aversion parameter η to match the mean risk-free rate of 0.0266 and the mean equity premium of 0.0575 from the data reported in Table I. We find $\beta = 0.975$ and $\eta = 8.864$.

There is no independent study of the magnitude of ambiguity aversion in the literature. To have a sense of our calibrated value, we conduct a thought experiment related to the Ellsberg paradox (Ellsberg (1961)) in a static setting. Suppose there are two urns. Subjects are told that there are 50 black and 50 white balls in urn 1. Urn 2 also contains 100 balls, but may contain either 100 black balls or 100 white balls. If a subject picks a black ball from an urn, he wins

a prize; otherwise, he does not win or lose anything. Experimental evidence reveals that most subjects prefer to bet on urn 1 rather than urn 2 (Camerer (1999) and Halevy (2007)). Paradoxically, if the subject is asked to pick a white ball, he still prefers to bet on urn 1. The standard expected utility model with any beliefs or any risk aversion level cannot explain this paradox. Our adopted smooth ambiguity model in the static setting (1) can explain this paradox whenever subjects display ambiguity aversion (i.e., v is more concave than u). Thus, ambiguity aversion and risk aversion have distinct behavioral meanings.

Formally, let w be a subject's wealth level and let d be the prize money. The subject knows that the distribution of black and white balls in urn 1 is $(1/2, 1/2)$. When he evaluates a bet on urn 1, his utility level in terms of certainty equivalent is equal to

$$(21) \qquad u^{-1}\left(\frac{1}{2}u(w+d) + \frac{1}{2}u(w)\right).$$

The subject believes that there are two possible equally likely distributions $(0, 1)$ and $(1, 0)$ in urn 2, and thus $\Pi = \{(0, 1), (1, 0)\}$ and $\mu = (1/2, 1/2)$. But he is not sure which one is the true distribution and is averse to this uncertainty. When he evaluates a bet on urn 2, his utility level in terms of certainty equivalent is equal to (1), where $C = w + d$ or w and $S = \{$black, white$\}$, that is,

$$(22) \qquad v^{-1}\left(\frac{1}{2}v(w+d) + \frac{1}{2}v(w)\right).$$

The expression in (21) is greater than that in (22) if v is more concave than u, causing the subject to bet on urn 1 rather than urn 2. The difference between the certainty equivalents in (21) and (22) is a measure of ambiguity premium. Given power functions of u and v, and fixing the risk aversion parameter, we can use the size of the ambiguity premium to gauge the magnitude of ambiguity aversion.[10] It is straightforward to compute that the ambiguity premium is equal to 1.7 percent of the expected prize value for our calibrated ambiguity aversion parameter $\eta = 8.864$, when we set $\gamma = 2$ and the prize–wealth ratio of 1 percent. Increasing the prize–wealth ratio raises the ambiguity premium. Camerer (1999) reported that the ambiguity premium is typically on the order of 10–20 percent of the expected value of a bet in the Ellsberg-style experiments. Given this evidence, our calibrated ambiguity aversion parameter seems small and reasonable. It is consistent with the experimental findings, though they are not the basis for our calibration.

---

[10]See Chen, Ju, and Miao (2009) for a more extensive discussion and an application of our generalized recursive ambiguity model to a portfolio choice problem.

4.2. *Unconditional Moments of Returns*

As a first check of the performance of our calibrated model, we compare our model's prediction of the volatility of the equity premium and the volatility of the risk-free rate with the data. Panel A of Table III reports model results. This table reveals that our model can match the volatility of the equity premium in the data quite closely (0.1826 versus 0.1902). However, our model generated volatility of the risk-free rate is lower than the data (0.0116 versus 0.0513). Campbell (1999) argued that the high volatility of the real risk-free rate in the century-long annual data could be due to large swings in inflation in the interwar period, particularly in 1919–1921. Much of this volatility is probably due to unanticipated inflation and does not reflect the volatility in the ex ante real interest rate. Campbell (1999) reported that the annualized volatility of the real return on Treasury bills is 0.018 using the U.S. postwar quarterly data.

To understand why our model is successful in matching unconditional moments of returns, we conduct a comparative statics analysis in panels B–E

TABLE III
UNCONDITIONAL MOMENTS AND COMPARATIVE STATISTICS[a]

η	r_f	$\sigma(r_f)$	r_e	$\sigma(r_e)$	μ_{eq}	$\sigma(\mu_{eq})$	$\frac{\sigma(M)}{E[M]}$
			Panel A (Baseline): $\rho = 1/1.5$, $\gamma = 2.0$				
8.864	2.66	1.16	8.41	17.98	5.75	18.26	0.60
			Panel B: $\rho = 1/1.5$, $\gamma = 2.0$				
2.0	3.61	0.63	4.60	14.49	0.98	14.50	0.09
3.0	3.55	0.70	4.81	14.72	1.26	14.74	0.12
8.0	2.88	1.10	7.44	17.28	4.56	17.51	0.47
15.0	−0.18	1.06	16.44	18.75	16.62	19.01	2.20
			Panel C: $\rho = 1/1.5$, $\gamma = 5.0$				
5.0	3.02	0.89	6.48	15.73	3.46	15.83	0.30
8.0	2.40	1.12	9.02	17.72	6.62	17.97	0.59
15.0	−0.87	0.97	17.87	18.47	18.74	18.68	2.42
			Panel D: $\rho = 1/2$, $\gamma = 2.0$				
2.0	3.30	0.51	4.31	14.57	1.02	14.58	0.09
3.0	3.23	0.58	4.53	14.82	1.30	14.85	0.12
8.0	2.52	1.00	7.30	17.59	4.78	17.82	0.48
15.0	−0.82	0.96	16.42	19.10	17.23	19.34	2.22
			Panel E: $\rho = 1/2$, $\gamma = 5.0$				
5.0	2.71	0.77	6.29	15.92	3.58	16.01	0.30
8.0	2.04	1.02	8.93	18.04	6.89	18.29	0.60
15.0	−1.52	0.87	17.88	18.80	19.39	19.00	2.45

[a]Except for the numbers in the first and the last columns, all other numbers are in percentage. The second through seventh columns present the means and standard deviations of the risk-free rate, the equity return, and the equity premium, respectively. $\sigma(M)/E[M]$ is the ratio of the standard deviation to the mean of the pricing kernel. We set $\beta = 0.975$ in all cases.

of Table III. The first row of each of these panels gives the result of benchmark model II with Epstein–Zin preferences under Bayesian learning. We first consider the effects of the three standard parameters (β, ρ, γ) familiar from the Epstein–Zin model. Equation (17) implies that the risk-free rate $R_{f,t+1} = 1/\mathbb{E}_t[M_{z_{t+1},t+1}]$. Because the pricing kernel $M_{z_{t+1},t+1}$ increases with the subjective discount factor β, a high value of β helps match the low risk-free rate. Table III reveals that an increase in EIS (or $1/\rho$) from 1.5 to 2.0 generally lowers the risk-free rate and stock returns due to the high intertemporal substitution effect. In addition, an increase in γ from 2.0 to 5.0 also lowers the risk-free rate and raises stock returns. These results follow from the usual intuition in the Epstein–Zin model.

Next, consider the role of ambiguity aversion, which is unique in our model. Table III reveals that an increase in the ambiguity aversion parameter η lowers the risk-free rate and raises stock returns. The intuition follows from the discussion in Section 3.2. An ambiguity-averse agent displays pessimistic behavior by attaching more weight to the worst state with low continuation utilities. Thus, he saves more for the future and invests less in the stock. In addition, as more weight is attached to the low-growth state, there is less variation of $\mathbb{E}_t[M_{z_{t+1},t+1}]$, and hence the risk-free rate $R_{f,t+1}$ is less volatile. By contrast, ambiguity aversion makes the pricing kernel $M_{z_{t+1},t+1}$ more volatile as revealed by the last term in (16), leading to high and volatile equity premium. It also generates a high market price of uncertainty defined by the ratio of the volatility of the pricing kernel and the mean of the pricing kernel (Hansen and Jagannathan (1991)). For our calibrated baseline parameter values, the market price of uncertainty is equal to 0.60, as reported in panel A of Table III. It is equal to 0.09 in benchmark model II with $\eta = \gamma$.

Finally, we analyze the role of learning under ambiguity. We decompose the risk-free r_f in our model into three components,

$$(23) \qquad r_f = r_f^* + (r_f^L - r_f^*) + (r_f - r_f^L),$$

where r_f^*, r_f^L, and r_f are the means of the risk-free rate delivered by benchmark model I, benchmark model II, and our ambiguity model, respectively. We do a similar decomposition for the mean stock returns and the volatility of the equity premium.[11] Table IV presents this decomposition.

Panel A of Table IV shows that under the baseline parameter values, benchmark model I with full information predicts that the mean risk-free rate $r_f^* = 0.0363$, the mean equity returns $r_e^* = 0.046$, and the volatility of equity premium $\sigma_{\text{eq}}^* = 0.1448$. For benchmark model II with hidden states, the stan-

---

[11]In a continuous-time multiple-priors model without learning, Chen and Epstein (2002) provided a similar decomposition and showed that equity premium reflects a premium for risk and a premium for ambiguity.

TABLE IV

DECOMPOSITION OF r_f, r_e AND σ_{eq}[a]

η	r_f^*	Δr_f^L	Δr_f	r_e^*	Δr_e^L	Δr_e	σ_{eq}^*	$\Delta\sigma_{\mathrm{eq}}^L$	$\Delta\sigma_{\mathrm{eq}}$
			Panel A (Baseline): $\rho = 1/1.5$, $\gamma = 2.0$						
8.864	3.63	−0.01	−0.96	4.60	0.01	3.80	14.48	0.02	3.76
			Panel B: $\rho = 1/1.5$, $\gamma = 2.0$						
2.0	3.62	−0.01	0.00	4.59	0.01	0.00	14.48	0.02	0.00
3.0	3.62	−0.01	−0.06	4.59	0.01	0.21	14.48	0.02	0.25
8.0	3.62	−0.01	−0.73	4.59	0.01	2.84	14.48	0.02	3.01
15.0	3.62	−0.01	−3.80	4.59	0.01	11.84	14.48	0.02	4.51
			Panel C: $\rho = 1/1.5$, $\gamma = 5.0$						
5.0	3.05	−0.03	0.00	6.45	0.03	0.00	15.78	0.05	0.00
8.0	3.05	−0.03	−0.61	6.45	0.03	2.54	15.78	0.05	2.14
15.0	3.05	−0.03	−3.89	6.45	0.03	11.40	15.78	0.05	2.85
			Panel D: $\rho = 1/2$, $\gamma = 2.0$						
2.0	3.30	−0.01	0.00	4.31	0.00	0.00	14.58	0.01	0.00
3.0	3.30	−0.01	−0.06	4.31	0.00	0.22	14.58	0.01	0.27
8.0	3.30	−0.01	−0.78	4.31	0.00	2.99	14.58	0.01	3.24
15.0	3.30	−0.01	−4.11	4.31	0.00	12.10	14.58	0.01	4.76
			Panel E: $\rho = 1/2$, $\gamma = 5.0$						
5.0	2.73	−0.02	0.00	6.26	0.02	0.00	15.99	0.03	0.00
8.0	2.73	−0.02	−0.67	6.26	0.02	2.64	15.99	0.03	2.27
15.0	2.73	−0.02	−4.22	6.26	0.02	11.59	15.99	0.03	2.99

[a]Except for the numbers in the first column, all numbers are in percentage. The variables r_f^*, r_e^*, and σ_{eq}^* are the mean risk-free rate, the mean stock return, and the equity premium volatility, respectively, for benchmark model I. The variables r_f^L, r_e^L, and σ_{eq}^L are the mean risk-free rate, the mean stock return, and the equity premium volatility, respectively, for benchmark model II. We denote by $\Delta r_f^L = r_f^L - r_f^*$ the change of the mean risk-free rate due to Bayesian learning and denote by $\Delta r_f = r_f - r_f^L$ the change of the mean risk-free rate due to learning under ambiguity. The other variables Δr_e^L, Δr_e, $\Delta\sigma_{\mathrm{eq}}^L$, and $\Delta\sigma_{\mathrm{eq}}$ are defined similarly. We set $\beta = 0.975$ in all cases.

dard Bayesian learning lowers the risk-free rate and raises the equity return and equity volatility, but by a negligible amount. By contrast, the component $(r_f - r_f^L)$ due to learning under ambiguity accounts for most of the decrease in the risk-free rate and the increase in the equity return and the volatility of the equity premium. In addition, the magnitude of this component is larger for a larger degree of ambiguity aversion. We find a similar result also for other values of the risk aversion parameter as presented in panels B and C. In particular, when the risk aversion parameter $\gamma = 2$ and 5, the corresponding effects of Bayesian learning are to lower the mean risk-free rate by 0.01 and 0.03 percent, to raise the mean stock return by 0.01 and 0.03 percent, and to raise the equity premium volatility by 0.02 and 0.05 percent. These effects are quantitatively negligible. Increasing EIS from 1.5 to 2.0 does not change this result much as revealed by panels D and E.

A surprising feature of benchmark model II with Bayesian learning is that the equity premium can become negative when risk aversion γ is sufficiently large in the special case of time-additive utility $\gamma = \rho$. Increasing risk aversion may worsen the equity premium puzzle. In a similar continuous-time model, Veronesi (2000) proved this result analytically. The intuition is that an increase in risk aversion raises the agent's hedging demand for the stock after bad news in dividends, thereby counterbalancing the negative pressure on prices due to the bad news in dividends. The former effect may dominate so that the pricing kernel and stock returns are positively correlated, resulting in a negative equity premium (see equation (19)). By contrast, in our model, an ambiguity-averse agent invests less in the stock, thereby counterbalancing the preceding hedging effect. In contrast to risk aversion, an increase in the degree of ambiguity aversion helps increase equity premium.

4.3. *Price–Consumption and Price–Dividend Ratios*

Panel A of Figure 1 presents the price–consumption ratio of the consumption claim as a function of the Bayesian posterior probabilities μ_t of the high-growth state for three values of η, holding other parameters fixed at the baseline values. It reveals two properties. First, the price–consumption ratio is increasing and convex. The intuition is similar to that described by Veronesi (1999), who analyzed the case with time-additive expected exponential utility. When times are good (μ_t is close to 1), a bad piece of news decreases μ_t, and hence decreases future expected consumption growth. But it also increases the agent's uncertainty about consumption growth since μ_t is now closer to 0.5, which gives approximately the maximal conditional volatility of the posterior probability of the high-growth state in the next period. Since the agent wants to be compensated for bearing more risk, he will require an additional discount on the price of consumption claims. Thus, the price reduction due to a bad piece of news in good times is higher than the reduction in expected future consumption. By contrast, suppose the agent believes times are bad and hence μ_t is close to zero. A good piece of news increases the expected future consumption growth, but also raises the agent's perceived uncertainty since it moves μ_t closer to 0.5. Thus, the price–consumption ratio increases, but not as much as it would in a present-value model.

The second property in panel A of Figure 1 is that an increase in the degree of ambiguity aversion lowers the price–consumption ratio because it induces the agent to invest less in the asset. In addition, an increase in the degree of ambiguity aversion raises the curvature of the price–consumption ratio function, thereby helping to increase the asset price volatility. In the special case of benchmark model II with $\eta = \gamma$, the price–consumption ratio is close to being
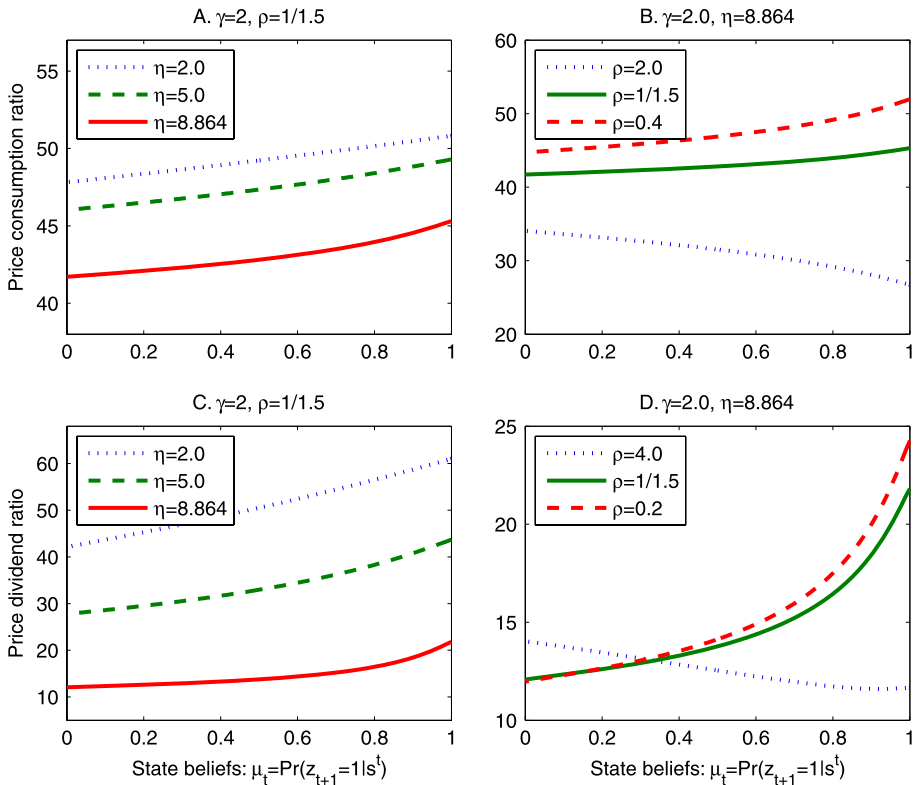
FIGURE 1.—Panels A and B plot the price–consumption ratio of the consumption claim as a function of the Bayesian posterior probabilities of the high-growth state. Panels C and D plot the price–dividend ratio of the dividend claim as a function of the posterior probabilities of the high-growth state. We set $\beta = 0.975$. The case with $\gamma = \eta$ corresponds to the Epstein–Zin model with Bayesian learning.

a linear function of the state beliefs.[12] Thus, this model cannot generate high asset price volatility.

Panel B of Figure 1 presents the price–consumption ratio function for various values of ρ, holding other parameters fixed at the baseline values. It reveals that the price–consumption ratio is an increasing function of μ_t when $\rho < 1$, while it is a decreasing function when $\rho > 1$. When $\rho = 1$, it is equal to $\beta/(1 - \beta)$ by (18) because wealth is equal to the present value of the consumption claim. These results follow from the usual intuition in the Epstein–Zin model (see Bansal and Yaron (2004)). When $\rho < 1$, EIS is greater than 1 and

---

[12]We can follow Veronesi (1999) to prove analytically that both the price–consumption and the price–dividend ratios are linear in state beliefs for time-additive expected utility.

hence the intertemporal substitution effect dominates the wealth effect. In response to good news of consumption growth, the agent buys more assets and hence the price–consumption ratio rises. The opposite result holds true when $\rho > 1$.

Panels C and D of Figure 1 present similar figures for the price–dividend ratio of the dividend claim. We find that the effects of η and ρ are similar. One difference is that the price–dividend ratio is not constant when $\rho = 1$ because dividends and aggregate consumption are not identical in our model. Due to leverage, we need a sufficiently small EIS (or a large ρ) to make the price–dividend ratio decrease with μ_t. Bansal and Yaron (2004) found a similar result in a full information model with Epstein–Zin preferences.

In summary, ambiguity aversion helps generate the variation in the price–consumption and price–dividend ratios. An EIS greater than 1 is important to generate procyclical price–consumption and price–dividend ratios. According to our baseline calibration, our model implied that correlation between consumption growth and the log price–dividend ratio is 0.21, close to the data (0.25) reported in panel C of Table I.

4.4. *Time-Varying Equity Premium and Equity Volatility*

Panel A of Figure 2 plots the conditional expected equity premium as a function of the posterior probability μ_t of the high-growth state for various values of η. We find that this function is hump-shaped and peaks when μ_t is around 0.6. This shape seems to suggest that a negative consumption shock can lead to either a higher or a lower equity premium, depending on whether μ_t is close to 0 or to 1. However, since the economy spends most of the time in the high-growth state, the steady-state distribution of the posterior is highly skewed. This implies that μ_t is close to 1 most of the time, leading to the pattern that equity premium rises following negative consumption shocks. As a result, our model can generate the countercyclical variation in equity premium observed in the data. Our model implies that the correlation between consumption growth and the next period equity premium is -0.14, close to the data -0.16 reported in panel C of Table I.

What is the role of ambiguity aversion? Panel A of Figure 2 shows that the curvature of the conditional expected equity premium function increases with η, implying that ambiguity aversion helps amplify the variation in equity premium (see equations (16) and (19)). In benchmark model II with Bayesian learning ($\eta = \gamma = 2.0$), the conditional expected equity premium is almost flat. Consequently, it cannot generate a highly time-varying expected equity premium. By contrast, when η is increased from 2 to 8.864, the conditional equity premium can rise from about 3 percent to 28 percent.

Panel B of Figure 2 plots the conditional volatility of equity premium as a function of μ_t for various values of η. This function is also hump-shaped, with the maximum attained at a value of μ_t close to 0.6. Following similar intuition
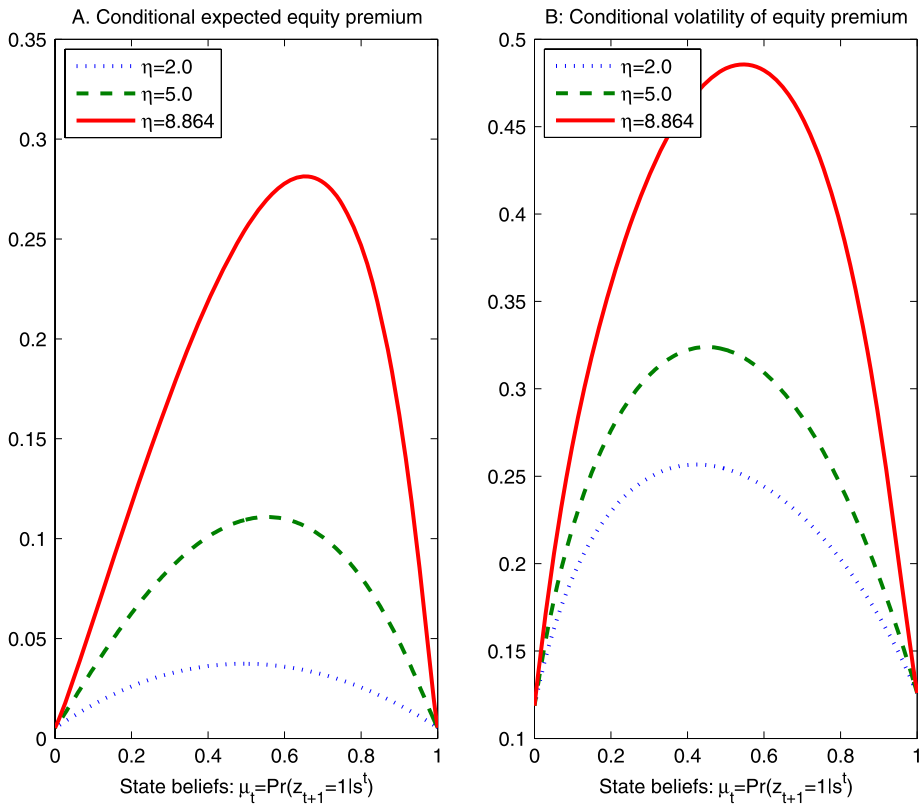
FIGURE 2.—The conditional expected equity premium and conditional volatility of equity premium as functions of the Bayesian posterior probabilities of the high-growth state. We set $\rho = 1/1.5$, $\beta = 0.975$, and $\gamma = 2$. The case with $\gamma = \eta = 2$ corresponds to the Epstein–Zin model with Bayesian learning.

discussed above, our model generates countercyclical variation in conditional volatility of the equity premium observed in the data. In addition, ambiguity aversion helps amplify this variation. Our model implies that the correlation between absolute excess returns and consumption growth is -0.10, the same sign as in the data (-0.28) reported in panel C of Table I, though it has a smaller magnitude. Our model implies that the correlation between the log price–dividend ratio and the next period absolute excess returns is -0.16, close to the data (-0.11). Thus, our model generates the leverage effect.

   Our model can also generate persistent changes in conditional volatility of the equity premium, documented by Bollerslev, Chou, and Kroner (1992). Our model implies that the autocorrelation of absolute excess returns is 0.17, close to the data (0.13) reported in panel C of Table I. The intuition is that the agent's beliefs are persistent in the sense that if he believes the high-growth
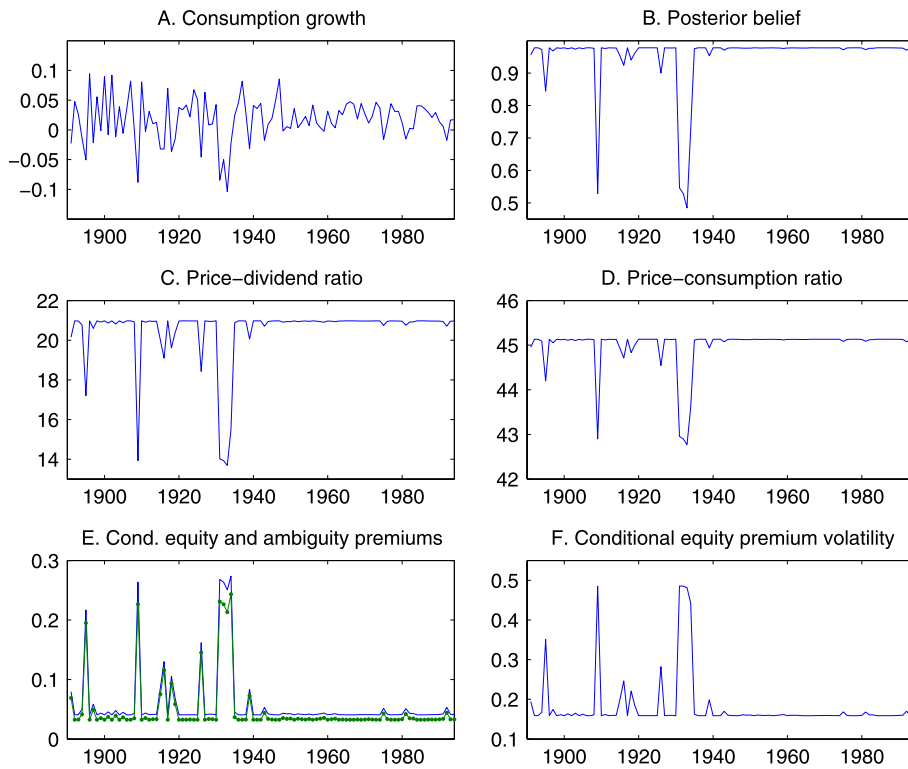
FIGURE 3.—Implied time series of some financial variables using historical consumption growth data. In panel E, the solid and dotted solid lines represent conditional equity and ambiguity premiums, respectively. We take baseline parameter values calibrated in Section 4.1.

state today has a high probability, then he expects the high-growth state tomorrow also has a high probability on average. The persistence of beliefs drives the persistence of the volatility of the equity premium.

Figure 3 illustrates the time-varying properties of some implied financial variables using the historical consumption growth data from 1890 to 1994. Panel A plots these data. Panel B plots the implied time series of the posterior probability of the high-growth state μ_t, computed using equation (15) in which we take the long-run stationary probability of the high-growth state as the initial value, $\mu_0 = 0.9565$. Panel B reveals that most of the time the agent believes that the economy is in the high-growth state in that μ_t is close to 1. During World War I and the Great Depression, the economy entered severe recessions and μ_t is close to 0.5. At this value, the agent's perceived uncertainty about the high-growth state is the highest. Using the time series of $\{\mu_t\}$, we can compute the implied time series of the price–dividend and price–consumption

TABLE V

PREDICTABILITY AND PERSISTENCE OF EXCESS RETURNS[a]

Horizon	Baseline Parameter Values			Benchmark Model I			Benchmark Model II		
	Slope	R^2	Variance Ratio	Slope	R^2	Variance Ratio	Slope	R^2	Variance Ratio
1	0.839	0.105	1.000	0.403	0.025	1.000	0.583	0.022	1.000
2	1.117	0.118	0.866	0.547	0.028	0.974	0.761	0.025	0.971
3	1.227	0.112	0.784	0.611	0.027	0.954	0.836	0.024	0.950
5	1.303	0.095	0.690	0.660	0.024	0.924	0.914	0.023	0.920
8	1.343	0.077	0.618	0.715	0.022	0.890	1.039	0.022	0.886

[a]The slope and R^2 are obtained from an ordinary least squares regression of the excess returns on the log dividend yield at different horizons. The variance ratio is computed in the same way as Cecchetti (1990, 2000). The reported numbers are the mean values of 10,000 Monte Carlo simulations, each consisting of 123 excess returns and dividend yields. Baseline parameter values are $\beta = 0.975$, $\rho = 1/1.5$, $\gamma = 2$, and $\eta = 8.864$ for our recursive ambiguity model. Benchmark model I is the Epstein–Zin model with full information. Benchmark model II is the Epstein–Zin model with Bayesian learning. In these two benchmark models, we set $\beta = 0.975$, $\rho = 1/1.5$, and $\gamma = 2$.

ratios, the conditional expected equity and ambiguity premiums,[13] and the conditional equity premium volatility. We plot these series in panels C–F of Figure 3, respectively. From these panels, we can visually see that (i) both the price–dividend and the price–consumption ratios are procyclical, and (ii) both the conditional equity premium volatility and the conditional expected equity premium are countercyclical. In addition, the countercyclical movements of the conditional equity premium are largely driven by the countercyclical movements of the ambiguity premium.

4.5. *Serial Correlation and Predictability of Returns*

To examine the ability of our model to generate the serial correlation and predictability of returns reported in Table II, we compare our model with benchmark models I and II. Table V reports the model implied values of the variance ratios, the regression slopes, and the R^2s at horizons of 1, 2, 3, 5, and 8 years based on the baseline parameter values given in Table III. To account for the small sample bias in these statistics, we generate them using 10,000 Monte Carlo experiments as described in Cecchetti, Lam, and Mark (2000).

From Table V, we observe that all three models can generate the pattern that variance ratios are less than 1 and decrease with the horizon, suggesting that excess returns are negatively serially correlated. In terms of predictive regressions, benchmark models I and II deliver very small R^2s, implying weak

---

[13]Conditional ambiguity premium is defined as the difference in the conditional equity premium delivered by our ambiguity model and by the benchmark model II.

predictability.[14] One may expect that learning should help generate return predictability. The intuition is that the change of state beliefs is persistent, and hence the price–dividend ratio is also persistent and positively serially correlated. However, Table V reports that benchmark model II with Bayesian learning helps little quantitatively. In a related model, Brandt, Zeng, and Zhang (2004) found a similar result.

We finally consider our model in which we introduce ambiguity aversion into benchmark model II. Table V reveals that while all three models can generate the pattern that the regression slopes increase with the horizon, our model with learning under ambiguity produces much higher R^2s. All three models cannot generate the pattern that R^2s increase with the horizon. The model predicted R^2s first increase with the horizon and then decrease with it after period 3. This could be due to the fact that the model generated price–dividend ratios are not persistent enough.[15] We should recognize that the predictability results in the empirical literature are quite sensitive to data sets, changing samples, and estimation techniques (Welch and Goyal (2008)). Thus, one should be cautious in interpreting empirical evidence on predictability.

5. CONCLUSION

In this paper, we propose a novel generalized recursive smooth ambiguity model which permits a three-way separation among risk aversion, ambiguity aversion, and intertemporal substitution. This model nests some utility models widely adopted in the literature as special cases. We also propose a tractable homothetic specification and apply this specification to asset pricing. When modeling consumption growth and dividend growth as regime-switching processes (nonlinear counterparts of the long-run risk processes as in Bansal and Yaron (2004)), our asset-pricing model can help explain a variety of asset-pricing puzzles. Our calibrated model can match the mean equity premium, the mean risk-free rate, and the volatility of the equity premium observed in the data. In addition, our model can generate a variety of dynamic asset-pricing phenomena, including the procyclical variation of price–dividend ratios, the countercyclical variation of equity premium and equity volatility, the leverage effect, and the mean reversion of excess returns.

We show that ambiguity aversion and learning under ambiguity play a key role in explaining asset-pricing puzzles. An ambiguity-averse agent displays pessimistic behavior in that he attaches more weight to the pricing kernel in bad times when his continuation values are low. This pessimistic behavior helps

---

[14]Beeler and Campbell (2009) and Garcia, Meddahi, and Tedongap (2008) reexamined the Bansal–Yaron model and found that it cannot match the predictability in the data, contrary to the finding of Bansal and Yaron (2004).

[15]See Campbell, Lo, and Mackinlay (1997, pp. 271–273) for a theoretical analysis of why R^2s may first increase and then decrease with the horizon.

propagate and amplify shocks to consumption growth, and generates the time-varying equity premium. We also find that Bayesian learning in the expected utility framework has a modest quantitative effect on asset returns, while learning under ambiguity is important to explain dynamic asset-pricing phenomena. One limitation of our model is that it cannot reproduce the predictability pattern in the data.

Other models may also simultaneously generate the unconditional moments and dynamics of asset returns observed in the data. For example, Campbell and Cochrane (1999) introduced a slow moving habit or time-varying subsistence level into a standard power utility function.[16] As a result, the agent's risk aversion is time varying. Bansal and Yaron (2004) applied the Epstein–Zin preferences, and incorporated fluctuating volatility and a persistent component in consumption growth.[17] Their calibrated risk aversion parameter is 10. Our model of consumption and dividend processes is similar to Bansal and Yaron (2004), but is much easier to estimate. We shut down exogenous fluctuations in consumption growth volatility and analyze how endogenous learning under ambiguity can generate a time-varying equity premium.

We view our model as a first step toward understanding the quantitative implications of learning under ambiguity for asset returns. We have shown that our model can go a long way to explain many asset-pricing puzzles. Much work still remains to be done. For example, how to empirically estimate parameters of ambiguity aversion, risk aversion, and intertemporal substitution would be an important future research topic. In addition, our proposed novel generalized recursive smooth ambiguity model can be applied to many other problems in economics.

APPENDIX: PROOFS OF RESULTS IN SECTION 3.2

We follow the method of Hansen, Heaton, and Li (2008) to derive the marginal utility of consumption and continuation value as

$$\text{MC}_t = \frac{\partial V_t(C)}{\partial C_t} = (1-\beta)V_t^\rho C_t^{-\rho},$$

$$\text{MV}_{z_{t+1},t+1} = \frac{\partial V_t(C)}{\partial V_{z_{t+1},t+1}}$$
$$= \beta V_t^\rho [\mathcal{R}_t(V_{t+1})]^{\eta-\rho}(\mathbb{E}_{z_{t+1},t}[V_{t+1}^{1-\gamma}])^{-(\eta-\gamma)/(1-\gamma)}V_{t+1}^{-\gamma},$$

---

[16]Ljungqvist and Uhlig (2009) showed that government interventions that occasionally destroy part of endowment can be welfare improving when endogenizing aggregate consumption choices in the Campbell–Cochrane habit formation model.

[17]See Beeler and Campbell (2009) for a critique of the Bansal–Yaron model.

where $V_{z_{t+1},t+1}$ denotes the continuation value $V_{t+1}(C)$ conditioned on the period $t + 1$ state being z_{t+1}. The pricing kernel is given by $M_{z_{t+1},t+1} = (\mathrm{MV}_{z_{t+1},t+1})(\mathrm{MC}_{t+1})/\mathrm{MC}_t$, which delivers (16).

We next use the dynamic programming method of Epstein and Zin (1989) to derive other results in Section 3.2. Suppose the agent trades N assets. The budget constraint is $W_{t+1} = (W_t - C_t)R_{w,t+1}$, where the return on the wealth portfolio $R_{w,t+1}$ is equal to $\sum_{k=1}^{N} \psi_{kt} R_{k,t+1}$, ψ_{kt} is the portfolio weight on asset k, and $R_{k,t+1}$ denotes its return. The value function $J_t(W_t, \mu_t)$ satisfies the Bellman equation

$$
(A.1) \quad J_t(W_t, \mu_t) = \max_{C_t, \{\psi_{kt}\}} \Big[(1 - \beta)C_t^{1-\rho}
$$
$$
+ \beta \big\{ \big(\mu_t \big(\mathbb{E}_{1,t}[J_t^{1-\gamma}(W_{t+1}, \mu_{t+1})] \big) \big)^{(1-\eta)/(1-\gamma)}
$$
$$
+ (1 - \mu_t)
$$
$$
\times \big(\mathbb{E}_{2,t}[J_t^{1-\gamma}(W_{t+1}, \mu_{t+1})] \big)^{(1-\eta)/(1-\gamma)} \big\}^{(1-\rho)/(1-\eta)} \Big]^{1/(1-\rho)}.
$$

Conjecture

$$
(A.2) \quad J_t(W_t, \mu_t) = A_t W_t \quad \text{and} \quad C_t = a_t W_t,
$$

where A_t and a_t are to be determined. Substituting (A.2) and the budget constraint into (A.1), we can then rewrite the Bellman equation as

$$
A_t = \max_{a_t, \{\psi_{kt}\}} \Big[(1 - \beta)a_t^{1-\rho}
$$
$$
+ (1 - a_t)^{1-\rho}
$$
$$
\times \beta \big(\mathbb{E}_{\mu_t} \big(\mathbb{E}_{z_{t+1},t}[(A_{t+1}R_{w,t+1})^{1-\gamma}] \big)^{(1-\eta)/(1-\gamma)} \big)^{(1-\rho)/(1-\eta)} \Big]^{1/(1-\rho)}.
$$

Use the first-order condition for consumption to derive

$$
(A.3) \quad \left(\frac{a_t}{1 - a_t} \right)^{-\rho} = \frac{\beta}{1 - \beta} \big(\mathbb{E}_{\mu_t} \big(\mathbb{E}_{z_{t+1},t}[(A_{t+1}R_{w,t+1})^{1-\gamma}] \big)^{(1-\eta)/(1-\gamma)} \big)^{(1-\rho)/(1-\eta)}.
$$

From the above two equations, we derive

$$
(A.4) \quad A_t = (1 - \beta)^{1/(1-\rho)} a_t^{-\rho/(1-\rho)} = (1 - \beta)^{1/(1-\rho)} \left(\frac{C_t}{W_t} \right)^{-\rho/(1-\rho)}.
$$

Substituting equation (A.4) into (A.2) yields equation (18).

REFERENCES

ABEL, A. B. (1999): "Risk Premia and Term Premia in General Equilibrium," *Journal of Monetary Economics*, 43, 3–33. [568,574]

——— (2002): "An Exploration of the Effects of Pessimism and Doubt on Asset Returns," *Journal of Economic Dynamics and Control*, 26, 1075–1092. [562]

ANDERSON, E. W., L. P. HANSEN, AND T. J. SARGENT (2003): "A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection," *Journal of the European Economic Association*, 1, 68–123. [561,566]

BACKUS, D., B. ROUTLEDGE, AND S. ZIN (2004): "Exotic Preferences for Macroeconomists," in *NBER Macroeconomics Annual 2004*, ed. by M. Gertler and K. Rogoff. Cambridge: MIT Press, 319–390. [561]

BANSAL, R., AND A. YARON (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59, 1481–1509. [567,568,572,574,580,581,585,586]

BARILLAS, F., L. P. HANSEN, AND T. J. SARGENT (2009): "Doubts or Variability?" *Journal of Economic Theory*, 144, 2388–2418. [566]

BEELER, J., AND J. Y. CAMPBELL (2009): "The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment," Working Paper, Harvard University. [585,586]

BLACK, F. (1976): "Studies of Stock Price Volatility Changes," in *Proceedings of the Business and Economic Statistics Section*. Washington: American Statistical Association. [573]

BOLLERSLEV, T., R. Y. CHOU, AND K. F. KRONER (1992): "ARCH Modeling in Finance: A Review of the Theory and Empirical Evidence," *Journal of Econometrics*, 52, 5–59. [560,573,582]

BRANDT, M. W., Q. ZENG, AND L. ZHANG (2004): "Equilibrium Stock Return Dynamics Under Alternative Rules of Learning About Hidden States," *Journal of Economic Dynamics and Control*, 28, 1925–1954. [561,562,585]

BRENNAN, M. J., AND Y. XIA (2001): "Stock Price Volatility and the Equity Premium," *Journal of Monetary Economics*, 47, 249–283. [561]

CAGETTI, M., L. P. HANSEN, T. SARGENT, AND N. WILLIAMS (2002): "Robustness and Pricing With Uncertain Growth," *Review of Financial Studies*, 15, 363–404. [561]

CAMERER, C. F. (1999): "Ambiguity-Aversion and Non-Additive Probability: Experimental Evidence, Models and Applications," in *Uncertain Decisions: Bridging Theory and Experiments*, ed. by L. Luini. Norwell, MA: Kluwer Academic Publishers, 53–80. [575]

CAMPBELL, J. Y. (1999): "Asset Prices, Consumption and the Business Cycle," in *Handbook of Macroeconomics*, Vol. 1, ed. by J. B. Taylor and M. Woodford. Amsterdam: Elsevier Science–North-Holland. [568,572,574,576]

CAMPBELL, J. Y., AND J. H. COCHRANE (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205–251. [570,572,574,586]

CAMPBELL, J. Y., AND R. J. SHILLER (1988a): "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1, 195–228. [560,572,573]

——— (1988b): "Stock Prices, Earnings, and Expected Dividends," *Journal of Finance*, 43, 661–676. [560,572,573]

CAMPBELL, J. Y., A. W. LO, AND A. C. MACKINLAY (1997): *The Econometrics of Financial Markets*. Princeton: Princeton University Press. [585]

CAO, H., T. WANG, AND H. ZHANG (2005): "Model Uncertainty, Limited Market Participation, and Asset Prices," *Review of Financial Studies*, 18, 1219–1251. [561]

CECCHETTI, S. G., P. LAM, AND N. C. MARK (1990): "Mean Reversion in Equilibrium Asset Prices," *American Economic Review*, 80, 398–418. [561,567,574,584]

——— (1993): "The Equity Premium and the Riskfree Rate: Matching the Moments," *Journal of Monetary Economics*, 31, 21–45. [567]

——— (2000): "Asset Pricing With Distorted Beliefs: Are Equity Returns Too Good to Be True," *American Economic Review*, 90, 787–805. [559,561,562,567,572-574,584]

CHEN, H., N. JU, AND J. MIAO (2009): "Dynamic Asset Allocation With Ambiguous Return Predictability," Working Paper, Boston University. [575]

CHEN, Z., AND L. G. EPSTEIN (2002): "Ambiguity, Risk and Asset Returns in Continuous Time," *Econometrica*, 70, 1403–1445. [561,562,577]

COLLARD, F., S. MUKERJI, K. SHEPPARD, AND J. TALLON (2009): "Ambiguity and Historical Equity Premium," Working Paper, Oxford University. [567]

DAVID, A. (1997): "Fluctuating Confidence in Stock Markets: Implications for Returns and Volatility," *Journal of Financial and Quantitative Analysis*, 32, 427–462. [561]

DETEMPLE, J. (1986): "Asset Pricing in a Production Economy With Incomplete Information," *Journal of Finance*, 41, 383–391. [561]

DOTHAN, M. U., AND D. FELDMAN (1986): "Equilibrium Interest Rates and Multiperiod Bonds in a Partially Observable Economy," *Journal of Finance*, 41, 369–382. [561]

DUFFIE, D., AND C. SKIADAS (1994): "Continuous-Time Security Pricing: A Utility Gradient Approach," *Journal of Mathematical Economics*, 23, 107–131. [569]

ELLSBERG, D. (1961): "Risk, Ambiguity and the Savage Axiom," *Quarterly Journal of Economics*, 75, 643–669. [560,574]

EPSTEIN, L. G. (1999): "A Definition of Uncertainty Aversion," *Review of Economic Studies*, 66, 579–608. [563]

EPSTEIN, L. G., AND J. MIAO (2003): "A Two-Person Dynamic Equilibrium Under Ambiguity," *Journal of Economic Dynamics and Control*, 27, 1253–1288. [561,570]

EPSTEIN, L. G., AND M. SCHNEIDER (2003): "Recursive Multiple-Priors," *Journal of Economic Theory*, 113, 1–31. [566]

——— (2007): "Learning Under Ambiguity," *Review of Economic Studies*, 74, 1275–1303. [562,566]

——— (2008): "Ambiguity, Information Quality and Asset Pricing," *Journal of Finance*, 63, 197–228. [562]

EPSTEIN, L. G., AND T. WANG (1994): "Intertemporal Asset Pricing Under Knightian Uncertainty," *Econometrica*, 62, 283–322. [561,566]

EPSTEIN, L. G., AND S. ZIN (1989): "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–969. [560,564,565,569,587]

ERGIN, H. I., AND F. GUL (2009): "A Theory of Subjective Compound Lotteries," *Journal of Economic Theory*, 144, 899–929. [562]

FAMA, E. F., AND K. R. FRENCH (1988a): "Dividend Yields and Expected Stock Returns," *Journal of Financial Economics*, 22, 3–25. [560,572]

——— (1988b): "Permanent and Temporary Components of Stock Prices," *Journal of Political Economy*, 96, 246–273. [560,572]

——— (1989): "Business Conditions and Expected Returns on Stocks and Bonds," *Journal of Financial Economics*, 25, 23–49. [560,573]

GARCIA, R., N. MEDDAHI, AND R. TEDONGAP (2008): "An Analytical Framework for Assessing Asset Pricing Models and Predictability," Working Paper, Edhec Business School. [568,585]

GARLAPPI, L., R. UPPAL, AND T. WANG (2007): "Portfolio Selection With Parameter and Model Uncertainty: A Multi-Prior Approach," *Review of Financial Studies*, 20, 41–81. [561]

GHIRARDATO, P., AND M. MARINACCI (2002): "Ambiguity Made Precise: A Comparative Foundation," *Journal of Economic Theory*, 102, 251–289. [563]

GILBOA, I., AND D. SCHMEIDLER (1989): "Maxmin Expected Utility With Non-Unique Priors," *Journal of Mathematical Economics*, 18, 141–153. [561,563]

GOLLIER, C. (2011): "Portfolio Choices and Asset Prices: The Comparative Statics of Ambiguity Aversion," *Review of Economic Studies*, 78, 1329–1344. [570]

HALEVY, Y. (2007): "Ellsberg Revisited: An Experimental Study," *Econometrica*, 75, 503–536. [563,575]

HAMILTON, J. D. (1989): "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 57, 357–384. [561,564]

HANSEN, L. P. (2007): "Beliefs, Doubts and Learning: The Valuation of Macroeconomic Risk," *American Economic Review*, 97, 1–30. [560-563,565-568]

HANSEN, L. P., AND R. JAGANNATHAN (1991): "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy*, 99, 225–262. [577]

HANSEN, L. P., AND T. J. SARGENT (2001): "Robust Control and Model Uncertainty," *American Economic Review*, 91, 60–66. [560,561,566]

——— (2007): "Recursive Robust Estimation and Control Without Commitment," *Journal of Economic Theory*, 136, 1–27. [560,562,566,567]

——— (2008): *Robustness*. Princeton: Princeton University Press. [560,561]

——— (2010): "Fragile Beliefs and the Price of Uncertainty," *Quantitative Economics*, 1, 129–162. [560,562,566]

HANSEN, L. P., J. C. HEATON, AND N. LI (2008): "Consumption Strikes Back?: Measuring Long-Run Risk," *Journal of Political Economy*, 116, 260–302. [569,586]

HANSEN, L. P., T. J. SARGENT, AND T. TALLARINI (1999): "Robust Permanent Income and Pricing," *Review of Economic Studies*, 66, 873–907. [561]

HAYASHI, T., AND J. MIAO (2011): "Intertemporal Substitution and Recursive Smooth Ambiguity Preferences," *Theoretical Economics*, 6, 423–475. [562]

JU, N., AND J. MIAO (2007): "Ambiguity, Learning, and Asset Returns," Working Paper, Boston University. [566,567,570,571]

——— (2012): "Supplement to 'Ambiguity, Learning, and Asset Returns': Technical Appendix," *Econometrica Supplemental Material*, 80, http://www.econometricsociety.org/ecta/Supmat/7618_miscellaneous.pdf; http://www.econometricsociety.org/ecta/Supmat/7618_data and programs.zip. [572]

JUDD, K. (1998): *Numerical Methods in Economics*. Cambridge, MA: MIT Press. [572]

KANDEL, S., AND R. F. STAMBAUGH (1991): "Asset Returns and Intertemporal Preferences," *Journal of Monetary Economics*, 27, 39–71. [567]

KEYNES, M. (1936): *The General Theory of Employment, Interest, and Money*. London: Macmillan. [560]

KLIBANOFF, P., M. MARINACCI, AND S. MUKERJI (2005): "A Smooth Model of Decision Making Under Ambiguity," *Econometrica*, 73, 1849–1892. [562,563,566]

——— (2009): "Recursive Smooth Ambiguity Preferences," *Journal of Economic Theory*, 144, 930–976. [560,562,566]

KNIGHT, F. (1921): *Risk, Uncertainty, and Profit*. Boston: Houghton Mifflin. [560]

KREPS, D. M., AND E. L. PORTEUS (1978): "Temporal Resolution of Uncertainty and Dynamic Choice," *Econometrica*, 46, 185–200. [564,565]

LEIPPOLD, M., F. TROJANI, AND P. VANINI (2008): "Learning and Asset Prices Under Ambiguous Information," *Review of Financial Studies*, 21, 2565–2597. [562,571]

LIU, J., J. PAN, AND T. WANG (2005): "An Equilibrium Model of Rare Event Premia and Its Implication for Option Smirks," *Review of Financial Studies*, 18, 131–164. [561]

LJUNQVIST, L., AND H. UHLIG (2009): "Optimal Endowment Destruction Under Campbell–Cochrane Habit Formation," Working Paper 14772, NBER. [586]

LUCAS, R. E. (1978): "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429–1445. [567]

MACCHERONI, F., M. MARINACCI, AND A. RUSTICHINI (2006): "Ambiguity Aversion, Robustness and the Variational Representation of Preferences," *Econometrica*, 74, 1447–1498. [561]

MAENHOUT, P. J. (2004): "Robust Portfolio Rules and Asset Pricing," *Review of Financial Studies*, 17, 951–983. [561]

MEHRA, R., AND E. C. PRESCOTT (1985): "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15, 145–161. [559,572]

MIAO, J. (2009): "Ambiguity, Risk and Portfolio Choice Under Incomplete Information," *Annals of Economics and Finance*, 10, 257–279. [561]

NAU, R. F. (2006): "Uncertainty Aversion With Second-Order Utilities and Probabilities," *Management Science*, 52, 136–145. [562]

POTERBA, J. M., AND L. H. SUMMERS (1988): "Mean Reversion in Stock Prices: Evidence and Implications," *Journal of Financial Economics*, 22, 27–59. [560,572]

ROUTLEDGE, B., AND S. ZIN (2009): "Model Uncertainty and Liquidity," *Journal: Review of Economic Dynamics*, 12, 543–566. [561]

SEGAL, U. (1987): "The Ellsberg Paradox and Risk Aversion: An Anticipated Utility Approach," *International Economic Review*, 28, 175–202. [563]

SEO, K. (2009): "Ambiguity and Second-Order Belief," *Econometrica*, 77, 1575–1605. [562,563]

SHILLER, R. J. (1981): "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?" *American Economic Review*, 71, 421–436. [559]

STRZALECKI, T. (2009): "Temporal Resolution of Uncertainty and Recursive Models of Ambiguity Aversion," Working Paper, Harvard University. [565]

TALLARINI, T. D., JR. (2000): "Risk-Sensitive Real Business Cycles," *Journal of Monetary Economics*, 45, 507–532. [560]

UPPAL, R., AND T. WANG (2003): "Model Misspecification and Underdiversification," *Journal of Finance*, 58, 2465–2486. [561]

VERONESI, P. (1999): "Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model," *Review of Financial Studies*, 12, 975–1007. [561,572,579,580]

——— (2000): "How Does Information Quality Affect Stock Returns?" *Journal of Finance*, 55, 807–837. [561,568,572,579]

WEIL, P. (1989): "The Equity Premium Puzzle and the Risk-Free Rate Puzzle," *Journal of Monetary Economics*, 24, 401–421. [559,565,569]

WEITZMAN, M. L. (2007): "Subjective Expectations and Asset-Return Puzzles," *American Economic Review*, 97, 1102–1130. [561]

WELCH, I., AND A. GOYAL (2008): "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction," *Review of Financial Studies*, 21, 1455–1508. [585]

Dept. of Finance, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong and SAIF, Shanghai Jiaotong University, China; nengjiu@ust.hk

and

Dept. of Economics, Boston University, 270 Bay State Road, Boston, MA 02215, U.S.A. and CEMA, Central University of Finance and Economics and AFR, Zhejiang University, China; miaoj@bu.edu.