Asset Pricing with Learning

Daniel Andrei* Michael Hasler† Alexandre Jeanneret‡

April 7, 2017

Abstract

The long-run risk literature assumes a highly persistent output growth to explain asset pricing puzzles. This assumption has been challenged. To help address this concern, we consider an equilibrium model in which the degree of persistence is unobservable, and let the data dictate its level and time-variation. Learning about persistence generates high and time-varying equity risk premia, volatility, and Sharpe ratio, although the degree of persistence is actually low. In contrast, these moments remain constant with a known degree of persistence or with learning about the level of output growth. The data lend support to the predictions of the model.

We are grateful to Michael Brennan for helpful comments, which have led to substantial improvements in the paper. We would like to thank Pat Akey, Jerome Detemple, Ian Dew-Becker (HEC-McGill Winter Finance Workshop discussant), Chanik Jo, Mariana Khapko, Gang Li, Chayawat Ornthanalai, Marcel Rindisbacher, and Gustavo Schwenkler for their suggestions. We also thank seminar/conference participants at Boston University, HEC-McGill Winter Finance Workshop, University of Technology Sydney, University of New South Wales, University of Sydney, and UCLA for comments and suggestions. We thank UCLA, the University of Toronto, and HEC Montréal for financial support. Hasler and Jeanneret are particularly grateful to the Montreal Institute of Structured Finance and Derivatives (IFSID) for its generous financial support.

*UCLA Anderson Graduate School of Management, 110 Westwood Plaza, Suite C420, Los Angeles, CA 90095, USA; +1 (310) 825-3544; daniel.andrei@anderson.ucla.edu; www.danielandrei.info.

†University of Toronto, Rotman School of Management, 105 St. George Street, Suite 431, Toronto, ON M5S 3E6, Canada, michael.hasler@rotman.utoronto.ca; www.rotman.utoronto.ca/faculty/hasler.

‡Department of Finance, HEC Montréal, 3000 Côte-Sainte-Catherine, Montréal, QC H3T 2A7, Canada; +1 (514) 589-6699; alexandre.jeanneret@hec.ca; www.alexandrejeanneret.com.

1 Introduction

The degree of persistence of the endowment process in asset pricing models has raised a debate. Bansal and Yaron (2004) argue that a high degree of persistence is important for the empirical performance of the long-run risk model. However, Beeler and Campbell (2012) reject this calibration with postwar consumption data.[1] Using data on S&P 500 options, van Binsbergen, Brandt, and Koijen (2012) provide evidence that equity risk premia and volatility decrease with maturity, at odds with the long-run risk model.[2] In addition, Epstein, Farhi, and Strzalecki (2014) show that the high degree of persistence required in Bansal and Yaron (2004) would inflate timing premia to seemingly implausible levels.

We propose a solution to this debate. We consider an equilibrium model in which the degree of persistence is unobservable, and let the data determine its level and time variation. Uncertainty about persistence helps the model explain the observed properties of asset pricing moments. The level and the dynamics of the equity risk premium, volatility, and Sharpe ratio are all in line with the data. Yet, the degree of persistence estimated using U.S. output growth data is actually low, but significantly time varying.

Our model considers a representative agent in an economy with imperfect information. The agent has Epstein-Zin preferences and favors early resolution of uncertainty. The expected growth of the output process mean-reverts, but the parameter that controls the speed of mean-reversion is unobservable. We also assume that the level of expected growth is unobservable, following the existing incomplete information literature.[3] The agent updates her beliefs about the level and the degree of persistence simultaneously. Both forms of learning are thus present in the economy. This general specification allows us to compare the implications of these two forms of learning. We show that all asset pricing moments are constant with learning about the level of expected growth, which is inconsistent with the data. Indeed, empirical findings show that the equity risk premium, volatility, and the Sharpe ratio are time-varying.[4] In contrast, learning about persistence generates significant time-variation in all asset pricing moments.

The theoretical mechanism is straightforward. When persistence is unobservable, rational learning generates an asymmetric effect on asset prices. In recessions, bad news about economic growth induces the agent to infer that there is more persistence. Higher persistence

[1]The same holds for output growth, dividends growth, or earnings growth. See also Belo, Collin-Dufresne, and Goldstein (2015).

[2]See also van Binsbergen, Hueskes, Koijen, and Vrugt (2013), van Binsbergen and Koijen (2016), and van Binsbergen and Koijen (2017).

[3]See David (1997), Veronesi (1999, 2000), Brennan and Xia (2001), David and Veronesi (2002) among many others, and excellent surveys by Pastor and Veronesi (2009) and Ziegler (2012).

[4]Schwert (1989); Ferson and Harvey (1991); Lettau and Ludvigson (2010); Lustig and Verdelhan (2012).

and a preference for early resolution of uncertainty makes the bad news worse. Similarly, good news imply less persistence, which makes the good news even better. The logic reverts during expansions. By decreasing the perceived level of persistence, bad news are not so bad, whereas good news increase persistence and are consequently not that good. Learning about persistence thus amplifies the response of asset prices to news during recessions, whereas it dampens the response in expansions. This asymmetry arises without assuming any exogenous fluctuations in economic uncertainty.[5]

We fit the model to real GDP growth and analyst forecast data over the period Q4:1968 to Q4:2015 by Maximum Likelihood (Hamilton, 1994). The estimation shows that the persistence is low and varies significantly over time. In comparison, the discrepancy between the expected output growth and the mean analyst forecast is zero, on average, and only weakly time varying. We also estimate two special cases of our general specification: i) a model with uncertainty about the level of expected growth only; and ii) a model with uncertainty about persistence only. The results show that the second model matches the U.S. data best. This finding may not be surprising, as survey forecasts available for the level of expected output growth tend to be accurate (e.g. Ang, Bekaert, and Wei, 2007). By contrast, we are not aware of any available survey forecast about the persistence of output growth. As a result, agents have an incentive to learn about persistence and this learning form appears to be fundamental to explain the dynamics of output growth and analyst forecast data.

We empirically test the asset-pricing predictions of our model with learning about persistence. First, we build model-implied time series for stock return volatility, the risk premium, and the Sharpe ratio over the Q4:1968–Q4:2015 period. As a preliminary test, we compare these model-implied quantities with their observed S&P 500 counterparts. The model-implied and observed time series are strongly and positively correlated, thereby indicating that our model of learning about persistence generates dynamics of asset pricing quantities that are in line with the data. Second, we directly test the predictions of our theory by regressing the observed asset prices on the model's state variables. In presence of uncertainty about persistence, all asset pricing quantities are expected to vary negatively with the business cycle and positively with the degree of persistence. Our empirical analysis provides empirical support for this prediction. Therefore, we are able to explain stylized facts in asset pricing with a low, but time-varying degree of persistence. Uncertainty about persistence thus helps address the critiques of the long-run risk model (Beeler and Campbell, 2012).

[5]The empirical importance of fluctuating economic uncertainty is emphasized in the Case II of Bansal and Yaron (2004), and by Bansal, Kiku, and Yaron (2012), Beeler and Campbell (2012), and Campbell, Giglio, Polk, and Turley (2017).

Another implication of the low persistence estimate is that the term structures of dividend strip volatility and risk premia implied by our model are almost flat. In contrast, they are markedly upward-sloping in the traditional long-run risk model (Bansal and Yaron, 2004), in the external habit formation model (Campbell and Cochrane, 1999), and in the time-varying risk of rare disaster model (Wachter, 2013). That is, our model helps to partially address the concerns raised by van Binsbergen et al. (2012, 2013) on the timing of volatility and risk premia implied by leading asset pricing models.

Collin-Dufresne, Johannes, and Lochstoer (2016) show that parameter uncertainty generates endogenous long-run risk. This implies a large equilibrium risk premium when the representative agent has a preference for early resolution of uncertainty. Our paper builds on this result. More precisely, we document an asymmetric response of asset prices to news, which arises when the uncertain parameter is the speed of mean-reversion, and we provide empirical support for this asymmetric response. Bidder and Dew-Becker (2016) study parameter uncertainty in a model with ambiguity aversion. Investors fear a "worst-case" model in which shocks to the consumption growth trend have a half-life of 70 years, which is more extreme than Bansal and Yaron (2004).[6] We offer an alternative view, in which the persistence is a latent state that needs to be filtered out and the representative agent features standard recursive preferences. Estimating this latent state using U.S. output growth data shows that the degree of persistence is much weaker than what is assumed in the long-run risk literature.

Our model assumes away any variation in the volatility of output growth. We do not assume time-varying disaster risk (Gabaix, 2012; Wachter, 2013). Nor do we consider any utility specification that induces time-varying preferences (Campbell and Cochrane, 1999). Yet, in our model, all asset pricing moments fluctuate endogenously over the business cycle, due to investors' learning about persistence.

Most studies in the literature on learning in financial markets assume that the unobservable dimension is the level of expected output growth.[7] None of our implications about the dynamics of asset prices obtain with this form of learning. Pakoš (2013) analyzes an economy where growth is a three-state Markov chain and in which a representative agent cannot distinguish between a mild recession and a "lost decade." This modeling choice exogenously

[6]In Bansal and Yaron (2004) the half-life of shocks is about three years.

[7]In Veronesi (1999), a representative agent learns about a discrete-state output growth trend. Volatility and risk premia are hump-shaped functions of the state of the economy. Ai (2010) analyzes learning in a production economy with Epstein-Zin preferences. The relation between information quality and the risk premium is negative. Croce, Lettau, and Ludvigson (2015) build a bounded rationality limited information model. They obtain a large risk premium and a downward-sloping term structure of risk. Johannes, Lochstoer, and Mou (2016) build an economy in which the agent has anticipated utility (Kreps, 1998; Cogley and Sargent, 2008). In their model, parameter uncertainty is not a priced risk factor.

introduces an asymmetry and a stronger response to news in bad times. In our case, the agent learns about persistence at all times, and the asymmetry arises endogenously.

Andrei, Carlin, and Hasler (2017) build an equilibrium model in which CRRA agents disagree about the length of business cycles. Disagreement about the length of business cycles yields time variation in the risk premium and dictates the strength of the risk-return relationship. In our model, a representative agent who prefers early resolution of uncertainty learns about both the persistence and the level of expected output growth. We show that a model of learning about persistence better fits U.S. output data than a model of learning about the level. Learning about persistence generates time-varying patterns of equity return volatility, risk premium, and Sharpe ratio that are supported by the data. Furthermore, we reconcile the debate about the high degree of persistence required by long-run risk models to solve asset pricing puzzles. Indeed, we show that the aforementioned moments are high, although the estimated persistence is actually low, precisely because the agent faces uncertainty about the degree of persistence and prefers early resolution of uncertainty.

The paper proceeds as follows. Section 2 introduces a model with learning about the persistence and the level of expected output growth. Section 3 calibrates the model and presents our theoretical predictions. Section 4 tests our main predictions. Section 5 concludes and offers directions for future research.

2 Model

In this section we first introduce the economic environment. We describe how an agent with recursive utility learns about both the level and the persistence of the expected output growth. We then solve the agent's learning problem and characterize equilibrium asset prices.

2.1 Environment

The economy is defined over a continuous-time horizon $[0, \infty)$. A representative agent derives utility from consumption. The agent has stochastic differential utility (Epstein and Zin, 1989) with subjective time preference rate β, relative risk aversion γ, and elasticity of intertemporal substitution ψ. The indirect utility function is given by

$$J_t = \mathbb{E}_t \left[\int_t^{\infty} h(C_s, J_s) ds \right], \tag{1}$$

where the aggregator h is defined as in Duffie and Epstein (1992):

$$h(C, J) = \frac{\beta}{1 - 1/\psi} \left(\frac{C^{1-1/\psi}}{[(1-\gamma)J]^{1/\phi - 1}} - (1-\gamma)J \right), \tag{2}$$

with $\phi \equiv \frac{1-\gamma}{1-1/\psi}$. Standard CRRA utility obtains if $\phi = 1$. Preference for early resolution of uncertainty obtains if $1 - \phi > 1$.

The agent can invest in a risk-free asset and a risky asset (the *stock*). The stock is a claim to the output stream:

$$\frac{d\delta_t}{\delta_t} = \mu_t dt + \sigma_\delta dW_t^\delta, \tag{3}$$

where W_t^δ is a standard Brownian motion.

The expected output growth rate, μ_t, is unobservable. The history of the output process (3) provides a signal about the expected growth rate. In addition to this signal, the agent continuously receives a forecast of μ_t from professional forecasters. The agent is aware that the forecast may be imperfect, but does not know the level of adjustment. The expected growth rate is then given by

$$\mu_t = f_t + l_t, \tag{4}$$

where f_t is the observed forecast and l_t is an adjustment in the forecasters' prediction. The expected output growth rate fluctuates around a long-term mean $\bar{f}$ with an unobservable mean-reversion parameter λ_t.

The forecast of expected output growth, f_t, its unobservable adjustment, l_t, and its unobservable mean-reversion speed, λ_t, follow

$$df_t = \lambda_t(\bar{f} - f_t)dt + \sigma_f dW_t^f, \tag{5}$$

$$dl_t = \sigma_l dW_t^l, \tag{6}$$

$$d\lambda_t = \sigma_\lambda dW_t^\lambda, \tag{7}$$

where the Brownians W_t^δ, W_t^f, W_t^l, and W_t^λ are independent and the parameters $\bar{f}$, σ_f, σ_l, and σ_λ are known/observed constants.[8]

The dynamics in (6) and (7) show that the adjustment l_t and the mean-reversion speed λ_t are unobservable constants perturbed by noise. We can alternatively assume that l_t and

[8] The assumption of independent shocks simplifies the description of the model without changing the main message. The model can be easily extended to allow for non-zero correlations between the four Brownians, but at the cost of making the estimation of the parameters (described in Section 3.1) significantly less stable.

λ_t are unobservable constants. Their posterior estimates would remain martingales through Bayesian learning (as they do in the next section), and all the asset pricing implications would hold (Collin-Dufresne et al., 2016).[9]

The economic environment described in (3)–(7) embeds two dimensions of uncertainty in a unified framework—uncertainty about the *level* of the expected growth rate and uncertainty about its *degree of persistence*. We will study separately two polar cases. First, fixing $\lambda_t = \bar{\lambda}\ \forall t$, where $\bar{\lambda}$ is a known constant, implies that the agent has perfect knowledge about the persistence of expected growth but faces uncertainty about its level. We call this case *learning about level*. Second, fixing $l_t = 0\ \forall t$ implies that the agent has perfect knowledge about the expected output growth but faces uncertainty about its persistence. We call this case *learning about persistence*. Finally, when both the adjustment l_t and the mean-reversion speed λ_t are unobservable, the agent learns simultaneously about the level of economic growth and about its persistence.[10]

2.2 Learning

The agent starts with the following prior distribution about the unobservable processes l_t and λ_t

$$\begin{bmatrix} l_0 \\ \lambda_0 \end{bmatrix} \sim N\left(\begin{bmatrix} \widehat{l}_0 \\ \widehat{\lambda}_0 \end{bmatrix}, \begin{bmatrix} \nu_{l,0} & 0 \\ 0 & \nu_{\lambda,0} \end{bmatrix}\right). \tag{8}$$

Define $\mathcal{F}_t$ the information set of the agent at time t, and denote by $\widehat{l}_t \equiv \mathbb{E}[l_t|\mathcal{F}_t]$ the estimated adjustment l_t and its posterior variance by $\nu_{l,t} \equiv \mathbb{E}[(l_t - \widehat{l}_t)^2|\mathcal{F}_t]$. Similarly, denote by $\widehat{\lambda}_t \equiv \mathbb{E}[\lambda_t|\mathcal{F}_t]$ the estimated mean-reversion speed λ_t and its posterior variance by $\nu_{\lambda,t} \equiv \mathbb{E}[(\lambda_t - \widehat{\lambda}_t)^2|\mathcal{F}_t]$. These estimates and posterior variances are such that

$$l_t \sim N(\widehat{l}_t, \nu_{l,t}), \qquad \lambda_t \sim N(\widehat{\lambda}_t, \nu_{\lambda,t}), \tag{9}$$

where $N(m, v)$ denotes the Normal distribution with mean m and variance v.

We denote the estimates $\widehat{l}_t$ and $\widehat{\lambda}_t$ as the *filters*, and the two posterior variances $\nu_{l,t}$ and $\nu_{\lambda,t}$ as the *uncertainties*. Following Liptser and Shiryayev (1977), the filters evolve according

[9] With constant l_t and λ_t, agent's learning would improve as time goes by, thereby changing our results quantitatively but not qualitatively. Alternatively, we can assume mean-reverting processes for l_t and λ_t. Such a specification would introduce additional parameters and thus decrease the stability of the parameter estimation. Notice, however, that the theoretical results are robust to this alternative specification.

[10] The feature of the model that preserves the linearity of the learning exercise in this latter case is the observability of the forecast f_t and its long-term mean $\bar{f}$.

to (see Appendix A):

$$\begin{bmatrix} d\widehat{l}_t \\ d\widehat{\lambda}_t \end{bmatrix} = \begin{bmatrix} \frac{\nu_{l,t}}{\sigma_\delta} & 0 \\ 0 & \frac{(\bar{f}-f_t)\nu_{\lambda,t}}{\sigma_f} \end{bmatrix} \begin{bmatrix} d\widehat{W}_t^\delta \\ d\widehat{W}_t^f \end{bmatrix}, \tag{10}$$

where

$$d\widehat{W}_t^\delta = \frac{1}{\sigma_\delta}\left(\frac{d\delta_t}{\delta_t} - [f_t + \widehat{l}_t]dt\right), \qquad d\widehat{W}_t^f = \frac{1}{\sigma_f}\left(df_t - \widehat{\lambda}_t(\bar{f} - f_t)dt\right), \tag{11}$$

are independent Brownian motions under the filtration $\mathcal{F}_t$. For clarity, we will hereafter use the term *output growth shocks* to refer to $d\widehat{W}_t^\delta$ innovations and the term *expected output growth shocks* to refer to $d\widehat{W}_t^f$ innovations.

The filter $\widehat{l}_t$ is perfectly and positively correlated with output δ. This implies that after positive (negative) output shocks, the agent revises the estimate of the expected growth rate upwards (downwards) (Brennan, 1998).

Learning about persistence induces a particular formation of beliefs. In this case, learning depends on the state of the economy, defined by the distance between the long-run growth forecast and the actual growth forecast, $(\bar{f} - f_t)$. We refer to this distance as the *output growth gap*. When the output growth gap is positive, the economy is in bad times. When the output growth gap is negative, the economy is in good times.

In good times, positive expected output growth shocks decrease the agent's estimate of λ_t. In bad times, negative expected output growth shocks decrease the agent's estimate of λ_t. In both situations—positive shocks in good times or negative shocks in bad times—the agent extrapolates that expected output growth becomes more persistent (i.e., lower mean-reversion speed λ_t). When $f_t = \bar{f}$ the agent is unable to learn about the mean-reversion speed. In this case, changes in the forecast f_t are uninformative about λ_t.

The dynamics of the uncertainties about l_t and λ_t are respectively given by

$$d\nu_{l,t} = \left[\sigma_l^2 - \frac{\nu_{l,t}^2}{\sigma_\delta^2}\right] dt, \tag{12}$$

$$d\nu_{\lambda,t} = \left[\sigma_\lambda^2 - \frac{(\bar{f} - f_t)^2\nu_{\lambda,t}^2}{\sigma_f^2}\right] dt. \tag{13}$$

The dynamics (12) imply that uncertainty about l_t converges to a constant, $\bar{\nu}_l$. We consequently follow Dumas, Kurshev, and Uppal (2009) and assume that uncertainty about

l_t has converged to its steady-state:[11]

$$\nu_{l,t} = \bar{\nu}_l = \sigma_\delta \sigma_l, \ \forall t. \tag{14}$$

The dynamics (13) imply that uncertainty about λ_t does not converge to a constant, due to the presence of the stochastic term $\bar{f} - f_t$. There are two terms in the dynamics of $\nu_{\lambda,t}$. The first term is the increase in uncertainty due to the variability of λ_t. The second term is the reduction in uncertainty due to learning. The magnitude of this second term depends on the output growth gap. A sizable output growth gap (positive or negative) makes changes in forecasts particularly informative about the mean-reversion speed.

It is worth noting that none of the uncertainties $\nu_{l,t}$ and $\nu_{\lambda,t}$ converge to zero. This is because l_t and λ_t are perturbed by noise, as opposed to being constant, which continuously regenerates learning. It implies that the dynamics of all state variables are non-degenerate. In contrast, if we assumed that l_t and λ_t were constants, both uncertainties would converge to zero in the long run.

2.3 Equilibrium

Solving for the equilibrium is standard (see Appendix B). It involves writing the HJB equation for problem (1):

$$\max_C \{h(C, J) + \mathcal{L}J\} = 0, \tag{15}$$

with the differential operator $\mathcal{L}J$ following from Itô's lemma. We guess the following value function (Benzoni, Collin-Dufresne, and Goldstein, 2011):

$$J(C, f, \widehat{l}, \widehat{\lambda}, \nu_\lambda) = \frac{C^{1-\gamma}}{1-\gamma} \left[\beta I(x)\right]^\phi, \tag{16}$$

where $I(x)$ is the price-dividend ratio, and $x \equiv [f \ \widehat{l} \ \widehat{\lambda} \ \nu_\lambda]^\top$ denotes the vector of four state variables, whose dynamics are given in (10)–(13).

Substituting the guess (16) in the HJB equation (15) and imposing the market clearing condition, $C_t = \delta_t$, yields a partial differential equation for the price-dividend ratio. We solve for this equation numerically using Chebyshev polynomials (Judd, 1998). The general PDE is presented in Equation (65) in Appendix B, while the PDEs corresponding to the two polar forms of learning are provided respectively in Equations (66) and (67).

[11] The steady-state uncertainty about l_t satisfies $\left.\frac{d\nu_{l,t}}{dt}\right|_{\nu_{l,t} \equiv \bar{\nu}_l} = 0$, which yields $\bar{\nu}_l = \sigma_\delta \sigma_l$. See Appendix A.

In order to characterize the effects of learning on equilibrium outcomes, we make the following conjecture.

Conjecture 1. *When $\gamma > 1 > 1/\psi$, we expect the partial derivatives of the price-dividend ratio with respect to the state variables to satisfy:*

$$I_f > 0, \; I_{\widehat{l}} > 0, \; I_{\widehat{\lambda}} > 0, \; I_{\nu_\lambda} < 0. \tag{17}$$

This conjecture holds for a wide range of parameters. In fact, several inequalities follow directly from the guess of the value function in (16). Taking the derivative of J with respect to any of the four state variables yields

$$J_{(\cdot)} = \phi J \frac{I_{(\cdot)}}{I}, \tag{18}$$

with the product ϕJ being positive when $\gamma > 1 > 1/\psi$. Due to non-satiation, expected lifetime utility must rise as investment opportunities improve, i.e., $J_f > 0$ and $J_{\widehat{l}} > 0$. Using (18), this reasoning yields the first two inequalities of Conjecture 1. Further, a risk-averse agent dislikes uncertainty, which implies that $J_{\nu_\lambda} < 0$. This yields the last inequality of Conjecture 1. The only inequality that needs numerical validation is $I_{\widehat{\lambda}} > 0$. Because the agent prefers early resolution of uncertainty, we expect that she prefers less persistence (i.e. higher mean-reversion speed), which yields $J_{\widehat{\lambda}} > 0$.[12]

Let $\sigma_I(x) \equiv [\sigma_{I1}(x) \; \sigma_{I2}(x)]$ be the diffusion vector of the price-dividend ratio. It has two elements, each of which loads on the Brownian motions defined in (11):

$$\sigma_{I1}(x_t) = \frac{\bar{\nu}_l}{\sigma_\delta} \frac{I_{\widehat{l}}}{I} \tag{19}$$

$$\sigma_{I2}(x_t) = \sigma_f \frac{I_f}{I} + \frac{(\bar{f} - f_t)\nu_{\lambda,t}}{\sigma_f} \frac{I_{\widehat{\lambda}}}{I}. \tag{20}$$

2.3.1 Risk-free rate and market price of risk

Following Duffie and Epstein (1992), the state-price density satisfies

$$\xi_t = \exp\left[\int_0^t h_J(C_s, J_s) ds\right] h_C(C_t, J_t) = \exp\left[\int_0^t \left(\frac{\phi - 1}{I(x_s)} - \beta\phi\right) ds\right] \beta^\phi C_t^{-\gamma} I(x_t)^{\phi - 1}. \tag{21}$$

[12] Appendix D reports a numerical analysis of $I_{\widehat{\lambda}}$ using a wide range of values for risk aversion, intertemporal of elasticity substitution, and the two state variables f_t and $\widehat{\lambda}_t$. We find that $I_{\widehat{\lambda}}$ is positive and large in all cases. We also elaborate on the parametrization for which this term can become negative (which can only happen outside the standard calibration of our model).

The dynamics of the state-price density follow

$$\frac{d\xi_t}{\xi_t} = -r_t dt - m_t^\top d\widehat{W}_t, \tag{22}$$

where r_t is the risk-free rate, m_t is the 2-dimensional market price of risk, and $\widehat{W} \equiv [\widehat{W}^\delta, \widehat{W}^f]^\top$ is the 2-dimensional standard Brownian motion defined in (11).

Applying Itô's lemma on the state-price density ξ_t defined in (21) yields the risk-free rate and the market price of risk:

$$r_t = \beta + \frac{1}{\psi}[f_t + \widehat{l}_t] - \frac{\gamma + \gamma\psi}{2\psi}\sigma_\delta^2 - (1-\phi)\left[\sigma_{I1}(x_t)\sigma_\delta + \frac{1}{2}\left(\sigma_{I1}^2(x_t) + \sigma_{I2}^2(x_t)\right)\right], \tag{23}$$

$$m_t = \begin{bmatrix} \gamma\sigma_\delta + (1-\phi)\sigma_{I1}(x_t) & (1-\phi)\sigma_{I2}(x_t) \end{bmatrix}^\top. \tag{24}$$

The second term of the equilibrium risk-free rate in (23) indicates that fluctuations in expected output growth generate a procyclical risk-free rate. Furthermore, when the agent prefers early resolution of uncertainty ($1-\phi > 0$), the risk-free rate contains an additional term due to variations in f, $\widehat{l}$, and $\widehat{\lambda}$. The resulting effect is a lower risk-free rate due to greater demand for the safe asset.

The market price of risk contains two elements, one for each of the two Brownians in the economy. The uncertainty about the expected growth rate, $\bar{\nu}_l$, increases the first component when the agent learns about the level of expected growth. The impact of learning about persistence is present in the second component. As Equation (20) shows, this component depends on the output growth gap. Following Conjecture 1, $I_{\widehat{\lambda}} > 0$ and the market price of risk thus increases in bad times and decreases in good times.

2.3.2 Stock market volatility

The diffusion of stock returns, σ_t, satisfies

$$\sigma_t = \begin{bmatrix} \sigma_\delta + \sigma_{I1}(x_t) & \sigma_{I2}(x_t) \end{bmatrix}, \tag{25}$$

which, after replacing (19)–(20), can be written as

$$\sigma_t = \begin{bmatrix} \sigma_\delta + \frac{\bar{\nu}_l}{\sigma_\delta}\frac{I_{\widehat{l}}}{I} & \sigma_f\frac{I_f}{I} + \frac{(\bar{f}-f_t)\nu_{\lambda,t}}{\sigma_f}\frac{I_{\widehat{\lambda}}}{I} \end{bmatrix}. \tag{26}$$

According to Conjecture 1, $I_{\widehat{l}} > 0$ and thus uncertainty about the level $\bar{\nu}_l$ increases the magnitude of the first diffusion component in (26). Learning about the level in output growth generates excess volatility in stock returns.

Learning about persistence creates an asymmetric stock market response to shocks. This is due to the presence of the output growth gap in the second diffusion component in (26). Consider a negative expected output growth shock, $d\widehat{W}_t^f < 0$. In bad times, the agent updates that economic growth is more persistent. The stock price drops not only because of the negative shock, but also because of a higher degree of persistence. The same intuition holds for a positive expected output growth shock in bad times. Therefore, stock returns react strongly to shocks when the economy is in bad times. In contrast, consider a negative expected output growth shock in good times. The agent perceives less persistence. This mitigates the initial effect of the negative shock. The same intuition holds for a positive expected output growth shock in good times. Therefore, learning about persistence attenuates (amplifies) the stock price response to shocks in good (bad) times.

Learning about level and learning about persistence have distinct effects on volatility. Learning about level helps generate a high level of stock return volatility. But in this case volatility does not vary with the business cycle.[13] In comparison, learning about persistence creates an asymmetric reaction of stock returns to news. The reaction is stronger in bad times, when the output growth gap is positive.

2.3.3 Equity risk premium

The risk premium in this economy is defined as $\mu_t - r_t = \sigma_t m_t$. Using expressions (24) and (26), we obtain

$$\mu_t - r_t = \left(\sigma_\delta + \frac{\bar{\nu}_l}{\sigma_\delta}\frac{I_{\hat{l}}}{I}\right)\left(\gamma\sigma_\delta + (1-\phi)\frac{\bar{\nu}_l}{\sigma_\delta}\frac{I_{\hat{l}}}{I}\right) + (1-\phi)\left(\sigma_f\frac{I_f}{I} + \frac{(\bar{f}-f_t)\nu_{\lambda,t}}{\sigma_f}\frac{I_{\hat{\lambda}}}{I}\right)^2. \quad (27)$$

The equity risk premium consists of two terms. Preference for early resolution of uncertainty $(1-\phi > 0)$ affects both terms. As shown by Collin-Dufresne et al. (2016), the risk premium increases when parameters are uncertain. Uncertainty about the level of expected growth enters in the first term. Uncertainty about persistence enters in the second term. The effect of uncertainty about persistence on the risk premium depends on a quadratic function in the output growth gap. This quadratic form is smallest when the growth forecast f_t is well above the long-term growth $\bar{f}$.[14] The risk premium is then higher when the economy is

[13] There are only two (indirect) channels through which volatility could potentially depend on f_t with learning about level: through the terms $I_{\hat{l}}/I$ and I_f/I in (26). With our calibration, the effect generated by these terms is negligible (more about this in Appendix E); in fact, this effect is zero with a standard log-linear approximation (Bansal and Yaron, 2004).

[14] Using a log-linear approximation of the price-dividend ratio, the minimum is attained when $f_t^{\min} = \bar{f} + \frac{I_f}{I_{\hat{\lambda}}}\frac{\sigma_f^2}{\nu_{\lambda,t}}$. This term is well above $\bar{f}$ with our calibration of Table 1: for $\nu_{\lambda,t} = 0.3$, we obtain $f_t^{\min} = 0.063$.

in bad times.

Our model thus implies that the risk premium fluctuates when persistence is uncertain. In contrast, the risk premium remains independent of the state of the economy if λ_t is observable, or if the agent only learns about the level of expected output growth.

3 Theoretical Predictions

In this section we first calibrate the model to U.S. output data. Second, we quantify and discuss how asset prices vary with the state variables. We show that learning about persistence generates time variation in asset-pricing moments, whereas a model with learning about the level only—or without learning—imply constant moments. Third, we show that, in the case of learning about persistence, the term structures of equity return volatility and risk premium are flatter than what the traditional long-run risk model predicts. Finally, we extend the model of learning about persistence by taking into account leverage. This feature helps better match the level of the return volatility and risk premium observed in the data.

3.1 Calibration

We now calibrate the model and determine which type of learning best explains the dynamics of the U.S. output growth rate and its forecasts. We use the mean analyst forecast on 1-quarter-ahead real GDP growth as a direct measure of f_t and the realized real GDP growth as a proxy for the growth rate of the output process δ_t. Data are obtained from the Federal Reserve Bank of Philadelphia and are available at quarterly frequency from Q4:1968 to Q4:2015.[15]

We use the dynamics of the filters $\widehat{l}_t$ and $\widehat{\lambda}_t$ from Equations (10), the dynamics of the uncertainty about λ from Equation (13), and the filtered Brownian shocks from Equation (11) to generate model-implied paths of the output growth and its forecast. We estimate the model by Maximum Likelihood (Hamilton, 1994) and determine the values of the parameters σ_δ, $\bar{f}$, σ_f, σ_l, σ_λ that provide the closest fit to realized observations. Table 1 reports the estimated parameters for different learning models. Note that the prior on uncertainty about the mean-reversion speed is set to $\nu_{\lambda,0} = \frac{\sigma_f \sigma_\lambda}{|\bar{f} - f_0|}$,[16] whereas the priors $\widehat{l}_0$ and $\widehat{\lambda}_0$ are estimated together with the parameters. Details are in Appendix C.

[15]Considering output rather than consumption data allows us exploiting a longer sample period, as the time series of consumption forecasts only starts in Q3:1981.

[16]We assume that the agent considers a (local) steady state when computing the prior on uncertainty about the mean-reversion speed. That is, the uncertainty about λ_0 solves $\frac{d\nu_{\lambda,t}}{dt} = 0$ with $t = 0$. It means that uncertainty initially starts at $\nu_{\lambda,0} = \frac{\sigma_f \sigma_\lambda}{|\bar{f} - f_0|}$ and then dynamically evolves according to Equation (13). The level of this prior has actually no impact on our results.

Parameter	Symbol	Full	Persistence	Level	No learning
Output growth volatility	σ_δ	0.0141***	0.0142***	0.0141***	0.0142***
		(0.001)	(0.001)	(0.001)	(0.001)
Forecast long-term mean	$\bar{f}$	0.0262***	0.0262***	0.0261***	0.0261***
		(0.003)	(0.003)	(0.004)	(0.004)
Forecast volatility	σ_f	0.0242***	0.0237***	0.0235***	0.0235***
		(0.001)	(0.001)	(0.001)	(0.001)
Constant mean-reversion speed	$\bar{\lambda}$			0.9538***	0.9539***
				(0.194)	(0.193)
Bias volatility	σ_l	0.0023**		0.0023**	
		(0.001)		(0.001)	
Mean-reversion speed volatility	σ_λ	0.2918***	0.1696*		
		(0.137)	(0.104)		
Bayesian information criterion	BIC	$-2,471.18$	$-2,475.43$	$-2,469.16$	$-2,473.29$
Akaike information criterion	AIC	$-2,487.39$	$-2,488.39$	$-2,485.36$	$-2,486.25$

Table 1: Parameter estimates
This table reports the estimates of the model parameters. We use the forecast on the 1-quarter-ahead real GDP growth and the realized real GDP growth as a proxy for the output growth. The estimates are obtained by Maximum Likelihood for the period Q4:1968 to Q4:2015, using data from the Federal Reserve Bank of Philadelphia. The table compares the estimation results of the full model with those of three special cases: i) learning about persistence only ($\widehat{l}_t = 0,\ \forall t$); ii) learning about the level only ($\widehat{\lambda}_t = \bar{\lambda},\ \forall t$); and iii) no learning ($\widehat{l}_t = 0, \widehat{\lambda}_t = \bar{\lambda},\ \forall t$). Standard errors are reported in brackets and statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

Figure 1 reports the time series of the main state variables for the different types of learning.[17] It is worth noting that that the estimated degree of persistence is much lower than what is typically considered in the long-run risk literature (e.g. Bansal et al., 2016), which is about 0.2. Further, the adjustment to output growth forecast is relatively small, on average. A plausible interpretation for this finding is that the forecasts available from professional surveys are of good quality (e.g. Ang et al., 2007). Thus, the main uncertainty that agents are facing is less about whether the economy will be in a recession or an expansion, but more about how persistent the current state of the economy is expected to be. For instance, it was pretty clear that the latest financial crisis of 2007-08 would induce a recession, but it was less clear whether the recession would be short-lived or rather long-lived and turn into a depression. This is the type of uncertainty that agents aim to resolve when learning about the degree of persistence.

The results in Table 1 confirm this intuition. We use the Bayesian and Akaike information criterions to compare the likelihood of observing different forms of learning: i) a model with learning about both the level and the persistence; ii) a model with learning about persistence only; iii) a model with learning about the level of economic growth only; and iv) a model

[17]See Table 4 in Appendix C.1 for the descriptive statistics and a discussion of these variables.

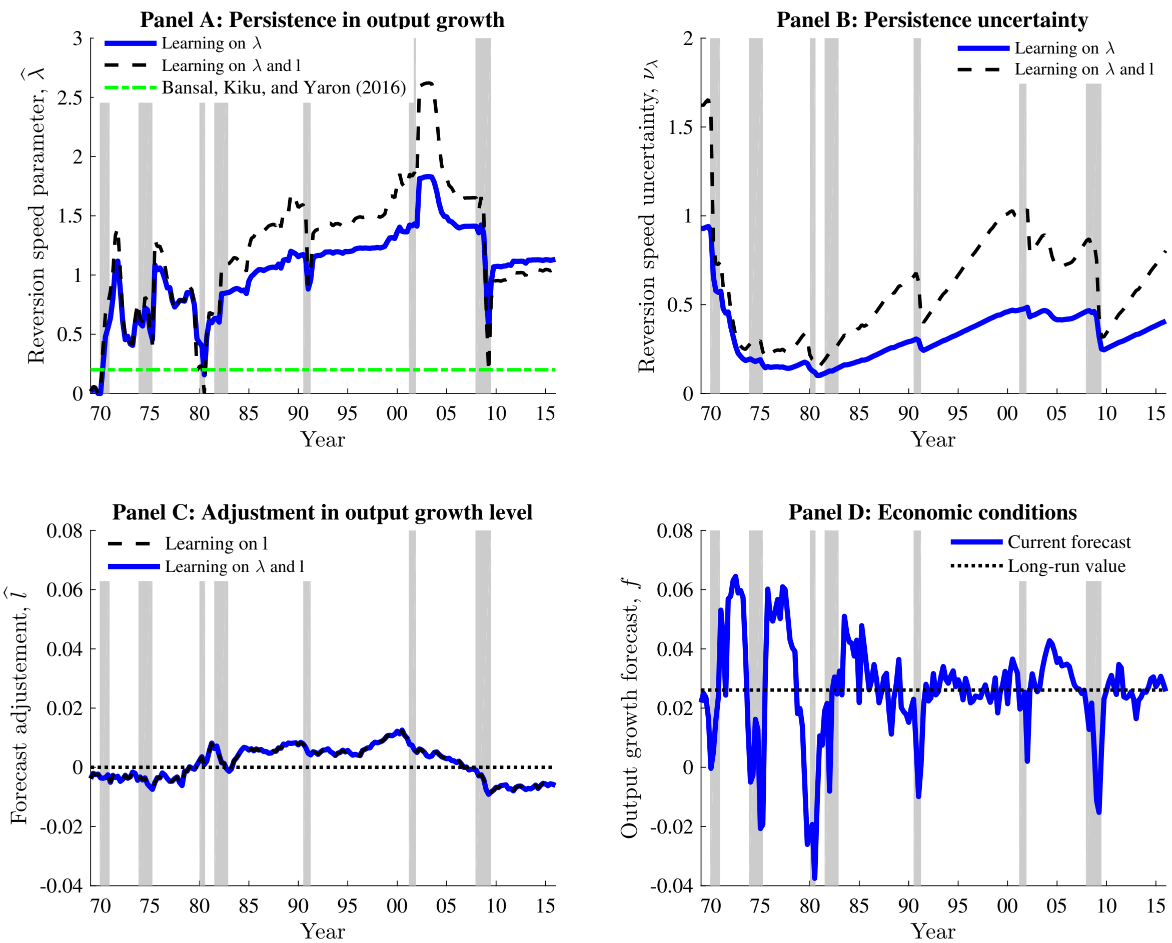

Figure 1: Historical path of the state variables.
This figure plots the time-series of the main state variables. Panel A reports the filtered persistence in output growth for different learning models. For comparison purposes, this panel also displays the average persistence level across the specifications used by Bansal, Kiku, and Yaron (2016). Panel B reports the filtered uncertainty about persistence, while Panel C shows the filtered adjustment in the output growth forecast in the case of learning about both the persistence and the level of expected output growth. Finally, Panel D shows the one-quarter ahead forecast of output growth, as reported by the Survey of Professional Forecasters. The sample spans the period Q4:1968 to Q4:2015.

without learning. We find that the model of learning about only the persistence of economic growth best fits historical data.[18]

In what follows, we will use these different calibrations to discuss the asset pricing implications of our model embedding different types of learning. We set the risk aversion to

[18]That is, the likelihood of the model combining both types of learning does not improve enough, compared to the model of learning about persistence, to dominate the penalty imposed by the additional number of parameters.

$\gamma = 10$, the elasticity of intertemporal substitution (EIS) to $\psi = 2$, and the subjective discount rate to $\beta = 0.04$.

3.2 Asset prices with learning about persistence

We first analyze a model with learning about persistence. According to our estimation, this model fits macroeconomic data particularly well. When learning about level is absent ($\widehat{l}_t = 0, \forall t$), the stock return volatility becomes

$$\|\sigma_t\| = \sqrt{\sigma_\delta^2 + \left(\sigma_f \frac{I_f}{I} + \frac{(\bar{f} - f_t)\nu_{\lambda,t}}{\sigma_f} \frac{I_{\widehat{\lambda}}}{I}\right)^2}, \tag{28}$$

whereas the risk premium is given by

$$\mu_t - r_t \equiv \sigma_t m_t = \gamma \sigma_\delta^2 + (1 - \phi) \left(\sigma_f \frac{I_f}{I} + \frac{(\bar{f} - f_t)\nu_{\lambda,t}}{\sigma_f} \frac{I_{\widehat{\lambda}}}{I}\right)^2. \tag{29}$$

Finally, the Sharpe ratio is defined as

$$SR_t \equiv \frac{\mu_t - r_t}{\|\sigma_t\|}. \tag{30}$$

Figure 2 displays these quantities. The left panels depict the relations with the forecast f_t, setting the filter $\widehat{\lambda}_t$ at its historical average. The right panels depict the relations with the filter $\widehat{\lambda}_t$, setting the forecast f_t at its long-term mean $\bar{f}$. All panels report values for various levels of uncertainty about persistence, $\nu_{\lambda,t}$. The limits for the X-axes in all panels correspond to the 5^{th}-percentiles and 95^{th}-percentiles from Table 4.[19]

Focusing first on the left panels, the main finding is that the asset-pricing moments are almost constant when $\nu_{\lambda,t} = 0$. In this case, there is no uncertainty about persistence and asset-pricing moments do not depend on the state of the economy. As Equations (28) and (29) show, the output growth gap does not drive the volatility or the risk premium when $\nu_{\lambda,t} = 0$. In contrast, when persistence is uncertain ($\nu_{\lambda,t} > 0$), the output growth gap becomes relevant, and the asset-pricing moments become negatively related to the forecast f_t.[20] Hence, the sensitivity of volatility, the risk premium, and the Sharpe ratio to the state of the economy depends on the degree of uncertainty about persistence. This is our main

[19] Our model assumes that output and dividends are the same. This is sufficient for our results on the dynamics of asset prices, but the model does not match the observed levels of volatility and risk premium. In Section 3.5, we distinguish between output and dividends (Bansal and Yaron, 2004; Beeler and Campbell, 2012). The stock market is a claim to dividends, which are more volatile than the output. This extension allows the model to better match the level of the S&P 500 return volatility and risk premium.

[20] The quadratic functions in Equations (28) and (29) reach a minimum when $f_t \gg \bar{f}$. See Footnote 14.

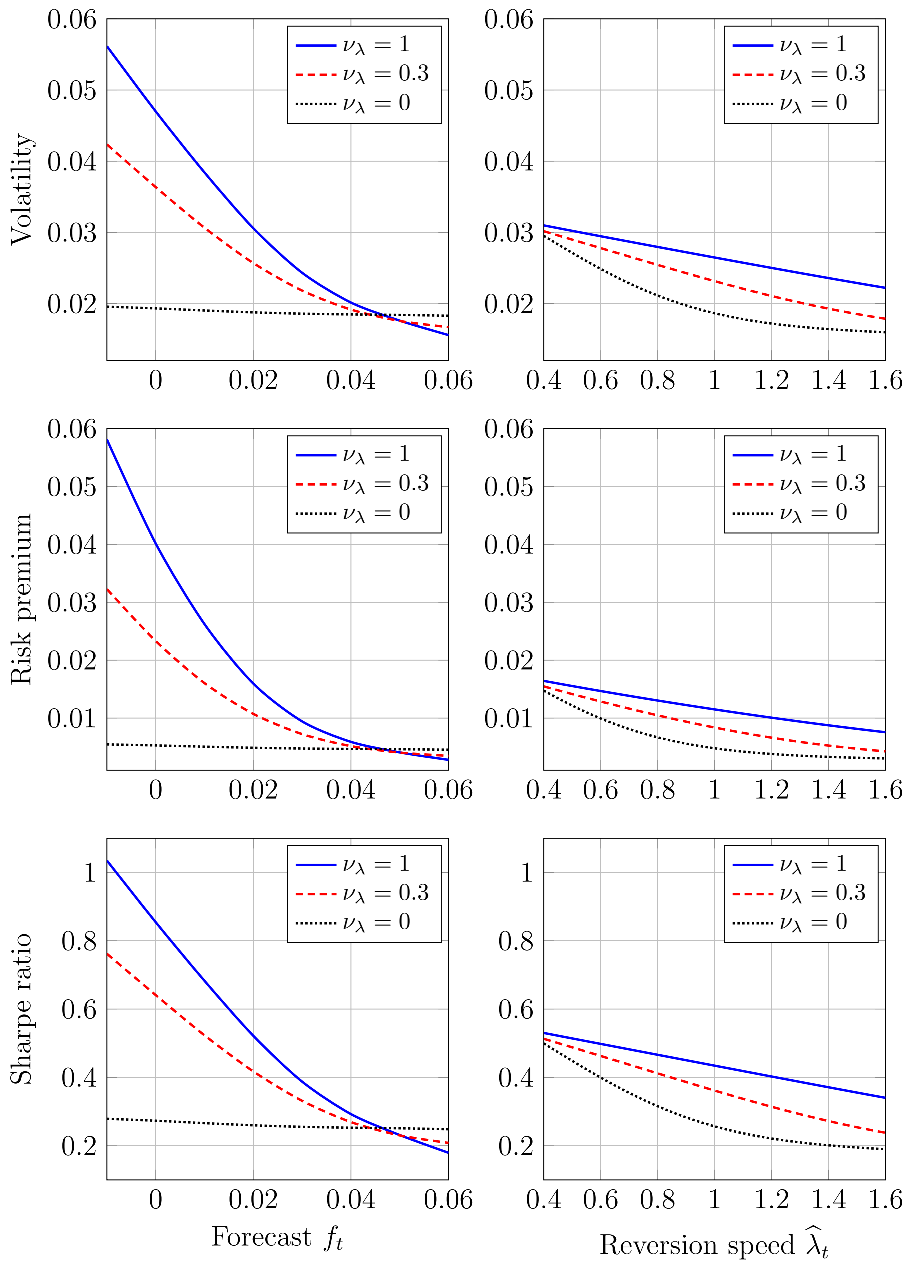

Figure 2: Stock return volatility, risk premium, and Sharpe ratio with learning about persistence.
This figure shows how the stock return volatility, the risk premium, and the Sharpe ratio vary with the state variables. For the left plots, we fix $\widehat{\lambda}_t = 1$. For the right plots, we fix $f_t = \bar{f}$. Unless otherwise specified, we consider the calibration provided in the second column of Table 1.

theoretical prediction.

The intuition behind this mechanism results from Equations (28)-(30). Consider a positive expected growth shock in bad times. This shock signals higher future output growth. In the same time, the agent expects less persistence, which is also good news because it means less risk in the long run. A negative expected growth shock does the opposite. It implies lower future output growth and more persistence (both effects are bad). Learning about persistence, thus, amplifies the response of asset returns to shocks—positive or negative—during bad times. Furthermore, agent's preferences imply that the coefficient $1 - \phi$ is positive and large (with our calibration, $1 - \phi = 19$). The risk premium then increases faster than the volatility as the economy enters bad times. This has an impact on the Sharpe ratio, which itself becomes higher in bad times.

In good times, learning about persistence has an opposite, dampening effect. Positive shocks signal higher future output growth, but also more persistence which is bad news because it means more risk in the long run. Negative shocks signal lower future output growth, but also less persistence. Agent's learning about persistence thus partially offsets the effect of the initial shock, reducing the volatility, the risk premium, and the Sharpe ratio in good times.[21]

The right panels of Figure 2 depict the impact of persistence on asset-pricing moments. The volatility, the risk premium, and the Sharpe ratio negatively depend on the mean-reversion speed $\widehat{\lambda}_t$. This arises because more persistence (lower $\widehat{\lambda}_t$) implies more risk in the long run. Notice that in the right panels we have fixed $f_t = \bar{f}$. When this equality holds, agents do not learn about the parameter λ_t (expected growth shocks are uninformative). Even in this case, the right panels show that uncertainty about persistence affects asset prices.

Overall, our model generates a set of new predictions. First, uncertainty about the degree of persistence creates a negative dependence of volatility, the risk premium, and the Sharpe ratio on the output growth forecast. Second, these asset-pricing moments increase with the persistence of expected output growth, i.e., they decrease with the estimated mean-reversion speed $\widehat{\lambda}_t$.

3.3 Comparison with learning about level and with no learning

We compare our model of learning about persistence with two other model specifications. First, we consider the case of learning about the level of expected output growth only ($\widehat{\lambda}_t =$

[21]The agent always dislikes an increase in persistence, in good or bad times. In good times, an increase in persistence signals indeed a longer expansion. But, due to preference for early resolution of uncertainty, the agent dislikes persistence because it increases risk in the long run. This effect dominates the good news of a longer expansion. Consequently, an increase in persistence is always bad news. Technically, the term $I_{\widehat{\lambda}}/I$ remains positive at all times. Table 5 in Appendix D clarifies this point.

$\bar{\lambda}\ \forall t$). In this case, the volatility of stock returns is

$$\|\sigma_t\| = \sqrt{\left(\sigma_\delta + \frac{\bar{\nu}_l}{\sigma_\delta}\frac{I_{\widehat{l}}}{I}\right)^2 + \left(\sigma_f \frac{I_f}{I}\right)^2}, \tag{31}$$

and the risk premium is

$$\mu_t - r_t \equiv \sigma_t m_t = \left(\sigma_\delta + \frac{\bar{\nu}_l}{\sigma_\delta}\frac{I_{\widehat{l}}}{I}\right)\left(\gamma\sigma_\delta + (1-\phi)\frac{\bar{\nu}_l}{\sigma_\delta}\frac{I_{\widehat{l}}}{I}\right) + (1-\phi)\left(\sigma_f \frac{I_f}{I}\right)^2. \tag{32}$$

Second, we consider an economy without learning ($\widehat{\lambda}_t = \bar{\lambda}, \widehat{l}_t = 0\ \forall t$). In this case, the volatility of stock returns is

$$\|\sigma_t\| = \sqrt{\sigma_\delta^2 + \left(\sigma_f \frac{I_f}{I}\right)^2}, \tag{33}$$

and the risk premium is

$$\mu_t - r_t \equiv \sigma_t m_t = \gamma\sigma_\delta^2 + (1-\phi)\left(\sigma_f \frac{I_f}{I}\right)^2. \tag{34}$$

In both cases, the Sharpe ratio is computed as in (30).

Equations (31)–(34) show that none of these models generates variations in volatility, the risk premium, and the Sharpe ratio, beyond the fluctuations that arise from the partial derivatives of the price-dividend ratio with respect to state variables. These fluctuations are weak; they are zero with a log-linear approximation of the log price-dividend ratio, which is the usual solution in the literature (Bansal and Yaron, 2004; Beeler and Campbell, 2012).

Figure 3 confirms this result. It compares the dependence of asset-pricing moments to state variables in three cases: (*i*) a model with learning about persistence, (*ii*) a model with learning about level, and (*iii*) a model without learning. Learning about level yields larger volatility, risk premium and Sharpe ratio than the other two models. However, the only model that generates variations in asset-pricing moments is the model of learning about persistence.

Appendix E provides additional results. It shows how the log price-dividend ratio and the risk-free rate vary with the state variables under different forms of learning. The log-price dividend ratio is pro-cyclical in all models (it increases with the growth forecast f_t), whereas the risk-free rate is counter-cyclical in all models. The counter-cyclicality of the risk-free rate is more pronounced in a model with learning about persistence.

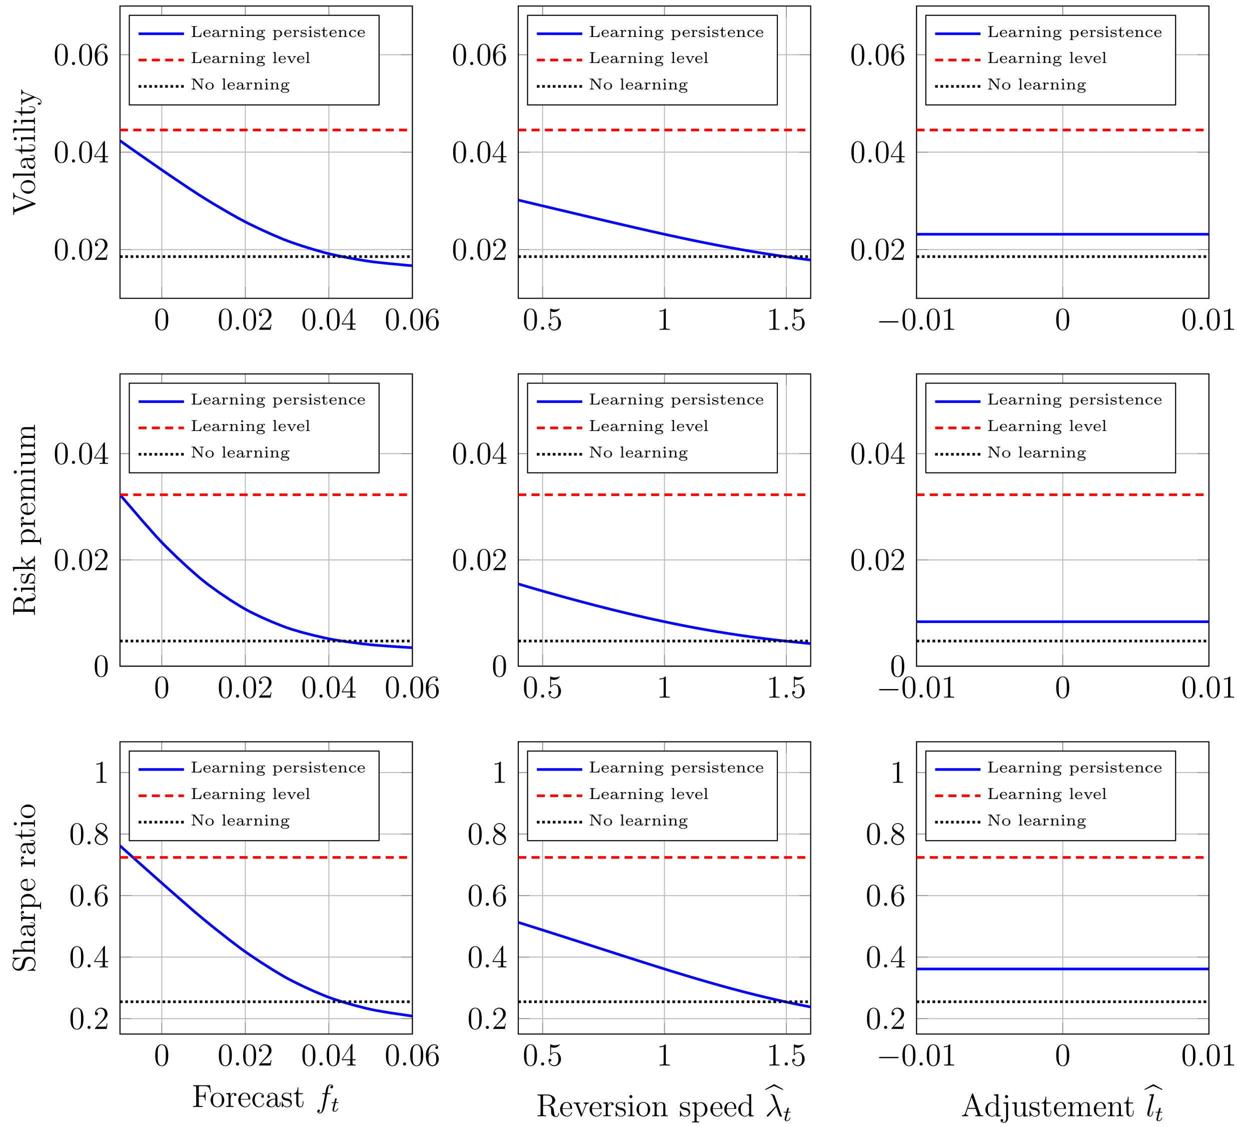

Figure 3: Stock return volatility, risk premium, and Sharpe ratio with and without learning.
This figure compares the asset-pricing implications in three cases: (i) a model with learning about persistence; (ii) a model with learning about level; and (iii) a model without learning. For the left plots, we fix $\widehat{l}_t = 0$ and $\widehat{\lambda}_t = 1$. For the middle plots, we fix $f_t = \bar{f}$ and $\widehat{l}_t = 0$. For the right plots, we fix $f_t = \bar{f}$ and $\widehat{\lambda}_t = 1$. We set $\nu_{\lambda,t} = 0.3$ in all plots. Unless otherwise specified, we consider the calibration provided in Table 1.

3.4 Term structure of equity volatility and risk premium

Recent evidence questions the validity of the calibration assumed by the long-run risk model. Using index options data, van Binsbergen et al. (2012) document that the term structures of equity volatility and equity risk premia are downward sloping. In contrast, the long-run risk model implies strongly upward-sloping term structures (Figure 5 in van Binsbergen and

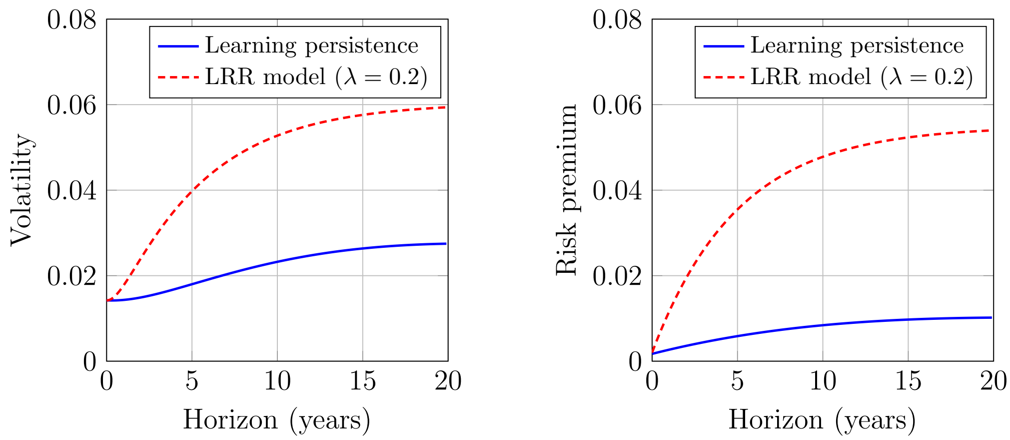

Figure 4: Term structures of volatility and risk premium.
This figure shows how the term structures of volatility and risk premium vary with learning about persistence. The solid lines correspond to a model with learning about persistence, in which we fix $f_t = \bar{f}$, $\widehat{\lambda}_t = 1$, and $\nu_{\lambda,t} = 0.3$. The dashed lines correspond to a model without learning, in which we fix $f_t = f$ and the mean-reverting parameter $\lambda = 0.2$. This value is in the middle range of the various calibrations considered in the long-run risk literature (Bansal et al., 2016). Unless otherwise specified, we consider the calibration provided in Table 1.

Koijen (2016)). This tension is mainly generated by the high degree of persistence assumed in the long-run risk model. Intuitively, high persistence increases the risk in the long run and thus generates a steeper term structure.

According to the estimation of Section 3.1, the degree of persistence in our model is much lower than in the long-run risk literature. It is then relevant to analyze the term structures generated by this new calibration. We do so by comparing two specifications. The first one is a model with learning about persistence. The second one is a model without learning in which we fix the mean-reverting parameter to $\lambda = 0.2$. This value is in the middle range of the recent calibrations considered in the long-run risk literature (see Bansal et al. (2016) and our discussion in Section 3.1). This second model without learning corresponds to the long-run risk specification of Bansal and Yaron (2004) with constant volatility of output growth.

We plot the term structures of volatility and risk premia for these two cases in Figure 4. The solid lines correspond to a model with learning about persistence in which the value of the filter is $\widehat{\lambda}_t = 1$ (this value represents the average of the model-implied filter $\widehat{\lambda}_t$). The dashed lines correspond to the long-run risk model. The two panels show that a model with learning about persistence implies term structures that are flatter than the ones obtained in the long-run risk model.[22] By calibrating the persistence parameter to the data, our model helps to partially address the concerns raised by van Binsbergen et al. (2012, 2013) on the timing of equity risk and risk premium implied by the long-run risk model.[23]

3.5 Dividend leverage

Figures 2 and 3 show that our model is able to generate significant fluctuations in volatility and the risk premium. The average levels of the model-implied volatility and risk premium shown in these figures, however, seem low when compared with the data (for the period 1968-2015, the annualized volatility of the S&P 500 is 15%; the annualized return in excess of the risk-free rate is 2.8%). The reason is that we do no distinguish between dividends and output in our baseline model to illustrate the effects of learning on asset prices in the simplest possible setting. In order to better match the level of volatility and risk premium observed in the data, we consider now an extension based on the notion of dividend leverage.

Abel (1999) was the first to introduce levered equity in an equilibrium model. Bansal and Yaron (2004) later adopted the same approach. They distinguish between output and dividends, and calibrate the leverage parameter to match the volatility of dividend growth, which is relatively higher.[24] In the context of our model, we assume that dividends are a levered claim of output:

$$D_t = \delta_t^\eta, \tag{35}$$

where $\eta \geq 1$ is the leverage parameter. The equity is then a claim to the dividend process

$$\frac{dD_t}{D_t} = \left[\eta\mu_t + \frac{1}{2}\eta(\eta-1)\sigma_\delta^2\right] dt + \eta\sigma_\delta dW_t^\delta, \tag{36}$$

which is obtained by applying Itô's lemma on (35). We are back to our baseline specification when $\eta = 1$.

[22]Learning about the level of expected output growth yields similar term structures to the ones obtained in the long-run risk model. Furthermore, a model with levered equity (see Section 3.5) yields a similar picture: the term structures are much flatter with learning about persistence than with the long-run risk specification.

[23]We acknowledge, however, that our model does not generate downward-sloping term structures. Other features are needed for that. See, for instance, Hasler and Marfe (2016) for an equilibrium model with rare disasters which are followed by recoveries. Disaster recovery helps explain the downward-sloping shape of the term structures of equity.

[24]The annualized volatility of the S&P 500 dividend growth from 1871 to 2015 (based on Robert Shiller's dataset) is about 5%. This is more than twice as large as the volatility of output or consumption growth. Other papers that distinguish between consumption and dividends are Bansal et al. (2012), Beeler and Campbell (2012), and Collin-Dufresne et al. (2016).

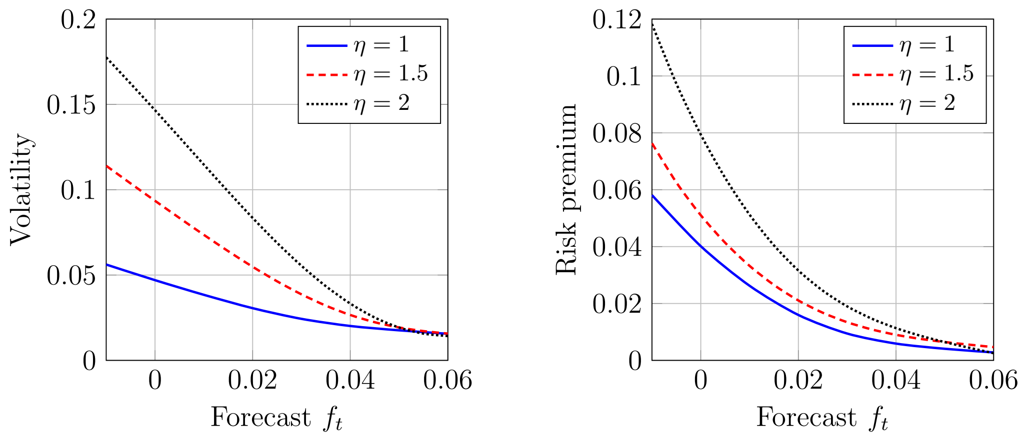

Figure 5: Levered equity volatility and risk premium with learning about persistence.
This figure shows how leverage affects the levels of asset-pricing moments. The two panels show the volatility and the risk premium of the equity, defined as the claim to the dividend process $D_t = \delta_t^{\eta}$, for three different values of the leverage parameter η. We obtain the baseline specification of Section 3.2 when $\eta = 1$. We set $\widehat{\lambda}_t = 1$ and $\nu_{\lambda,t} = 0.3$. Unless otherwise specified, we consider the calibration provided in the second column of Table 1.

Importantly, this extension does not change the learning problem of the agent. Since the dividend is a monotonic function of output, it does not provide additional information about the level of expected growth. Thus, the learning exercise of Section 2.2 still holds. But the parameter η will now change the volatility and the risk premium of the stock market. To see this, define V_t as the equity in the new model (the claim to dividends):

$$V_t = D_t I^D(x_t) = \mathbb{E}_t \left[\int_t^{\infty} \frac{\xi_s}{\xi_t} D_s ds \right], \tag{37}$$

where $I^D(x_t)$ is the price-dividend ratio (with this new specification, $I(x)$ becomes the wealth-consumption ratio). The pricing kernel ξ_t is the same as before. Finding the price-dividend ratio $I^D(x_t)$ involves writing a new partial differential equation, which we solve numerically (details are in Appendix B.1).

Figure 5 presents the results for three different values of the leverage parameter η. It shows that both the volatility and the risk premium are amplified with leverage, reaching values that are close to their empirical counterparts. Interestingly, the amplification occurs mainly in normal and bad times.

4 Empirical Tests

In this section, we empirically evaluate the predictions of the model with learning about persistence. We start by examining how the model-implied risk premium, Sharpe ratio, stock return volatility, price-dividend ratio, and the risk-free rate compare to their empirical counterparts. Then, we formally test our theory, according to which learning about persistence generates two main testable theoretical predictions for asset prices. First, the equity risk premium, the Sharpe ratio, and stock return volatility increase with the agent's estimate of the persistence in output growth. Hence, these moments are negatively related to the filtered mean-reversion speed, $\widehat{\lambda}_t$. Second, these asset pricing moments are expected to be more sensitive to changes in expected growth forecasts when there is greater uncertainty about the persistence level. The product of the output growth gap and the uncertainty about persistence, $(\bar{f} - f_t)\nu_{\lambda,t}$, should then have a positive impact on the aforementioned moments. These are the theoretical predictions that we are going to test in the data.

Our empirical analysis uses quarterly U.S. data over the period Q4:1968–Q4:2015. Following the calibration, we consider the mean analyst forecast on 1-quarter-ahead real GDP growth as a measure of f_t. The estimation performed in Section 3.1 provides time series of the filtered mean-reversion speed $\widehat{\lambda}_t$ and of the corresponding uncertainty $\nu_{\lambda,t}$. These state variables allow us constructing model-implied time series for the risk premium, Sharpe ratio, stock return volatility, price-dividend ratio, and the risk-free rate defined in Section 2.3. The empirical counterparts of the equity-related moments are computed using real data on the S&P 500 index, whereas we proxy for the risk-free rate with the real 10-year Treasury rate. The data come from Robert J. Shiller's website. Our proxy for the equity risk premium is the fitted value obtained by regressing one-quarter ahead S&P 500 excess returns on current price-dividend ratios, following the return predictability literature (Cochrane, 2008; van Binsbergen and Koijen, 2010). We then use the residuals of this regression to compute the return volatility on the S&P 500. Our estimation is based on the EGARCH(1,1,1) to account for potential asymmetric volatility responses to positive and negative shocks.[25]

In a first analysis, we compare the dynamics of the model-implied equity risk premium, Sharpe ratio, and stock return volatility with their empirical counterparts. Table 2 shows that the model-implied moments, which we derive in equilibrium, help understand the variations in the moments observed in the data. The correlations are all positive and highly statistically significant. These tests confirm that the dynamics of stock returns obtained

[25] This choice of specification is motivated by the leverage (or asymmetric) effect documented by Black (1976), French, Schwert, and Stambaugh (1987), Schwert (1989), Nelson (1991), and Glosten, Jagannathan, and Runkle (1993), among others. Note that using a GARCH model or changing the number of lags does not qualitatively alter our results.

	RP	SR	Vol	PD	r
Correlation	0.322***	0.173**	0.399***	0.576***	0.197***
$p - value$	0.000	0.017	0.000	0.000	0.007
N	189	189	189	189	189

Table 2: Empirical vs. model-implied asset pricing quantities
This table reports the correlations between our model-implied asset pricing quantities and their empirical counterparts. The relation using the equity risk premium is reported in the first column, whereas the second column considers the Sharpe ratio, the third column considers the stock return volatility, the fourth column considers the price-dividend ratio, and the last column reports the results for the risk-free rate. The second line reports the p-values, while N is the number of observations. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

with our model align well with what is observed in the data. We also obtain strong results for the price-dividend ratio and the risk-free rate.

We now discuss the relations between the empirical asset pricing moments and the state variables. The results for the equity risk premium, the Sharpe ratio, and stock return volatility are presented in panels A, B, and C of Table 3, respectively. The data confirm the theoretical predictions that these moments decrease with the filtered mean-reversion speed $\widehat{\lambda}_t$ (see Column 1) but increase with the term $(\bar{f} - f_t)\nu_{\lambda,t}$ (see Column 2). This term appears in (28)-(30) and combines the effect of the output growth gap and uncertainty about persistence. Columns 3 and 4 repeat the previous analysis when we control for both the output growth forecast and the degree of uncertainty regarding the degree of persistence. Our conclusion remains unchanged with these alternative specifications. This is important because it indicates that the degree of persistence is not a mere proxy of the current economic conditions. Moreover, it suggests that neither the growth forecast, nor the uncertainty about persistence alone drive asset prices, but it is the interaction between the two that matters. As a last robustness check, Column 5 reports the results when all potential explanatory variables are considered in the regression.

These results are in favor of our theoretical result that stock return volatility, the equity risk premium, and the Sharpe ratio all increase when the agent believes that there is higher persistence in the economy. These moments are also more sensitive to changes in growth forecasts when there is greater uncertainty about persistence.

Overall, the data lend support to the predictions of a model with learning about persistence. That is, our theory successfully pins down key determinants of the dynamics of the empirical risk premia, stock return volatility, and the risk-return tradeoff. It is thus important to consider learning and uncertainty about the persistence in expected output growth for understanding the dynamics of asset prices.

Panel A: Equity risk premium

	(1)	(2)	(3)	(4)	(5)
$\widehat{\lambda}$	-0.042^{***}		-0.039^{***}		-0.038^{***}
	(0.004)		(0.003)		(0.003)
$(\bar{f}-f)\nu_\lambda$		0.628^{**}		1.877^{**}	0.741
		(0.305)		(0.909)	(0.462)
Constant	0.072^{***}	0.028^{***}	0.009^{***}	0.056^{***}	0.009^{***}
	(0.004)	(0.002)	(0.004)	(0.005)	(0.004)
R^2	0.444	0.017	0.660	0.343	0.666
N	189	189	189	189	189

Panel B: Sharpe ratio

	(1)	(2)	(3)	(4)	(5)
$\widehat{\lambda}$	-0.247^{***}		-0.234^{***}		-0.227^{***}
	(0.029)		(0.021)		(0.021)
$(\bar{f}-f)\nu_\lambda$		1.072		12.11^{**}	5.243^{*}
		(1.917)		(5.739)	(3.058)
Constant	0.447^{***}	0.189^{***}	0.590^{***}	0.381^{***}	0.590^{***}
	(0.031)	(0.011)	(0.026)	(0.031)	(0.025)
R^2	0.331	0.001	0.580	0.337	0.587
N	189	189	189	189	189

Panel C: Stock return volatility

	(1)	(2)	(3)	(4)	(5)
$\widehat{\lambda}$	-0.027^{***}		-0.018^{***}		-0.015^{*}
	(0.009)		(0.008)		(0.008)
$(\bar{f}-f)\nu_\lambda$		3.563^{***}		2.995^{**}	2.545^{*}
		(0.842)		(1.289)	(1.353)
Constant	0.179^{***}	0.150^{***}	0.169^{***}	0.155^{***}	0.169^{***}
	(0.010)	(0.003)	(0.001)	(0.007)	(0.001)
R^2	0.056	0.159	0.161	0.166	0.181
N	189	189	189	189	189

Common specification across panels (control variables)

$\bar{f}-f$	No	No	Yes	Yes	Yes
ν_λ	No	No	Yes	Yes	Yes

Table 3: Asset pricing with learning about persistence
This table reports the relations between the asset pricing moments and the state variables obtained with learning about the persistence in output growth. Panel A reports the results for the equity risk premium, Panel B for the Sharpe ratio, and Panel C for the stock return volatility. N is the number of observations. Newey and West (1987) standard errors are in brackets and statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

5 Conclusion

This paper shows that a model of learning about the persistence of expected economic growth fits U.S. output growth data better than a model of learning about the level of expected growth. A plausible interpretation for this finding is that investors can easily obtain analyst forecasts on economic growth from professional surveys, but do not have access to forecasts for the persistence of the growth rate.

Learning about persistence appears to be determinant for financial markets. In equilibrium, this type of learning implies a negative relation between asset pricing moments (equity return volatility, equity risk premium, and the Sharpe ratio) and economic conditions. This relation disappears when there is no uncertainty about the persistence, or with learning about the level of expected growth. In addition, our model suggests that asset pricing moments increase with the degree of persistence in expected growth. The data provide empirical support for our theoretical predictions.

Our analysis offers several extensions for future research. First, a theory in which agents endogenously choose which dimension to learn (level or persistence) would be useful to investigate whether learning about persistence would be a rational response of investors, as well as to compute the welfare gains of a publicly available indicator of the persistence in economic growth. Second, extending our model to an economy with multiple risky assets offers the possibility to investigate the implications of learning about persistence for the cross-section of asset returns. Third, since the persistence in economic growth remains unobservable and can only be estimated using a long history of data, agents are very likely to disagree about it (Andrei et al., 2017). Future work could then construct a measure of disagreement about the persistence of economic growth by exploiting the cross-section of analyst forecasts on future economic growth. Such disagreement measure might help better predict future market returns and their volatility than the existing measures of disagreement about the level of economic growth.

References

Abel, A. B. (1999). Risk premia and term premia in general equilibrium. *Journal of Monetary Economics 43*(1), 3–33.

Ai, H. (2010). Information quality and long-run risk: Asset pricing implications. *The Journal of Finance 65*(4), 1333–1367.

Andrei, D., B. Carlin, and M. Hasler (2017). Asset pricing with structural uncertainty and structural disagreement. *Working Paper*.

Ang, A., G. Bekaert, and M. Wei (2007). Do macro variables, asset markets, or surveys forecast inflation better? *Journal of Monetary Economics 54*(4), 1163 – 1212.

Bansal, R., D. Kiku, and A. Yaron (2012). An empirical evaluation of the long-run risks model for asset prices. *Critical Finance Review 1*(1), 183–221.

Bansal, R., D. Kiku, and A. Yaron (2016). Risks for the long run: Estimation with time aggregation. *Journal of Monetary Economics 82*, 52–69.

Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance 59*(4), 1481–1509.

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

Belo, F., P. Collin-Dufresne, and R. S. Goldstein (2015). Dividend dynamics and the term structure of dividend strips. *The Journal of Finance 70*(3), 1115–1160.

Benzoni, L., P. Collin-Dufresne, and R. S. Goldstein (2011). Explaining asset pricing puzzles associated with the 1987 market crash. *Journal of Financial Economics 101*(3), 552–573.

Bidder, R. and I. Dew-Becker (2016). Long-run risk is the worst-case scenario. *The American Economic Review 106*(9), 2494–2527.

Black, F. (1976). Studies of stock price volatility changes. *American Statistical Association, Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economic Statistics Section*, 177–181.

Brennan, M. J. (1998). The role of learning in dynamic portfolio decisions. *European Finance Review 1*(3), 295–306.

Brennan, M. J. and Y. Xia (2001). Stock price volatility and equity premium. *Journal of Monetary Economics 47*(2), 249–283.

Campbell, J. Y. and J. H. Cochrane (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *The Journal of Political Economy 107*(2), 205–251.

Campbell, J. Y., S. Giglio, C. Polk, and R. Turley (2017). An intertemporal capm with stochastic volatility. *Journal of Financial Economics*, Forthcoming.

Cochrane, J. H. (2008). The dog that did not bark: a defense of return predictability. *Review of Financial Studies 21*(4), 1533–1575.

Cogley, T. and T. J. Sargent (2008). The market price of risk and the equity premium: A legacy of the great depression? *Journal of Monetary Economics 55*(3), 454 – 476.

Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer (2016). Parameter learning in general equilibrium: The asset pricing implications. *The American Economic Review 106*(3), 664–698.

Croce, M. M., M. Lettau, and S. C. Ludvigson (2015). Investor information, long-run risk, and the term structure of equity. *Review of Financial Studies 28*(3), 706–742.

David, A. (1997). Fluctuating confidence in stock markets: Implications for returns and volatility. *Journal of Financial and Quantitative Analysis 32*(04), 427–462.

David, A. and P. Veronesi (2002). Option prices with uncertain fundamentals. *Working Paper*.

Duffie, D. and L. G. Epstein (1992). Asset pricing with stochastic differential utility. *Review of Financial Studies 5*(3), 411–436.

Dumas, B., A. Kurshev, and R. Uppal (2009). Equilibrium portfolio strategies in the presence of sentiment risk and excess volatility. *The Journal of Finance 64*(2), 579–629.

Epstein, L. G., E. Farhi, and T. Strzalecki (2014). How much would you pay to resolve long-run risk? *American Economic Review 104*(9), 2680–97.

Epstein, L. G. and S. E. Zin (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica 57*(4), 937–969.

Ferson, W. E. and C. R. Harvey (1991). The variation of economic risk premiums. *Journal of Political Economy 99*(2), 385–415.

French, K. R., G. W. Schwert, and R. F. Stambaugh (1987). Expected stock returns and volatility. *Journal of Financial Economics 19*(1), 3–29.

Gabaix, X. (2012). Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *The Quarterly Journal of Economics 127*(2), 645–700.

Glosten, L. R., R. Jagannathan, and D. E. Runkle (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance 48*(5), 1779–1801.

Hamilton, J. D. (1994). *Time Series Analysis*. Princeton: Princeton University Press.

Hasler, M. and R. Marfe (2016). Disaster recovery and the term structure of dividend strips. *Journal of Financial Economics 122*(1), 116–134.

Johannes, M., L. A. Lochstoer, and Y. Mou (2016). Learning about consumption dynamics. *The Journal of Finance 71*(2), 551–600.

Judd, K. L. (1998). *Numerical methods in economics*. MIT press.

Kreps, D. M. (1998). Anticipated utility and dynamic choice. *Frontiers of research in economic theory*, 242–274. In: Frontiers of research in economic theory.

Lettau, M. and S. C. Ludvigson (2010). Measuring and modelling variation in the risk-return trade-off. *Handbook of Financial Econometrics 1*, 617–690.

Liptser, R. and A. Shiryayev (1977). Statistics of random processes. Springer-Verlag, Berlin, New York.

Lustig, H. and A. Verdelhan (2012). Business cycle variation in the risk-return trade-off. *Journal of Monetary Economics 59*, S35–S49.

Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica 59*(2), 347–70.

Newey, W. K. and K. D. West (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica 55*(3), 703–08.

Pakoš, M. (2013). Long-run risk and hidden growth persistence. *Journal of Economic Dynamics and Control 37*(9), 1911–1928.

Pastor, L. and P. Veronesi (2009). Learning in financial markets. *Annual Review of Financial Economics 1*(1), 361–381.

Schwert, G. W. (1989). Why does stock market volatility change over time? *Journal of Finance 44*(5), 1115–53.

van Binsbergen, J., M. Brandt, and R. Koijen (2012). On the timing and pricing of dividends. *American Economic Review 102*(4), 1596–1618.

van Binsbergen, J., W. Hueskes, R. Koijen, and E. Vrugt (2013). Equity yields. *Journal of Financial Economics 110*(3), 503 – 519.

van Binsbergen, J. H. and R. S. Koijen (2017). The term structure of returns: Facts and theory. *Journal of Financial Economics*, Forthcoming.

van Binsbergen, J. H. and R. S. J. Koijen (2010). Predictive regressions: a present-value approach. *Journal of Finance 65*(4), 1439–1471.

van Binsbergen, J. H. and R. S. J. Koijen (2016). On the timing and pricing of dividends: Reply. *American Economic Review 106*(10), 3224–37.

Veronesi, P. (1999). Stock market overreactions to bad news in good times: a rational expectations equilibrium model. *Review of Financial Studies 12*(5), 975–1007.

Veronesi, P. (2000). How does information quality affect stock returns? *The Journal of Finance 55*(2), 807–837.

Wachter, J. A. (2013). Can time-varying risk of rare disasters explain aggregate stock market volatility? *The Journal of Finance 68*(3), 987–1035.

Ziegler, A. C. (2012). *Incomplete information and heterogeneous beliefs in continuous-time finance*. Springer Science & Business Media.

A Learning

Theorem 1. *(Liptser and Shiryayev, 1977) Consider an unobservable process u_t and an observable process s_t with dynamics given by*

$$du_t = [a_0(t,s_t) + a_1(t,s_t)u_t]\,dt + b_1(t,s_t)dZ_t^u + b_2(t,s_t)dZ_t^s \tag{38}$$
$$ds_t = [A_0(t,s_t) + A_1(t,s_t)u_t]\,dt + B_1(t,s_t)dZ_t^u + B_2(t,s_t)dZ_t^s. \tag{39}$$

All the parameters can be functions of time and of the observable process. Liptser and Shiryayev (1977) show that the filter evolves according to (we drop the dependence of coefficients on t and s_t for notational convenience):

$$d\widehat{u}_t = (a_0 + a_1\widehat{u}_t)dt + [(b\circ B) + \nu_t A_1^\top](B\circ B)^{-1}[ds_t - (A_0 + A_1\widehat{u}_t)dt] \tag{40}$$
$$\frac{d\nu_t}{dt} = a_1\nu_t + \nu_t a_1^\top + (b\circ b) - [(b\circ B) + \nu_t A_1^\top](B\circ B)^{-1}[(b\circ B) + \nu_t A_1^\top]^\top, \tag{41}$$

where

$$b\circ b = b_1 b_1^\top + b_2 b_2^\top \tag{42}$$
$$B\circ B = B_1 B_1^\top + B_2 B_2^\top \tag{43}$$
$$b\circ B = b_1 B_1^\top + b_2 B_2^\top. \tag{44}$$

Write the dynamics of the observable variables

$$\begin{bmatrix} d\log\delta_t \\ df_t \end{bmatrix} = \underbrace{\begin{bmatrix} f_t - \frac{1}{2}\sigma_\delta^2 \\ 0 \end{bmatrix}}_{A_0} + \underbrace{\begin{bmatrix} 1 & 0 \\ 0 & (\bar{f} - f_t) \end{bmatrix}}_{A_1} \begin{bmatrix} l_t \\ \lambda_t \end{bmatrix} + \underbrace{\begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}}_{B_1} \begin{bmatrix} dW_t^l \\ dW_t^\lambda \end{bmatrix} + \underbrace{\begin{bmatrix} \sigma_\delta & 0 \\ 0 & \sigma_f \end{bmatrix}}_{B_2} \begin{bmatrix} dW_t^\delta \\ dW_t^f \end{bmatrix} \tag{45}$$

and unobservable variables

$$\begin{bmatrix} dl_t \\ d\lambda_t \end{bmatrix} = \underbrace{\begin{bmatrix} 0 \\ 0 \end{bmatrix}}_{a_0} + \underbrace{\begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}}_{a_1} \begin{bmatrix} l_t \\ \lambda_t \end{bmatrix} + \underbrace{\begin{bmatrix} \sigma_l & 0 \\ 0 & \sigma_\lambda \end{bmatrix}}_{b_1} \begin{bmatrix} dW_t^l \\ dW_t^\lambda \end{bmatrix} + \underbrace{\begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}}_{B_2} \begin{bmatrix} dW_t^\delta \\ dW_t^f \end{bmatrix}. \tag{46}$$

Then,

$$\begin{bmatrix} d\widehat{l}_t \\ d\widehat{\lambda}_t \end{bmatrix} = \begin{bmatrix} \frac{\nu_{l,t}}{\sigma_\delta} & 0 \\ 0 & \frac{(\bar{f}-f_t)\nu_{\lambda,t}}{\sigma_f} \end{bmatrix} \begin{bmatrix} d\widehat{W}_t^\delta \\ d\widehat{W}_t^f \end{bmatrix}, \tag{47}$$

where the independent Brownian motions $\widehat{W}_t^\delta$ and $\widehat{W}_t^f$ are such that

$$\frac{d\delta_t}{\delta_t} = [f_t + \widehat{l}_t]dt + \sigma_\delta d\widehat{W}_t^\delta, \tag{48}$$
$$df_t = \widehat{\lambda}_t(\bar{f} - f_t)dt + \sigma_f d\widehat{W}_t^f. \tag{49}$$

The posterior uncertainties about l_t and λ_t evolve according to

$$d\nu_{l,t} = \left[\sigma_l^2 - \frac{\nu_{l,t}^2}{\sigma_\delta^2}\right] dt, \tag{50}$$

$$d\nu_{\lambda,t} = \left[\sigma_\lambda^2 - \frac{(\bar{f} - f_t)^2 \nu_{\lambda,t}^2}{\sigma_f^2}\right] dt. \tag{51}$$

Therefore, the uncertainty about l_t admits a constant steady-state solution (i.e. $\left.\frac{d\nu_{l,t}}{dt}\right|_{\nu_{l,t} \equiv \bar{\nu}_l} = 0$)

$$\bar{\nu}_l = \sigma_\delta \sigma_l, \tag{52}$$

but this is not the case for the uncertainty about λ_t because of the term $(\bar{f} - f_t)$.

B Equilibrium

The dynamics of the vector of state variables are

$$\begin{bmatrix} d\delta_t \\ df_t \\ d\widehat{l}_t \\ d\widehat{\lambda}_t \\ d\nu_{\lambda,t} \end{bmatrix} = \begin{bmatrix} \delta_t[f_t + \widehat{l}_t] \\ \widehat{\lambda}_t(\bar{f} - f_t) \\ 0 \\ 0 \\ \sigma_\lambda^2 - \frac{(\bar{f} - f_t)^2 \nu_{\lambda,t}^2}{\sigma_f^2} \end{bmatrix} dt + \begin{bmatrix} \delta_t \sigma_\delta & 0 \\ 0 & \sigma_f \\ \frac{\nu_{l,t}}{\sigma_\delta} & 0 \\ 0 & \frac{(\bar{f} - f_t)\nu_{\lambda,t}}{\sigma_f} \\ 0 & 0 \end{bmatrix} \begin{bmatrix} d\widehat{W}_t^\delta \\ d\widehat{W}_t^f \end{bmatrix}. \tag{53}$$

Proof that $I(x_t)$ is the price-dividend ratio. The following relationship results directly from replacing the conjectured form of the value function J in $h(C, J)$:

$$\frac{h(C, J)}{J} = \frac{\phi}{I(x_t)} - \beta\phi. \tag{54}$$

Define

$$S_t = C_t I(x_t), \tag{55}$$

and replace (21) in the product $\xi_t S_t$ to get

$$\xi_t S_t = (1 - \gamma) \exp\left(\int_0^t \left(\frac{\phi - 1}{I(x_s)} - \beta\phi\right) ds\right) \frac{C_t^{1-\gamma}}{1 - \gamma} [\beta I(x_t)]^\phi \tag{56}$$

$$= (1 - \gamma) \exp\left(\int_0^t \left(\frac{\phi - 1}{I(x_s)} - \beta\phi\right) ds\right) J. \tag{57}$$

This is a function of J and of time. Applying Itô's lemma yields:

$$d(\xi_t S_t) = (1 - \gamma) \exp\left(\int_0^t \left(\frac{\phi - 1}{I(x_s)} - \beta\phi\right) ds\right) \left[dJ - J\left(\beta\phi - \frac{\phi - 1}{I(x_t)}\right) dt\right]. \tag{58}$$

We also know that

$$
\begin{aligned}
dJ &= -h(C,J)dt + dM_t && (59)\\
&= J\left(\beta\phi - \frac{\phi}{I(x_t)}\right)dt + dM_t, && (60)
\end{aligned}
$$

where M_t is a martingale. The second equality follows from (54). Replace this in (58):

$$
\begin{aligned}
d(\xi_t S_t) &= (1-\gamma)\exp\left(\int_0^t\left(\frac{\phi-1}{I(x_s)} - \beta\phi\right)ds\right)\left[J\left(\beta\phi - \frac{\phi}{I(x_t)}\right)dt + dM_t - J\left(\beta\phi - \frac{\phi-1}{I(x_t)}\right)dt\right] && (61)\\
&= -(1-\gamma)\exp\left(\int_0^t\left(\frac{\phi-1}{I(x_s)} - \beta\phi\right)ds\right)\frac{J}{I(x_t)}dt + d\widetilde{M}_t && (62)\\
&= -\xi_t C_t dt + d\widetilde{M}_t, && (63)
\end{aligned}
$$

where $d\widetilde{M}_t$ is a martingale. The third equality follows from replacing the conjectured form of the value function. The last equation can be integrated. Then, taking expectation and assuming that the transversality condition holds yields the total wealth (claim to all future output):

$$
S_t = \mathbb{E}_t\left[\int_t^\infty \frac{\xi_s}{\xi_t} C_s ds\right] \tag{64}
$$

which proves that $I(x_t)$ is the price-dividend ratio.

Partial differential equation for the price-dividend ratio. Substituting the guess (16) in the HJB Equation (15) and imposing the market clearing condition, $C = \delta$, yields the following PDE for the log price-dividend ratio $i \equiv \log I$

$$
\begin{aligned}
0 = {} & (\gamma-1)\left[-f - l + \frac{1}{2}\gamma\sigma_\delta^2\right] - \beta\phi + \phi e^{-i}\\
& + \phi(\bar{f}-f)\widehat{\lambda} i_f + \phi\left[\sigma_\lambda^2 - \frac{(\bar{f}-f)^2\nu_\lambda^2}{\sigma_f^2}\right] i_{\nu_\lambda} - \phi^2(1-1/\psi)\bar{\nu}_l i_{\widehat{l}}\\
& + \phi\frac{\sigma_f^2}{2} i_{ff} + \phi\frac{\bar{\nu}_l^2}{2\sigma_\delta^2} i_{\widehat{l}\widehat{l}} + \phi\frac{(\bar{f}-f)^2\nu_\lambda^2}{2\sigma_f^2} i_{\widehat{\lambda}\widehat{\lambda}} + \phi\nu_\lambda(\bar{f}-f) i_{f\widehat{\lambda}}\\
& + \phi^2\frac{\sigma_f^2}{2} i_f^2 + \phi^2(\bar{f}-f)\nu_\lambda i_f i_{\widehat{\lambda}} + \phi^2\frac{\bar{\nu}_l^2}{2\sigma_\delta^2} i_{\widehat{l}}^2 + \phi^2\frac{(\bar{f}-f)^2\nu_\lambda^2}{2\sigma_f^2} i_{\widehat{\lambda}}^2
\end{aligned} \tag{65}
$$

Setting $\widehat{\lambda}_t = \bar{\lambda}$, $\forall t$ in Equation (65) yields

$$
\begin{aligned}
0 = {} & (\gamma-1)\left[-f - l + \frac{1}{2}\gamma\sigma_\delta^2\right] - \beta\phi + \phi e^{-i}\\
& + \phi(\bar{f}-f)\bar{\lambda} i_f - \phi^2\bar{\nu}_l(1-1/\psi) i_{\widehat{l}}\\
& + \phi\frac{\sigma_f^2}{2} i_{ff} + \phi\frac{\bar{\nu}_l^2}{2\sigma_\delta^2} i_{\widehat{l}\widehat{l}} + \phi^2\frac{\sigma_f^2}{2} i_f^2 + \phi^2\frac{\bar{\nu}_l^2}{2\sigma_\delta^2} i_{\widehat{l}}^2.
\end{aligned} \tag{66}
$$

Setting $\widehat{l}_t = 0,\ \forall t$ in Equation (65) yields

$$
\begin{aligned}
0 = & -(\gamma-1)f + \gamma(\gamma-1)\frac{\sigma_\delta^2}{2} - \beta\phi + \phi e^{-i} \\
& + \phi\left[\widehat{\lambda}(\bar{f}-f)i_f + \left(\sigma_\lambda^2 - \frac{(\bar{f}-f)^2\nu_\lambda^2}{\sigma_f^2}\right)i_{\nu_\lambda}\right] \\
& + \phi\left[\frac{\sigma_f^2}{2}i_{ff} + \frac{(\bar{f}-f)^2\nu_\lambda^2}{2\sigma_f^2}i_{\widehat{\lambda}\widehat{\lambda}} + \nu_\lambda(\bar{f}-f)i_{f\widehat{\lambda}}\right] + \phi^2\left[\frac{\sigma_f^2}{2}i_f^2 + \frac{(\bar{f}-f)^2\nu_\lambda^2}{2\sigma_f^2}i_{\widehat{\lambda}}^2 + \nu_\lambda(\bar{f}-f)i_f i_{\widehat{\lambda}}\right].
\end{aligned}
\tag{67}
$$

B.1 Levered equity

Define

$$V_t = D_t I^D(x_t) = C_t^\eta I^D(x_t). \tag{68}$$

In this section $I(x_t)$ becomes the wealth-consumption ratio, whereas $I^D(x_t)$ becomes the price-dividend ratio which we need to solve for. Let us compute

$$\xi_t V_t = (1-\gamma)\exp\left(\int_0^t \left(\frac{\phi-1}{I(x_s)} - \beta\phi\right)ds\right) J C_t^{\eta-1}\frac{I^D(x_t)}{I(x_t)}. \tag{69}$$

One can clearly see that if $\eta = 1$, then $C_t^{\eta-1}$ drops out and the last fraction equals one, which brings us back to (57). The case of interest is $\eta > 1$. Define

$$K(C_t, x_t) = C_t^{\eta-1}\frac{I^D(x_t)}{I(x_t)} \tag{70}$$

and thus

$$
\begin{aligned}
d(\xi_t V_t) &= (1-\gamma)\exp\left(\int_0^t \left(\frac{\phi-1}{I(x_s)} - \beta\phi\right)ds\right)\left[-\frac{KJ}{I(x_t)} + KdM_t + JdK + (dJ)(dK)\right] & (71)\\
&= -K\xi_t C_t dt + (1-\gamma)\exp\left(\int_0^t \left(\frac{\phi-1}{I(x_s)} - \beta\phi\right)ds\right)\left[JdK + (dJ)(dK) + KdM_t\right]. & (72)
\end{aligned}
$$

We know that if V_t is the stock price, then we should also have:

$$d(\xi_t V_t) = -\xi_t C_t^\eta dt + d\widehat{M}_t, \tag{73}$$

where $d\widehat{M}_t$ is a martingale. This means that the drifts in (72) and (73) have to be equal. This yields a partial differential equation to be solved by $I^D(x_t)$. Finally, replacing $I^D = \exp(j)$ results in a partial differential equation for the log price-dividend ratio j:

$$
\begin{aligned}
0 = e^{-j} &+ h(f,\beta,\sigma_\delta^2) + \frac{1-\phi}{2\sigma_f^2}\left[\sigma_f^2 i_f + (\bar{f}-f)\nu_\lambda i_{\widehat{\lambda}}\right]^2 \\
&+ \left[\widehat{\lambda}(\bar{f}-f) - (1-\phi)\sigma_f^2 i_f - (1-\phi)\nu_\lambda(\bar{f}-f)i_{\widehat{\lambda}}\right] j_f \\
&+ \left[-(1-\phi)\nu_\lambda(\bar{f}-f)i_f - (1-\phi)\frac{(\bar{f}-f)^2\nu_\lambda^2}{\sigma_f^2} i_{\widehat{\lambda}}\right] j_{\widehat{\lambda}} \\
&+ \left[\sigma_\lambda^2 - \frac{(\bar{f}-f)^2\nu_\lambda^2}{\sigma_f^2}\right] j_{\nu_\lambda} \\
&+ \frac{\sigma_f^2}{2} j_{ff} + \frac{(\bar{f}-f)^2\nu_\lambda^2}{2\sigma_f^2} j_{\widehat{\lambda}\widehat{\lambda}} + \nu_\lambda(\bar{f}-f) j_{f\widehat{\lambda}} \\
&+ \frac{\sigma_f^2}{2} j_f^2 + \frac{(\bar{f}-f)^2\nu_\lambda^2}{2\sigma_f^2} j_{\widehat{\lambda}}^2 + \nu_\lambda(\bar{f}-f) j_f j_{\widehat{\lambda}}
\end{aligned} \tag{74}
$$

This equation has a similar structure as (65), except that it also involves the log wealth-consumption ratio i. We consequently need to first solve (65) to get i, and then solve (74) to get j.

C Estimation procedure

To fit our continuous-time model to the data, we first discretize the filtered dynamics in Equations (10), (11), and (13) using the following approximations

$$\log\left(\delta_{t+\Delta}/\delta_t\right) = \left(f_t + \widehat{l}_t - \frac{1}{2}\sigma_\delta^2\right)\Delta + \sigma_\delta\sqrt{\Delta}\epsilon_{1,t+\Delta}, \tag{75}$$

$$f_{t+\Delta} = e^{-\widehat{\lambda}_t\Delta} f_t + \left(1 - e^{-\widehat{\lambda}_t\Delta}\right)\bar{f} + \sigma_f\sqrt{\frac{1-e^{-2\widehat{\lambda}_t\Delta}}{2\widehat{\lambda}_t}}\epsilon_{2,t+\Delta}, \tag{76}$$

$$\widehat{l}_{t+\Delta} = \widehat{l}_t + \frac{\bar{\nu}_l}{\sigma_\delta}\sqrt{\Delta}\epsilon_{1,t+\Delta}, \tag{77}$$

$$\widehat{\lambda}_{t+\Delta} = \widehat{\lambda}_t + \frac{(\bar{f}-f_t)\nu_{\lambda,t}}{\sigma_f}\sqrt{\Delta}\epsilon_{2,t+\Delta} \tag{78}$$

$$\nu_{\lambda,t+\Delta} = \nu_{\lambda,t} + \left[\sigma_\lambda^2 - \left(\frac{(\bar{f}-f_t)\nu_{\lambda,t}}{\sigma_f}\right)^2\right]\Delta, \tag{79}$$

where $\epsilon_{1,t}$, $\epsilon_{2,t}$ are independent normally distributed random variables with mean 0 and variance 1.

We use the mean analyst forecast on the 1-quarter-ahead real GDP growth as a proxy for the expected growth rate f_t and the realized real GDP growth as a proxy for the output growth $\log(\delta_{t+\Delta}/\delta_t)$. The time interval is $\Delta = 1/4$. The system above shows that, conditional on knowing the parameters of the model and the prior means and uncertainty $(\widehat{l}_0, \widehat{\lambda}_0, \nu_{\lambda,0})$, the time series of the GDP growth forecast and realized GDP growth allow us to sequentially back out the time series of the posterior means and uncertainty $(\widehat{l}_t, \widehat{\lambda}_t, \nu_{\lambda,t})$ as well as the noises $(\epsilon_{1,t+\Delta}, \epsilon_{1,t+\Delta})$ for $t = \Delta, 2\Delta, 3\Delta \ldots$

	Mean	Standard deviation	5-percentile	95-percentile
f_t	0.0261	0.0170	-0.0099	0.0569
$\widehat{\lambda}_t$	1.0434	0.3618	0.4180	1.5321
$\nu_{\lambda,t}$	0.3227	0.1568	0.1327	0.4895
$\widehat{l}_t$	0.0013	0.0056	-0.0072	0.0097

Table 4: Descriptive statistics of the main variables.
This table reports the descriptive statistics of the state variables in the economy: the forecast f_t, the estimated adjustment $\widehat{l}_t$, the estimated mean-reversion speed $\widehat{\lambda}_t$, and the uncertainty about the mean-reversion speed $\nu_{\lambda,t}$. The statistics for $\widehat{\lambda}_t$ and $\nu_{\lambda,t}$ are obtained using the model of learning about persistence. The statistics for $\widehat{l}_t$ are obtained using the model of learning about level.

Maximizing the log-likelihood function L

$$L(\Theta; u_\Delta, \ldots, u_{N\Delta}) = \sum_{i=1}^{N} \log\left(\frac{1}{2\pi\sqrt{|\Sigma_{(i-1)\Delta}|}}\right) - \frac{1}{2} u_{i\Delta}^\top \Sigma_{i\Delta}^{-1} u_{i\Delta}, \tag{80}$$

where $\Theta \equiv (\Theta_1, \widehat{l}_0, \widehat{\lambda}_0)^\top$, N is the number of observations, $\top$ is the transpose operator, and $|.|$ is the determinant operator yields the vector of parameters $\Theta_1 \equiv (\sigma_\delta, \bar{f}, \sigma_f, \sigma_l, \sigma_\lambda)$. The prior means $\widehat{l}_0$ and $\widehat{\lambda}_0$ are estimated together with the vector of parameters Θ_1. The 2-dimensional vector u satisfies

$$u_{t+\Delta} \equiv \begin{pmatrix} u_{1,t+\Delta} \\ u_{2,t+\Delta} \end{pmatrix} = \begin{pmatrix} \log(\delta_{t+\Delta}/\delta_t) - \left(f_t + \widehat{l}_t - \frac{1}{2}\sigma_\delta^2\right)\Delta \\ f_{t+\Delta} - e^{-\widehat{\lambda}_t \Delta} f_t - \left(1 - e^{-\widehat{\lambda}_t \Delta}\right)\bar{f} \end{pmatrix} = \begin{pmatrix} \sigma_\delta \sqrt{\Delta}\epsilon_{1,t+\Delta} \\ \sigma_f \sqrt{\frac{1-e^{-2\widehat{\lambda}_t\Delta}}{2\widehat{\lambda}_t}}\epsilon_{2,t+\Delta} \end{pmatrix}. \tag{81}$$

Therefore, the conditional expectation and conditional variance-covariance matrix of $u_{t+\Delta}$ are

$$\mathbb{E}_t(u_{t+\Delta}) = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \qquad \text{Var}_t(u_{t+\Delta}) \equiv \Sigma_t = \begin{pmatrix} \sigma_\delta^2 \Delta & 0 \\ 0 & \sigma_f^2 \frac{1-e^{-2\widehat{\lambda}_t\Delta}}{2\widehat{\lambda}_t} \end{pmatrix}. \tag{82}$$

C.1 Descriptive statistics of the state variables

Table 4 reports statistics describing the level, time-variation, and range of the state variables. The growth rate forecast f_t is 2.6% on average and fluctuates mostly between -1% to 6%. The estimated mean-reversion speed $\widehat{\lambda}_t$ varies strongly over time, fluctuating mostly between 0.4 and 1.6. Finally, uncertainty about the mean-reversion speed $\nu_{\lambda,t}$ varies also substantially, fluctuating mostly between 0.1 and 0.5. Overall, these results suggest that persistence clearly fluctuates over time.

The estimated adjustment in the forecast $\widehat{l}_t$ is close to zero, on average, thereby confirming the view that professional forecasters provide accurate forecasts. This result explains why learning about the level of the expected growth rate may be irrelevant when such forecasts are available to investors.

D Numerical evaluation of $I_{\widehat{\lambda}}/I$

According to our discussion of Conjecture 1, it is not clear whether the sign of $I_{\widehat{\lambda}}/I$ is always positive. The sgn may depend on the utility parameters γ and ψ, as well as on the value of the state variables f_t and $\widehat{\lambda}_t$. In this section we numerically evaluate $I_{\widehat{\lambda}}/I$ and discuss the results.

(a) $f_t = -1\%, \widehat{\lambda}_t = 1$

γ/ψ	1.5	2	2.5
8	1.2	2.1	2.7
10	1.4	2.3	2.9
12	1.3	2.1	2.6

(c) $f_t = 2.6\%, \widehat{\lambda}_t = 1$

γ/ψ	1.5	2	2.5
8	0.8	1.4	1.9
10	1.1	1.8	2.3
12	1.1	1.8	2.3

(e) $f_t = 6\%, \widehat{\lambda}_t = 1$

γ/ψ	1.5	2	2.5
8	0.3	0.6	0.9
10	0.6	1.1	1.4
12	0.7	1.2	1.5

(b) $f_t = 2.6\%, \widehat{\lambda}_t = 0.4$

γ/ψ	1.5	2	2.5
8	-0.1	0.0	0.1
10	0.3	0.6	0.8
12	0.5	0.8	0.9

(d) $f_t = 2.6\%, \widehat{\lambda}_t = 1$

γ/ψ	1.5	2	2.5
8	0.8	1.4	1.9
10	1.1	1.8	2.3
12	1.1	1.8	2.3

(f) $f_t = 2.6\%, \widehat{\lambda}_t = 1.6$

γ/ψ	1.5	2	2.5
8	0.8	1.5	2.0
10	1.3	2.2	3.0
12	1.5	2.5	3.2

Table 5: Values of the coefficient $I_{\widehat{\lambda}}/I$
This table reports a numerical evaluation of $I_{\widehat{\lambda}}/I$. We consider different levels of risk aversion, $\gamma \in \{8, 10, 12\}$, and different levels of the elasticity of intertemporal substitution, $\psi \in \{1.5, 2, 2.5\}$. There are six panels. The upper panels keep $\widehat{\lambda}_t = 1$ but use different values for f_t. The lower panels keep $f_t = \bar{f} = 2.6\%$ but consider different values $\widehat{\lambda}_t$. For all tables, the uncertainty is $\nu_{\lambda,t} = 0.3$ (the effect of a change in uncertainty on $I_{\widehat{\lambda}}/I$ is relatively weak). Unless otherwise specified, we consider the calibration provided in the second column of Table 1.

Table 5 reports the values of the coefficient $I_{\widehat{\lambda}}/I$ in several situations. The different panels of Table 5 correspond to various levels of the forecast f_t and mean-reversion speed $\widehat{\lambda}$. Within each panel, we compute the value of $I_{\widehat{\lambda}}/I$ for different preference parameters γ and ψ.

The results indicate that the coefficient becomes smaller for lower values of risk aversion and elasticity of intertemporal substitution. It essentially remains positive, unless the risk aversion, the EIS, and the mean-reversion speed are all sufficiently small. Importantly, $I_{\widehat{\lambda}}/I$ is always positive with our preference parameters ($\gamma = 10, \psi = 2$), thus confirming our Conjecture 1 numerically.

The coefficient $I_{\widehat{\lambda}}/I$ tends to increase in bad times (i.e. lower f_t) and when the speed of mean-reversion is higher. The fact that $I_{\widehat{\lambda}}/I$ becomes smaller in good times is related to the effect discussed in Section 3, Footnote 21. More precisely, positive shocks in good times do not only signal higher persistence (which is bad for the agent), but also a longer economic boom (which is good). However, because the term $I_{\widehat{\lambda}}/I$ remains positive in good times, the second effect appears to be small.

E Price-dividend ratio and risk-free rate

This Appendix provides additional results for Sections 3.2 and 3.3. Figure 6 plots the log price-dividend ratio and the equilibrium risk-free rate. The log price-dividend ratio increases with the

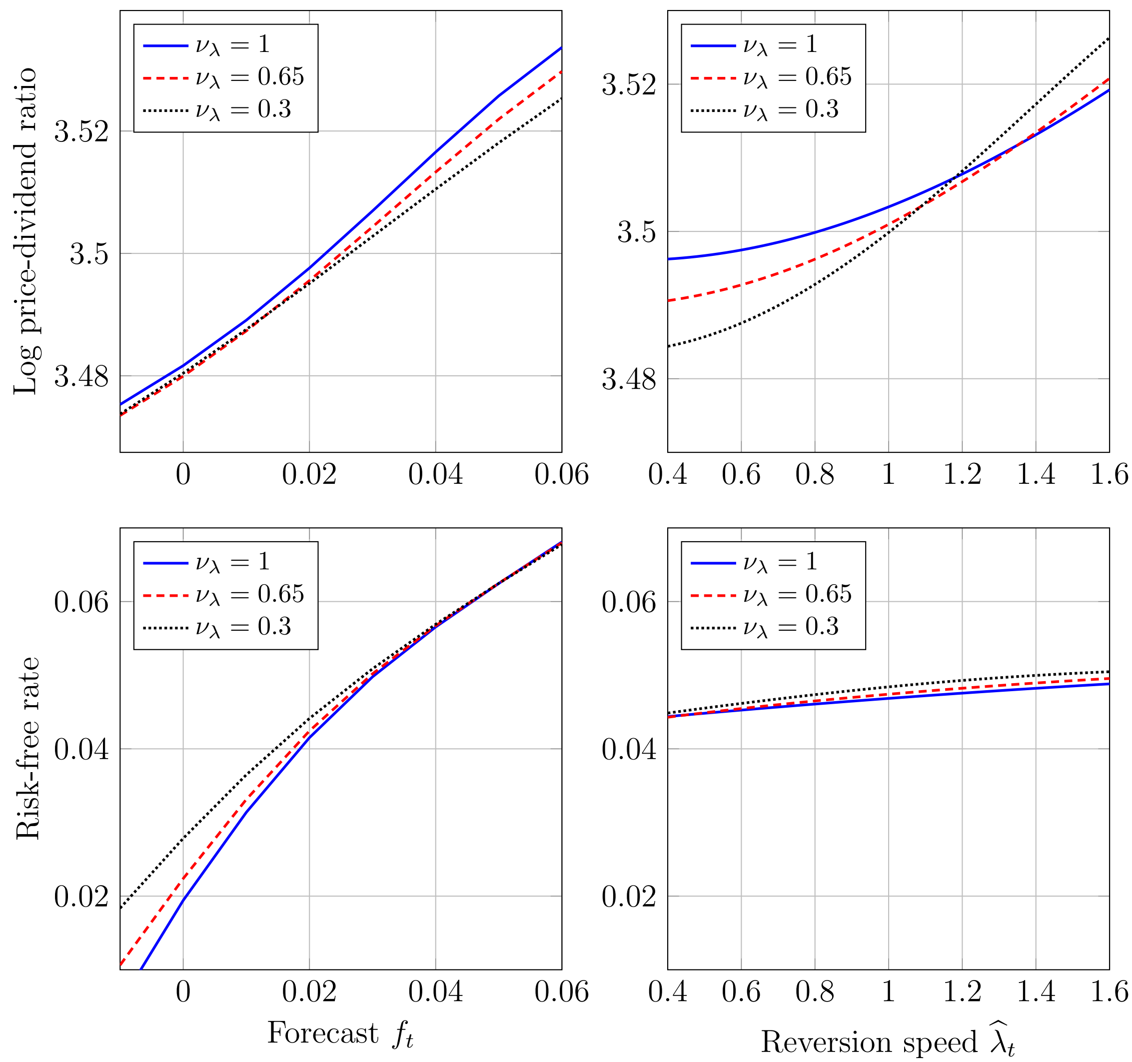

Figure 6: Behavior of the price-dividend ratio and risk-free rate with learning about persistence.
This figure shows how the price-dividend ratio and the equilibrium risk-free rate vary with the state variables. For the left graph, we fix $\widehat{\lambda}_t = 1$. For the right graph, we fix $f_t = \bar{f}$. Unless otherwise specified, we consider the calibration provided in the second column of Table 1.

output growth forecast f_t (upper left panel). The relationship is almost linear, implying that I_f/I is positive and close to being a constant. In the upper right panel, the log price-dividend ratio increases with the filter of the persistence parameter $\widehat{\lambda}_t$. This implies $I_{\widehat{\lambda}}/I > 0$.

In a model with learning about persistence, the risk-free rate satisfies

$$r_t = \beta + \frac{1}{\psi} f_t - \frac{\gamma(1 + 1/\psi)}{2}\sigma_\delta^2 - \frac{1}{2}(1-\phi)\left(\sigma_f \frac{I_f}{I} + \frac{(\bar{f} - f_t)\nu_{\lambda,t}}{\sigma_f}\frac{I_{\widehat{\lambda}}}{I}\right)^2. \quad (83)$$

The lower panels of Figure 6 depict the behavior of the equilibrium risk-free rate. The risk-free rate increases with growth forecast f_t (lower left panel) and with the filter $\widehat{\lambda}_t$ (lower right panel).

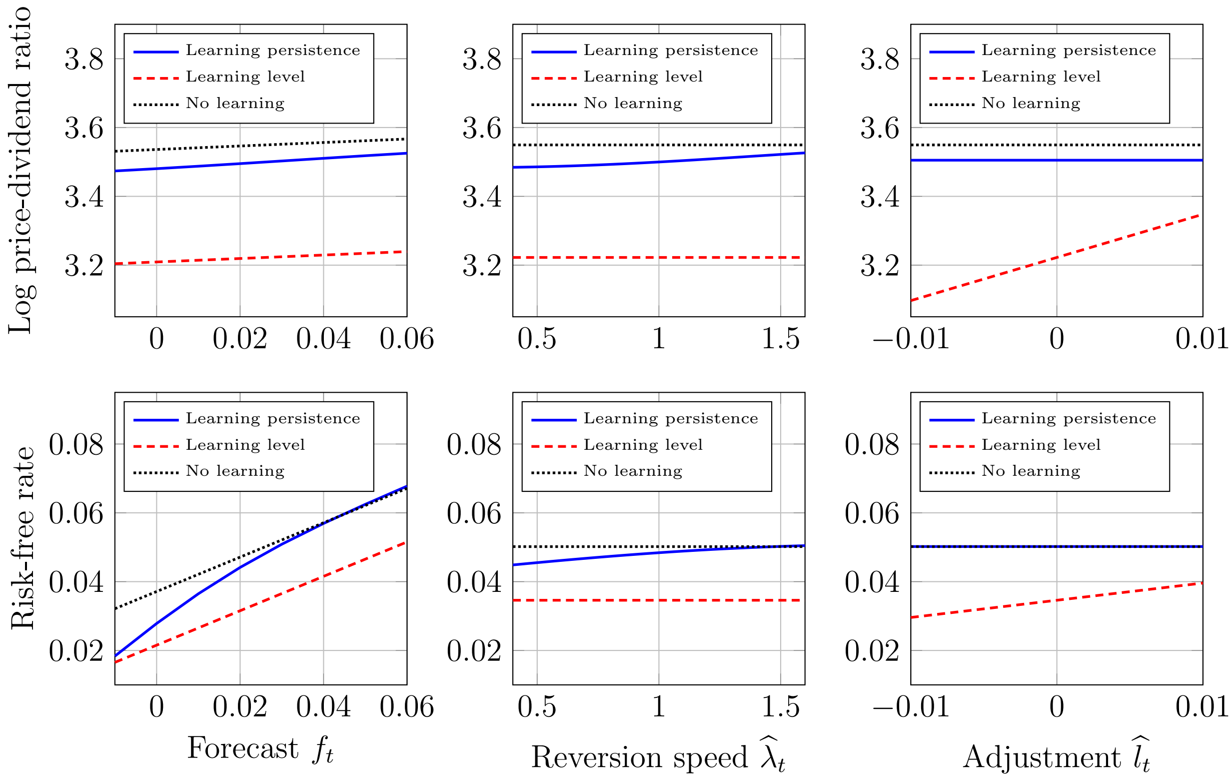

Figure 7: Price-dividend ratio and risk-free rate with and without learning. This figure compares the price-dividend ratio and the equilibrium risk-free rate in three cases: (i) a model with learning about persistence; (ii) a model with learning about level; and (iii) a model without learning. For the left plots, we fix $\widehat{l}_t = 0$ and $\widehat{\lambda}_t = 1$. For the middle plots, we fix $f_t = \bar{f}$ and $\widehat{l}_t = 0$. For the right plots, we fix $f_t = \bar{f}$ and $\widehat{\lambda}_t = 1$. We set $\nu_{\lambda,t} = 0.3$ in all plots. Unless otherwise specified, we consider the calibration provided in Table 1.

Uncertainty about persistence decreases the risk-free rate but its impact is weak.

Figure 7 compares the dependence of the log price-dividend ratio and the risk-free rate on state variables in three cases: (i) a model with learning about persistence, (ii) a model with learning about level, and (iii) a model without learning. All types of learning imply that the risk-free rate and price-dividend ratio increase with the forecast f_t. Unsurprisingly, both the risk-free rate and the price-dividend ratio increase with the estimated mean-reversion speed $\widehat{\lambda}_t$ and with the level adjustment $\widehat{l}_t$ when the agent learns about persistence and the level, respectively.