

Asset Pricing with Persistence Risk

Daniel Andrei* Michael Hasler[†] Alexandre Jeanneret[‡]

September 5, 2017

Abstract

Persistence risk is a form of uncertainty that arises when a rational agent learns about the length of business cycles. Persistence risk is positive during recessions and negative during expansions. This property, which results exclusively from learning about persistence, generates high and time-varying volatility, risk premia, and Sharpe ratios, although the estimated degree of persistence in the data is actually low. Persistence risk is a good predictor of future excess returns, both in the equilibrium model and in the data. Unlike long-run risk models, our framework delivers high predictability of equity returns, while at the same time maintaining low predictability of future consumption growth.

We are grateful to Michael Brennan, Mike Chernov, Lars Lochstoer, Avanihar Subrahmanyam, an anonymous referee, and Ron Kaniel (editor) 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), Andrey Ermolov (EFA 2017 discussant), Chanik Jo, Mariana Khapko, Gang Li, Erwan Morellec, Chayawat Ornthanalai, Marcel Rindisbacher, Gustavo Schwenkler, and Alexandre Ziegler 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, UCLA, the Einaudi Junior Conference, HEC Montréal, University of Zurich, HEC Lausanne, and the EFA 2017 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 long-run risk model developed by [Bansal and Yaron \(BY, 2004\)](#) postulates that investors demand a high risk premium in response to persistent fluctuations in aggregate consumption growth. This theory has intuitive appeal and is increasingly popular, but is also controversial on empirical grounds. The reason is that a high degree of persistence in consumption growth is necessary for the empirical performance of the long-run risk model. [Beeler and Campbell \(BC, 2012\)](#), however, challenge both the calibration of the model and its asset pricing implications. They find weaker persistence using postwar consumption data,¹ and demonstrate that a high degree of persistence leads to counterfactual predictions: in the long-run risk model stock prices predict future consumption growth but not future excess returns, whereas the data suggest the opposite.²

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. We show that a model in which a representative agent faces uncertainty about persistence (i.e., *persistence risk*) deals favorably with the above critiques. The level and the dynamics of the equity risk premium, volatility, and Sharpe ratio generated by the model are all in line with the data. Yet, the degree of persistence estimated using U.S. data is low. The model endogenously generates countercyclical time-variation in asset pricing moments over the business cycle,³ without assuming any exogenous fluctuations in output growth volatility, and overcomes the predictability challenges of the long-run risk theory.

We build a representative agent economy with imperfect information. As in the long-run risk model, the agent has Epstein-Zin preferences, favors early resolution of uncertainty, and the expected growth of the endowment process is persistent. In contrast with the long-run risk model, we assume that the persistence parameter is unobservable. To keep our setup close to the incomplete information literature, we also assume that the level of expected growth is unobservable. Thus, in our model the agent updates her beliefs simultaneously about the degree of persistence and the level of expected output growth. We show that all asset pricing moments are constant with learning about the level of expected growth,

¹Similar results are obtained with output growth, dividends growth, or earnings growth. See also [Belo, Collin-Dufresne, and Goldstein \(2015\)](#).

²There are other critiques of the long-run risk model. First, using data on S&P 500 options, [van Binsbergen, Brandt, and Kojen \(2012\)](#) provide evidence that equity risk premia and volatility decrease with maturity, at odds with the long-run risk model (See also [van Binsbergen, Hueskes, Kojen, and Vrugt \(2013\)](#), [van Binsbergen and Kojen \(2016\)](#), and [van Binsbergen and Kojen \(2017\)](#)). Second, [Epstein, Farhi, and Strzalecki \(2014\)](#) show that the high degree of persistence required in BY would inflate timing premia to seemingly implausible levels.

³In line with the empirical evidence of [Schwert \(1989\)](#), [Ferson and Harvey \(1991\)](#), [Lettau and Ludvigson \(2010\)](#), [Lustig and Verdelhan \(2012\)](#), among others.

but become strongly countercyclical with learning about persistence. This arises exclusively because the degree of persistence is unobservable—none of the above predictions arise in an alternative model in which the degree of persistence is observable, yet time-varying. Hence, uncertainty about persistence is not only sufficient, but necessary for the above results.

The theoretical mechanism through which asset pricing moments vary over the business cycle results from rational learning about the degree of persistence. During recessions, bad news about economic growth increases agent’s perception that the economic growth is more persistent. The agent demands a high risk premium to bear this *positive* persistence risk. The logic reverts during expansions, when the effect of news is mitigated by learning about persistence. Persistence risk is now *negative*, and the agent demands a smaller risk premium. Thus, learning creates an asymmetry through which risk premia and asset price fluctuations are amplified during recessions.

This asymmetry arises without assuming any exogenous fluctuations in the volatility of economic growth. In fact, we intentionally construct a model with no built-in asymmetry and with constant output growth volatility. In this sense, our model corresponds to the Case I of [BY](#), in which asset pricing moments are constant. To obtain time variation in asset pricing moments, [BY](#), and subsequently [Bansal, Kiku, and Yaron \(BKY, 2012\)](#), resort to “driving-process assumptions” ([Cochrane, 2017](#)), i.e., exogenous and persistent variation in the volatility of output growth.⁴ In our case, persistence risk is a type of uncertainty that arises purely from rational learning.

We fit the model to real GDP growth and analyst forecast data over the period Q4:1968 to Q4:2016 by Maximum Likelihood ([Hamilton, 1994](#)). The estimation shows that the degree of persistence is much lower than in the long-run risk literature, but varies significantly over time. A model with learning about persistence fits the data well and is statistically close to alternative setups with incomplete information, such as a model with learning about the level of expected growth, or a model with learning on both dimensions. Our estimation further shows that survey forecasts are reasonably accurate in predicting the expected output growth. By contrast, we are not aware of any available survey forecast about the persistence of expected output growth. Hence, asset prices should depend on how agents learn about this dimension.

We implement several empirical tests of our model with learning about persistence. We start by generating model-implied time series of stock return volatility, risk premium, and Sharpe ratio over the sample period. We compare these model-implied quantities with their observed counterparts, using market data from CRSP. The model-implied and observed

⁴The role of fluctuating economic uncertainty is analyzed in the Case II of [BY](#) and discussed in [BC, Campbell, Giglio, Polk, and Turley \(2017\)](#), and [Cochrane \(2017\)](#).

time series are well aligned. The model matches the levels and the variation of asset pricing moments, although we do not use financial market data in our calibration.

We derive testable predictions that are specific to our theory. First, the equity risk premium, the Sharpe ratio and the stock return volatility increase non-linearly with persistence risk. We directly test these positive and non-linear relationships and find supporting evidence. Second, because greater persistence risk commands a higher risk premium, it should positively predict future returns. The data lend robust support to this prediction, even after controlling for the log price-dividend ratio, the stock return volatility, the default premium, and the macro uncertainty index built by [Jurado, Ludvigson, and Ng \(2015\)](#). This suggests that persistence risk captures a form of uncertainty that is distinct from other measures proposed in the literature. Third, our model implies that the predictive power of persistence risk should be concentrated during times when uncertainty about persistence matters most. These are times when the expected growth forecast is away from its long-term mean. When splitting our sample into “high informative times” and “low informative times,” we find that the return predictability in the data is entirely concentrated during high informative times, as our model predicts.

Our framework also addresses two predictability concerns of the long-run risk model. First, a large literature shows that the price-dividend ratio predicts excess returns (e.g., [Fama and French, 1988](#); [Hodrick, 1992](#); [Cochrane, 2008](#)). Yet, as [BC](#) point out, this relationship is absent or weaker in the long-run risk model. In our setup, expected returns are high during recessions and when the perceived degree of persistence is high. These are periods in which the price-dividend ratio is depressed. Consequently, the price-dividend ratio negatively predicts future excess returns, reconciling theory and evidence. The second predictability concern of the long-run risk model is the excessive predictability of future consumption growth by the price-dividend ratio. Our model features a considerably lower degree of persistence, which helps maintain a low predictability of future consumption growth, in line with the data.

The asymmetric response of asset prices to news relates our work to the habit literature. [Campbell and Cochrane \(1999\)](#) reverse-engineer preferences in order to capture a large and countercyclical equity premium. In our case, a large risk premium arises endogenously during recessions (when consumption is low relative to its recent past) through the learning problem of the agent, rather than through the preference structure. This feature also constitutes an improvement of the long-run risk model, in which countercyclical risk premia arise merely through a coincidence of large fear of long-run risk when current consumption goes down (point extensively discussed by [Cochrane, 2017](#)). We validate this argument by showing that our model-implied price-dividend ratio and equity premium endogenously become history

dependent, as in the habit formation model.

Most of the literature on learning assumes that the unobservable dimension is the level of expected output growth.⁵ None of our implications about the dynamics of asset prices obtain with this form of learning. [Pakoš \(2013\)](#) analyzes an economy where growth follows 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 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.

[Collin-Dufresne, Johannes, and Lochstoer \(2016\)](#) show that parameter uncertainty and preference for early resolution of uncertainty generate endogenous long-run risk and a large equilibrium risk premium. Parameter uncertainty plays a key role in our model as well. But the focus of our paper is different in three respects. First, we investigate a particular form of parameter uncertainty, i.e. persistence risk. Second, we estimate the parameters of the output process using growth forecasts and find a considerable lower degree of persistence than in the long-run risk literature. This lower degree of persistence helps our model to successfully overcome the empirical hurdles of the long-run risk theory. Third, we develop, test, and find empirical support for a set of predictions specific to our model with persistence risk.

Our theory also connects with the ambiguity aversion and robust-control literature (see [Hansen, 2014](#), for a comprehensive review). In [Hansen and Sargent \(2010\)](#), a representative agent with *concerns about robustness* assigns positive probabilities to an i.i.d. or a long-run risk framework. The agent deals with model uncertainty by always picking the worst-case scenario. As a result, the price of risk becomes countercyclical and history-dependent. In our case, this does not arise because the agent chooses to be pessimistic at all times, but because the agent learns about persistence.⁶

⁵See [David \(1997\)](#), [Veronesi \(1999, 2000\)](#), [Brennan and Xia \(2001\)](#), [David and Veronesi \(2002\)](#) among many others, and comprehensive surveys by [Pastor and Veronesi \(2009\)](#) and [Ziegler \(2012\)](#). [Veronesi \(1999\)](#) shows that volatility and risk premia are hump-shaped functions of the state of the economy when a representative agent learns about a discrete-state output growth. [Ai \(2010\)](#) analyzes learning in a production economy with Epstein-Zin preferences and finds a negative relation between information quality and the risk premium. [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.

⁶Along the same lines, [Drechsler \(2013\)](#) shows that model uncertainty can explain the variance risk premium. In [Bidder and Dew-Becker \(2016\)](#), ambiguity-averse investors fear a “worst-case” model in which shocks to expected consumption growth have a half-life of 70 years, much longer than in [BY](#) (about 3 years). Our estimation shows that the persistence is much weaker than what is assumed in the long-run risk literature. Furthermore, in our case the agent has standard preferences and updates the economic model through Bayesian learning.

In recent work, [Ghosh and Constantinides \(2016\)](#) argue that consumers learn about the economic regime not only from the history of consumption, but also from other signals. Similarly, in our model the representative agent observes not only the output growth, but also survey forecasts about the expected output growth.

A growing literature uses analyst forecast dispersion to investigate the impact of uncertainty on asset prices.⁷ [Bansal and Shaliastovich \(2010\)](#) and [Drechsler \(2013\)](#) show that exogenous fluctuations in uncertainty can explain large movements in asset prices and the variance risk premium. For uncertainty to have a sizable impact on asset prices, [Bansal and Shaliastovich \(2010\)](#) rely on recency-biased learning (as opposed to Bayesian learning) and [Drechsler \(2013\)](#) relies on fear of model misspecification. We also conclude that fluctuations in uncertainty are important for asset prices. Our approach is distinct, however, because we provide a theoretical foundation for fluctuations in uncertainty. Furthermore, behavior biases or fear of model uncertainty are not necessary for our results.

[Andrei, Carlin, and Hasler \(2017\)](#) build an equilibrium model in which CRRA agents disagree about the length of business cycles. Disagreement yields time variation in risk premia and dictates the strength of the risk-return relationship. In our model, a representative agent learns about both the persistence and the level of expected output growth. We show that the equity return volatility, risk premium and Sharpe ratio are high and time-varying, although the estimated persistence is low. This arises because the agent faces uncertainty about the degree of persistence and prefers early resolution of uncertainty.

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 predictions. Section 5 concludes and offers directions for future research.

2 Model

In this section we introduce the economic model. We keep our formulation intentionally close to the long-run risk model of [BY](#). The main difference is that the representative agent operates under incomplete information. We solve the agent's learning problem and characterize the equilibrium asset prices.

⁷Other measures of uncertainty have been proposed in the literature. Two prominent examples are [Bloom \(2009\)](#) and [Jurado et al. \(2015\)](#). Because our measure of persistence risk is not based on financial market data, it is closer to the measure proposed by [Jurado et al. \(2015\)](#), with the main difference being that our measure results from a theoretical model of learning.

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 discount rate β , relative risk aversion γ , and elasticity of intertemporal substitution ψ . The indirect utility function is given by

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

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

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

with $\phi \equiv \frac{1-\gamma}{1-1/\psi}$. Standard CRRA utility obtains if $\phi = 1$. The agent prefers early resolution of uncertainty when $1 - \phi > 0$.

The agent observes two processes. First, δ_t , represents the process of the aggregate output in the economy. Second, f_t , represents the outcome of a survey about the expected output growth from professional forecasters:

$$d\delta_t/\delta_t = \mu_t dt + \sigma_\delta dW_t^\delta \quad (3)$$

$$df_t = \theta_t(\bar{\mu} - f_t)dt + \sigma_f dW_t^f, \quad (4)$$

where W_t^δ and W_t^f are standard Brownian motions. The expected output growth rate μ_t depends on the survey forecast f_t in a way that we will describe below.

As in the long-run risk model, the representative agent observes the output process and a time-varying component of the expected output growth rate. In contrast with the long-run risk model, however, we are assuming away exogenous fluctuations in output growth volatility. That is, σ_δ and σ_f are constants, whereas they are time-varying in BY and BKY. This modeling choice helps us isolate the effect of learning on asset prices.

We further depart from the long-run risk model by assuming that the agent operates under incomplete information. Specifically, we consider two dimensions of uncertainty. First, the expected output growth rate, μ_t , is unobservable by the agent. Instead, the agent observes the expected growth forecast provided by the survey f_t , which is given by

$$f_t = \mathbb{E}[\mu_t \mid \text{Information set of forecasters}]. \quad (5)$$

This formulation can be interpreted as follows: professional forecasters have access to pro-

proprietary information and build the best posterior estimate of μ_t given their information set. The true expected growth rate, μ_t , is then equal to their estimate plus an unobservable, independent innovation:⁸

$$\mu_t = \underbrace{\mathbb{E}[\mu_t \mid \text{Information set of forecasters}]_{f_t}} + \epsilon_t. \quad (6)$$

For the rest of the paper, we will refer to the innovation ϵ_t as the *survey error*. To ensure that this survey error is unbiased in the long term, we let it fluctuate around zero:

$$d\epsilon_t = -\varphi\epsilon_t dt + \sigma_\epsilon dW_t^\epsilon. \quad (7)$$

The second dimension of uncertainty is about the persistence of the expected output growth. We assume that the mean-reversion speed θ_t is time-varying and unobservable. It has two components: an observable long-term average and an unobservable, time-varying noise with zero mean. Thus, we define $\theta_t = \bar{\theta} + \lambda_t$, where λ_t follows

$$d\lambda_t = -\kappa\lambda_t dt + \sigma_\lambda dW_t^\lambda. \quad (8)$$

The parameters $\bar{\mu}$, $\bar{\theta}$, φ , κ , σ_δ , σ_f , σ_ϵ , and σ_λ are positive constants known by the agent and the four Brownians W_t^δ , W_t^f , W_t^ϵ , and W_t^λ are independent.⁹

We can alternatively assume that ϵ_t and θ_t are unobservable constants. Their posterior estimates would then become martingales through Bayesian learning (as shown by Collin-Dufresne et al., 2016), and all the asset pricing implications would hold. However, this alternative is less desirable for our purposes because it makes the model non-stationary. Non-stationarity precludes a proper empirical assessment of the model.

The economic environment described above 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.¹⁰ In what follows, we refer to agent’s updating about ϵ_t as *learning about level* and to agent’s updating about λ_t as *learning about persistence*. We can exclude any of these two dimensions of uncertainty by assuming $\sigma_\epsilon = 0$ and/or $\sigma_\lambda = 0$. This

⁸Notice that this formulation is slightly different from the standard approach adopted in the literature, where the forecast is modeled as truth plus noise (from the perspective of the economic agent). This is because in our model f_t results from the *learning exercise of professional forecasters*, and thus it is a posterior estimate that must satisfy Eq. (5).

⁹The assumption of independent shocks simplifies the description of the model without changing the main message. The model can be extended to allow for non-zero correlations between the four Brownians, at the cost of adding parameters and making the estimation (described in Section 3.1) less stable.

¹⁰The feature of the model that preserves the linearity of the learning exercise with both dimensions of uncertainty is the observability of the survey forecast f_t .

allows us to disentangle the impact of each type of learning on asset prices.

2.2 Learning

Taken together, the dynamics described in (3)-(8) imply

$$d\delta_t/\delta_t = (f_t + \epsilon_t)dt + \sigma_\delta dW_t^\delta \quad (9)$$

$$df_t = (\bar{\theta} + \lambda_t)(\bar{\mu} - f_t)dt + \sigma_f dW_t^f, \quad (10)$$

with the dynamics of ϵ_t and λ_t (both unobservable) provided respectively in (7) and (8).

Denote by \mathcal{F}_t the information set of the agent at time t , by $\hat{\epsilon}_t \equiv \mathbb{E}[\epsilon_t|\mathcal{F}_t]$ the estimated survey error, and its posterior variance by $\nu_{\epsilon,t} \equiv \mathbb{E}[(\epsilon_t - \hat{\epsilon}_t)^2|\mathcal{F}_t]$. Similarly, denote by $\hat{\lambda}_t \equiv \mathbb{E}[\lambda_t|\mathcal{F}_t]$ the estimated unobservable component of the mean-reversion speed and its posterior variance by $\nu_{\lambda,t} \equiv \mathbb{E}[(\lambda_t - \hat{\lambda}_t)^2|\mathcal{F}_t]$. Thus,

$$\epsilon_t \sim N(\hat{\epsilon}_t, \nu_{\epsilon,t}), \quad \lambda_t \sim N(\hat{\lambda}_t, \nu_{\lambda,t}), \quad (11)$$

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

We refer to the estimates $\hat{\epsilon}_t$ and $\hat{\lambda}_t$ as the *filters*, and to the two posterior variances $\nu_{\epsilon,t}$ and $\nu_{\lambda,t}$ as the *uncertainties*. The filters evolve according to (Liptser and Shirayayev, 1977):

$$\begin{bmatrix} d\hat{\epsilon}_t \\ d\hat{\lambda}_t \end{bmatrix} = \begin{bmatrix} -\varphi & 0 \\ 0 & -\kappa \end{bmatrix} \begin{bmatrix} \hat{\epsilon}_t \\ \hat{\lambda}_t \end{bmatrix} dt + \begin{bmatrix} \frac{\nu_{\epsilon,t}}{\sigma_\delta} & 0 \\ 0 & \frac{(\bar{\mu} - f_t)\nu_{\lambda,t}}{\sigma_f} \end{bmatrix} \begin{bmatrix} d\widehat{W}_t^\delta \\ d\widehat{W}_t^f \end{bmatrix}, \quad (12)$$

where

$$d\widehat{W}_t^\delta = \frac{1}{\sigma_\delta} \left(\frac{d\delta_t}{\delta_t} - (f_t + \hat{\epsilon}_t)dt \right), \quad d\widehat{W}_t^f = \frac{1}{\sigma_f} \left(df_t - (\bar{\theta} + \hat{\lambda}_t)(\bar{\mu} - f_t)dt \right), \quad (13)$$

are independent Brownian motions under the filtration \mathcal{F}_t (see Appendix A).¹¹ For the sake of brevity, we will hereafter use the term *output growth shocks* to refer to $d\widehat{W}_t^\delta$ innovations and *expected output growth shocks* to refer to $d\widehat{W}_t^f$ innovations.

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

¹¹Note that the mean-reversion speed, $\theta_t = \bar{\theta} + \lambda_t$, can theoretically become negative because it is an Ornstein-Uhlenbeck process. With our calibration, the unconditional probability of a negative mean-reversion speed is 1.7%. Since the filtered mean-reversion speed, $\hat{\theta}_t = \bar{\theta} + \hat{\lambda}_t$, is a projection of θ_t on the observation filtration \mathcal{F}_t , the probability that $\hat{\theta}_t$ becomes negative is even smaller. In fact, $\hat{\theta}_t$ is never negative over the sample Q4:1968-Q4:2016, which we use for our calibration in Section 3.1 (see Figure 1).

Learning about persistence induces a particular formation of beliefs. In this case, agent's updating depends on the state of the economy, defined by the distance between the long-run output growth and the actual growth forecast, $(\bar{\mu} - f_t)$. We refer to this distance as the *output growth gap*. In good times, when the output growth gap is below zero, positive expected output growth shocks decrease the agent's estimate of λ_t . In bad times, when the output growth gap is above zero, 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). When $f_t = \bar{\mu}$, changes in the forecast f_t are uninformative about λ_t and the agent is unable to learn about the mean-reversion speed.

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

$$\frac{d\nu_{\epsilon,t}}{dt} = \sigma_\epsilon^2 - \left(2\varphi\nu_{\epsilon,t} + \frac{\nu_{\epsilon,t}^2}{\sigma_\delta^2} \right) \quad (14)$$

$$\frac{d\nu_{\lambda,t}}{dt} = \sigma_\lambda^2 - \left(2\kappa\nu_{\lambda,t} + \frac{(\bar{\mu} - f_t)^2\nu_{\lambda,t}^2}{\sigma_f^2} \right). \quad (15)$$

Eq. (14) implies that uncertainty about ϵ_t converges to a constant. We consequently assume that it has converged to its steady-state (e.g., [Dumas, Kurshev, and Uppal, 2009](#)) (see Appendix A):

$$\bar{\nu}_\epsilon \equiv \sigma_\delta \left(\sqrt{\varphi^2\sigma_\delta^2 + \sigma_\epsilon^2} - \varphi\sigma_\delta \right). \quad (16)$$

In contrast, Eq. (15) implies that uncertainty about λ_t does not converge to a constant, due to the presence of the output growth gap, which is stochastic.¹²

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 the forecast particularly informative about the unobservable degree of persistence.

The dynamics of the two variables that result from learning (the filter $\hat{\lambda}_t$ and the uncertainty $\nu_{\lambda,t}$) depend on the product $(\bar{\mu} - f_t)\nu_{\lambda,t}$. Because this product is directly driven by

¹²It is worth noting that none of the uncertainties $\nu_{\epsilon,t}$ and $\nu_{\lambda,t}$ converge to zero. This is because ϵ_t and λ_t are perturbed by noise, as opposed to being constants, which continuously regenerates learning. It implies that the dynamics of all state variables are non-degenerate (in contrast, both uncertainties would converge to zero in the long run if we were to assume that ϵ_t and λ_t are constants).

the uncertainty about persistence, we refer to it as *persistence risk* in the rest of the paper:

$$\text{Persistence risk} \equiv (\bar{\mu} - f_t)\nu_{\lambda,t}. \quad (17)$$

Persistence risk is positive in bad times and negative in good times.

2.3 Equilibrium asset prices

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

$$\max_C \{h(C, J) + \mathcal{L}J\} = 0, \quad (18)$$

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

$$J(C, f, \hat{\epsilon}, \hat{\lambda}, \nu_{\lambda}) = \frac{C^{1-\gamma}}{1-\gamma} [\beta I(x)]^{\phi}, \quad (19)$$

where $I(x)$ is the wealth-consumption ratio, and $x \equiv [f \hat{\epsilon} \hat{\lambda} \nu_{\lambda}]^{\top}$ denotes the vector of four state variables,¹³ whose dynamics are given in (12)–(15).

Substituting the guess (19) in the HJB equation (18) and imposing the market clearing condition $C_t = \delta_t$ yields a partial differential equation (PDE) for the wealth-consumption ratio. We solve this equation numerically using Chebyshev polynomials (Judd, 1998). We refer the reader to Eq. (77) in Appendix B for the PDE and for more details about our numerical procedure.

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 wealth-consumption ratio with respect to the state variables to satisfy:*

$$I_f > 0, \quad I_{\hat{\epsilon}} > 0, \quad I_{\hat{\lambda}} > 0, \quad I_{\nu_{\lambda}} < 0. \quad (20)$$

This conjecture holds for a reasonable range of parameters. In fact, several inequalities follow directly from the guess of the value function in (19). Taking the derivative of J with

¹³We can equivalently use $\hat{\mu}_t = f_t + \hat{\epsilon}_t$ (i.e., agent's expected growth rate) as a state variable, instead of $\hat{\epsilon}_t$. However, it is notationally convenient to work with the filtered survey error.

respect to any of the four state variables yields

$$J_{(\cdot)} = \phi J \frac{I_{(\cdot)}}{I}, \quad (21)$$

with the product ϕJ being positive when $\gamma > 1 > 1/\psi$. Due to non-satiation, expected lifetime utility must rise as investment opportunities improve, and thus $J_f > 0$ and $J_{\hat{\epsilon}} > 0$ (because agent's expected output growth is the sum of f_t and $\hat{\epsilon}_t$: $\hat{\mu}_t = f_t + \hat{\epsilon}_t$). Using (21), this reasoning yields the first two inequalities of Conjecture 1. Further, the representative agent dislikes uncertainty when the coefficient of risk aversion is higher than one (Ziegler, 2003; Ai, 2010), which implies that $J_{\nu_\lambda} < 0$. This yields the last inequality of Conjecture 1. The only inequality that needs numerical validation is $I_{\hat{\lambda}} > 0$. Assuming that the agent prefers early resolution of uncertainty, we expect that she prefers less persistence (i.e. higher mean-reversion speed), which yields $J_{\hat{\lambda}} > 0$.¹⁴ Using (21), this yields $I_{\hat{\lambda}} > 0$.

Let $\sigma_I(x) \equiv [\sigma_{I1}(x) \ \sigma_{I2}(x)]$ be the diffusion vector of the wealth-consumption ratio. It has two elements:

$$\sigma_{I1}(x_t) = \frac{\bar{\nu}_\epsilon I_{\hat{\epsilon}}}{\sigma_\delta I} \quad (22)$$

$$\sigma_{I2}(x_t) = \sigma_f \frac{I_f}{I} + \frac{(\bar{\mu} - f_t)\nu_{\lambda,t} I_{\hat{\lambda}}}{\sigma_f I}. \quad (23)$$

A key implication of our setup is that the volatility of the wealth-consumption ratio is directly driven by persistence risk and therefore becomes endogenously stochastic (Eq. 23). In contrast, the long-run risk models of BY and BKY generate stochastic volatility of the wealth-consumption ratio by assuming exogenous time-varying fluctuations in output growth volatility. In our setup, stochastic volatility arises endogenously from learning. Furthermore, Conjecture 1 implies that learning increases the volatility of the wealth-consumption ratio during bad times, when persistence risk is positive.

2.3.1 Risk-free rate and market price of risk

Following Duffie and Epstein (1992), the state-price density is given by

$$\xi_t = \exp \left[\int_0^t h_J(C_s, J_s) ds \right] h_C(C_t, J_t) = \exp \left[\int_0^t \left(\frac{\phi - 1}{I(x_s)} - \beta\phi \right) ds \right] \beta^\phi C_t^{-\gamma} I(x_t)^{\phi-1}. \quad (24)$$

¹⁴Appendix B.2 reports a numerical analysis of $I_{\hat{\lambda}}$ for several values of risk aversion, intertemporal of elasticity substitution, and of the two state variables, f_t and $\hat{\lambda}_t$. We find that $I_{\hat{\lambda}}$ is positive and large in all cases. We also discuss the parametrization for which this term can become negative (which can only happen outside the standard calibration of our model).

The risk-free rate, $r_{f,t}$, and the 2-dimensional market price of risk, m_t , follow directly from the dynamics of the state price density,

$$\frac{d\xi_t}{\xi_t} = -r_{f,t}dt - m_t^\top d\widehat{W}_t, \quad (25)$$

where $\widehat{W} \equiv [\widehat{W}^\delta, \widehat{W}^f]^\top$ is defined in (13). Itô's lemma yields

$$r_{f,t} = \beta + \frac{1}{\psi}(f_t + \widehat{\epsilon}_t) - \frac{\gamma + \gamma\psi}{2\psi}\sigma_\delta^2 - (1 - \phi) \left[\sigma_{I1}(x_t)\sigma_\delta + \frac{1}{2}(\sigma_{I1}^2(x_t) + \sigma_{I2}^2(x_t)) \right] \quad (26)$$

$$m_t = \left[\gamma\sigma_\delta + (1 - \phi)\sigma_{I1}(x_t) \quad (1 - \phi)\sigma_{I2}(x_t) \right]^\top. \quad (27)$$

Fluctuations in expected output growth generate a procyclical risk-free rate, as observed from the second term in (26). Furthermore, when the agent prefers early resolution of uncertainty ($1 - \phi > 0$), variations in f_t , $\widehat{\epsilon}_t$, and $\widehat{\lambda}_t$ yield the last term in (26). The resulting effect is a lower risk-free rate due to greater demand for the safe asset.

The market price of risk has two components (Eq. 27). The uncertainty $\bar{\nu}_\epsilon$ 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. Following Conjecture 1 and Eq. (23), $I_{\widehat{\lambda}} > 0$ and the market price of risk increases in bad times.

2.3.2 Asset prices

We assume that dividends are a levered claim of output (Abel, 1999):

$$D_t = e^{-\beta_d t} \delta_t^\eta, \quad (28)$$

where $\eta \geq 1$ is the leverage parameter and $\beta_d > 0$ is a parameter that determines the growth rate of dividends. Leverage is motivated by the observation that the volatility of dividend growth is larger than the volatility of output growth in the data.¹⁵ It also makes our results directly comparable with BY and BKY, who adopt the same approach (see also Benzoni et al. (2011), BC, among others). This specification does not change the learning problem of the agent, because the process (28) does not bring any additional information.

¹⁵Using data from January 1969 to December 2016, the annualized CRSP dividend growth volatility is about 19%. BC report values between 11% and 27%. See also Drechsler (2013) for numbers of similar magnitude. In comparison, the annualized output growth volatility is 1.4%.

Given (28), the stock price is defined as a claim to the dividend process

$$\frac{dD_t}{D_t} = \left(-\beta_d + \eta\mu_t + \frac{1}{2}\eta(\eta-1)\sigma_\delta^2 \right) dt + \eta\sigma_\delta dW_t^\delta. \quad (29)$$

Define the price-dividend ratio as $\Pi(x_t)$. Its diffusion vector has two components:

$$\sigma_{\Pi 1}(x_t) = \frac{\bar{\nu}_\epsilon}{\sigma_\delta} \frac{\Pi_{\hat{\epsilon}}}{\Pi} \quad (30)$$

$$\sigma_{\Pi 2}(x_t) = \sigma_f \frac{\Pi_f}{\Pi} + \frac{(\bar{\mu} - f_t)\nu_{\lambda,t}}{\sigma_f} \frac{\Pi_{\hat{\lambda}}}{\Pi}. \quad (31)$$

Without leverage ($\eta = 1$, $\beta_d = 0$), these two components coincide with $\sigma_{I1}(x_t)$ and $\sigma_{I2}(x_t)$ from (22)-(23). With leverage, we expect (and verify in Appendix B.2) the partial derivatives of Π to satisfy the same inequalities of Conjecture 1. The partial differential equation to be solved by $\Pi(x_t)$ is provided in Appendix B.1.

2.3.3 Stock market volatility

The diffusion of stock returns, which we denote by σ_t , satisfies

$$\sigma_t = \begin{bmatrix} \eta\sigma_\delta + \sigma_{\Pi 1}(x_t) & \sigma_{\Pi 2}(x_t) \end{bmatrix} = \begin{bmatrix} \eta\sigma_\delta + \frac{\bar{\nu}_\epsilon}{\sigma_\delta} \frac{\Pi_{\hat{\epsilon}}}{\Pi} & \sigma_f \frac{\Pi_f}{\Pi} + \frac{(\bar{\mu} - f_t)\nu_{\lambda,t}}{\sigma_f} \frac{\Pi_{\hat{\lambda}}}{\Pi} \end{bmatrix}. \quad (32)$$

The leverage parameter directly increases the volatility of stock returns through the multiplication with σ_δ . Furthermore, according to Conjecture 1, we expect $\Pi_{\hat{\epsilon}} > 0$ and therefore uncertainty about the level $\bar{\nu}_\epsilon$ increases the magnitude of the first diffusion component in (32). Thus, uncertainty about the level of expected output growth can generate excess volatility in stock returns.

The process of learning about persistence described in Section 2.2 creates an asymmetric stock market response to shocks. This is due to the presence of persistence risk in the second diffusion component in (32). Stock returns react strongly to shocks when the economy is in bad times and persistence risk is positive. In contrast, when the economy is in good times, persistence risk is negative and attenuates the sensitivity of stock prices to news.

2.3.4 Equity risk premium

The equity risk premium is defined as $RP_t \equiv \sigma_t m_t$. Using Eqs. (27) and (32), we obtain

$$RP_t = [\gamma\sigma_\delta + (1 - \phi)\sigma_{I1}(x_t)] \left(\eta\sigma_\delta + \frac{\bar{\nu}_\epsilon}{\sigma_\delta} \frac{\Pi_{\hat{\epsilon}}}{\Pi} \right) + (1 - \phi)\sigma_{I2}(x_t) \left(\sigma_f \frac{\Pi_f}{\Pi} + \frac{(\bar{\mu} - f_t)\nu_{\lambda,t}}{\sigma_f} \frac{\Pi_{\hat{\lambda}}}{\Pi} \right), \quad (33)$$

where $\sigma_{I1}(x_t)$ and $\sigma_{I2}(x_t)$ are defined in (22)-(23).

The equity risk premium consists of two terms. The first term, which pertains to risk generated by output growth shocks, is higher in presence of uncertainty $\bar{\nu}_\epsilon$ about the level of expected growth (see also Ai, 2010). This arises both because the agent has a preference for early resolution of uncertainty ($1 - \phi > 0$) and because the volatility of the price-dividend ratio $\Pi(x_t)$ is amplified by $\bar{\nu}_\epsilon$.

The second term of the equity risk premium is directly driven by persistence risk and thus is specific to our setup. The model implies that the risk premium fluctuates when persistence is uncertain. Furthermore, the risk premium is higher in bad times, when persistence risk is positive, because both the market price of risk (Eq. 27) and the volatility of the price-dividend ratio (Eq. 32) are amplified when $\bar{\mu} - f_t > 0$.

3 Theoretical Predictions

We now calibrate the model to U.S. output data and evaluate its quantitative implications. We do not use asset prices in the calibration—our aim is to investigate whether a model with learning about persistence is able to produce realistic asset pricing implications without being calibrated to do so. We then quantify and discuss how asset prices vary with the state variables. Learning about persistence generates significant time variation in asset pricing moments, whereas a model with learning about the level only (or without learning) implies constant moments. Finally, we show that the time-variation in asset pricing moments is negligible in a model in which λ_t is time-varying and observable.

3.1 Calibration

We use the mean analyst forecast on 1-quarter-ahead real GDP growth as a measure of f_t and the realized real GDP growth rate 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 Q4:2016.¹⁶

We use the dynamics of the filters $\hat{\epsilon}_t$ and $\hat{\lambda}_t$ from Eq. (12), the dynamics of the uncertainty about λ from Eq. (15), and the filtered Brownian shocks from Eq. (13) 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 , $\bar{\mu}$, σ_f , σ_ϵ , σ_λ , $\bar{\theta}$, κ , and φ that provide the closest fit to realized observations. Note that the prior on

¹⁶Considering output rather than consumption data allows us to exploit a longer sample period (the time series of consumption forecasts starts only in Q3:1981).

Parameter	Symbol	Full	Persistence	Level	No learning
Vol. of output growth	σ_δ	0.0140*** (24.60)	0.0141*** (25.04)	0.0140*** (23.70)	0.0141*** (24.00)
L-T growth rate	$\bar{\mu}$	0.0254*** (9.32)	0.0255*** (10.35)	0.0256*** (6.84)	0.0256*** (7.33)
Volatility of growth forecast	σ_f	0.0237*** (28.77)	0.0238*** (23.43)	0.0230*** (21.84)	0.0230*** (29.75)
L-T reversion speed	$\bar{\theta}$	1.3134*** (6.90)	1.3467*** (6.70)	0.9601*** (5.20)	0.9601*** (7.23)
Vol. of reversion speed	σ_λ	0.2674*** (4.93)	0.2790*** (6.37)		
Reversion speed of reversion speed	κ	0.0855** (2.71)	0.0954*** (4.53)		
Volatility of survey error	σ_ϵ	0.0034*** (5.16)		0.0034*** (3.59)	
Reversion speed of survey error	φ	0.1460 (0.88)		0.1462*** (3.02)	
Likelihood ratio statistic	LR	-0.8409	-	-0.3628	0.1805
	p-value	0.200	-	0.358	0.572

Table 1: Parameter estimates

This table reports the estimates of the model parameters, obtained by Maximum Likelihood for the period Q4:1968 to Q4:2016. The table compares the estimation results of the full model with those of three special cases: i) learning about persistence only ($\hat{\epsilon}_t = 0, \forall t$); ii) learning about the level only ($\hat{\lambda}_t = 0, \hat{\theta}_t = \bar{\theta}, \forall t$); and iii) no learning ($\hat{\epsilon}_t = 0, \hat{\lambda}_t = 0, \hat{\theta}_t = \bar{\theta}, \forall t$). t -statistics are reported in brackets and statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

uncertainty about the mean-reversion speed is set to its steady-state value,¹⁷ whereas the priors $\hat{\epsilon}_0$ and $\hat{\lambda}_0$ are set to their long-term means, which are zero. Details about the estimation are provided in Appendix C.

Table 1 reports the estimated parameters for different learning models and Figure 1 displays the time series of the state variables.¹⁸ It is worth noting that the estimated average degree of persistence is much lower than what is typically considered in the long-run risk literature. The long-term mean of the mean-reversion speed is $\bar{\theta} = 1.31$ in the full model, whereas it is about 0.2 in the long-run risk models (e.g. [Bansal et al., 2016](#)). Overall, the estimation indicates that the degree of persistence estimated using U.S. data is low on average, but significantly time varying due to the high value of σ_λ (Figure 1, panel A).

In contrast, the volatility of the survey error, σ_ϵ , is relatively small (see also Figure 1,

¹⁷We 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_t at time $t = 0$ is the positive root of the polynomial obtained when $\frac{d\nu_{\lambda,t}}{dt} = 0$. Uncertainty about persistence initially starts at this level, $\nu_{\lambda,0}$, and then dynamically evolves according to Eq. (15).

¹⁸See also Table 13 in Appendix C.1 for the descriptive statistics and a discussion of these variables.

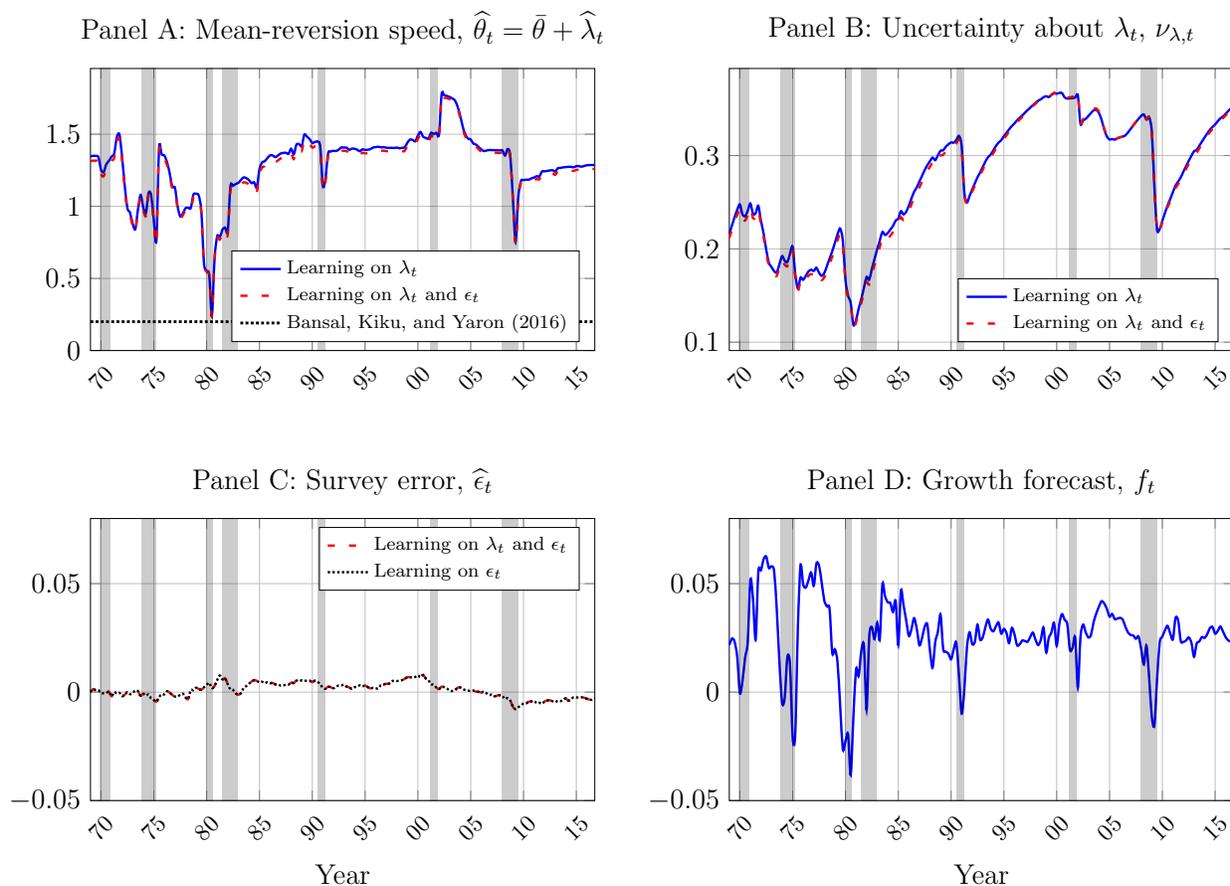


Figure 1: Historical path of the state variables.

This figure plots the time series of the main state variables. Panel A reports the filtered mean-reversion speed of expected 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 the mean-reversion speed, while Panel C shows the filtered survey error when learning about 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-Q4:2016.

panel C). A plausible interpretation for this finding is that the forecasts available from professional surveys are of good quality (e.g. [Ang, Bekaert, and Wei, 2007](#)). Thus, the main uncertainty that agents are face is less whether the economy will be in a recession or an expansion, but rather 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- or long-lived. This is the type of uncertainty that agents aim to resolve when learning about the degree of persistence.

To compare the likelihood of the different models, we follow [Vuong \(1989\)](#) and perform likelihood-ratio (LR) tests for each pair of models (see Appendix C). We choose the model

with learning about persistence (second column) as the benchmark and compare it to the alternative models. The last two rows of Table 1 provide the LR statistics and their p-values. None of the LR statistics are significantly different from zero, which shows that all models are statistically equivalent in terms of fitting observed output growth data.

We generate theoretical predictions using the calibration provided in Table 1. To keep our analysis as close as possible to [BY](#), we set the risk aversion to $\gamma = 10$ and the elasticity of intertemporal substitution (EIS) to $\psi = 1.5$. The leverage parameter is set to $\eta = 7$, the subjective discount rate to $\beta = 0.02$, and $\beta_d = 0.145$. This implies a dividend growth volatility of 10%, which is a lower bound of what is typically considered in the literature (the dividend growth volatility is 12% in [BY](#) and 16% in [BKY](#), while the CRSP dividend growth volatility is 19%). The parameters β_d and β are chosen to obtain reasonable values of the average wealth-consumption ratio ([Lustig, Van Nieuwerburgh, and Verdelhan, 2013](#)), price-dividend ratio, and dividend growth rate, which are about 80, 35, and 3.5%, respectively, in our model.

3.2 Asset prices with learning about persistence

We analyze the implications of a model with learning about persistence (whose calibration is given in the second column of Table 1). We exclude learning about level in order to isolate the contribution of learning about persistence for the dynamics of asset prices. We discuss other types of learning in the next section.

In a model with learning about persistence only, the stock return volatility is

$$\|\sigma_t\| = \sqrt{\eta^2 \sigma_\delta^2 + \left(\sigma_f \frac{\Pi_f}{\Pi} + \frac{(\bar{\mu} - f_t) \nu_{\lambda,t} \Pi_{\hat{\lambda}}}{\sigma_f \Pi} \right)^2}, \quad (34)$$

whereas the risk premium is

$$RP_t = \gamma \eta \sigma_\delta^2 + (1 - \phi) \left(\sigma_f \frac{I_f}{I} + \frac{(\bar{\mu} - f_t) \nu_{\lambda,t} I_{\hat{\lambda}}}{\sigma_f I} \right) \left(\sigma_f \frac{\Pi_f}{\Pi} + \frac{(\bar{\mu} - f_t) \nu_{\lambda,t} \Pi_{\hat{\lambda}}}{\sigma_f \Pi} \right). \quad (35)$$

Finally, the Sharpe ratio is defined as

$$SR_t \equiv \frac{RP_t}{\|\sigma_t\|}. \quad (36)$$

Figure 2 displays these quantities. The left panels depict the relations with the forecast f_t , setting the filter $\hat{\lambda}_t$ at zero (i.e., the mean-reversion speed at $\hat{\theta}_t = \bar{\theta}$). The right panels depict the relations with the mean-reversion speed $\hat{\theta}_t = \bar{\theta} + \hat{\lambda}_t$, setting the forecast f_t at its long-term

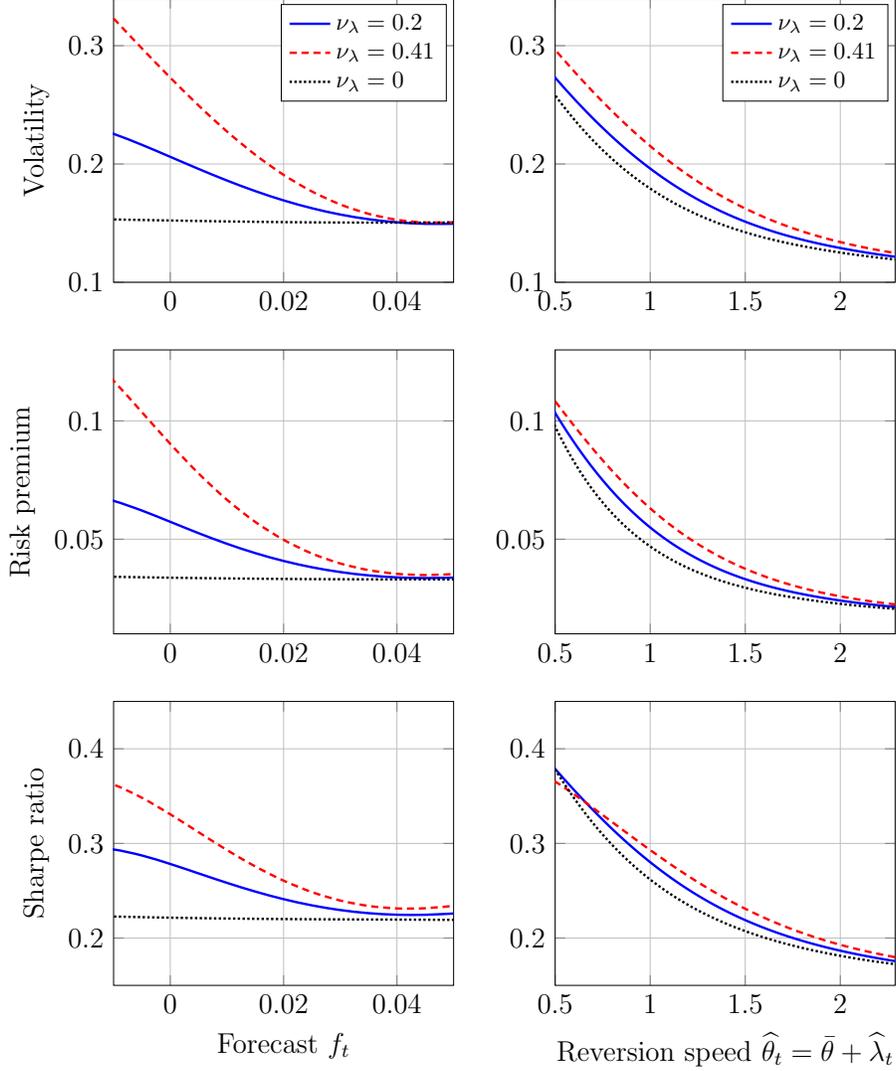


Figure 2: Stock return volatility, risk premium, and Sharpe ratio with learning about persistence.

This figure shows how the stock return volatility, the risk premium, and the Sharpe ratio vary with the state variables in a model of learning about persistence. For the left plots, we fix $\hat{\lambda}_t = 0$. For the right plots, we fix $f_t = \bar{\mu} = 0.0254$. Unless otherwise specified, we consider the calibration provided in the second column of Table 1.

mean $\bar{\mu}$. All panels report values for various levels of uncertainty about persistence, $\nu_{\lambda,t}$.

Focusing on the left panels, we observe that 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. In contrast, in presence of uncertainty about persistence ($\nu_{\lambda,t} > 0$), asset pricing moments decrease with the forecast f_t . This can be seen from Eqs. (34) and (35), where the volatility and the risk premium increase nonlinearly in persistence risk $(\bar{\mu} - f_t)\nu_{\lambda,t}$, and thus are higher in bad times. Furthermore, since

the second component of the market price of risk (Eq. 27) also increases in bad times, the Sharpe ratio becomes countercyclical.

The economic mechanism behind these results comes from learning about persistence. In recessions, bad news increases agent's perception that economic growth is more persistent. Persistence risk is positive and amplifies all asset pricing moments. In expansions, the effect of news is mitigated by learning about persistence. Persistence risk is now negative and thus dampens asset pricing moments. Therefore, learning about persistence creates the asymmetry depicted in the left panels of Figure 2.

The right panels of Figure 2 depict the impact of persistence on asset pricing moments. The volatility, the risk premium, and the Sharpe ratio depend negatively on $\widehat{\theta}_t$. This arises because more persistence (lower $\widehat{\theta}_t$) implies more risk in the long run. Note that in the right panels we fix $f_t = \bar{\mu}$. When this equality holds, the agent cannot learn about the mean-reversion speed λ_t because expected growth shocks are uninformative. Yet, the plots show that uncertainty about persistence still affects asset prices, through the partial derivatives of the price-dividend ratio.

3.3 Comparison with alternative models

We compare our model with learning about persistence against alternative setups. In Section 3.3.1, we consider two alternatives, one in which the agent learns about level and one in which there is no incomplete information. None of these alternatives generate time variation in asset pricing moments. In Section 3.3.2, we consider a model in which the degree of persistence is time-varying, but observable, and show that it does not deliver substantial variation in asset pricing moments. We conclude that uncertainty about persistence is not only sufficient, but also necessary to generate plausible asset price dynamics.

3.3.1 Other types of learning

We analyze two alternative specifications. First, we consider the case of learning about the level of expected output growth only ($\widehat{\lambda}_t = 0, \forall t$). Uncertainty about the level of expected growth is the premise of a large incomplete information literature and therefore constitutes an important benchmark (e.g., [Detemple, 1986](#); [Veronesi, 1999, 2000](#); [Scheinkman and Xiong, 2003](#); [Dumas et al., 2009](#)). With learning about level only, the volatility of stock returns is given by

$$\|\sigma_t\| = \sqrt{\left(\eta\sigma_\delta + \frac{\bar{\nu}_\epsilon \Pi_{\widehat{\epsilon}}}{\sigma_\delta \Pi}\right)^2 + \left(\sigma_f \frac{\Pi_f}{\Pi}\right)^2}, \quad (37)$$

whereas the risk premium is

$$RP_t = \left(\eta\sigma_\delta + \frac{\bar{\nu}_\epsilon \Pi_{\hat{\epsilon}}}{\sigma_\delta \Pi} \right) \left(\gamma\sigma_\delta + (1 - \phi) \frac{\bar{\nu}_\epsilon I_{\hat{\epsilon}}}{\sigma_\delta I} \right) + (1 - \phi) \sigma_f^2 \frac{I_f \Pi_f}{I \Pi}. \quad (38)$$

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

$$\|\sigma_t\| = \sqrt{\sigma_\delta^2 + \left(\sigma_f \frac{\Pi_f}{\Pi} \right)^2}, \quad (39)$$

whereas the risk premium is given by

$$RP_t = \gamma\eta\sigma_\delta^2 + (1 - \phi) \sigma_f^2 \frac{I_f \Pi_f}{I \Pi}. \quad (40)$$

In both alternative models, the Sharpe ratio is computed as in (36).

Eqs. (37)–(40) 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 and the wealth-consumption ratio. We expect these fluctuations to be relatively weak—they are actually zero with a log-linear approximation of the log price-dividend ratio.

Figure 3 confirms this result. It compares the sensitivity of asset pricing moments to the 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, which corresponds to Case I of BY. In each case, we use the parameters from the corresponding column in Table 1. The plots show that learning about level yields larger volatility, risk premium and Sharpe ratio than a model with no learning, confirming previous results (e.g., Ai, 2010). However, the only model that generates variations in asset pricing moments is the one with learning about persistence.

3.3.2 Time-varying, but observable persistence

We now analyze an alternative model in which the persistence is time-varying, but *observable*. Besides the output process δ_t and the growth forecast f_t (Eqs. (3)–(4), with $\mu_t = f_t$), this alternative model features an additional observable state variable λ_t :

$$d\lambda_t = -\kappa\lambda_t dt + \sigma_\lambda \rho dW_t^f + \sigma_\lambda \sqrt{1 - \rho^2} dW_t^\lambda. \quad (41)$$

We allow for an exogenous correlation ρ between the expected growth and its persistence.

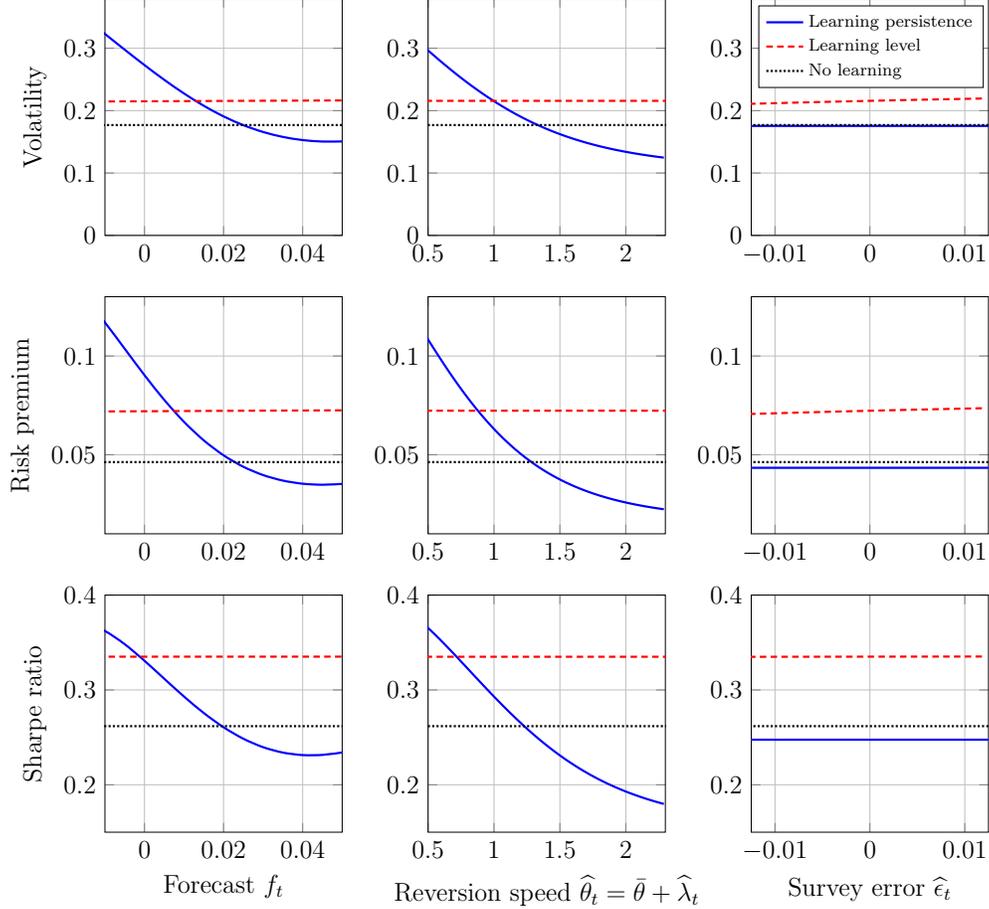


Figure 3: Stock return volatility, risk premium, and Sharpe ratio with and without learning.

This figure compares the asset pricing implications of three models: (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 $\hat{\epsilon}_t = 0$ and $\hat{\lambda}_t = 0$. For the middle plots, we fix $f_t = \bar{\mu}$ and $\hat{\epsilon}_t = 0$. For the right plots, we fix $f_t = \bar{\mu}$ and $\hat{\lambda}_t = 0$. We set $\nu_{\lambda,t} = 0.41$ in all plots. Unless otherwise specified, we use the calibration provided in Table 1.

We consider a range of positive values for this correlation (i.e., persistence increases or, equivalently, the mean-reversion speed decreases after negative shocks). This model is therefore *designed* to generate strong persistence in bad times. Our aim is to evaluate whether it is able to generate results that are comparable with those obtained by a model with learning about persistence.

The price-dividend ratio in this economy, $\Pi(f_t, \lambda_t)$, depends on two state variables, f_t and λ_t . Because these variables are driven by the Brownians W_t^f and W_t^λ , the price-dividend ratio has two diffusion components (we relegate all technical details regarding the equilibrium

to Appendix B.4):

$$\sigma_{\Pi 1} = \sigma_f \frac{\Pi_f}{\Pi} + \rho \sigma_\lambda \frac{\Pi_\lambda}{\Pi} \quad (42)$$

$$\sigma_{\Pi 2} = \sqrt{1 - \rho^2} \sigma_\lambda \frac{\Pi_\lambda}{\Pi}. \quad (43)$$

Since we assume ρ to be positive, the first diffusion term increases due to fluctuations in λ_t . However, direct comparison with its counterpart in presence of learning about persistence (Eq. 31) reveals that our key mechanism is absent when λ_t becomes observable. When persistence is uncertain, the output growth gap directly enters in the diffusion of the price-dividend ratio in Eq. (31) and magnifies it during bad times. In contrast, when λ_t is observable the amplification occurs at all times. Consequently, this model does not generate an asymmetry in asset pricing moments.

It might nevertheless be the case that an (indirect) asymmetry arises through the partial derivatives of the price-dividend ratio in (42)-(43). We investigate this by solving the model for two different values of the correlation parameter, $\rho \in \{0.5, 1\}$. To keep the results comparable with a model with learning about persistence, we use the parameters from column 2 of Table 1. We derive the volatility and the risk premium in Appendix B.4.

Figure 4 depicts the results. The two panels in column (a) show the volatility and the risk premium in an economy with learning about persistence. There are three lines, each one corresponding to a different value of the filter $\hat{\lambda}_t \in \{-0.3, 0, 0.3\}$. Both panels show significant variation in asset pricing moments. In contrast, columns (b) and (c) show weaker effects when the persistence is observable, without any significant asymmetry on the expected growth dimension. It is therefore the learning exercise of the agent that induces strong variations in equilibrium risk premia and return volatility, and not the fact that the degree of persistence is time varying.

4 Evidence

In this section, we empirically evaluate the predictions of the model with learning about persistence. We start by examining how the model-implied asset prices compare to their empirical counterparts. The model matches well both the level and variation of asset pricing moments. We then formally test our theory, which predicts that the risk premium, the stock return volatility, and the Sharpe ratio increase non-linearly with persistence risk. We find support in the data for the increasing and non-linear relationships. Furthermore, persistence risk has strong predictive power for future excess returns. We show that that the predictabil-

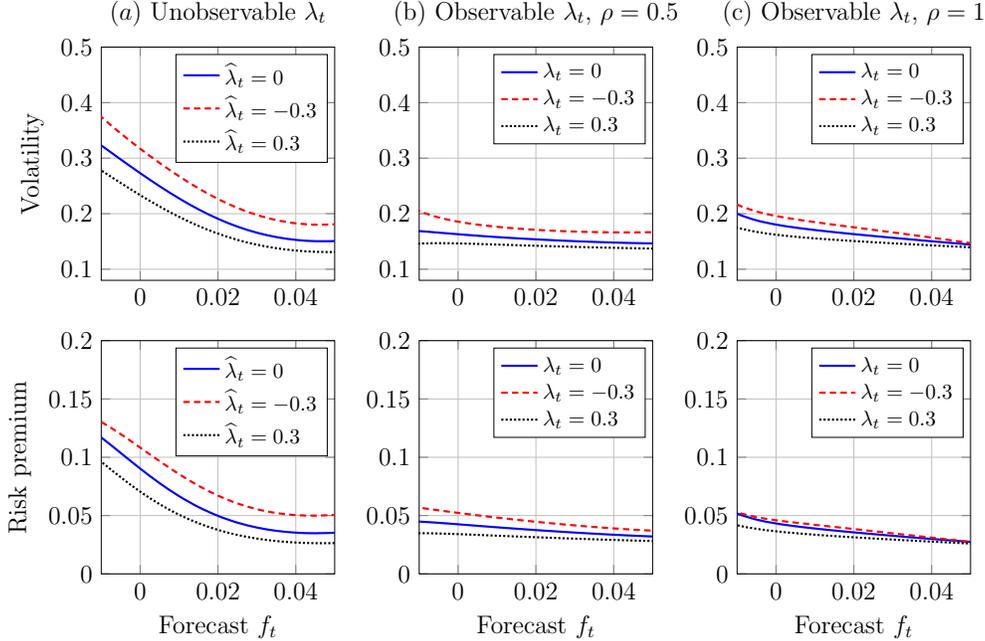


Figure 4: Comparison with an alternative model in which the mean-reversion speed θ_t is observable and time-varying.

This figure compares the asset pricing implications of two models: (i) a model with learning about persistence (column *a*), and (ii) a model without learning but with time-varying and observable persistence (columns *b* – *c*). In the left panels, we fix $\nu_{\lambda,t} = 0.41$. In the right panels, we set the correlation coefficient to $\rho \in \{0.5, 1\}$. Unless otherwise specified, we use the calibration provided in the second column of Table 1.

ity is concentrated when news are particularly informative about the degree of persistence, i.e., when the forecast f_t is away from the long-term growth level $\bar{\mu}$.

We also investigate how the presence of uncertainty about persistence overcomes the predictability challenges facing the long-run risk model. Specifically, the BY and BKY calibrations imply that stock prices have predictive power for future consumption growth but not for future excess returns, whereas the data suggest the opposite (BC). Our model generates predictions that are in line with the data, thus addressing these critiques.

Finally, we show how the countercyclicality in asset prices, driven by persistence risk in our setup, relates to the habit formation model of Campbell and Cochrane (1999).

4.1 Data

Our empirical analysis is based on quarterly U.S. data over the period Q1:1969–Q4:2016. The estimation performed in Section 3.1 provides time series of the filtered mean-reversion speed $\hat{\lambda}_t$ and persistence risk, defined by $(\bar{\mu} - f_t)\nu_{\lambda,t}$. The mean analyst forecast on 1-quarter-ahead real GDP growth is our measure of f_t . Using these state variables, we construct model-

implied time series for the risk premium, Sharpe ratio, stock return volatility, price-dividend ratio, and the risk-free rate, all defined in Section 2.3.

The empirical counterparts of these asset pricing quantities are constructed as follows. We compute quarterly stock returns and dividend growth from the value-weighted CRSP index, which covers NYSE, Amex, and Nasdaq data, and convert them to real terms using the consumer price index (CPI). We create a proxy for the ex-ante risk-free rate by forecasting the ex-post quarterly real return on three-month Treasury bills with past one-year inflation and the most recent available three-month nominal bill yield. This procedure is equivalent to forecasting inflation and subtracting the inflation forecast from the nominal bill yield (BC). The price-dividend ratio is the price in the last month of the quarter divided by the sum of dividends paid in the last twelve months. Our proxy for the risk premium is the fitted value obtained by regressing excess stock returns on the lagged dividend yield (inverse of the price-dividend ratio), the lagged default premium (Baa yield minus ten-year government bond yield), and stock return volatility. This choice is based on the empirical evidence that the dividend yield (Fama and French, 1988), the default premium (Fama and French, 1989), and the level of stock market volatility (French, Schwert, and Stambaugh, 1987) have predictive power for stock market returns. Stock return volatility is the conditional volatility of real stock returns estimated with a GARCH(1,1). The data construction is described in detail in Appendix D.

4.2 Descriptive analysis of asset pricing moments

In a first analysis, we compare the model-implied asset pricing moments with their empirical counterparts. We consider the risk premium, the stock return volatility, the Sharpe ratio, the log price-dividend ratio, and the risk-free interest rate. All variables are measured in real terms. Table 2 starts by showing descriptive statistics of these unconditional asset pricing moments in the model and in the data.

The model is successful in matching most properties of the data. In particular, the model generates a risk premium of 5% and a volatility of 18%, which are close to their empirical estimates. With the exception of the risk-free rate (for which the model delivers the volatility but not the level), and the price-dividend ratio (for which the model delivers the level but not the volatility), these numbers show that the asset pricing moments implied by our model align reasonably well with what is observed in the data. It is worth noting that the standard deviation of the log price-dividend ratio is similar to the one obtained in BY and BKY, although we do not introduce time-variation in output volatility. While the variability in stock prices results from exogenous fluctuations in volatility in the standard long-run risk

Variable	Mean	Std dev.	Median	5-percentile	95-percentile
Risk premium					
Data	0.0447	0.0706	0.0317	-0.0462	0.1785
Model	0.0499	0.0320	0.0410	0.0309	0.1057
Stock return volatility					
Data	0.1751	0.0428	0.1636	0.1289	0.2654
Model	0.1819	0.0499	0.1694	0.1431	0.2834
Sharpe ratio					
Data	0.1913	0.3211	0.2062	-0.3482	0.7251
Model	0.2591	0.0564	0.2437	0.2160	0.3753
Log (P/D)					
Data	3.6390	0.4021	3.6304	3.0266	4.3087
Model	3.5035	0.1433	3.5248	3.2076	3.6746
Real risk-free rate					
Data	0.0071	0.0094	0.0073	-0.0231	0.0393
Model	0.0345	0.0116	0.0348	0.0091	0.0544

Table 2: Unconditional asset pricing moments.

This table reports the unconditional asset pricing moments in the model and in the data. The statistics are based on quarterly data and are annualized. All values are in real terms. The construction of the empirical moments is detailed in Appendix D. Stock prices represent the value-weighted CRSP index and the sample spans the period Q1:1969-Q4:2016.

model, the numbers in Table 2 show that learning about persistence generates endogenous fluctuations of similar magnitude.

Figure 5 plots the dynamics of the main asset pricing quantities, in the model and in the data. We notice that the model generates substantial time-variation in the risk premium and stock market volatility (panels A and B). Comparing the time series with their empirical counterparts, the spikes in the risk premium and volatility are generally well aligned. Panel C shows that the price-dividend ratio decreases during recessions. Panel D shows that the model delivers significant fluctuations in the risk-free rate, comparable with the variation in the data, although the model does not match the downward trend in the risk-free rate towards the end of the sample.

Finally, Table 3 reports the relations between the model-implied asset pricing quantities and their empirical counterparts. Observed moments are regressed onto model-implied moments. With the exception of the risk-free rate, the slopes of the regressions are all positive and statistically significant. In particular, two of the slope coefficients are close to one (for the risk premium and the log price-dividend ratio).

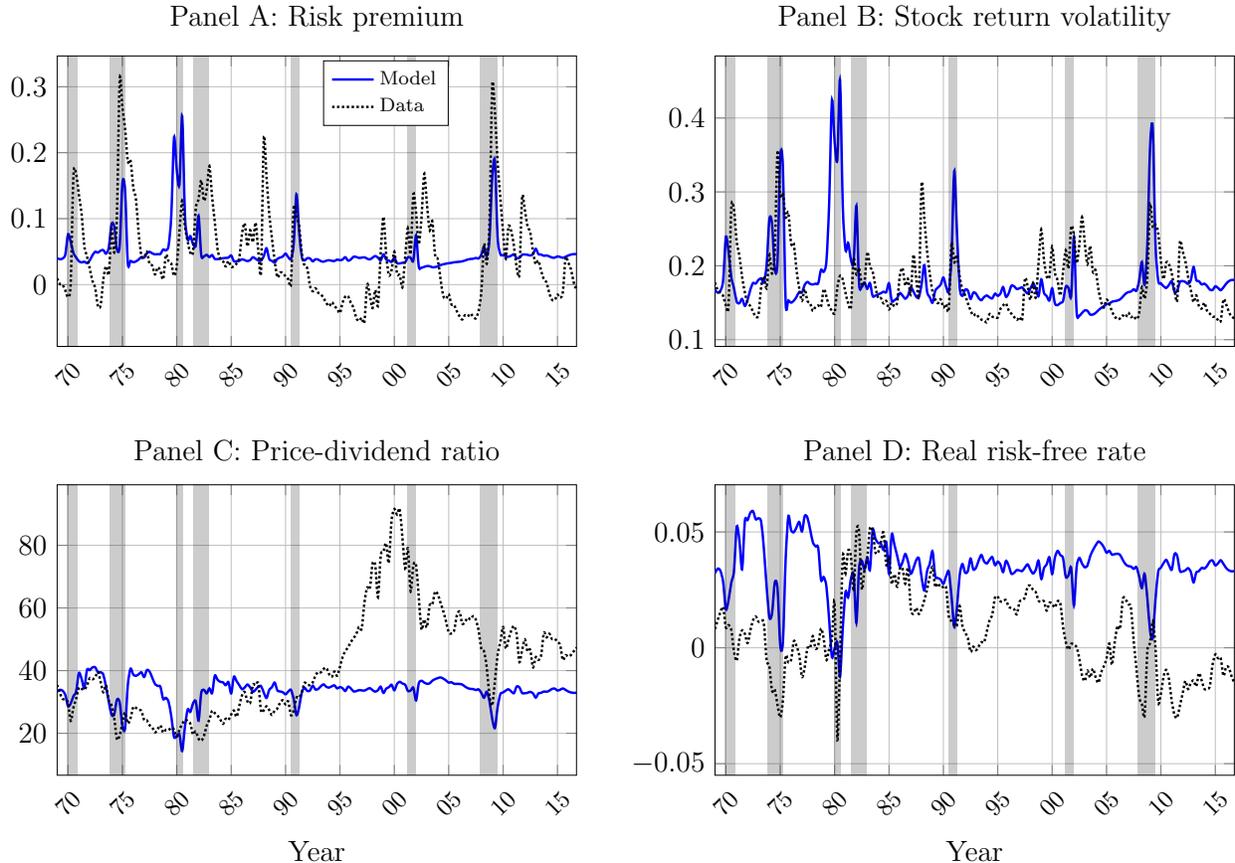


Figure 5: Empirical vs. model-implied asset pricing quantities.

This figure plots the time-series of the main asset pricing quantities in the model and in the data. Panel A reports the risk premium, panel B reports the stock return volatility, panel C reports the price-dividend ratio, while panel D reports the risk-free rate. All values are in real terms and annualized. The construction of the empirical moments is detailed in Appendix D. Stock prices represent the value-weighted CRSP index and the sample spans the period Q1:1969-Q4:2016.

Overall, Tables 2 and 3 and Figure 5 show that the dynamics of asset returns obtained in our framework align well with what is observed in the data. The model-implied moments thus help understand the variations in the empirical moments. It is important to emphasize that we do not use financial market data in the estimation of our model (and thus in generating model-implied asset pricing moments).

4.3 Asset prices with persistence risk

We now perform a number of analyses that test the quantitative implications of learning about persistence. The model predicts that the risk premium, the stock return volatility, and the Sharpe ratio increase with persistence risk $(\bar{\mu} - f_t)\nu_{\lambda,t}$ (see Figure 2). Our first

	Risk premium	Stock return volatility	Sharpe ratio	Log (P/D)	Risk-free rate
Constant	0.005	0.143***	-0.336***	0.567	0.001
t-stat	0.423	9.682	-2.829	1.157	0.592
Slope	0.792***	0.178**	2.035***	0.877***	0.032
t-stat	3.170	2.149	4.500	6.176	1.048
R-squared	0.129	0.043	0.128	0.098	0.006
N	192	192	192	192	192

Table 3: Empirical vs. model-implied asset pricing quantities.

This table reports the relations between the model-implied asset pricing quantities and their empirical counterparts. Observed moments are regressed onto model-implied moments. The t -statistics are computed with Newey and West (1987) standard errors. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The construction of the empirical moments is discussed in Appendix D. Stock prices represent the value-weighted CRSP index and the sample spans the period Q1:1969-Q4:2016.

objective is to determine whether the empirical asset pricing moments are indeed higher when persistence risk is more severe.

As a preliminary exercise, we split the sample into two parts according to the level of persistence risk, above and below the median. Table 4 reports conditional asset pricing moments, both in the model and in the data. Consistent with the theory, we find evidence that the risk premium, the Sharpe ratio, and the stock return volatility are greater when persistence risk is higher, while the price-dividend ratio and the risk-free rate are lower. The differences are all statistically significant, except for the empirical risk-free rate. Figure 6 provides additional support by showing that the relations between these asset pricing moments and persistence risk are positive over the period Q1:1969–Q4:2016, both in the model (panels A-C) and in the data (panels D-F). These results confirm that asset pricing moments increase with persistence risk. In addition, Table 4 shows that the variance of asset pricing moments is higher when persistence risk is higher, both in the model and in the data.

The model also suggests a strong non-linearity in the relation between asset pricing moments and persistence risk. This can be seen in Eqs (32) and (33), where the stock return variance and the risk premium depend on the square of $(\bar{\mu} - f_t)\nu_{\lambda,t}$. For the risk premium, this quadratic term arises from the multiplication of the market price risk (Eq. 27) with the quantity of risk (Eq. 32). These non-linearities in the model are illustrated in panels A-C of Figure 6.

We provide empirical evidence in support of this prediction. Figure 6 plots the predictive

Variable	High persistence risk		Low persistence risk		High-minus-low	
	Mean	Std	Mean	Std	Mean diff.	Std diff.
Risk premium						
Data	0.0620	0.0783	0.0273	0.0572	0.0347***	0.0211**
Model	0.0619	0.0416	0.0379	0.0056	0.0239***	0.0361***
Stock return volatility						
Data	0.1845	0.0482	0.1656	0.0344	0.0189***	0.0138***
Model	0.2043	0.0622	0.1595	0.0112	0.0449***	0.0510***
Sharpe ratio						
Data	0.2691	0.3269	0.1135	0.2970	0.1556***	0.0299
Model	0.2814	0.0710	0.2367	0.0187	0.0447***	0.0522***
Log (P/D)						
Data	3.5867	0.4125	3.6914	0.3865	-0.1048*	0.0260
Model	3.4241	0.1604	3.5829	0.0535	-0.1588***	0.1069***
Risk-free rate						
Data	0.0061	0.0100	0.0082	0.0089	-0.0021	0.0005
Model	0.0269	0.0104	0.0421	0.0067	-0.0152***	0.0037***

Table 4: Conditional asset pricing moments.

This table reports the asset pricing moments conditional on persistence risk. The sample is split into two parts based on the median of $(\bar{\mu} - f_t)\nu_{\lambda,t}$. High values correspond to periods of high persistence risk, whereas low values reflect times of low persistence risk. The means and standard deviations (Std) are based on quarterly data and are annualized. For the last two columns, statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. Stock prices represent the value-weighted CRSP index and the sample spans the period Q1:1969-Q4:2016.

relation between persistence risk and the risk premium (panel D), the volatility of stock returns (panel E), and the Sharpe ratio (panel F) over the period Q1:1969–Q4:2016, using a quadratic (OLS) regression.¹⁹ The empirical relations are indeed non-linear. Asset pricing moments are more sensitive to persistence risk when it is high, which corresponds to worse economic times and/or greater uncertainty about persistence. This means that learning about persistence amplifies the sensitivity of asset prices to news in bad times. Hence, fluctuations in the risk premium, the level of stock return volatility, and the Sharpe ratio are countercyclical due to uncertainty about persistence.

We formally test the relations between asset pricing moments and persistence risk. The results for the risk premium, the stock return volatility, and the Sharpe ratio are presented in Table 5. Persistence risk explains a large fraction of the time variation in asset pricing moments. Both the linear and the quadratic terms are statistically significant, in the model and in the data, thus providing evidence of the asymmetric relations. These results are in

¹⁹Although not reported, we also consider alternative specifications as robustness checks. The shape of the relation is similar with least absolute deviations and spline regressions.

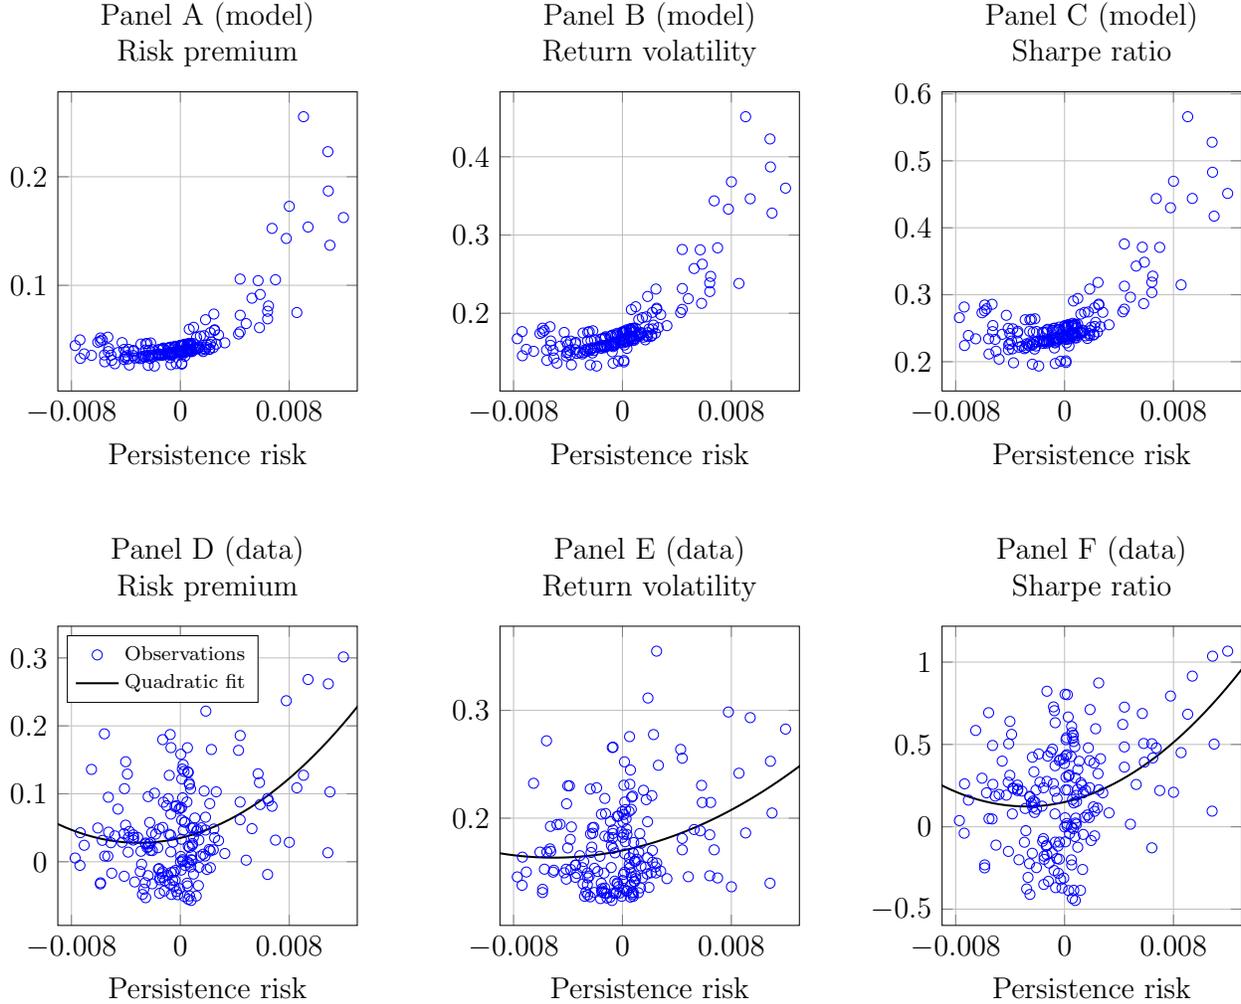


Figure 6: Relation between asset pricing moments and persistence risk.

This figure plots the relations between asset pricing moments and persistence risk. The left panels display the relation with the risk premium, the middle panels display the relation with the level of stock return volatility, while the right panels show the relation with the Sharpe ratio. The upper panels show the model-implied moments, while the lower panels show their empirical counterparts. The construction of the empirical moments is discussed in Appendix D. Stock prices represent the value-weighted CRSP index and the sample spans the period Q1:1969-Q4:2016.

favor of our theoretical prediction that stock return volatility, the risk premium, and the Sharpe ratio increase in times of greater persistence risk.

Furthermore, consistent with the long-run risk theory, asset pricing moments increase when the agent views economic growth as more persistent (see Figure 2). Table 6 confirms that the empirical moments decrease with the estimated mean-reversion speed $\hat{\lambda}_t$, which relates negatively with the degree of persistence. That is, the relation between asset pricing moments and persistence is positive and non-linear in the data, as implied by the model.

	Risk premium		Return volatility		Sharpe ratio	
	Model	Data	Model	Data	Model	Data
$(\bar{\mu} - f_t)\nu_{\lambda,t}$	4.364***	4.736***	8.016***	2.632***	7.576***	18.717***
t-stat	10.094	3.437	14.134	2.950	10.115	3.253
$((\bar{\mu} - f_t)\nu_{\lambda,t})^2$	783.658***	778.242***	1099.877***	260.292*	1386.833***	3319.759***
t-stat	7.694	2.858	8.633	1.722	8.662	3.243
R-squared	0.802	0.176	0.845	0.095	0.793	0.144
N	192	192	192	192	192	192

Table 5: Asset pricing with learning about persistence — persistence risk.

This table reports the relations between the asset pricing moments and persistence risk, which is defined by $(\bar{\mu} - f_t)\nu_{\lambda,t}$. N is the number of observations. The t -statistics are computed with Newey and West (1987) standard errors. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The construction of the empirical moments is discussed in Appendix D. Stock prices represent the value-weighted CRSP index and the sample spans the period Q1:1969-Q4:2016.

	Risk premium		Return volatility		Sharpe ratio	
	Model	Data	Model	Data	Model	Data
$\hat{\lambda}_t$	-0.387***	-0.234**	-0.511***	-0.130**	-0.552***	-1.219***
t-stat	-10.424	-2.123	-10.180	-2.002	-10.694	-2.617
$\hat{\lambda}_t^2$	0.126***	0.072	0.154***	0.057**	0.157***	0.364*
t-stat	8.835	1.492	7.830	2.042	7.790	1.777
R-squared	0.678	0.064	0.638	0.017	0.717	0.091
N	192	192	192	192	192	192

Table 6: Asset pricing with learning about persistence — level of persistence.

This table reports the relations between the asset pricing moments and the degree of persistence. Persistence (or the lack of) is determined by the mean-reversion speed, $\hat{\lambda}_t$. N is the number of observations. The t -statistics are computed with Newey and West (1987) standard errors. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The construction of the empirical moments is discussed in Appendix D. Stock prices represent the value-weighted CRSP index and the sample spans the period Q1:1969-Q4:2016.

Overall, the data validate the theoretical relations of our model and show that incomplete information about the degree of persistence in expected output growth helps explain the

time-variation in asset pricing moments.

4.4 Predictability

This section shows that persistence risk predicts future excess stock returns. We find evidence that its predictive power is concentrated during times when news are particularly informative about the degree of persistence. We also show that our model addresses the critique that the long-run risk model typically fails at predicting future returns with the price-dividend ratio, while displaying excessive predictive power for future consumption growth.

4.4.1 Return predictability with persistence risk

We have thus far provided evidence that expected excess returns are higher when there is greater persistence risk. Persistence risk should then also help predict future excess returns. This is a prediction that is unique to our theory. We test this prediction with the following regression specification, at quarterly frequency:

$$\sum_{k=1}^K (r_{t+k} - r_{f,t+k}) = a_K + b_K(\bar{\mu} - f_t)\nu_{\lambda,t} + \mathbf{c}_K X_t + \epsilon_{t+K}, \quad (44)$$

where r_{t+k} and $r_{f,t+k}$ are the log real return and real risk-free rate for quarter $t+k$. We consider different horizons: 1 year, 3 years, 5 years, and 7 years. We denote by X_t a vector of control variables that are expected to have predictive power in the data. Consistent with our construction of the risk premium in Section 4.1, we control for the price-dividend ratio (Fama and French, 1988), the stock market volatility (French et al., 1987), and the default premium (Fama and French, 1989). Because persistence risk is a measure of economic uncertainty, we also control for the macro uncertainty index of Jurado et al. (2015). These controls allow us to determine whether persistence risk has predictive power beyond that of existing predictors.

Table 7 shows support for the return predictability of persistence risk. This variable is highly statistically significant in the model and in the data, particularly at the 5 and 7 year horizon. The results become even stronger when we include controls, as we now observe significant predictability over all horizons. This means that persistence risk contains additional information for explaining future excess returns, beyond what is already embedded in the price-dividend ratio, the current level of stock return volatility, or in the level of macroeconomic uncertainty. Importantly, our measure of persistence risk is only driven by economic shocks (on real GDP growth rate and output growth forecasts) and does not incorporate any market prices. Moreover, this predictability is completely driven by learning about persis-

Panel A: Predictability in the model					Panel B: Predictability in the data				
	Excess return					Excess return			
	1Y	3Y	5Y	7Y		1Y	3Y	5Y	7Y
$(\bar{\mu} - f_t)\nu_{\lambda,t}$	10.603***	45.126***	39.952***	46.533***	$(\bar{\mu} - f_t)\nu_{\lambda,t}$	6.287*	17.095*	23.641**	29.325***
t-stat	4.138	3.892	3.518	5.518	t-stat	1.799	1.920	2.501	4.402
R-squared	0.024	0.211	0.139	0.136	R-squared	0.016	0.052	0.069	0.120
N	188	180	172	164	N	188	180	172	164
Control: log (P/D)					Control: log (P/D)				
$(\bar{\mu} - f_t)\nu_{\lambda,t}$	9.861***	42.113***	34.285***	36.159***	$(\bar{\mu} - f_t)\nu_{\lambda,t}$	5.014	14.749	20.016**	24.495***
t-stat	3.527	3.373	3.326	4.361	t-stat	1.254	1.532	2.105	3.327
Control: stock return volatility					Control: stock return volatility				
$(\bar{\mu} - f_t)\nu_{\lambda,t}$	3.988	40.238***	41.310***	50.608***	$(\bar{\mu} - f_t)\nu_{\lambda,t}$	5.110	17.941**	24.686***	34.749***
t-stat	1.189	3.654	3.473	5.156	t-stat	1.600	2.032	2.950	4.475
Control: default premium					Control: default premium				
$(\bar{\mu} - f_t)\nu_{\lambda,t}$	7.803**	44.717***	41.951***	52.685***	$(\bar{\mu} - f_t)\nu_{\lambda,t}$	4.429	14.446	17.127***	26.829***
t-stat	2.274	3.852	3.288	4.994	t-stat	1.391	1.633	3.358	4.737
Control: macro uncertainty					Control: macro uncertainty				
$(\bar{\mu} - f_t)\nu_{\lambda,t}$	15.058***	46.276***	42.397**	57.263***	$(\bar{\mu} - f_t)\nu_{\lambda,t}$	11.774***	24.460**	22.867**	27.735***
t-stat	2.785	3.102	2.147	3.955	t-stat	2.680	2.407	2.390	2.809
All controls					All controls				
$(\bar{\mu} - f_t)\nu_{\lambda,t}$	10.476**	42.271***	43.184***	59.557***	$(\bar{\mu} - f_t)\nu_{\lambda,t}$	11.183*	26.002***	29.795***	35.655***
t-stat	2.426	3.132	2.746	6.656	t-stat	1.849	2.708	4.150	3.399

Table 7: Predictability of excess returns with persistence risk.

This table shows the predictability of cumulative future excess stock returns with persistence risk. Panel A reports the regression estimates in the model, while panel B reports the empirical counterparts. The top panels report the results for the univariate regressions and the lower panels report results when controlling for the log price-dividend ratio (Fama and French, 1988), the level of stock market volatility (French et al., 1987), the default premium (Fama and French, 1989), the macro uncertainty index of Jurado et al. (2015), and all controls together. Each column represents a different forecast horizon K and N is the number of observations. t -statistics are computed with Newey and West (1987) standard errors with $2(K - 1)$ lags. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The construction of the empirical moments is discussed in Appendix D. Stock prices represent the value-weighted CRSP index and the sample spans the period Q1:1969-Q4:2016.

tence. With an observable persistence level, as assumed in the standard long-run risk theory, $\nu_{\lambda,t} = 0$ and the role of persistence risk would disappear. To sum up, these results show that investors are compensated for the risk premium that they demand for bearing persistence risk.

4.4.2 Asymmetric predictability: high vs low informative times

Learning about persistence fundamentally depends on news about expected output growth. Yet not all news are viewed as equal. When $f_t = \bar{\mu}$, changes in the forecast f_t are uninformative because the agent is unable to learn about the mean-reversion speed. However, news become particularly informative for lower and higher values of f_t . This is when learning about persistence matters for asset prices. Consequently, the predictability of future excess returns should be greater during times when the expected growth forecast is far away from its long-term mean.

We test this prediction by decomposing the sample into two subsamples. We first consider the observations when f_t is far from the long-term average, as determined by the bottom and the top quartiles of f_t . We refer to such period as "high informative times." Then, we consider the remaining observations when f_t is relatively close to the long-term average. We call this period "low informative times".²⁰ Based on this sample split, we estimate the following regression specification:

$$\sum_{k=1}^K (r_{t+k} - r_{f,t+k}) = a_K + b_K \mathbb{1}_{HIT}(\bar{\mu} - f_t)\nu_{\lambda,t} + c_K \mathbb{1}_{LIT}(\bar{\mu} - f_t)\nu_{\lambda,t} + \mathbf{d}_K X_t + \epsilon_{t+K}, \quad (45)$$

where $\mathbb{1}_{HIT}$ and $\mathbb{1}_{LIT}$ are dummies equal to one for observations during high (HIT) and low (LIT) informative times, respectively, and zero otherwise. We include a vector of control variables X_t , which consist of the log price-dividend ratio, the level of stock market volatility, the default premium, and the macro uncertainty index of [Jurado et al. \(2015\)](#).

Table 8 shows that persistence risk is statistically significant for predicting future stock returns within the model, but only during high informative times. This is when f_t is far from $\bar{\mu}$, either above or below. Panel A shows that this finding is robust to all forecast horizons, ranging from 1 year to 7 years. This prediction is confirmed by the data (Panel B), in particular when we add controls.

Further, we separately estimate the regressions during either high or low informative times and compare the R^2 statistics, following the predictability literature (e.g., [Rapach, Strauss, and Zhou, 2010](#); [Henkel, Martin, and Nardari, 2011](#); [Dangl and Halling, 2012](#)). Table 9 confirms that the predictive power is concentrated when economic news are informative for updating the degree of persistence.

We validate our return predictability results through finite-sample simulations of the model with learning about persistence. These simulations are equal in length to the dataset

²⁰We also considered alternatives, such as separating observations by low and high values of $(\bar{\mu} - f_t)^2$, by the median or quartiles. Results remain very similar.

Panel A: Predictability in the model					Panel B: Predictability in the data				
	Excess return					Excess return			
	1Y	3Y	5Y	7Y		1Y	3Y	5Y	7Y
Low informative times (f_t close to $\bar{\mu}$)					Low informative times (f_t close to $\bar{\mu}$)				
$(\bar{\mu} - f_t)\nu_{\lambda,t}$	0.514	21.705	5.506	54.127	$(\bar{\mu} - f_t)\nu_{\lambda,t}$	2.054	-1.981	-78.264*	-23.871
t-stat	0.035	0.920	0.235	1.255	t-stat	0.158	-0.066	-1.654	-0.563
High informative times (f_t far from $\bar{\mu}$)					High informative times (f_t far from $\bar{\mu}$)				
$(\bar{\mu} - f_t)\nu_{\lambda,t}$	10.998***	45.972***	41.145***	46.289***	$(\bar{\mu} - f_t)\nu_{\lambda,t}$	6.453*	17.783*	27.169***	31.038***
t-stat	4.329	4.007	3.579	5.565	t-stat	1.685	1.905	2.786	4.605
R-squared	0.025	0.213	0.142	0.136	R-squared	0.017	0.054	0.112	0.132
N	188	180	172	164	N	188	180	172	164
With controls					With controls				
Low informative times (f_t close to $\bar{\mu}$)					Low informative times (f_t close to $\bar{\mu}$)				
$(\bar{\mu} - f_t)\nu_{\lambda,t}$	4.153	14.242	-9.429	27.413	$(\bar{\mu} - f_t)\nu_{\lambda,t}$	10.177	3.483	-51.469	-23.520
t-stat	0.190	0.501	-0.602	1.050	t-stat	0.893	0.151	-1.460	-0.713
High informative times (f_t far from $\bar{\mu}$)					High informative times (f_t far from $\bar{\mu}$)				
$(\bar{\mu} - f_t)\nu_{\lambda,t}$	10.740***	43.470***	44.940**	60.600***	$(\bar{\mu} - f_t)\nu_{\lambda,t}$	11.225*	26.965***	32.508***	37.576***
t-stat	2.739	3.267	2.597	6.537	t-stat	1.782	2.624	3.640	3.515
R-squared	0.197	0.285	0.250	0.438	R-squared	0.222	0.280	0.405	0.422
N	188	180	172	164	N	188	180	172	164

Table 8: Asymmetric excess return predictability with persistence risk.

This table reports the return predictability of persistence risk during high vs. low informative times. High informative times correspond to observations when the growth forecast f_t falls into its bottom or top quartiles. Low informative times include the remaining observations. Panel A reports the results for the model and Panel B for the data. The lower panels report results when controlling for the log price-dividend ratio (Fama and French, 1988), the level of stock market volatility (French et al., 1987), the default premium (Fama and French, 1989), and the macro uncertainty index (Jurado et al., 2015). Each column represents a different forecast horizon K and N is the number of observations. t -statistics are computed with Newey and West (1987) standard errors with $2(K - 1)$ lags. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The construction of the empirical variables is discussed in Appendix D. Stock prices represent the value-weighted CRSP index and the sample spans the period Q1:1969-Q4:2016.

Q1:1969-Q4:2016 (192 quarters). We simulate 100,000 time series at quarterly frequency using the discretization shown in Appendix C. We then calculate for each simulated point the wealth-consumption ratio $I(f_t, \hat{\lambda}_t, \nu_{\lambda,t})$, the price-dividend ratio $\Pi(f_t, \hat{\lambda}_t, \nu_{\lambda,t})$, and the risk-free rate $r_{f,t}$ (Eq. 26). These numbers allow us to generate time series of quarterly excess returns and to run the regressions (44) and (45) on simulated data. We only present results at the 5-year horizon for space considerations.

Figure 7 shows the simulation results. The first row (panels (a) to (c)) corresponds to the

Panel A: Predictive power in the model

	1Y	3Y	5Y	7Y
High informative times (f_t far from $\bar{\mu}$)	0.021	0.311	0.225	0.263
Low informative times (f_t close to $\bar{\mu}$)	0.035	0.061	0.006	0.001
Unconditional	0.024	0.211	0.139	0.136

Panel B: Predictive power in the data

	1Y	3Y	5Y	7Y
High informative times (f_t far from $\bar{\mu}$)	0.050	0.138	0.203	0.310
Low informative times (f_t close to $\bar{\mu}$)	0.023	0.004	0.009	0.028
Unconditional	0.016	0.052	0.069	0.120

Table 9: Asymmetry in predictive power of persistence risk for excess returns.

This table reports the conditional predictive power of persistence risk. Predictive power is measured with the R^2 of the regression of future excess returns on persistence risk, defined by $(\bar{\mu} - f_t)\nu_{\lambda,t}$. We run separate regression for high informative times, which correspond to times when f_t is far from $\bar{\mu}$ (bottom and top quartiles), and low informative times, which correspond to times when f_t is close to $\bar{\mu}$. Panel A reports the results for the model and Panel B in the data. Each column represents a different forecast horizon K . The construction of excess returns is discussed in Appendix D. Stock prices represent the value-weighted CRSP index and the sample spans the period Q1:1969-Q4:2016.

regression specification (44). The plots show the distributions of the slope coefficients in panel (a), the corresponding t-stats in panel (b), and R^2 coefficients in panel (c). The dashed lines indicate medians across simulations, whereas the solid lines indicate the regression results in the data, without controls (upper right panel of Table 7, column labeled “5Y”). The medians of the simulations are very close to what we observe in the data. The plots show clear evidence that the return predictability that we obtain from persistence risk is consistent with the data.

The panels (d) to (h) in Figure 7 show the results of the asymmetric regression specification (45), based on the same simulated data. Panels (d) and (e) show the slope coefficients and their t-stats for high informative times, panels (g) and (h) show coefficients and t-stats for low informative times, and finally panel (f) shows the R^2 of the regressions. The dashed lines are the medians across simulations, to be compared with the solid lines, which correspond to the upper right panel of Table 8, column labeled “5Y.” The simulations confirm the existence of an asymmetric return predictability that arises in the presence of persistence risk.

Overall, Figure 7 endorses our theoretical message that persistence risk is an important determinant of future excess returns. During high informative times—when f_t is in the

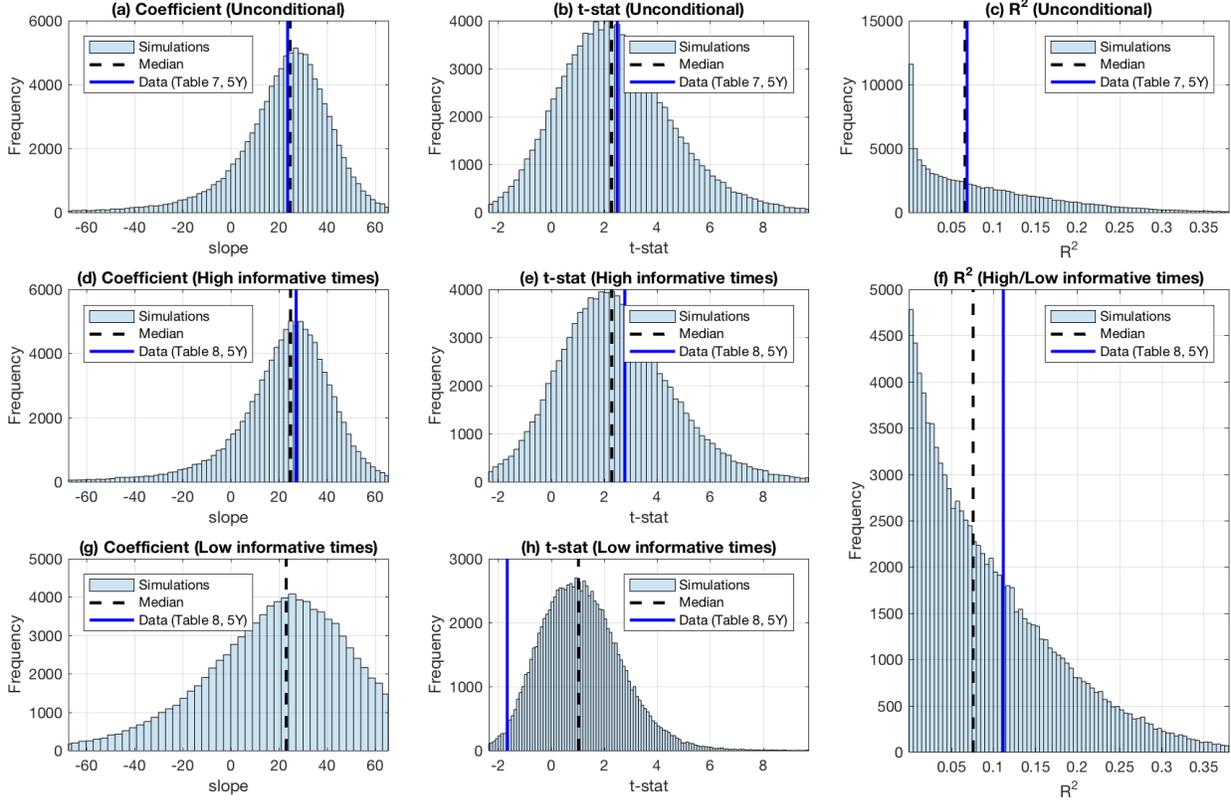


Figure 7: Return predictability with persistence risk (simulations).

Panels (a)-(c) show the results of the regression specification (44), based on 100,000 simulations. Each simulation has 192 quarters, as in our empirical analysis. The dashed lines are the medians from the simulations. The solid lines correspond to the regression results in the data (upper right panel of Table 7, column labeled “5Y”). The plots in the first row show 99% of the observations. Panels (d)-(h) do the same for the asymmetric regression specification (45). The solid lines correspond to the column labeled “5Y” in the upper right panel of Table 8. For ease of comparison, we maintain the x-axis limits of the plots from the first row on the rows below.

bottom or the top quartile—persistence risk predicts future excess returns better than during times when f_t is close to $\bar{\mu}$.

4.4.3 Return predictability with the price-dividend ratio

There is strong empirical evidence on the predictability of excess stock returns with the price-dividend ratio (e.g., Fama and French, 1988; Hodrick, 1992; Cochrane, 2008). For example, BC show that the logarithm of the price-dividend ratio predicts excess returns negatively with an economically important R^2 . We confirm this predictability over our more recent sample, which spans the period Q1:1969–Q4:2016. We consider the following regression at

Panel A: Predictability in the model					Panel B: Predictability in the data				
	Excess return					Excess return			
	1Y	3Y	5Y	7Y		1Y	3Y	5Y	7Y
Log (P/D)	-0.418***	-1.158***	-1.269***	-1.464***	Log (P/D)	-0.073***	-0.146**	-0.216**	-0.290***
t-stat	-6.190	-5.250	-6.261	-6.349	t-stat	-2.953	-2.001	-2.367	-3.346
R-squared	0.060	0.222	0.224	0.216	R-squared	0.027	0.046	0.070	0.140
N	188	180	172	164	N	188	180	172	164
With controls					With controls				
Log (P/D)	-0.778***	-1.370***	-1.763***	-2.328***	Log (P/D)	-0.163***	-0.274***	-0.364***	-0.369***
t-stat	-6.429	-3.458	-3.244	-3.937	t-stat	-5.289	-3.362	-3.854	-3.369
R-squared	0.275	0.288	0.264	0.312	R-squared	0.184	0.189	0.302	0.283
N	188	180	172	164	N	188	180	172	164

Table 10: Predictability of excess returns with the price-dividend ratio.

This table reports the predictability of excess stock returns with the log price-dividend ratio. Panel A reports the results for the model and Panel B in the data. The upper panels report the results for the univariate regressions and the lower panels report results when controlling for the level of stock market volatility (French et al., 1987), the default premium (Fama and French, 1989), and the macro uncertainty index of Jurado et al. (2015). Each column represents a different forecast horizon K and N is the number of observations. t -statistics are computed with Newey and West (1987) standard errors with $2(K - 1)$ lags. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The construction of the price-dividend ratio is discussed in Appendix D. Stock prices represent the value-weighted CRSP index and the sample spans the period Q1:1969-Q4:2016.

quarterly frequency:

$$\sum_{k=1}^K (r_{t+k} - r_{f,t+k}) = a_K + b_K pd_t + \mathbf{c}_K X_t + \epsilon_{t+K}, \quad (46)$$

where pd_t denotes the logarithm of the price-dividend (P/D) ratio and X_t is a vector of control variables, which consists of the level of stock market volatility, the default premium, and the macro uncertainty index of Jurado et al. (2015). Panel B of Table 10 provides support for the strong predictability of excess returns with the price-dividend ratio at any horizon.

We show that such predictability is also present in our model. Panel A of Table 10 reproduces the panel B but with model-implied asset prices. Consistent with the data, the relations are statistically significant at any horizon (1 year to 7 years). Moreover, we find strong predictive power, as the R^2 rises to 21.6%. Such predictive power is in contrast with that obtained in long-run risk models. For example, BC find that R^2 statistics are typically close to zero when simulating the BY and BKY models.

In our model, the price-dividend ratio predicts future returns for the following reason. A decrease in the growth forecast f_t and/or its perceived mean-reverting speed $\hat{\theta}_t$ reduces

the price-dividend ratio (Figure 11). This leads to greater expected returns (Figure 2). As a result, the price-dividend ratio *negatively* predicts future excess returns.

The channel through which the price-dividend ratio predicts future excess returns is fundamentally different from the one in long-run risk models. In [BY](#) and [BKY](#), return predictability is generated by persistent variation in the volatility of consumption growth—a feature which is absent in our case. As [BC](#) show, however, this channel is not strong enough to produce the return predictability observed in the data (the [BKY](#) calibration fares better by assuming *extremely persistent* consumption volatility, but even in this case the model is only in rough concordance with the data). We propose here an alternative predictability channel, which arises from learning about persistence. A distinction between our return predictability channel and the one analyzed by [BY](#) and [BKY](#) is that ours results endogenously from the rational response of economic agents to incomplete information about the degree of persistence in economic growth.

To better understand how predictability arises in our model, we once again use simulations. We generate 100,000 samples of 192 quarterly observations and then run the regression specification (46) on each simulated sample. We report results for the 5-year horizon and compare these results with the predictability found in the data (upper right panel of Table 10, column labeled “5Y”).

Figure 8 shows the simulation results with four rows of plots. The first row provides results for the model with learning about persistence; the second row depicts results from simulations of a model with learning about level; the third row simulates a model without learning; and the fourth row simulates a model with time-varying but observable persistence parameter λ_t , as considered in Section 3.3.2, with $\rho = 1$. Rows two to four in the figure adopt the x-axis limits of row one. These limits include 99% of the occurrences in the model with learning about persistence.

Panels (a) to (c) in the first row of Figure 8 show that the model with learning about persistence generates predictability that is in line with the data. With the exception of the slope coefficient, which is stronger than in the data, the median t-stat and the median R^2 are close to their empirical counterparts. The rest of the rows (panels (d) to (l)) show that none of the alternative models can generate the same degree of predictability. The only model that comes slightly close is the one with learning about level (second row), although the median t-stat is about 1.5, below standard significance levels. Overall, the R^2 coefficients in panels (f), (i) and (l) show weak evidence of return predictability in alternative model specifications.

These simulations emphasize the importance of incomplete information *about the degree of persistence* for generating sizable return predictability. Instead of assuming that return

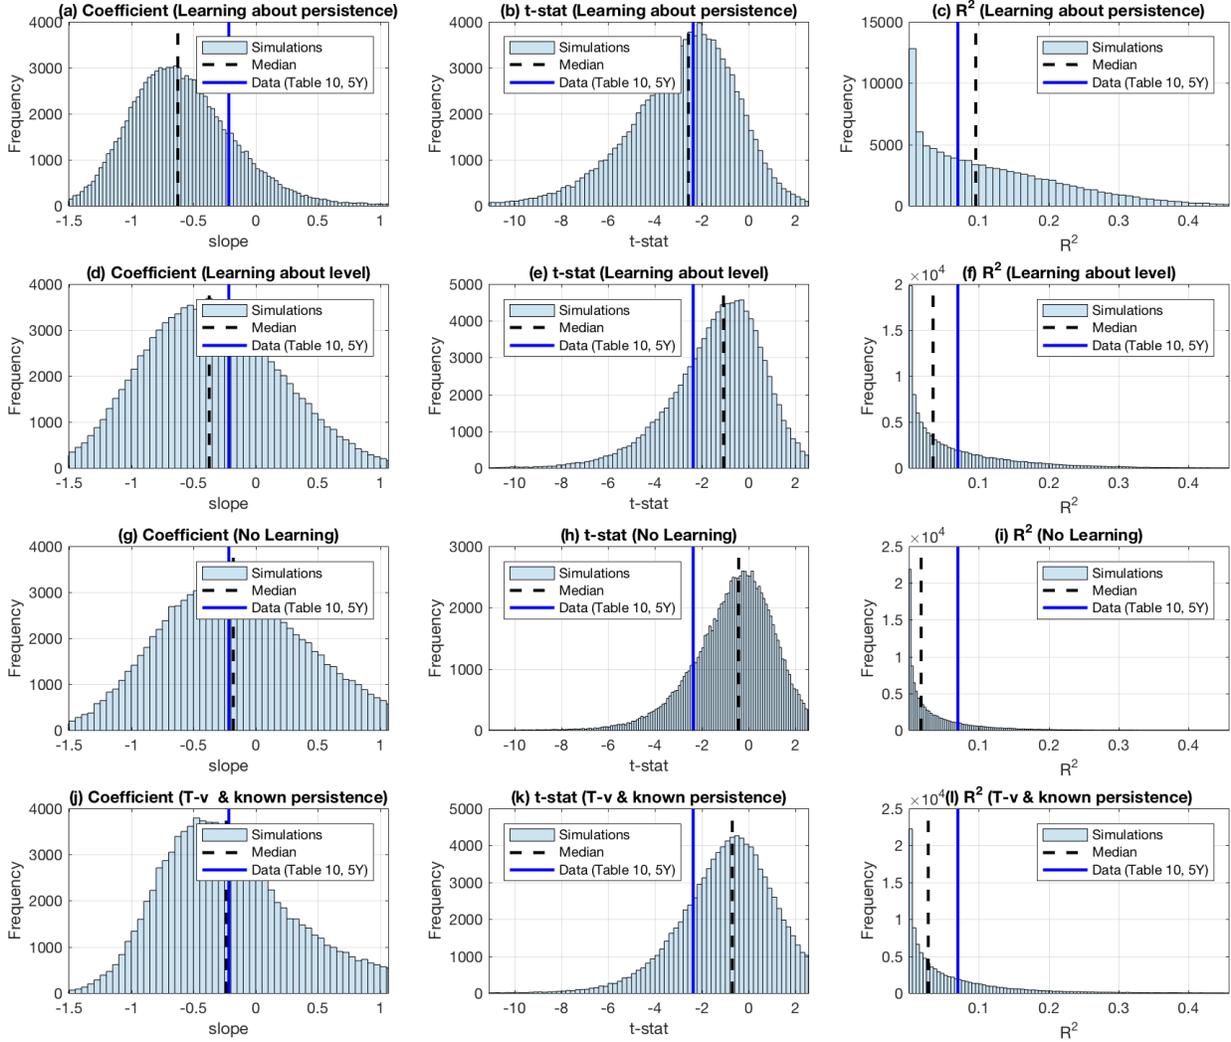


Figure 8: Return predictability in different models (simulations).

This figure shows the results of the regression specification (46), based on 100,000 simulations. Each simulation has 192 quarters, as in our empirical analysis. Each row corresponds to a different model: the first row to a model with learning about persistence; the second row to a model with learning about level; the third row to a model without learning; the fourth row to a model with time-varying and known persistence (Section 3.3.2, with $\rho = 1$). The dashed lines are the medians from the simulations. The solid lines are the numbers from the data (upper right panel of Table 10, column labeled “5Y”). The plots in the first row show 99% of the observations. For ease of comparison, we maintain the x-axis limits of the plots from the first row on the rows below.

predictability arises due to fluctuations in the volatility of consumption growth, we propose a stronger predictability channel that arises when agents are uncertain about the degree of persistence in economic growth. As the above results show, when parameters are carefully calibrated to fit U.S. output data, this learning channel generates strong return predictability.

4.4.4 Consumption growth predictability with the price-dividend ratio

The long-run risk model implies that the price-dividend ratio should predict future consumption growth. This predictability, however, is at odds with the data (BC). Using time series from the Bureau of Economic Analysis on U.S. real nondurables and services consumption per capita, we estimate the following regression specification at quarterly frequency from Q2:1947 to Q4:2016 (the longest possible postwar dataset):

$$\sum_{k=1}^K (\Delta c_{t+k}) = a_K + b_K p d_t + \epsilon_{t+K}. \quad (47)$$

Panel B of Table 11 shows the estimation results. The coefficients are negative and not statistically different from zero, except at the 7-year horizon (if we end our sample in Q4:2008—the exact sample analyzed by BC—all coefficients lose statistical significance and the R^2 are much lower). The conclusion is that the price-dividend ratio does *not* predict future consumption growth.

This represents a challenge for long-run risk models. Intuitively, if economic agents perceive a high degree of persistence in consumption, then it should influence asset prices, which in turn should predict future consumption growth. Thus, the feature that generates “excessive” predictability of consumption growth in long-run risk models is the assumed high degree of persistence in consumption.

We expect our model to overcome this challenge. As the estimation in Section 3.1 shows, the degree of persistence in consumption growth is much lower than in long-run risk models (for comparison, the parameter $\bar{\theta}$ would be 0.255 with the BY calibration; in our case, $\bar{\theta} = 1.35$). We therefore run the regression specification (47) on model-implied data and show the results in panel A of Table 11 (notice that our sample size is smaller, because of the availability of forecast data). The regression coefficients are all close to zero, as in the data. They are negative for longer horizons, although only weakly statistically significant. Importantly, all R^2 coefficients are low, as in the data.

We remark that the observed price-dividend does not predict the 1-year-ahead consumption growth in the data, whereas the model-implied price-dividend ratio does. The reason is that the price-dividend ratio in the model varies positively with the growth rate forecast (see Figure 11), which is an accurate predictor of the realized growth rate for the near future. The current price-dividend ratio thus positively predicts consumption growth over the short term in our model. Notice, however that the R^2 is relatively low. This means that the price-dividend ratio is able to explain well only a small fraction of the variation in the 1-year-ahead consumption growth.

Panel A: Predictability in the model

Consumption growth				
Period: 1969Q1–2016Q4				
	1Y	3Y	5Y	7Y
Log (P/D)	0.028***	0.005	-0.032*	-0.050*
t-stat	7.125	0.405	-1.903	-1.703
R-squared	0.095	0.001	0.014	0.025
N	188	180	172	164

Panel B: Predictability in the data

Consumption growth				
Period: 1947Q2–2016Q4				
	1Y	3Y	5Y	7Y
Log (P/D)	-0.002	-0.016	-0.029	-0.045*
t-stat	-0.638	-1.382	-1.594	-1.792
R-squared	0.005	0.054	0.090	0.139
N	274	266	258	250

Table 11: Predictability of consumption growth with the price-dividend ratio.

This table reports the predictability of consumption growth with the log price-dividend ratio. Panel A reports the results using the model-implied price-dividend ratio, while Panel B reports the results with the empirical counterpart. Each column represents a different forecast horizon K and N is the number of observations. t -statistics are computed with Newey and West (1987) standard errors with $2(K - 1)$ lags. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The construction of the price-dividend ratio is discussed in Appendix D. Stock prices represent the value-weighted CRSP index. The samples are Q1:1969-Q4:2016 for Panel A and Q2:1947-Q4:2016 for Panel B.

We once again consider simulations to investigate the predictability of long-run consumption growth. In panels (a)-(c) of Figure 9, we draw 100,000 samples of quarterly data from our model with learning about persistence. The panels show distributions of the slope coefficients, the t -stats, and the R^2 coefficients. Dashed lines indicate medians over simulations and solid lines the numbers from Panel B of Table 11, column labeled “5Y.” The median t -stat is below two and, most important, the median R^2 is low as in the data. This confirms that our model maintains a low predictability of future consumption growth over the long run.

Panels (d)-(f) of Figure 9 perform the same exercise, but with the calibration of the standard long-run risk model (BY, Case I without fluctuating consumption growth volatility). The parameters of the BY model (adapted to a continuous time setting) are provided in the caption of the figure. These panels show strong evidence of predictability of future consumption growth. First, the median of the slope coefficients quadruples, the median t -stat is now close to four, and the median R^2 is much higher than in the data. We conclude that the excessive predictability of consumption growth in long-run risk models is therefore caused by the high degree of persistence considered in their calibration. Our model maintains a relatively low predictability of future consumption growth by considering a lower degree of persistence.

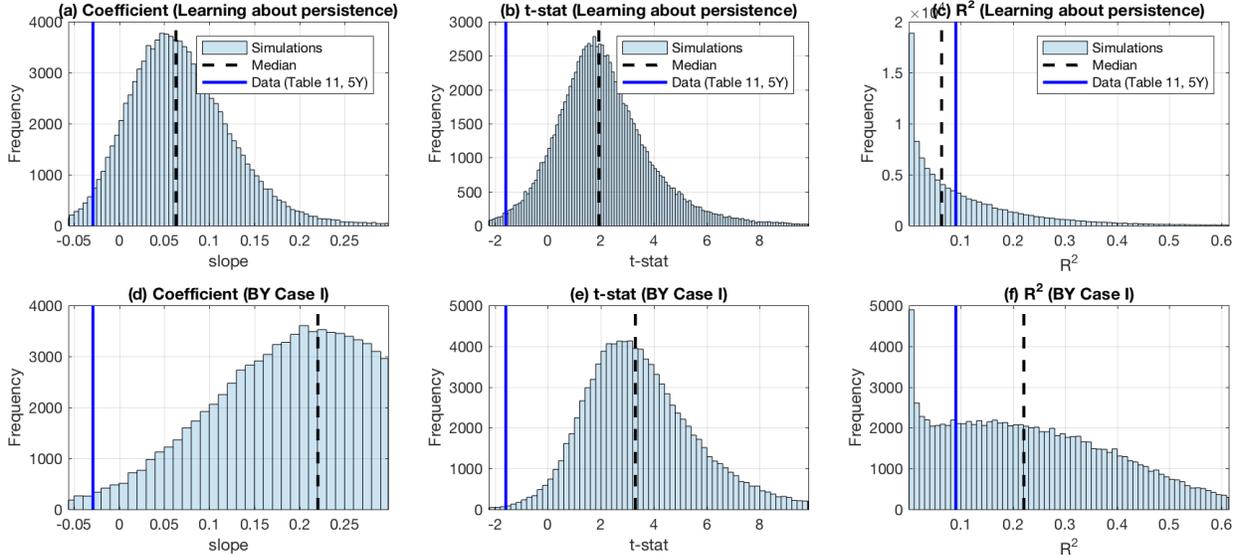


Figure 9: Consumption growth predictability (simulations).

This figure shows the consumption growth predictability based on simulations. We compare a model with learning about persistence (panels (a)-(c)) with the long-run risk model (BY, Case I) (panels (d)-(f)). The panels show the results of the regression (47) on 100,000 samples of simulated data of 192 quarters. The dashed lines are the medians from the simulations. The solid lines are the numbers from the data (panel B of Table 11, column labeled “5Y”). The plots in the first row show 99% of the observations. For ease of comparison, we maintain the x-axis limits of the plots from the first row on the row below. The calibration used for the long-run risk model is: $\sigma_\delta = 0.0270$, $\bar{\mu} = 0.0184$, $\sigma_f = 0.0144$, $\bar{\theta} = 0.2547$, $\beta = 0.0240$, $\beta_d = 0.0630$, $\psi = 1.5$, $\gamma = 10$, and $\eta = 4.5$.

4.5 Relation with the habit literature

Learning about persistence generates predictions that are related to those of the habit formation literature. One feature that both models capture is the dependence of stock prices and expected returns on past consumption.

In Campbell and Cochrane (1999), the price-dividend ratio falls and expected returns rise when consumption approaches a slow-moving habit level. To illustrate this, assume that the habit level, X_t , is a simple moving average of past consumption:

$$X_t = \phi X_{t-1} + (1 - \phi)C_t \tag{48}$$

where ϕ is arbitrarily fixed at 0.9.²¹ Figure 10 compares the time series of consumption minus habit, $C_t - X_t$, with the price-dividend ratio (left panels) and with the risk premium

²¹See Cochrane (2017, Figure 3), for a similar exercise. The results are (qualitatively) insensitive to the choice of ϕ .

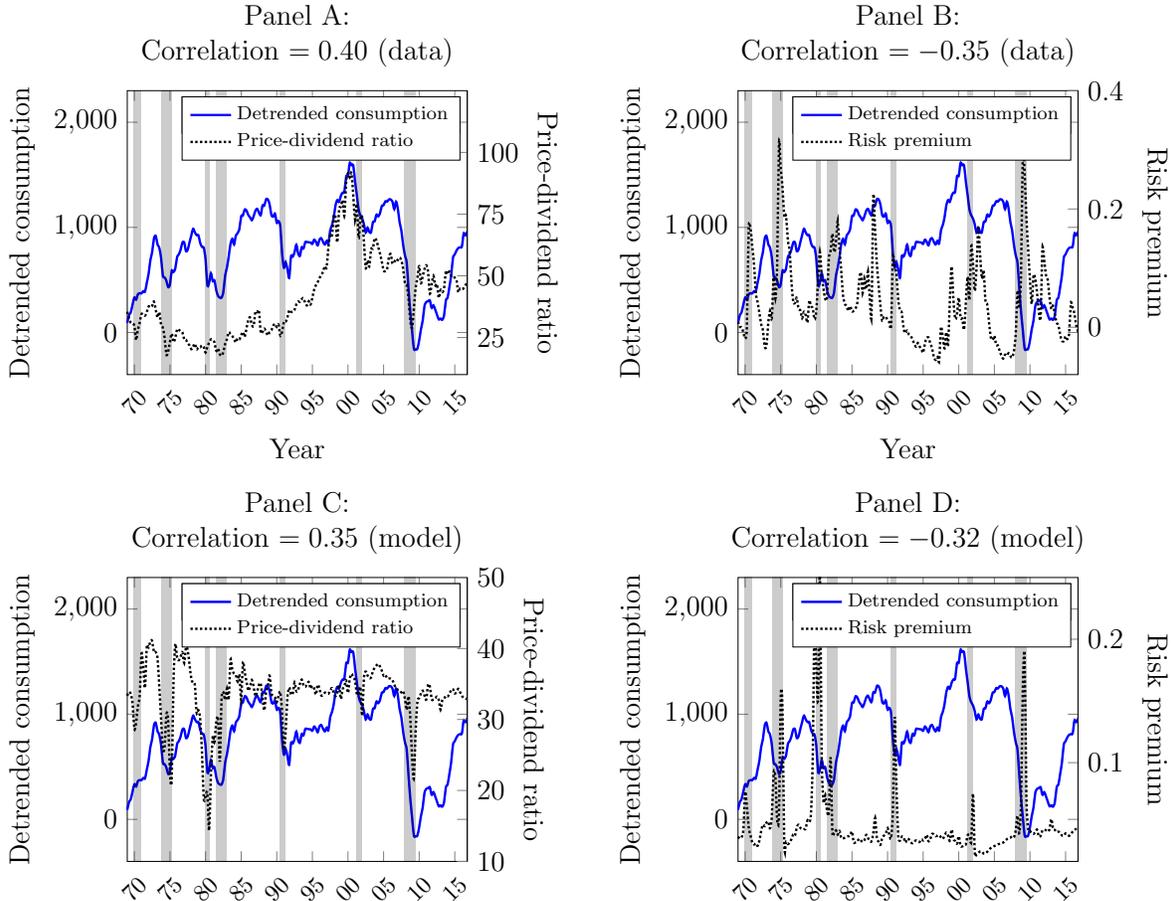


Figure 10: Asset prices and detrended consumption: link with habit formation.

The left panels depict the relations between the price-dividend ratio and the detrended consumption, whereas the right panels depict the relations between the risk premium and the detrended consumption. Detrended consumption is computed as the consumption minus habit, $C_t - X_t$, following Cochrane (2017). The upper panels report the empirical relations, and the lower panels report the model-implied relations. The construction of the price-dividend ratio and the risk premium is discussed in Appendix D. Stock prices represent the value-weighted CRSP index and the sample spans the period Q1:1969-Q4:2016.

(right panels). The correlations between the time series are given in the titles of the plots. In line with the prediction of the habit model, stock prices are low and expected returns are high in bad times, when consumption is low relative to the habit level. This occurs both in the data (upper panels) and in the model with learning about persistence (lower panels).

The habit formation model delivers this mechanism through reverse-engineering of the preference structure. By contrast, in the long-run risk model, the fear about long-run consumption growth—which derives from a preference for early resolution of uncertainty and is responsible for the high risk premium—is by construction unrelated to current macroeconomic conditions or to past consumption. Critics of the long-run risk model argue that

it can deliver the relationships depicted in the upper panels of Figure 10 only through an exogenously assumed “coincidence” (Cochrane, 2017), whereby negative consumption shocks are correlated with bad news about long-run future consumption growth.

This is not the case in a model with uncertainty about persistence. Bad news about long-run future consumption growth amplifies fluctuations precisely during bad times, when the output growth gap is above zero. This arises endogenously through the optimal learning problem of the agent. Most important, because growth forecasts are presumably positively related to the consumption minus habit $C_t - X_t$ variable,²² both the price dividend ratio and the risk premium implied by our model become endogenously linked to past consumption. We illustrate this in the lower panels of Figure 10, which show a positive correlation between consumption minus habit $C_t - X_t$ and the model-implied price-dividend ratio (panel C), and a negative correlation with the risk premium (Panel D). The correlations are similar when using model-implied or empirical asset prices, thus providing evidence that our model is able to capture a salient feature of the data.

As such, incomplete information about the degree of persistence helps the long-run risk model generate depressed prices and high expected returns during recessions, when consumption is low relative to its recent past. It relates two competing theories, habit and long-run risk, in an interesting—and non-exclusive—way. The habit model postulates an increase in risk aversion as consumption declines, whereas incomplete information about persistence postulates a fear of longer business cycles when expected growth is relatively low.

5 Conclusion

This paper shows that learning about the degree of persistence of expected economic growth can be a determinant force in financial markets. In equilibrium, this type of learning generates a form of uncertainty that we call *persistence risk*, which is positive during recessions and negative during expansions. This property of persistence risk implies a negative relation between asset pricing moments (equity return volatility, equity risk premium, and the Sharpe ratio) and economic conditions. The relation disappears when there is no uncertainty about persistence, or when economic agents learn about the level of expected growth.

Our model produces several unique predictions for which the data provide empirical support. It implies an increasing, non-linear relationship between persistence risk and asset pricing moments. In the model, as in the data, persistence risk is a good predictor of future excess returns, but only during times when uncertainty about persistence matters most (i.e.,

²²A regression of $C_t - X_t$ onto f_t yields a positive and statistically significant slope with a t -stat of 3.8; the correlation between the two time series is 27%.

when the expected economic growth is away from its long-term mean). We also contribute to the debate that long-run risk models struggle to explain the high predictability of equity returns, while at the same time maintaining low predictability of consumption growth. In our model with learning about persistence, the price-dividend ratio strongly predicts future excess returns, but only weakly predicts future consumption growth.

Our analysis offers several directions for future research. One direction is to build a theory in which agents endogenously choose on which dimension to learn (level or persistence). This theory would help investigate at least two questions. First, if investors' attention is costly, when is it optimal to be more attentive to the level versus the persistence? Second, what are the welfare gains of a publicly available indicator about the persistence of economic growth?

Another direction for future research is to study an economy with multiple risky assets. This extension would help investigate the implications of learning about persistence for the cross-section of asset returns.

Finally, 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 data. Such disagreement measure might help better predict future market returns and their volatility than 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 disagreement and uncertainty about the length of business cycles. *Management Science*.
- 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 I. Shaliastovich (2010, May). Confidence Risk and Asset Prices. *American Economic Review* 100(2), 537–541.
- 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.
- Bloom, N. (2009). The impact of uncertainty shocks. *econometrica* 77(3), 623–685.
- 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.
- Cochrane, J. H. (2017). Macro-finance. *Review of Finance* 21(3), 945–985.
- 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.
- Dangl, T. and M. Halling (2012). Predictive regressions with time-varying coefficients. *Journal of Financial Economics* 106(1), 157–181.
- 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*.
- Detemple, J. (1986). Asset pricing in a production economy with incomplete information. *The Journal of Finance* 41(2), pp. 383–391.
- Drechsler, I. (2013). Uncertainty, time-varying fear, and asset prices. *The Journal of Finance* 68(5), 1843–1889.
- 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.

- Fama, E. F. and K. R. French (1988). Dividend yields and expected stock returns. *Journal of financial economics* 22(1), 3–25.
- Fama, E. F. and K. R. French (1989). Business conditions and expected returns on stocks and bonds. *Journal of Financial Economics* 25(1), 23–49.
- 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.
- Ghosh, A. and G. M. Constantinides (2016). What information drives asset prices? *Working Paper*.
- Hamilton, J. D. (1994). *Time Series Analysis*. Princeton: Princeton University Press.
- Hansen, L. P. (2014). Uncertainty outside and inside economic models. *Journal of Political Economy* 122(5), 945–87.
- Hansen, L. P. and T. J. Sargent (2010). Fragile beliefs and the price of uncertainty. *Quantitative Economics* 1(1), 129–162.
- Henkel, S. J., J. S. Martin, and F. Nardari (2011). Time-varying short-horizon predictability. *Journal of Financial Economics* 99(3), 560–580.
- Hodrick, R. J. (1992). Dividend yields and expected stock returns: Alternative procedures for inference and measurement. *The Review of Financial Studies* 5(3), 357–386.
- 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.
- Jurado, K., S. C. Ludvigson, and S. Ng (2015). Measuring uncertainty. *The American Economic Review* 105(3), 1177–1216.
- 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. Shirayev (1977). *Statistics of random processes*. Springer-Verlag, Berlin, New York.

- Lustig, H., S. Van Nieuwerburgh, and A. Verdelhan (2013). The wealth-consumption ratio. *Review of Asset Pricing Studies* 3(1), 38–94.
- Lustig, H. and A. Verdelhan (2012). Business cycle variation in the risk-return trade-off. *Journal of Monetary Economics* 59, S35–S49.
- 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.
- Rapach, D. E., J. K. Strauss, and G. Zhou (2010). Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. *The Review of Financial Studies* 23(2), 821–862.
- Scheinkman, J. A. and W. Xiong (2003). Overconfidence and speculative bubbles. *The Journal of Political Economy* 111(6), 1183–1219.
- 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 (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.
- Vuong, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica* 57(2), 307–333.

Ziegler, A. C. (2003). *Incomplete Information and Heterogeneous Beliefs in Continuous-time Finance*. Springer-Verlag Berlin Heidelberg.

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

A Appendix: Learning

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

$$du_t = [a_0(t, s_t) + a_1(t, s_t)u_t] dt + b_1(t, s_t)dZ_t^u + b_2(t, s_t)dZ_t^s \quad (49)$$

$$ds_t = [A_0(t, s_t) + A_1(t, s_t)u_t] dt + B_1(t, s_t)dZ_t^u + B_2(t, s_t)dZ_t^s. \quad (50)$$

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

$$d\hat{u}_t = (a_0 + a_1\hat{u}_t)dt + [(b \circ B) + \nu_t A_1^\top](B \circ B)^{-1}[ds_t - (A_0 + A_1\hat{u}_t)dt] \quad (51)$$

$$\frac{d\nu_t}{dt} = a_1\nu_t + \nu_t a_1^\top + (b \circ b) - [(b \circ B) + \nu_t A_1^\top](B \circ B)^{-1}[(b \circ B) + \nu_t A_1^\top]^\top, \quad (52)$$

where

$$b \circ b = b_1 b_1^\top + b_2 b_2^\top \quad (53)$$

$$B \circ B = B_1 B_1^\top + B_2 B_2^\top \quad (54)$$

$$b \circ B = b_1 B_1^\top + b_2 B_2^\top. \quad (55)$$

Write the dynamics of the observable variables

$$\begin{bmatrix} d \log \delta_t \\ df_t \end{bmatrix} = \left(\underbrace{\begin{bmatrix} f_t - \frac{1}{2}\sigma_\delta^2 \\ \bar{\theta}(\bar{\mu} - f_t) \end{bmatrix}}_{A_0} + \underbrace{\begin{bmatrix} 1 & 0 \\ 0 & (\bar{\mu} - f_t) \end{bmatrix}}_{A_1} \begin{bmatrix} \epsilon_t \\ \lambda_t \end{bmatrix} \right) dt + \underbrace{\begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}}_{B_1} \begin{bmatrix} dW_t^\epsilon \\ dW_t^\lambda \end{bmatrix} + \underbrace{\begin{bmatrix} \sigma_\delta & 0 \\ 0 & \sigma_f \end{bmatrix}}_{B_2} \begin{bmatrix} dW_t^\delta \\ dW_t^f \end{bmatrix} \quad (56)$$

and unobservable variables

$$\begin{bmatrix} d\epsilon_t \\ d\lambda_t \end{bmatrix} = \left(\underbrace{\begin{bmatrix} 0 \\ 0 \end{bmatrix}}_{a_0} + \underbrace{\begin{bmatrix} -\varphi & 0 \\ 0 & -\kappa \end{bmatrix}}_{a_1} \begin{bmatrix} \epsilon_t \\ \lambda_t \end{bmatrix} \right) dt + \underbrace{\begin{bmatrix} \sigma_\epsilon & 0 \\ 0 & \sigma_\lambda \end{bmatrix}}_{b_1} \begin{bmatrix} dW_t^\epsilon \\ dW_t^\lambda \end{bmatrix} + \underbrace{\begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}}_{b_2} \begin{bmatrix} dW_t^\delta \\ dW_t^f \end{bmatrix}. \quad (57)$$

Then,

$$\begin{bmatrix} d\hat{\epsilon}_t \\ d\hat{\lambda}_t \end{bmatrix} = \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix} + \begin{bmatrix} -\varphi & 0 \\ 0 & -\kappa \end{bmatrix} \begin{bmatrix} \hat{\epsilon}_t \\ \hat{\lambda}_t \end{bmatrix} \right) dt + \begin{bmatrix} \frac{\nu_{\epsilon,t}}{\sigma_\delta} & 0 \\ 0 & \frac{(\bar{\mu} - f_t)\nu_{\lambda,t}}{\sigma_f} \end{bmatrix} \begin{bmatrix} d\widehat{W}_t^\delta \\ d\widehat{W}_t^f \end{bmatrix}, \quad (58)$$

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

$$\frac{d\delta_t}{\delta_t} = (f_t + \hat{\epsilon}_t)dt + \sigma_\delta d\widehat{W}_t^\delta \quad (59)$$

$$df_t = (\bar{\theta} + \hat{\lambda}_t)(\bar{\mu} - f_t)dt + \sigma_f d\widehat{W}_t^f. \quad (60)$$

The posterior uncertainties about ϵ_t and λ_t evolve according to

$$\frac{d\nu_{\epsilon,t}}{dt} = \sigma_\epsilon^2 - 2\varphi\nu_{\epsilon,t} - \frac{\nu_{\epsilon,t}^2}{\sigma_\delta^2} \quad (61)$$

$$\frac{d\nu_{\lambda,t}}{dt} = \sigma_\lambda^2 - 2\kappa\nu_{\lambda,t} - \frac{(\bar{\mu} - f_t)^2\nu_{\lambda,t}^2}{\sigma_f^2}. \quad (62)$$

The uncertainty about ϵ_t admits a constant steady-state solution (i.e. $\left. \frac{d\nu_{\epsilon,t}}{dt} \right|_{\nu_{\epsilon,t} \equiv \bar{\nu}_\epsilon} = 0$):

$$\bar{\nu}_\epsilon = \sigma_\delta \left(\sqrt{\varphi^2\sigma_\delta^2 + \sigma_\epsilon^2} - \varphi\sigma_\delta \right), \quad (63)$$

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

B Appendix: Equilibrium

The dynamics of the vector of state variables are

$$\begin{bmatrix} d\delta_t \\ df_t \\ d\hat{\epsilon}_t \\ d\hat{\lambda}_t \\ d\nu_{\lambda,t} \end{bmatrix} = \begin{bmatrix} \delta_t(f_t + \hat{\epsilon}_t) \\ (\bar{\theta} + \hat{\lambda}_t)(\bar{\mu} - f_t) \\ -\varphi\hat{\epsilon}_t \\ -\kappa\hat{\lambda}_t \\ \sigma_\lambda^2 - 2\kappa\nu_{\lambda,t} - \frac{(\bar{\mu} - f_t)^2\nu_{\lambda,t}^2}{\sigma_f^2} \end{bmatrix} dt + \begin{bmatrix} \delta_t\sigma_\delta & 0 \\ 0 & \sigma_f \\ \frac{\bar{\nu}_\epsilon}{\sigma_\delta} & 0 \\ 0 & \frac{(\bar{\mu} - f_t)\nu_{\lambda,t}}{\sigma_f} \\ 0 & 0 \end{bmatrix} \begin{bmatrix} d\widehat{W}_t^\delta \\ d\widehat{W}_t^f \end{bmatrix}. \quad (64)$$

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

$$\frac{h(C, J)}{J} = \frac{\phi}{I(x_t)} - \beta\phi. \quad (65)$$

Define

$$W_t = C_t I(x_t), \quad (66)$$

and replace (24) in the product $\xi_t W_t$ to get

$$\xi_t W_t = (1 - \gamma) \exp \left(\int_0^t \left(\frac{\phi - 1}{I(x_s)} - \beta\phi \right) ds \right) \frac{C_t^{1-\gamma}}{1 - \gamma} [\beta I(x_t)]^\phi \quad (67)$$

$$= (1 - \gamma) \exp \left(\int_0^t \left(\frac{\phi - 1}{I(x_s)} - \beta\phi \right) ds \right) J. \quad (68)$$

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

$$d(\xi_t W_t) = (1 - \gamma) \exp \left(\int_0^t \left(\frac{\phi - 1}{I(x_s)} - \beta\phi \right) ds \right) \left[dJ - J \left(\beta\phi - \frac{\phi - 1}{I(x_t)} \right) dt \right]. \quad (69)$$

We also know that

$$dJ = -h(C, J)dt + dM_t \quad (70)$$

$$= J \left(\beta\phi - \frac{\phi}{I(x_t)} \right) dt + dM_t, \quad (71)$$

where M_t is a martingale. The second equality follows from (65). Replace this in (69):

$$d(\xi_t W_t) = (1 - \gamma) \exp \left(\int_0^t \left(\frac{\phi - 1}{I(x_s)} - \beta\phi \right) ds \right) \left[J \left(\beta\phi - \frac{\phi}{I(x_t)} \right) dt + dM_t - J \left(\beta\phi - \frac{\phi - 1}{I(x_t)} \right) dt \right] \quad (72)$$

$$= -(1 - \gamma) \exp \left(\int_0^t \left(\frac{\phi - 1}{I(x_s)} - \beta\phi \right) ds \right) \frac{J}{I(x_t)} dt + d\widetilde{M}_t \quad (73)$$

$$= -\xi_t C_t dt + d\widetilde{M}_t, \quad (74)$$

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 on $[t, \infty)$. Then, taking expectation and assuming that the transversality condition holds yields the total wealth (claim to all future output):

$$W_t = \mathbb{E}_t \left[\int_t^\infty \frac{\xi_s}{\xi_t} C_s ds \right] \quad (75)$$

which proves that $I(x_t)$ is indeed the wealth-consumption ratio.

Partial differential equation for the wealth-consumption ratio. Define the log wealth-consumption ratio:

$$i \equiv \log I. \quad (76)$$

Substituting the guess (19) in the HJB Eq. (18) and imposing the market clearing condition, $C = \delta$, yields the following PDE for the log wealth consumption ratio:

$$\begin{aligned} 0 = & \frac{\gamma - 1}{\phi} \left[-f - \epsilon + \frac{1}{2} \gamma \sigma_\delta^2 \right] - \beta + e^{-i} \\ & + (\bar{\theta} + \widehat{\lambda})(\bar{\mu} - f)i_f - [(\gamma - 1)\bar{\nu}_\epsilon + \widehat{\epsilon}\varphi]i_{\widehat{\epsilon}} - \kappa\widehat{\lambda}i_{\widehat{\lambda}} + \left[\sigma_\lambda^2 - 2\kappa\nu_\lambda - \frac{(\bar{\mu} - f)^2\nu_\lambda^2}{\sigma_f^2} \right] i_{\nu_\lambda} \\ & + \frac{\sigma_f^2}{2} i_{ff} + \frac{\sigma_\epsilon^2 - 2\varphi\bar{\nu}_\epsilon}{2} i_{\widehat{\epsilon}\widehat{\epsilon}} + \frac{(\bar{\mu} - f)^2\nu_\lambda^2}{2\sigma_f^2} i_{\widehat{\lambda}\widehat{\lambda}} + \nu_\lambda(\bar{\mu} - f)i_{f\widehat{\lambda}} \\ & + \phi \frac{\sigma_f^2}{2} i_f^2 + \phi(\bar{\mu} - f)\nu_\lambda i_f i_{\widehat{\lambda}} + \phi \frac{\sigma_\epsilon^2 - 2\varphi\bar{\nu}_\epsilon}{2} i_{\widehat{\epsilon}}^2 + \phi \frac{(\bar{\mu} - f)^2\nu_\lambda^2}{2\sigma_f^2} i_{\widehat{\lambda}}^2. \end{aligned} \quad (77)$$

B.1 Levered equity

Define

$$P_t = D_t \Pi(x_t) = e^{-\beta a t} C_t^\eta \Pi(x_t). \quad (78)$$

Compute

$$\xi_t P_t = \underbrace{(1 - \gamma) \exp \left(\int_0^t \left(\frac{\phi - 1}{I(x_s)} - \beta \phi \right) ds - \beta_d t \right)}_{\equiv \Delta(t)} J C_t^{\eta-1} \frac{\Pi(x_t)}{I(x_t)}. \quad (79)$$

One can clearly see that if $\eta = 1$ and $\beta_d = 0$, then $C_t^{\eta-1}$ drops out and the last fraction equals one, which brings us back to (68). The case of interest is $\eta > 1$. Define

$$K(C_t, x_t) = C_t^{\eta-1} \frac{\Pi(x_t)}{I(x_t)} \quad (80)$$

and thus

$$d(\xi_t P_t) = \Delta(t) \left[-KJ \left(\frac{1}{I(x_t)} + \beta_d \right) + K dM_t + J dK + (dJ)(dK) \right]. \quad (81)$$

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

$$d(\xi_t P_t) = -\xi_t e^{-\beta_d t} C_t^\eta dt + d\widehat{M}_t, \quad (82)$$

where $d\widehat{M}_t$ is a martingale. This means that the drifts in (81) and (82) have to be equal. This yields a partial differential equation to be solved by $\Pi(x_t)$. Replacing $j(x_t) \equiv \ln \Pi(x_t)$ results in the following partial differential equation:

$$\begin{aligned} 0 = & e^{-j} - \beta - \beta_d - \frac{\gamma - 1 - \phi(\eta - 1)}{\phi} (f + \widehat{\epsilon}) + \frac{\sigma_\delta^2}{2} \left[\frac{\gamma(\gamma - 1)}{\phi} + (1 - 2\gamma)(\eta - 1) + (\eta - 1)^2 \right] \\ & - (1 - \phi)(\eta - 1) \bar{\nu}_\epsilon i_{\widehat{\epsilon}} + \frac{1 - \phi}{2} (\sigma_\epsilon^2 - 2\varphi \bar{\nu}_\epsilon) i_{\widehat{\epsilon}}^2 + \frac{1 - \phi}{2\sigma_f^2} [\sigma_f^2 i_f + (\bar{\mu} - f) \nu_\lambda i_{\widehat{\lambda}}]^2 \\ & + \left[(\bar{\theta} + \widehat{\lambda})(\bar{\mu} - f) - (1 - \phi)(\sigma_f^2 i_f + (\bar{\mu} - f) \nu_\lambda i_{\widehat{\lambda}}) \right] j_f \\ & - [(\gamma - \eta) \bar{\nu}_\epsilon + \varphi \widehat{\epsilon} + (1 - \phi)(\sigma_\epsilon^2 - 2\varphi \bar{\nu}_\epsilon)] j_{\widehat{\epsilon}} \\ & - \left[\kappa \widehat{\lambda} + (1 - \phi)(\bar{\mu} - f) \nu_\lambda \left(i_f + \frac{(\bar{\mu} - f) \nu_\lambda}{\sigma_f^2} i_{\widehat{\lambda}} \right) \right] j_{\widehat{\lambda}} \\ & + \left[\sigma_\lambda^2 - 2\kappa \nu_\lambda - \frac{(\bar{\mu} - f)^2 \nu_\lambda^2}{\sigma_f^2} \right] j_{\nu_\lambda} \\ & + \frac{\sigma_f^2}{2} j_{ff} + \frac{1}{2} (\sigma_\epsilon^2 - 2\varphi \bar{\nu}_\epsilon) j_{\widehat{\epsilon}\widehat{\epsilon}} + \frac{(\bar{\mu} - f)^2 \nu_\lambda^2}{2\sigma_f^2} j_{\widehat{\lambda}\widehat{\lambda}} + \nu_\lambda (\bar{\mu} - f) j_{f\widehat{\lambda}} \\ & + \frac{\sigma_f^2}{2} j_f^2 + \frac{1}{2} (\sigma_\epsilon^2 - 2\varphi \bar{\nu}_\epsilon) j_{\widehat{\epsilon}}^2 + \frac{(\bar{\mu} - f)^2 \nu_\lambda^2}{2\sigma_f^2} j_{\widehat{\lambda}}^2 + \nu_\lambda (\bar{\mu} - f) j_f j_{\widehat{\lambda}}. \end{aligned} \quad (83)$$

This equation has a similar structure with (77), except that it also involves the log wealth-consumption ratio i . It is a matter of algebra to verify that replacing $\eta = 1$ and $\beta_d = 0$ in (83) gives exactly (77).

Numerical procedure The PDE for $i(f, \hat{\epsilon}, \hat{\lambda}, \nu_\lambda)$ is solved numerically using the Chebyshev collocation method (Judd, 1998). That is, we approximate the function $i(f, \hat{\epsilon}, \hat{\lambda}, \nu_\lambda)$ as follows:

$$i(f, \hat{\epsilon}, \hat{\lambda}, \nu_\lambda) \approx P(\hat{\mu}, y, V, \nu) = \sum_{i=0}^I \sum_{j=0}^J \sum_{k=0}^K \sum_{l=0}^L a_{i,j,k,l} T_i[f] T_j[\hat{\epsilon}] T_k[\hat{\lambda}] T_l[\nu_\lambda],$$

where $T_m[\cdot]$ is the Chebyshev polynomial of order m . The interpolation nodes are obtained by meshing the scaled roots of the Chebyshev polynomials of order $I + 1$, $J + 1$, $K + 1$, and $L + 1$. We scale the roots of the Chebyshev polynomials of order $I + 1$, $J + 1$, $K + 1$, and $L + 1$ such that they cover the intervals (these bounds cover about 95% of the unconditional distributions of the 4 state variables, with the calibration provided in Table 1):

$$f \in [-0.01, 0.06] \tag{84}$$

$$\hat{\epsilon} \in [-0.013, 0.013] \tag{85}$$

$$\hat{\lambda} \in [-1.3, 1.3] \tag{86}$$

$$\nu_\lambda \in [0.07, 0.42]. \tag{87}$$

The polynomial $P(\hat{\mu}, y, V, \nu)$ and its partial derivatives are then substituted into the PDE, and the resulting expression is evaluated at the interpolation nodes. This yields a system of $(I + 1) \times (J + 1) \times (K + 1) \times (L + 1)$ equations with $(I + 1) \times (J + 1) \times (K + 1) \times (L + 1)$ unknowns (the coefficients $a_{i,j,k,l}$). This system of equations is solved numerically.

Once we solve for the wealth-consumption ratio i , we replace it in Eq (83), then we solve for the price-dividend ratio using the same procedure.

We solve separately the various cases in the paper. For our main model of interest (model of learning about persistence) there are only three state variables: f , $\hat{\lambda}$, and ν_λ . We generate a grid of 10^3 points. The mean squared PDE residuals over the set of 1,000 interpolation nodes is of order 10^{-8} . That is, the Chebyshev collocation method yields an accurate solution to the PDE.

B.2 Numerical evaluation of $I_{\hat{\lambda}}/I$ and $\Pi_{\hat{\lambda}}/\Pi$

We verify numerically that the signs of $I_{\hat{\lambda}}/I$ and $\Pi_{\hat{\lambda}}/\Pi$ are always positive with our calibration. Because these signs may depend on the utility parameters γ and ψ , as well as on the value of the state variables f_t and $\hat{\theta}_t = \hat{\theta} + \hat{\lambda}_t$, we tabulate results for different values of the utility parameters and of the two state variables.

Table 12 reports the values of the coefficient $I_{\hat{\lambda}}/I$ in several situations. The different panels of Table 12 correspond to various levels of the forecast f_t and mean-reversion speed $\hat{\theta}_t$. Within each panel, we compute the value of $I_{\hat{\lambda}}/I$ for different preference parameters γ and ψ .

The results indicate that the coefficient increases with the risk aversion and with the elasticity of intertemporal substitution. It is essentially positive, unless the risk aversion and the EIS are sufficiently small, and the mean-reversion speed is sufficiently large. Importantly, $I_{\hat{\lambda}}/I$ is always positive with our preference parameters ($\gamma = 10, \psi = 1.5$), thus confirming numerically our Conjecture 1.

The coefficient $I_{\hat{\lambda}}/I$ tends to increase in bad times (i.e. lower f_t) and when the persistence is stronger (i.e. lower $\hat{\theta}_t$). The fact that $I_{\hat{\lambda}}/I$ becomes smaller in good times is related to the following effect: 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 $I_{\hat{\lambda}}/I$ remains positive in good times, the second effect appears to be small.

(a) $f_t = 0\%, \hat{\theta}_t = 1.35$				(c) $f_t = 2.5\%, \hat{\theta}_t = 1.35$				(e) $f_t = 5\%, \hat{\theta}_t = 1.35$			
γ/ψ	1.5	2	2.5	γ/ψ	1.5	2	2.5	γ/ψ	1.5	2	2.5
10	0.81	1.78	2.90	10	0.46	1.01	1.67	10	0.07	0.17	0.31
12	0.89	1.91	3.05	12	0.55	1.20	1.94	12	0.17	0.39	0.66
14	0.99	2.09	3.28	14	0.66	1.42	2.24	14	0.28	0.63	1.02

(b) $f_t = 2.5\%, \hat{\theta}_t = 0.4$				(d) $f_t = 2.5\%, \hat{\theta}_t = 1.35$				(f) $f_t = 2.5\%, \hat{\theta}_t = 2.30$			
γ/ψ	1.5	2	2.5	γ/ψ	1.5	2	2.5	γ/ψ	1.5	2	2.5
10	0.15	0.30	0.49	10	0.46	1.01	1.67	10	0.31	0.70	1.17
12	0.21	0.43	0.69	12	0.55	1.20	1.94	12	0.37	0.83	1.35
14	0.29	0.61	0.96	14	0.66	1.42	2.24	14	0.45	0.98	1.56

Table 12: Values of the coefficient $I_{\hat{\lambda}}/I$

This table reports a numerical evaluation of $I_{\hat{\lambda}}/I$. We consider different levels of risk aversion, $\gamma \in \{10, 12, 14\}$, and different levels of the elasticity of intertemporal substitution, $\psi \in \{1.5, 2, 2.5\}$. There are six panels. The upper panels keep $\hat{\theta}_t = \bar{\theta} = 1.35$ but use different values for f_t . The lower panels keep $f_t = \bar{\mu} = 2.5\%$ but consider different values $\hat{\lambda}_t$. For all tables, the uncertainty is $\nu_{\lambda,t} = 0.2$ (the effect of a change in uncertainty on $I_{\hat{\lambda}}/I$ is relatively weak). Unless otherwise specified, we consider the calibration provided in the second column of Table 1.

For the price-dividend ratio Π , we find coefficients $\Pi_{\hat{\lambda}}/\Pi$ that are an order of magnitude higher than $I_{\hat{\lambda}}/I$. The coefficients are positive at all times, decrease in γ , decrease in ψ , increase in bad times and when the persistence is stronger. We therefore do not tabulate them here.

B.3 Price-dividend ratio and risk-free rate

This Appendix provides additional results for Section 3.2. Figure 11 plots the log price-dividend ratio and the equilibrium risk-free rate. The log price-dividend ratio increases with the output growth forecast f_t (upper left panel). The relationship is almost linear, implying that Π_f/Π 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 $\hat{\theta}_t$. This implies $\Pi_{\hat{\lambda}}/\Pi > 0$.

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

$$r_{f,t} = \beta + \frac{1}{\psi} f_t - \frac{\gamma + \gamma\psi}{2\psi} \sigma_\delta^2 - \frac{1}{2}(1 - \phi) \left(\sigma_f \frac{I_f}{I} + \frac{(\bar{\mu} - f_t)\nu_{\lambda,t}}{\sigma_f} \frac{I_{\hat{\lambda}}}{I} \right)^2. \quad (88)$$

The lower panels of Figure 11 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 perceived persistence $\hat{\theta}_t$ (lower right panel), although the latter effect is relatively weaker. Uncertainty about persistence decreases the risk-free rate but its impact is weak.

B.4 A model with time-varying, but observable persistence

The solution method follows the same steps as before (Appendix B). There are three state variables, δ_t , f_t , and λ_t , whose dynamics are given in (3), (4), and respectively (41). The partial differential

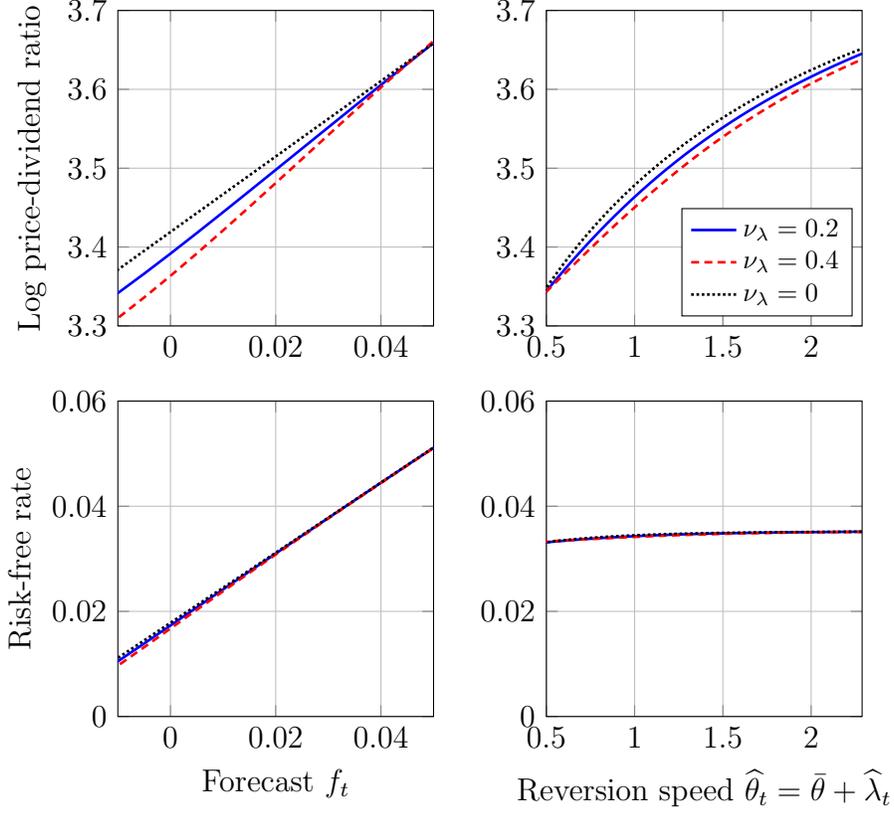


Figure 11: Behavior of the price-dividend ratio and risk-free rate with learning about persistence.

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 $\hat{\lambda}_t = 0$. For the right graph, we fix $f_t = \bar{\mu}$. Unless otherwise specified, we consider the calibration provided in the second column of Table 1.

equation for the wealth-consumption ratio (which is a function of f_t and λ_t) writes:

$$\begin{aligned}
0 = & \frac{\gamma - 1}{\phi} \left[-f - \epsilon + \frac{1}{2} \gamma \sigma_\delta^2 \right] - \beta + e^{-i} + (\bar{\theta} + \lambda)(\bar{\mu} - f)i_f - \kappa \lambda i_\lambda \\
& + \frac{\sigma_f^2}{2} i_{ff} + \frac{\sigma_\lambda^2}{2} i_{\lambda\lambda} + \rho \sigma_f \sigma_\lambda i_{f\lambda} + \phi \frac{\sigma_f^2}{2} i_f^2 + \phi \rho \sigma_f \sigma_\lambda i_f i_\lambda + \phi \frac{\sigma_\lambda^2}{2} i_\lambda^2.
\end{aligned} \tag{89}$$

The partial differential equation for the price-dividend ratio then follows as in Appendix B.1. We solve these two equations numerically as before. In order to keep the results comparable, we use the same number of Chebyshev nodes.

In this model, the stock return volatility is

$$\|\sigma_t\| = \sqrt{\eta^2 \sigma_\delta^2 + \left(\sigma_f \frac{\Pi_f}{\Pi} + \rho \sigma_\lambda \frac{\Pi_\lambda}{\Pi} \right)^2 + (1 - \rho^2) \sigma_\lambda^2 \left(\frac{\Pi_\lambda}{\Pi} \right)^2}, \tag{90}$$

whereas the risk premium is given by:

$$RP_t = \gamma\eta\sigma_\delta^2 + (1 - \phi) \left(\sigma_f \frac{I_f}{I} + \rho\sigma_\lambda \frac{I_\lambda}{I} \right) \left(\sigma_f \frac{\Pi_f}{\Pi} + \rho\sigma_\lambda \frac{\Pi_\lambda}{\Pi} \right) + (1 - \phi)(1 - \rho^2)\sigma_\lambda^2 \frac{I_\lambda}{I} \frac{\Pi_\lambda}{\Pi}. \quad (91)$$

C Appendix: Estimation

To fit our continuous-time model to the data, we first discretize the filtered dynamics in Equations (58), (59), (60), and (62) using the following approximations

$$\log(\delta_{t+\Delta}/\delta_t) = \left(f_t + \hat{\epsilon}_t - \frac{1}{2}\sigma_\delta^2 \right) \Delta + \sigma_\delta \sqrt{\Delta} v_{1,t+\Delta}, \quad (92)$$

$$f_{t+\Delta} = e^{-\hat{\theta}_t \Delta} f_t + \left(1 - e^{-\hat{\theta}_t \Delta} \right) \bar{\mu} + \sigma_f \sqrt{\frac{1 - e^{-2\hat{\theta}_t \Delta}}{2\hat{\theta}_t}} v_{2,t+\Delta}, \quad (93)$$

$$\hat{\epsilon}_{t+\Delta} = e^{-\varphi \Delta} \hat{\epsilon}_t + \frac{\bar{\nu}_\epsilon}{\sigma_\delta} \sqrt{\frac{1 - e^{-2\varphi \Delta}}{2\varphi}} v_{1,t+\Delta}, \quad (94)$$

$$\hat{\lambda}_{t+\Delta} = e^{-\kappa \Delta} \hat{\lambda}_t + \frac{(\bar{\mu} - f_t) \nu_{\lambda,t}}{\sigma_f} \sqrt{\frac{1 - e^{-2\kappa \Delta}}{2\kappa}} v_{2,t+\Delta} \quad (95)$$

$$\nu_{\lambda,t+\Delta} = \nu_{\lambda,t} + \left[\sigma_\lambda^2 - 2\kappa \nu_{\lambda,t} - \left(\frac{(\bar{\mu} - f_t) \nu_{\lambda,t}}{\sigma_f} \right)^2 \right] \Delta, \quad (96)$$

where $\hat{\theta}_t = \bar{\theta} + \hat{\lambda}_t$ and $v_{1,t}$, $v_{2,t}$ are independent normally distributed random variables with mean 0 and variance 1. The time interval is $\Delta = 1/4$. 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 system above shows that, conditional on knowing the parameters of the model and the priors $(\hat{\epsilon}_0, \hat{\lambda}_0, \nu_{\lambda,0})$, the time series of the GDP growth forecast and realized GDP growth allow us to sequentially back out the time series of the posteriors $(\hat{\epsilon}_t, \hat{\lambda}_t, \nu_{\lambda,t})$ as well as the noises $(v_{1,t+\Delta}, v_{2,t+\Delta})$ for $t = \Delta, 2\Delta, 3\Delta \dots$. For the initial values, we set $\hat{\epsilon}_0$ and $\hat{\lambda}_0$ to zero, which corresponds to their long-term mean, while $\nu_{\lambda,0}$ is set to the positive root of the polynomial obtained when $\frac{d\nu_{\lambda,t}}{dt} = 0$, which defines a local steady state.

The objective is to maximize the log-likelihood function L

$$L(\Theta; u_\Delta, \dots, u_{N\Delta}) = \sum_{i=1}^N \log \left(\frac{1}{2\pi \sqrt{|\Sigma_{(i-1)\Delta}|}} \right) - \frac{1}{2} u_{i\Delta}^\top \Sigma_{i\Delta}^{-1} u_{i\Delta}, \quad (97)$$

where $\Theta \equiv (\sigma_\delta, \bar{\mu}, \sigma_f, \sigma_\epsilon, \sigma_\lambda, \bar{\theta}, \varphi, \kappa)^\top$, N is the number of observations, \top is the transpose operator, and $|\cdot|$ is the determinant operator. The 2-dimensional vector u satisfies

$$u_{t+\Delta} \equiv \begin{pmatrix} u_{1,t+\Delta} \\ u_{2,t+\Delta} \end{pmatrix} = \begin{pmatrix} \log(\delta_{t+\Delta}/\delta_t) - (f_t + \hat{\epsilon}_t - \frac{1}{2}\sigma_\delta^2) \Delta \\ f_{t+\Delta} - e^{-\hat{\theta}_t \Delta} f_t - (1 - e^{-\hat{\theta}_t \Delta}) \bar{\mu} \end{pmatrix} = \begin{pmatrix} \sigma_\delta \sqrt{\Delta} v_{1,t+\Delta} \\ \sigma_f \sqrt{\frac{1 - e^{-2\hat{\theta}_t \Delta}}{2\hat{\theta}_t}} v_{2,t+\Delta} \end{pmatrix}. \quad (98)$$

	Mean	Standard deviation	5-percentile	95-percentile
f_t	0.0256	0.0166	-0.0098	0.0553
$\hat{\theta}_t$	1.2771	0.2342	0.8269	1.5151
$\nu_{\lambda,t}$	0.2737	0.0669	0.1668	0.3624
$\hat{\epsilon}_t$	0.0009	0.0034	-0.0047	0.0062

Table 13: 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 survey error $\hat{\epsilon}_t$, the estimated mean-reversion speed $\hat{\theta}_t = \bar{\theta} + \hat{\lambda}_t$, and the uncertainty about the mean-reversion speed $\nu_{\lambda,t}$. The statistics for $\hat{\theta}_t$ and $\nu_{\lambda,t}$ are obtained using the model of learning about persistence. The statistics for $\hat{\epsilon}_t$ are obtained using the model of learning about level.

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

$$\mathbb{E}_t(u_{t+\Delta}) = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \quad \Sigma_t \equiv \text{Var}_t(u_{t+\Delta}) = \begin{pmatrix} \sigma_\delta^2 \Delta & 0 \\ 0 & \sigma_f^2 \frac{1 - e^{-2\hat{\theta}_t \Delta}}{2\hat{\theta}_t} \end{pmatrix}. \quad (99)$$

Following Vuong (1989), the likelihood-ratio test of two (non-nested) models is based on the statistic:

$$LR(M_1, M_2) \equiv n^{-1/2} \frac{\sum_{i=1}^n \log \frac{h^{M_1}(y_{t+1}|\Theta^{M_1})}{h^{M_2}(y_{t+1}|\Theta^{M_2})}}{\sqrt{\frac{1}{n} \sum_{i=1}^n \left[\log \frac{h^{M_1}(y_{t+1}|\Theta^{M_1})}{h^{M_2}(y_{t+1}|\Theta^{M_2})} \right]^2 - \left[\frac{1}{n} \sum_{i=1}^n \log \frac{h^{M_1}(y_{t+1}|\Theta^{M_1})}{h^{M_2}(y_{t+1}|\Theta^{M_2})} \right]^2}} \quad (100)$$

$$\stackrel{d}{\rightarrow} N(0, 1), \quad (101)$$

where N denotes the normal distribution, n is the number of observations, y_t is the 2-dimensional vector of observations at time t , h^{M_i} is the density associated with model M_i , and Θ^{M_i} is the vector of parameters associated with model M_i , $i \in \{1, 2\}$. If $LR > \mathcal{N}^{-1}(p)$, where \mathcal{N} is the normal CDF, the econometrician rejects the null that both models are equivalent in favor of model M_1 at the $p\%$ confidence level. If $LR < \mathcal{N}^{-1}(1-p)$, the econometrician rejects the null that both models are equivalent in favor of model M_2 at the $p\%$ confidence level.

C.1 Descriptive statistics of the state variables

Table 13 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 mean-reversion speed $\hat{\theta}_t$ varies strongly over time, fluctuating mostly between 0.8 and 1.5. Finally, uncertainty about the mean-reversion speed $\nu_{\lambda,t}$ varies also substantially, fluctuating mostly between 0.15 and 0.4. Overall, these results suggest that persistence clearly fluctuates over time.

The estimated survey error $\hat{\epsilon}_t$ 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 Appendix: Data description

Real GDP growth rate and forecast data We proxy the output process with the realized Gross Domestic Product (GDP). We compute the log growth rate of the real quarterly GDP over the period 1968Q4–2016Q4. We consider the mean real GDP growth forecast for the next quarter from the Survey of Professional Forecasters as a measure of expected real GDP growth. The first forecast observation consists of the expected real GDP growth rate for 1969Q1, as released in 1968Q4. The reported growth forecasts are annualized.

Realized and forecasted GDP data are seasonally adjusted. Real GDP data are from the Bureau of Economic Analysis and available from the Federal Reserve Bank of St. Louis, while forecast data are obtained from the Federal Reserve Bank of Philadelphia. These series can be retrieved using the following links:

- Real GDP : <https://fred.stlouisfed.org/series/GDPC1>
- Real GDP growth forecasts : <https://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/data-files/rgdp>

Consumption We use real nondurable consumption per capita and real services consumption per capita. The growth rate of consumption in quarter $t + 1$ is

$$\Delta c_{t+1} = \ln \left(\frac{C_{nd,t+1} + C_{s,t+1}}{C_{nd,t} + C_{s,t}} \right) \quad (102)$$

where $C_{nd,t}$ denotes real nondurables and $C_{s,t}$ denotes real services consumption per capita. The series are at the quarterly frequency and seasonally adjusted. The data are from the Bureau of Economic Analysis but can be retrieved from the Federal Reserve Bank of St. Louis, using the following links:

- Real nondurables consumption : <https://fred.stlouisfed.org/series/A796RX0Q048SBEA>
- Real services consumption : <https://fred.stlouisfed.org/series/A797RX0Q048SBEA>

Real risk-free rate We compute the real risk-free rate as the three month nominal yield adjusted by the expected inflation rate over the next three months. As in Beeler and Campbell (2012), we first take the nominal yield on a three month Treasury bill $y_{3,t}$ in month t and subtract the three month inflation $\pi_{t,t+3}$ from period t to $t + 3$ to form a measure of the ex post real three month interest rate. This is the dependent variable in the predictive regression below:

$$y_{3,t} - \pi_{t,t+3} = \beta_0 + \beta_1 y_{3,t} + \beta_2 \pi_{t-12,t} + \epsilon_{t+3} \quad (103)$$

where the independent variables are the inflation over the previous year $\pi_{t-12,t}$ divided by four and the three month nominal yield $y_{3,t}$. The predicted value for the regression in month t gives the ex ante risk free rate for month $t + 1$. Our quarterly measure of the real risk free rate is the annualized value at the beginning of the quarter, which we denote by $r_{f,t}$.

The nominal yield is the three-month Treasury Bill secondary market rate, which we continuously compound as follows: $y_{3,t} = \ln(1 + y_{3,t,obs}/100)/4$. Inflation is computed as the monthly log growth rate of the Consumer Price Index (CPI) from the Bureau of Labor Statistics, which is seasonally adjusted. Both series are at the monthly frequency. The data can be retrieved from the Federal Reserve Bank of St. Louis, using the following links:

- Three-month Treasury Bill rate : <https://fred.stlouisfed.org/series/TB3MS>
- Consumer Price Index : <https://fred.stlouisfed.org/series/CPIAUCSL>

Stock prices and dividends We compute the stock market price index and extract the dividends using CRSP data. The stock price index in month t is constructed as:

$$P_t = P_{t-1} (1 + R_{noD,t}) \quad (104)$$

where $R_{x,t}$ denotes the return of a value-weighted index excluding distributions in month t .

The monthly dividend is given by

$$D_t = P_t \left(\frac{1 + R_{D,t}}{1 + R_{noD,t}} - 1 \right) \quad (105)$$

where $R_{d,t}$ denotes the return of a value-weighted index including distributions in month t .

Quarterly dividend is the sum of dividends within a quarter, which are not seasonally adjusted. We then calculate the log quarter over quarter growth rate in dividends. Dividend growth is converted from nominal to real terms using the CPI. We thus subtract log inflation to form real growth rates.

The price-dividend ratio is the price in the last month of the quarter divided by the sum of dividends paid in the last twelve months. We use the series of the value-weighted index including distributions (VWRETD) and the value-weighted index excluding distributions (VWRETX) from CRSP, which cover NYSE, Amex, and Nasdaq data.

Stock return volatility Stock return volatility is the volatility of real stock returns computed at the quarterly frequency. We first fit an AR(1) process on the quarterly log return of the stock price index and take the residuals. We then obtain the conditional volatility estimate, denoted by $Vol_{R,t}$, with a GARCH(1,1). We finally annualize the series. We use the quarterly value-weighted market price index excluding distributions from CRSP.

Realized and expected excess stock returns We first compute the quarterly real excess stock returns by subtracting the real risk-free rate from real returns. The real return is the log return of the market price index deflated by the CPI, whereas the real risk-free rate is constructed as in section D. The realized real excess stock return $R_{X,t}$ in quarter t is thus given by:

$$R_{X,t} = \ln \left(\frac{P_t}{P_{t-1}} \right) - \pi_{t-1,t} - r_{f,t} \quad (106)$$

where quarterly inflation $\pi_{t-1,t}$ is the log growth rate of the CPI in the final month of the current quarter over the final month in the previous quarter. We use the Consumer Price Index from the Bureau of Labor Statistics, which is seasonally adjusted.

To compute the expected excess returns, we regress the returns $R_{X,t}$ on the lagged dividend yield (measured at time $t - 1$), the lagged default premium (Baa yield minus ten-year government bond yield), and stock return volatility. The estimated expected real excess return in quarter t is the fitted value at time t , $\widehat{R}_{X,t}$. We then annualize the series. This approach follows Fama and French (1989)'s measurement procedure for estimating expected returns.

The Sharpe ratio is given by the expected real excess stock returns $\widehat{R}_{X,t}$ divided by the volatility of real stock returns $Vol_{R,t}$. Both series are at the quarterly frequency.

Regarding the data, the market price index is the value-weighted index excluding distributions from CRSP. The dividend yield is the inverse of the price-dividend ratio defined in Section D. The default premium is defined as the Moody's seasoned Baa corporate bond yield relative to yield on 10-year Treasury constant maturity, as available from the Federal Reserve Bank of St. Louis (<https://fred.stlouisfed.org/series/BAA10YM>).