Learning and Consumption Dynamics

Daniel Andrei* Steffen Hitzemann† Alexandre Jeanneret‡

August 31, 2016

Abstract

In a production economy, uncertainty and learning about persistence generates counter-cyclical consumption volatility. When the persistence of productivity is unobservable, consumption responds differently to long-run productivity shocks, depending on the state of the economy. In bad economic times, a negative shock prompts agents to extrapolate that productivity becomes more persistent, which reinforces consumption's response to bad news. The opposite extrapolation occurs in good times, when productivity becomes less persistent after a negative shock, partially offsetting the bad news. This asymmetric response to productivity shocks amplifies the volatility of consumption in bad times, but attenuates it in good times. In contrast, other types of learning generate constant or procyclical volatility of consumption, at odds with empirical evidence.

*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; danielandrei.info.

†Fisher College of Business, The Ohio State University, 2100 Neil Avenue, Columbus, OH 43210, USA; +1 (614) 688-2181; hitzemann.6@osu.edu; steffenhitzemann.com.

‡Department of Finance, HEC Montréal, 3000 Ch. de la Côte-Sainte-Catherine, Montréal, Québec, CANADA; +1 (514) 589-6699; alexandre.jeanneret@hec.ca; alexandrejeanneret.com.

1 Introduction

Standard asset pricing theory predicts that assets that covary positively with consumption should earn a positive risk premium. This compensation for risk depends on the level of consumption uncertainty, which is often assumed to be constant. However, empirical evidence suggests that consumption volatility varies over time and, in particular, becomes higher during recessions.[1] Countercyclical volatility of consumption should then provide an additional risk premium, as investors dislike assets that pay poorly not only when their marginal utility of consumption is low but also when consumption becomes more uncertain.

Based on this intuition, consumption volatility risk has recently attracted much interest among economists and is now viewed as an important ingredient for asset pricing. For instance, Bansal and Yaron (2004) show that fluctuations in consumption volatility generate a time-varying and countercyclical risk premium.[2] Moreover, this channel is important for understanding the negative relation between returns and return volatility, also known as the volatility feedback effect (French, Schwert, and Stambaugh, 1987; Campbell and Hentschel, 1992). In addition, the exposure to fluctuating consumption volatility explains the cross section of stock returns (Boguth and Kuehn, 2013; Bansal, Kiku, Shaliastovich, and Yaron, 2014; Tédongap, 2015). It is therefore crucial to understand the fundamental mechanism behind consumption volatility risk. Our paper takes one step in this direction.

We show that countercyclical consumption volatility can arise endogenously in a general equilibrium production economy with incomplete information. The economy is populated by a representative agent with Epstein and Zin (1989) lifetime utility function and a preference for early resolution of uncertainty (Bansal and Yaron, 2004). In this economy, the productivity is mean-reverting and observable, but the agent has incomplete information about its persistence. That is, the agent observes the level of productivity but does not know its mean-reverting speed. We derive the equilibrium implications of this learning exercise and show that it generates time-varying and countercyclical consumption volatility.

The mechanism behind this prediction becomes clear when we disentangle good economic times (i.e., times during which the productivity is above its long-run mean) from bad economic times. In good times, a positive productivity shock has two opposing effects on consumption. First, higher productivity means better investment opportunities, which induces the agent to invest more and thus to reduce today's consumption.[3] Second, the agent

[1] See, for example, Kandel and Stambaugh (1990), Kim and Nelson (1999), Bansal and Yaron (2004), Bansal, Khatchatrian, and Yaron (2005), Boguth and Kuehn (2013), and Bekaert and Engstrom (2015).

[2] Bansal and Yaron (2004) show that both the persistence in the consumption growth rate and in the volatility of consumption growth are necessary to replicate key asset pricing moments. See also Bansal, Kiku, and Yaron (2012) and Beeler and Campbell (2012).

[3] We assume that the parameters of relative risk aversion (RRA) and of intertemporal elasticity of sub-

uses this new information to update her estimate of the persistence parameter; the positive productivity shock then becomes bad news, as the agent now extrapolates that productivity is more persistent, a feature which she dislikes given the preference for early resolution of uncertainty. These two forces partially offset each other, thereby dampening fluctuations in consumption. As a consequence, we observe a low level of consumption volatility in good times. In contrast, a positive productivity shock in bad times is always perceived as good news: the agent becomes wealthier and, furthermore, extrapolates that the economy is recovering faster, which implies less persistent productivity. Both forces influence consumption in the same direction, thus yielding a high level of consumption volatility in bad times. Given this, learning about persistence creates an asymmetry in consumption volatility.

This countercyclicality does not arise in a model where the agent learns about the *level* of productivity (Ai, 2010) instead of its persistence. To demonstrate this, we build a model which nests both alternatives and thus allows us to show that only uncertainty about the degree of persistence is able to generate countercyclical consumption volatility. Instead, learning about the level of productivity growth helps better match asset pricing moments, in comparison to an economy with perfect information, but implies no variation or even procyclical consumption volatility, at odds with the data. This key difference arises because uncertainty about the level of productivity impacts the sensitivity of consumption to short-run shocks, whereas uncertainty about its persistence impacts the sensitivity of consumption to long-run shocks. This also emphasizes the importance of long-run productivity shocks for macroeconomics and finance considerations (Favilukis and Lin, 2013; Croce, 2014).

The properties of consumption volatility that we generate endogenously from the learning process have several critical implications. First, the persistence in productivity translates into a persistent conditional consumption volatility in equilibrium. Because the volatility of asset returns is a linear affine function of the conditional variance of consumption growth (Bansal and Yaron, 2004), the model would help explain the autoregressive property of market return volatility, as demonstrated by the (G)ARCH processes (Engle, 1982; Bollerslev, 1986). Second, this relation implies that both market return volatility and the equity risk premium rise during recessions, consistent with empirical evidence (Fama and French, 1989; Schwert, 1989).[4] Finally, time-variation in consumption volatility also plays a key role in the pricing of corporate claims and in capital structure decisions, which thereby contribute to

stitution (IES) are higher than one. With an IES higher than one, the intertemporal substitution effect dominates the income effect, which implies that the agent consumes a lower proportion of her wealth when the expected return on the technology increases (Ai, 2010).

[4]This would also help explain why expected market volatility rises when economic conditions deteriorates, as highlighted with the CBOEs VIX index. The rationale is that the VIX is positively related with the market return volatility, and thus with consumption volatility, in addition to containing a variance risk premium (Carr and Wu, 2009; Bollerslev, Gibson, and Zhou, 2011).

solving the credit spread puzzle (Bhamra, Kuehn, and Strebulaev, 2010; Chen, 2010). Our theory thus provides a fundamental explanation for the dynamics of consumption, which has been shown to help explain several salient observations in finance.

This paper is related to several strands of literature. First, our approach builds on the growing literature analyzing equilibrium conditions and asset-price properties in production-based economies. The most closely related paper to ours is Ai (2010), who describes the effect of learning on the unconditional moments of the wealth-consumption ratio and of the return on aggregate wealth. Our paper focuses on conditional moments and shows that fluctuations in consumption volatility arise endogenously from learning about persistence of productivity, rather than about its level. Hirshleifer, Li, and Yu (2015) analyze extrapolative expectations about productivity growth. They focus on the unconditional moments of consumption and show that consumption growth becomes more predictable and volatile if one assumes a greater extrapolative bias. The contribution of our paper is to show that agents form extrapolative expectations rationally when updating their beliefs and that this behavior yields time-variation in consumption volatility.[5]

Second, our analysis of the asset pricing implications builds on the literature analyzing the role of consumption volatility as an exogenous source of risk. Bansal et al. (2005) find that time-varying consumption volatility predicts aggregate asset valuation ratios, which is then confirmed by Lettau, Ludvigson, and Wachter (2008). Bansal and Yaron (2004) demonstrate that both long-run risk and fluctuations in consumption volatility help match key asset pricing moments with a reasonable parameterization of recursive preferences. Boguth and Kuehn (2013) and Tédongap (2015) suggest that the volatility of consumption growth contributes to explaining individual expected stock returns, while Bekaert and Engstrom (2015) conclude that conditional volatility, skewness, and kurtosis of consumption growth matter for asset pricing.

Finally, our paper also relates to Kaltenbrunner and Lochstoer (2010) and Croce (2014), who obtain long-run consumption risk endogenously in a production economy. Although our model generates an endogenous persistence in the consumption growth rate, which arises from learning, we focus instead on the dynamics of consumption volatility. Furthermore, and in contrast to these papers, our results obtain without any investment friction.

The remainder of the article is organized as follows. Section 2 provides empirical evidence

[5]The paper also complements a recent set of studies highlighting the asset pricing implications of learning about consumption dynamics (Johannes, Lochstoer, and Mou, 2016; Collin-Dufresne, Johannes, and Lochstoer, 2016). Another related study is Brennan and Xia (2001), which shows that learning on the level of dividend growth rate can explain high levels of stock price volatility and equity premium. The main difference with these articles is that we explore learning about the dynamics of productivity and its endogenous effect on consumption.

on time-varying and countercyclical consumption volatility. Section 3 introduces the model, while Section 4 presents the theoretical predictions. Section 5 outlines the calibration of the model parameters and discusses the results of our numerical analysis. The final section offers some concluding remarks.

2 Empirical Properties of Consumption Volatility

We first discuss the empirical properties of consumption volatility in the U.S. Our analysis uses data on quarterly per capita real consumption expenditures on nondurable goods and services from the Bureau of Economic Analysis. We examine the 1950Q1-2015Q4 period, based on the observation that U.S. consumption behaved quite specifically in the first years following World War II (Yogo, 2006; Lettau et al., 2008; Boguth and Kuehn, 2013). For comparison, we also report results using annual consumption data from 1929 until 2015.

We first display the conditional volatility of consumption growth, which we estimate using a GARCH(1,1) model.[6] Figure 1 (top panel) shows that conditional consumption volatility presents persistence and fluctuates over time at the business cycle frequency. Furthermore, consumption volatility appears to be countercyclical, as it tends to increase in bad times and to decrease in good times.

We demonstrate this pattern in two different ways. First, consumption volatility tends to be higher during official NBER recessions than during normal times. The average conditional annualized volatility is respectively 1.32% and 1.07%, as reported in Table 1 (Panel A).

Second, we compare the levels of consumption volatility across different states of the economy estimated directly from the consumption dynamics. We thus exploit business cycle information that likely differs from the NBER classification of recession/expansion periods. Specifically, we estimate a Markov regime-switching model with two regimes on U.S. consumption growth. The transition probability matrix is obtained by maximum likelihood using the Hamilton (1989)'s approach. The middle and bottom panels of Figure 1 indicate that consumption volatility increases (decreases) with the probability of being in bad (good) economic times. We find that the average consumption volatility is respectively 1.57% and 0.66% during the bad and the good state (see Table 1, Panel B). The conclusion of our analysis remains qualitatively similar, although quantitatively amplified, if we instead consider annual consumption data over the 1929-2015 period (see Table 1).

Overall, these results provide empirical support that consumption volatility is countercyclical and that it presents an asymmetry that strongly depends on the state of the economy. We now propose a theory that will help us understand these empirical observations.

[6]Data are first demeaned to remove MA(1) variation in the conditional mean of consumption growth.

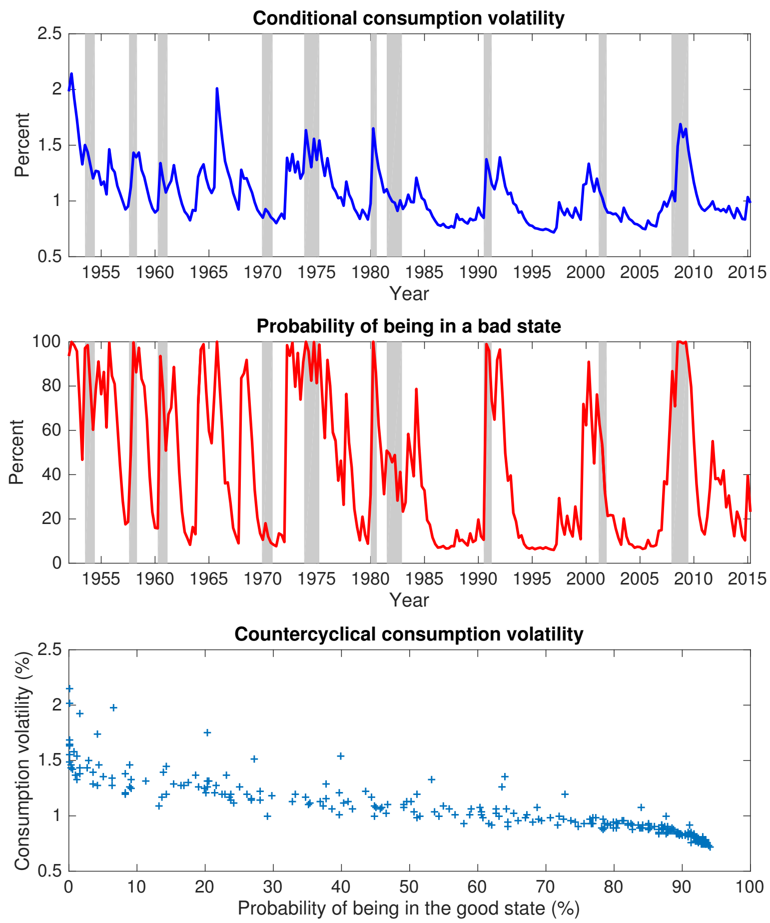

Figure 1: **Conditional variation in consumption volatility.** The top panel displays the time-variation in conditional consumption volatility in the U.S., estimated with a GARCH(1,1). The middle panel reports the probability of being in bad economic times, estimated with a two-state Markov-regime switching model (Hamilton, 1989). The bottom panel presents the relation between conditional consumption volatility and the probability of being in good economic times. We use data on real non-durables goods plus service consumption expenditures over the period 1950Q1-2015Q4 from the Bureau of Economic Analysis.

3 Model

3.1 Economic Environment

Consider an economy populated by a representative agent, who derives utility from consumption. The agent has stochastic differential utility with subjective time preference rate β, relative risk aversion γ, and elasticity of intertemporal substitution ψ. The indirect utility

Table 1: **Conditional moments of consumption volatility** This table reports the conditional moments of consumption volatility in the U.S. Panel A uses official NBER dates to characterize periods of expansions and recessions, whereas Panel B considers a two-state Markov-regime switching model based on Hamilton (1989). For each method, we consider quarterly consumption data over the period 1950Q1-2015Q4 and annual consumption data over the period 1929-2015. We measure consumption using data on real non-durables goods plus service consumption expenditures from the Bureau of Economic Analysis.

	Recession (1)	Expansion (2)	(1) - (2)
Panel A : NBER classification			
Quarterly data (1950Q1-2015Q4)			
Consumption volatility (%)	1.323	1.069	0.254
Frequency	0.164	0.836	
Annual data (1929-2015)			
Consumption volatility (%)	3.349	1.592	1.757
Frequency	0.164	0.836	
Panel B : Markov-regime switching model			
Quarterly data (1950Q1-2015Q4)			
Consumption volatility (%)	1.566	0.660	0.906
Frequency	0.443	0.557	
Annual data (1929-2015)			
Consumption volatility (%)	4.313	1.208	3.105
Frequency	0.145	0.855	

function is given by

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

where the aggregator f writes (Campbell, Chacko, Rodriguez, and Viceira, 2004):

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

In this economy, the capital, K_t, and the productivity, a_t, evolve according to:

$$dK_t = K_t \left\{ [a_t + (1-\theta)\pi]dt + \sigma_K dB_{K,t} \right\} - C_t dt \tag{3}$$

$$da_t = [\theta\lambda + (1-\theta)\bar{\lambda}](\bar{a} - a_t)dt + \sigma_a dB_{a,t}, \tag{4}$$

where we assume, for simplicity, that the two Brownians in (3)-(4) are uncorrelated. Considering a non-zero correlation between these two Brownians would make the notation heavier without changing the main message of the paper.

The agent observes the system (3)-(4) but has only limited information about the constants π and λ. All other parameters are known. If the parameter θ is either 0 or 1, only one of two constants remains unobservable. In the case $\theta = 0$, the agent knows the mean-reverting speed of the productivity, $\bar{\lambda}$, but does not know its overall level, $a_t + \pi$ (we call this case *learning about level*). By contrast, the case $\theta = 1$ indicates that the agent perfectly observes the level of productivity, a_t, but does not know its mean-reversion speed, λ (we call this case *learning about persistence*).[7] As such, the parameter θ defines the learning exercise of the representative agent.

Although the main focus of our paper is on the case of learning about persistence ($\theta = 1$), we keep the parameter θ for expositional purposes. This allows us to clearly distinguish between the different equilibrium implications of learning about the level versus learning about the persistence of productivity.

3.2 Learning

We now explain how the representative agent updates her beliefs about each of the two parameters, using the information provided by the different types of shocks. The agent starts with the following priors about π and λ:

$$\pi \sim \mathcal{N}(0, \nu_{\pi,0}) \tag{5}$$

$$\lambda \sim \mathcal{N}(\bar{\lambda}, \nu_{\lambda,0}). \tag{6}$$

There is no correlation between these two priors (i.e., $\nu_{\pi\lambda,0} = 0$). The value of θ affects neither the agent's prior about the long-term level of productivity, which always equals $\bar{a}$, nor the agent's prior about the mean-reverting speed of productivity, which always equals

[7]In fact, within our model, the parameter θ can take any value between 0 and 1, in which case the agent learns simultaneously about level *and* persistence. Intuitively, if the parameter θ approaches 0, the agent is more concerned with learning about the level of productivity, whereas she rather focuses on learning about persistence when the parameter θ approaches 1. A worthwhile research question is to use macroeconomic and asset price data to structurally estimate the parameter θ.

$\bar{\lambda}$. That is, starting with a θ of zero or one generates exactly the same priors; the *only* difference is that in one case the agent learns about the level, whereas in the other case the agent learns about persistence.

Let us denote by $\widehat{\pi}_t \equiv \mathbb{E}[\pi|\mathcal{F}_t]$ the estimated parameter π and its posterior variance by $\nu_{\pi,t} \equiv \mathbb{E}[(\pi-\widehat{\pi}_t)^2|\mathcal{F}_t]$. Similarly, let us denote by $\widehat{\lambda}_t \equiv \mathbb{E}[\lambda|\mathcal{F}_t]$ the estimated parameter λ and its posterior variance by $\nu_{\lambda,t} \equiv \mathbb{E}[(\lambda - \widehat{\lambda}_t)^2|\mathcal{F}_t]$. The estimates and the posterior variances are such that

$$\pi \sim N(\widehat{\pi}_t, \nu_{\pi,t}) \tag{7}$$

$$\lambda \sim N(\widehat{\lambda}_t, \nu_{\lambda,t}), \tag{8}$$

where $N(m, v)$ denotes the Normal distribution with mean m and variance v. Henceforth, we refer to either of the two estimated predictive coefficients $\widehat{\pi}_t$ and $\widehat{\lambda}_t$ as the *filter* and to either of the two posterior variances $\nu_{\pi,t}$ and $\nu_{\lambda,t}$ as the *uncertainty*. Note that learning is standard in our setup (Liptser and Shiryayev, 1977) and the reader can refer to Appendix A for the technical details.

The filters evolve according to

$$\begin{bmatrix} d\widehat{\pi} \\ d\widehat{\lambda} \end{bmatrix} = \begin{bmatrix} \frac{(1-\theta)\nu_{\pi,t}}{\sigma_K} & 0 \\ 0 & \frac{\theta(\bar{a}-a_t)\nu_{\lambda,t}}{\sigma_a} \end{bmatrix} \begin{bmatrix} d\widehat{B}_{K,t} \\ d\widehat{B}_{a,t} \end{bmatrix}, \tag{9}$$

where $d\widehat{B}_{K,t}$ and $d\widehat{B}_{a,t}$ are independent Brownian motions coming from the filtration of the agent. We define these Brownian motions in Appendix A. One can observe from (9) that the agent does not update λ if $\theta = 0$ and thus sticks with the prior $\bar{\lambda}$ as the mean-reverting speed. If $\theta = 1$, the agent does not learn about π and thus maintains the prior $\bar{a}$ as the long-term productivity in the economy.

The agent's estimate of π and the capital K are perfectly and positively correlated. As a result, the agent's expectation formation is "extrapolative", as in Brennan (1998): when capital increases, the agent revises her estimate upwards; when capital decreases, she revises her estimate downward. This extrapolative expectation formation is also valid when learning about λ, but the effect now depends on the state of the economy, which is characterized by the difference $\bar{a}-a_t$. In good times, when $\bar{a}-a_t < 0$, any *positive* productivity shock decreases the agent's estimate of λ, whereas in bad times, when $\bar{a} - a_t > 0$, any *negative* productivity shock decreases the agent's estimate of λ. Both situations (i.e., positive shocks in good times or negative shocks in bad times) induce the agent to extrapolate that productivity becomes more persistent. As we will show in the next section, this extrapolative expectation formation plays a critical role in the agent's optimal consumption decision.[8]

The Bayesian uncertainties about π and λ evolve according to

$$d\nu_{\pi,t} = -\frac{(1-\theta)^2\nu_{\pi,t}^2}{\sigma_K^2}dt \tag{10}$$

$$d\nu_{\lambda,t} = -\frac{\theta^2(\bar{a}-a_t)^2\nu_{\lambda,t}^2}{\sigma_a^2}dt. \tag{11}$$

Because the agent is learning about constants, both uncertainties converge to zero. It is easy to generate positive steady-state uncertainties by assuming that learning is regenerated, i.e., that π and λ are not constants but rather move over time. Yet this would unnecessarily complicate the setup without affecting our qualitative implications.[9]

The uncertainty about π has a closed-form solution and is a deterministic function of time:

$$\nu_{\pi,t} = \frac{1}{\frac{(1-\theta)^2}{\sigma_K^2}t + \frac{1}{\nu_{\pi,0}}}, \tag{12}$$

where the term $(1-\theta)^2/\sigma_K^2$ represents the speed of learning, which decreases with θ and σ_K. If we consider time $t = 0$, we obtain the prior $\nu_{\pi,0}$. Although there is no straightforward solution for the uncertainty about λ because of the $(\bar{a}-a_t)$ term in (11), it is clear, however, that this uncertainty also decreases with time. Notably, the convergence to zero is quicker when a_t is away from $\bar{a}$, as news become more informative. The agent thus learns faster in either very good or very bad times, in comparison to productivity being close to its long-term mean. To see it differently, new observations that are further away from a regression line are more informative about the slope coefficient.

To summarize, the dynamics of the capital and other state variables generated by this learning exercise are given respectively by

$$dK_t = K_t\left\{[a_t + (1-\theta)\widehat{\pi}_t]dt + \sigma_K d\widehat{B}_{K,t}\right\} - C_t dt \tag{13}$$

$$da_t = [\theta\widehat{\lambda} + (1-\theta)\bar{\lambda}](\bar{a}-a_t)dt + \sigma_a d\widehat{B}_{a,t} \tag{14}$$

[8] The extrapolative nature of learning is different from the "extrapolation bias," a pervasive phenomenon in human judgement and decisions. Extrapolation bias refers to the tendency to overweight recent events when making decisions about the future. In our case, the agent does not overweight recent events, but applies standard Bayesian updating rules. See Hirshleifer et al. (2015) for a rigorous analysis of extrapolation bias in a production economy.

[9] Collin-Dufresne et al. (2016) show that parameter learning—even about a constant—is important when the representative agent has a preference for early resolution of uncertainty, which will be the case in our setup.

and Equations (9), (10), and (11).

3.3 Equilibrium

The equilibrium is standard and thus technical details are relegated to Appendix B. In this economy, solving for equilibrium involves writing the HJB-equation for problem (1) as

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

with the differential operator $\mathcal{L}J$ following from Itô's lemma, and defined in Equation (52) of Appendix B. We then obtain from (15) the first order condition for consumption:

$$C = \beta^{\psi} J_K^{-\psi} [(1-\gamma)J]^{\frac{1-\gamma\psi}{1-\gamma}}, \tag{16}$$

which we insert into the HJB-equation (52) and guess the following form for the value function:

$$J(K, a, \widehat{\pi}, \widehat{\lambda}, \nu_{\pi}, \nu_{\lambda}) = \frac{K^{1-\gamma}}{1-\gamma} \beta^{\phi} x(a, \widehat{\pi}, \widehat{\lambda}, \nu_{\pi}, \nu_{\lambda})^{-\frac{1-\gamma}{1-\psi}}, \tag{17}$$

where ϕ is a constant defined as $\phi \equiv \frac{1-\gamma}{1-1/\psi}$. This constant equals one in the CRRA case and is lower than one when $\gamma > 1 > 1/\psi$. The function $x(a, \widehat{\pi}, \widehat{\lambda}, \nu_{\pi}, \nu_{\lambda})$ represents the wealth-consumption ratio in the economy:

$$x(a, \widehat{\pi}, \widehat{\lambda}, \nu_{\pi}, \nu_{\lambda}) = \frac{K}{C}. \tag{18}$$

The HJB (52) can then be written in terms of the wealth-consumption ratio x to obtain the partial differential equation (56) in Appendix B. We solve this equation in two special cases: $\theta = 0$ (stated in Equation (59) of Appendix B.1) and $\theta = 1$ (stated in Equation (60) of Appendix B.2). In each one fo these two cases, the partial differential equation is solved using Chebyshev polynomials (Judd, 1998).

4 Theoretical Predictions

This section presents and discusses the main predictions of the paper. We show that productivity shocks prompt the agent to extrapolate about the perceived persistence of productivity. This endogenous response reduces consumption volatility in good times but increases it in bad times. Hence, consumption volatility becomes countercyclical, consistent with the empirical evidence of Section 2. We contrast this result with the case of learning about level, which can potentially yield procyclical volatility of consumption.

4.1 Consumption Dynamics

To understand how learning affects consumption volatility, it is useful to first analyze the endogenous dynamics of consumption in our economy. Starting from (18), the consumption evolves according to:

$$\frac{dC_t}{C_t} = \mu_t^C dt + \begin{bmatrix} \sigma_{K,t}^C & \sigma_{a,t}^C \end{bmatrix} \begin{bmatrix} d\widehat{B}_{K,t} \\ d\widehat{B}_{a,t} \end{bmatrix}, \tag{19}$$

with the diffusion terms $\sigma_{K,t}^C$ and $\sigma_{a,t}^C$ satisfying

$$\sigma_{K,t}^C \equiv \sigma_K - (1-\theta)\frac{\nu_{\pi,t}}{\sigma_K}\frac{x_{\widehat{\pi}}}{x} \tag{20}$$

$$\sigma_{a,t}^C \equiv -\sigma_a \frac{x_a}{x} - \theta\frac{(\bar{a} - a_t)\nu_{\lambda,t}}{\sigma_a}\frac{x_{\widehat{\lambda}}}{x}, \tag{21}$$

where x_y denotes the partial derivative of the wealth-consumption ratio with respect to the state variable y.

Our focus is on the impact of learning on the volatility of consumption.[10] The two diffusion terms in (20) and (21) are primarily driven by three state variables: the level of productivity a_t, the uncertainty about the level of productivity $\nu_{\pi,t}$, and the uncertainty about the mean-reversion speed $\nu_{\lambda,t}$. Based on these diffusion terms, our model indicates that these two types of learning affect the volatility of consumption differently. On the one hand, if the agent learns about the level of productivity ($\theta = 0$), (20) is affected by uncertainty about π, while (21) is not. On the other hand, if the agent learns about persistence ($\theta = 1$), only (21) depends on the uncertainty about λ. Thus, an important conclusion is that uncertainty about the level of productivity affects consumption's response to short-run shocks $d\widehat{B}_K$, whereas uncertainty about the persistence of productivity affects consumption's response to long-run shocks $d\widehat{B}_a$.

Note that these relations depend on the wealth-consumption ratio terms x_a/x, $x_{\widehat{\pi}}/x$, and $x_{\widehat{\lambda}}/x$. We now discuss the signs of these terms, which will be useful to understand the overall impact of the state variables on the volatility of consumption.

When the agent has relative risk aversion γ and intertemporal elasticity of substitution ψ higher than one, the wealth-consumption increases with good news about either the

[10] See Appendix B, Equation (58), for the drift of consumption, which has been analyzed in detail by Kaltenbrunner and Lochstoer (2010) in production economies.

transitory component or the long-run level of productivity. Intuitively, this prediction arises because the intertemporal substitution effect dominates the income effect, which implies that the agent consumes a lower proportion of her wealth as she becomes wealthier (Ai, 2010). Hence, we have that $x_a/x > 0$ and $x_{\widehat{\pi}}/x > 0$. In addition, the agent has preference for early resolution of uncertainty and thus dislikes long-run risk (Bansal and Yaron, 2004), which is captured by a lower persistence parameter $\widehat{\lambda}$. The wealth-consumption is then increasing in $\widehat{\lambda}$, which yields $x_{\widehat{\lambda}}/x > 0$.[11] To summarize, preferences satisfying $\gamma > 1 > 1/\psi$ imply that

$$x_a/x > 0, \quad x_{\widehat{\pi}}/x > 0, \quad x_{\widehat{\lambda}}/x > 0. \tag{22}$$

We now analyze each form of learning in detail and compare their implications for consumption volatility.

4.2 Learning about level

When the representative agent learns about the level of productivity, uncertainty about π affects consumption's response to short-run shocks and decreases the diffusion term $\sigma^C_{K,t}$, as it can be seen from (20). This type of learning decreases the volatility of consumption, as noted by Ai (2010). Intuitively, a positive capital shock, $d\widehat{B}_{K,t} > 0$, makes the agent wealthier, which allows her to increase consumption. At the same time, the agent uses this innovation to update her estimate about the level of productivity. She now believes that the expected return on the technology becomes higher, which encourages her to invest more in order to increase future consumption. As these forces offset each other, the presence of learning therefore creates an intertemporal substitution effect that reduces the response of consumption to capital shocks when the IES is higher than one.[12] As highlighted in Ai (2010), learning about the level of productivity therefore decreases the consumption's response to short-run shocks and hence its volatility.

[11] These conclusions can be reached by starting directly from the specification of the value function (17): when $\gamma > 1 > 1/\psi$, the value function is increasing in the wealth-consumption ratio. Furthermore, we know that the agent prefers higher a and higher $\widehat{\pi}$ (due to non-satiation, expected lifetime utility must rise as investment opportunities improve), i.e. $J_a > 0$ and $J_{\widehat{\pi}} > 0$. We also know that the agent prefers early resolution of uncertainty, i.e. $J_{\widehat{\lambda}} > 0$. Finally, we know that the agent dislikes uncertainty, i.e. $J_{\nu_\pi} < 0$ and $J_{\nu_\lambda} < 0$. We then obtain (22).

[12] Note that the second effect can actually become stronger than the first effect, leading to a decrease in consumption after a positive capital shock.

4.3 Learning about persistence

Learning about persistence impacts consumption's response to long-run shocks. This impact arises asymmetrically through the presence of the difference $\bar{a} - a_t$ in $\sigma^C_{a,t}$, as indicated in (21). Current economic conditions thus determine whether learning attenuates or rather amplifies the direct effect of productivity shocks on consumption. To understand the role of this asymmetry, we examine how consumption responds to a positive productivity shock, $d\widehat{B}_{a,t} > 0$, by comparing good and bad economic times.

Productivity a_t above long-term mean, $a_t > \bar{a}$. In good economic times, a positive productivity shock increases a_t but decreases $\widehat{\lambda}_t$. This second effect arises because learning is extrapolative—the agent believes that good times will last longer. The increase in productivity increases the wealth-consumption ratio, whereas the decrease in the filter $\widehat{\lambda}$ reduces this ratio. Hence, consumption decreases with the first term through a, but increases with the second term through $\widehat{\lambda}$, as illustrated below:

$$\sigma^C_{a,t} = \underbrace{-\sigma_a \frac{x_a}{x}}_{\text{Increase in } a\ (-)} \underbrace{-\theta \frac{(\bar{a} - a_t)\nu_{\lambda,t}}{\sigma_a} \frac{x_\lambda}{x}}_{\text{Decrease in } \widehat{\lambda}\ (+)} \tag{23}$$

In other words, a positive productivity shock in good times is good news because of higher productivity (increase in a), but is bad news because it also implies more persistence (decrease in $\widehat{\lambda}$). Because these two forces have opposing effects on consumption, fluctuations in consumption are *dampened* in good times, resulting in relatively low consumption volatility.

Productivity a_t below long-term mean, $a_t < \bar{a}$. In bad economic times, a positive productivity shock increases both a_t and $\widehat{\lambda}_t$. The increase in the filter arises again because of learning—with a positive shock, the agent believes that the economy is recovering faster. Hence, the increase in both productivity a and the filter $\widehat{\lambda}$ lead to a higher wealth-consumption ratio and to lower consumption, as we illustrate below

$$\sigma^C_{a,t} = \underbrace{-\sigma_a \frac{x_a}{x}}_{\text{Increase in } a\ (-)} \underbrace{-\theta \frac{(\bar{a} - a_t)\nu_{\lambda,t}}{\sigma_a} \frac{x_\lambda}{x}}_{\text{Increase in } \widehat{\lambda}\ (-)} \tag{24}$$

Hence, a positive productivity shock in bad times is good news because of higher productivity (increase in a), and is also good news because it implies less persistence (increase in $\widehat{\lambda}$). Thus, consumption goes down in both cases. These shocks therefore work together to *increase* the volatility of consumption in bad times.[13]

Overall, news about productivity, or long-run news, have two effects on consumption. The first effect is that higher productivity stimulates investment and decreases consumption. It increases the opportunity cost of consumption and, when the substitution effect dominates the income effect, the agent decides to invest more and to consume less today. Importantly, the sign of this effect does not vary with the state of the economy (i.e., whether the economy is in good or bad times), but only depends on the parameters of the utility function.

The second effect depends on the state of the economy, thereby introducing an asymmetry in consumption's response to productivity shocks. This arises because the agent prefers early resolution of uncertainty and thus dislikes persistence in the productivity. If the economy is in good times, a positive productivity shock signals *more persistence*—the agent then chooses to decrease investment and to consume more; if the economy is in bad times, a positive productivity shock signals *less persistence*—the agent then chooses to increase investment and to consume less.

These two effects have opposing signs in good times, but the same sign in bad times. Because consumption becomes more sensitive to long-run productivity shocks in bad times, this asymmetry generates countercyclical consumption volatility.

5 Numerical Illustration

We now offer a quantitative analysis of the role of learning in consumption volatility. We illustrate first the implications of learning about persistence. Then, we compare our findings with the case of learning about the level of productivity ($\theta = 0$).

5.1 Calibration

In our baseline calibration, we set preferences such that $\beta = 0.014$, $\gamma = 3$, and $\psi = 1.5$. The volatility of the aggregate wealth is set to $\sigma_K = 0.05$. The productivity process assumes that $\bar{a} = 0.035$ and $\sigma_a = 0.005$, as in Ai (2010). We specify the following prior distribution for the mean-reverting parameter λ: $\bar{\lambda} = 0.05$ and $\nu_{\lambda,0} = 0.000625$ (in comparison Ai (2010) chooses a mean-reverting parameter of 0.027, which is within our specified prior distribution). Finally, regarding the parameter π, we specify a zero prior mean and $\nu_\pi = \sigma_a^2/(2\bar{\lambda}) = 0.00025$, which is the unconditional variance of the productivity process a_t.

[13]The intuition for a negative productivity shock is similar: in good times the shock is not that bad, because it signals less long-run risk; in bad times the shock is very bad, because it signals more long-run risk.

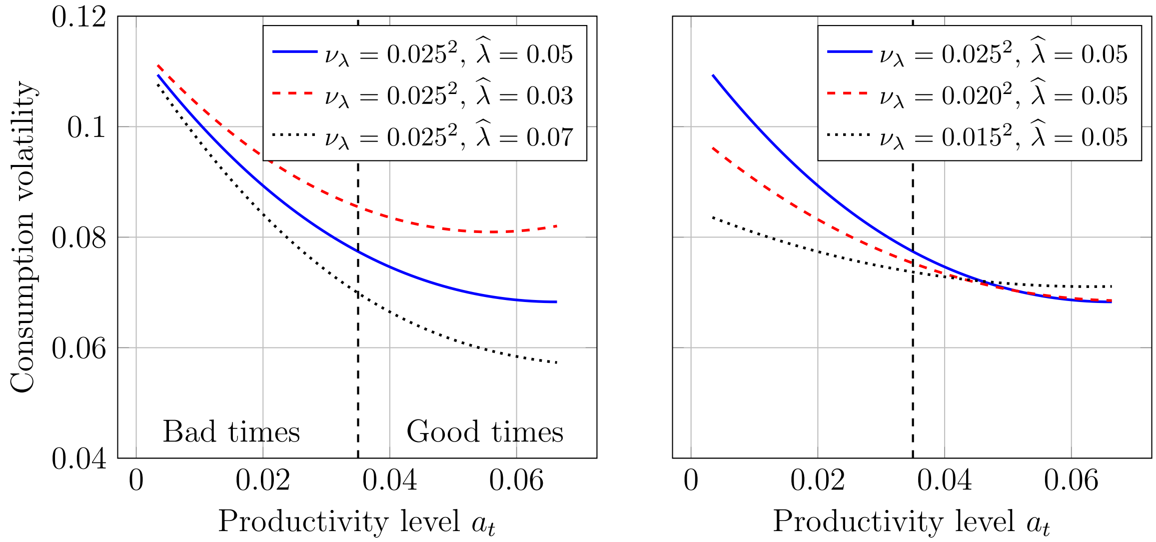

Figure 2: **Consumption volatility with learning about productivity persistence.** The figure depicts the volatility of consumption when the agent learns about the persistence in productivity. The left panel compares the model predictions for different values of the estimated mean-reversion speed in productivity $\widehat{\lambda}$, while the right panel illustrates the role of the uncertainty about this parameter, which corresponds to ν_λ. Consumption volatility is reported as a function of the productivity a_t. Low and high productivity levels a_t capture bad and good economic times, respectively. The long-term mean in productivity is given by $\bar{a} = 0.035$. The calibration is discussed in Section 5.1.

5.2 Learning and Conditional Consumption Volatility

We start by analyzing the case in which $\theta = 1$ (i.e., learning about persistence), which implies that the wealth-consumption ratio depends on only three state variables: a_t, $\widehat{\lambda}_t$, and $\nu_{\lambda,t}$. Figure 2 plots the instantaneous volatility of consumption as a function of productivity. The left panel shows that a lower value of $\widehat{\lambda}$ translates into more persistence and thus increases the volatility of consumption. In the right panel, we fix the filter $\widehat{\lambda}$ to its prior value ($\bar{\lambda} = 0.05$) and vary the uncertainty about the mean-reversion speed ν_λ. This panel shows that the relation between the level of consumption volatility and the uncertainty about persistence ν_λ is ambiguous; it varies with the state of the economy, as determined by the level of productivity a_t. With our calibration, uncertainty amplifies the volatility of consumption in bad economic times, whereas the effect of uncertainty on the volatility of consumption is weaker and ambiguous in good economic times.

The key implication of this numerical analysis is that consumption volatility is countercyclical. In particular, we show that consumption volatility is more sensitive to long-run productivity shocks for higher estimates of the mean-reversion speed $\widehat{\lambda}$. In addition, uncertainty about persistence ν_λ amplifies the countercyclical property of consumption volatility.

We now turn to learning about the level of productivity and contrast the results obtained with learning about persistence. The effect of learning about the level of productivity arises through the diffusion of consumption multiplying the short-run shock $d\widehat{B}_K$, which we restate here for convenience from Equation (20) in the case $\theta = 0$:

$$\sigma_{K,t}^C = \underbrace{\sigma_K}_{(+)} - \underbrace{\frac{\nu_{\pi,t}}{\sigma_K}\frac{x_{\widehat{\pi}}}{x}}_{(+)}. \tag{25}$$

Figure 3 depicts the level of consumption volatility for different values of the filter $\widehat{\pi}$ (left panel) and for different values of the uncertainty ν_π (right panel). The level of consumption volatility decreases with the estimate of long-term productivity $\widehat{\pi}$ (left panel) and increases with the uncertainty about the level of productivity $\nu_{\pi,t}$ (right panel). Hence, learning about the level of productivity reduces the volatility of consumption, as predicted in Ai (2010).[14]

Notably, we find that consumption volatility increases with the level of productivity a_t, which means that it becomes procyclical in this learning environment. Although (25) does not depend directly on the level of productivity, as it is the case with learning about persistence, this diffusion term actually varies with productivity through the partial derivative of the wealth-consumption ratio, $x_{\widehat{\pi}}/x$.

In conclusion, our numerical illustration depicts an important difference between learning about persistence and learning about level. We show that learning about persistence generates countercyclical consumption volatility, whereas learning about the level of productivity can potentially generate procyclical consumption volatility, as it is the case under our calibration. The key difference arises because learning about level impacts consumption's response to short-run shocks, whereas learning about persistence impacts consumption's response to long-run shocks. In bad economic times, uncertainty about persistence makes consumption more sensitive to long-run productivity shocks, which in turn generates generates countercyclical consumption volatility.

5.3 The relation between investment and consumption

In this section, we analyze the dynamic properties of investment and its covariance with consumption. For the sake of brievity, we only consider the case in which the agent learns

[14]Note that Ai (2010) reports results using a log-linearization, and thus $x_{\widehat{\pi}}/x$ becomes a constant in his setup. In our case, an exponential linear approximation of the wealth-consumption ratio would also generate a constant $x_{\widehat{\pi}}/x$, and thus both panels of Figure 3 would only show a single horizontal line. Furthermore, depending on the relative magnitude of the two terms in (25), the overall diffusion term can take positive or negative values. With our calibration, $\sigma_{K,t}^C$ becomes negative, but with the calibration from Ai (2010) it stays positive.

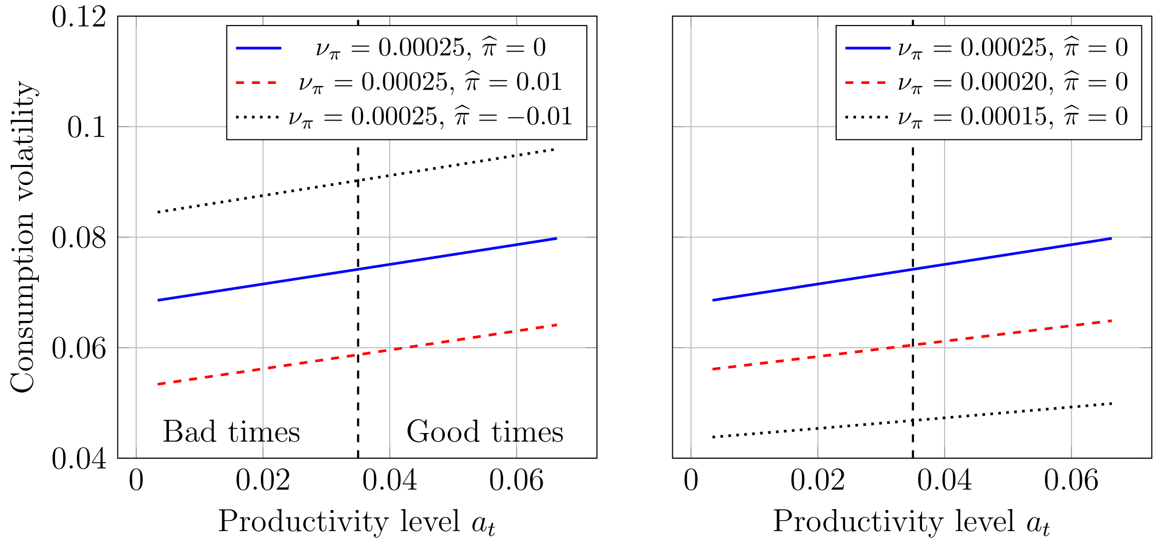

Figure 3: **Consumption volatility with learning about productivity level.** The figure depicts the volatility of consumption when the agent learns about the level of productivity. The left panel compares the model predictions for different estimates of long-term productivity $\widehat{\pi}$, while the right panel illustrates the role of the uncertainty about this parameter, which corresponds to ν_π. Consumption volatility is reported as a function of productivity in our economy. Low and high productivity levels a_t capture bad and good economic times, respectively. The long-term mean in productivity is given by $\bar{a} = 0.035$. The calibration is discussed in Section 5.1.

about persistence ($\theta = 1$). The optimal level of investment is given by

$$I_t = K_t[a_t + (1-\theta)\widehat{\pi}_t] - C_t. \tag{26}$$

Applying Ito's lemma on (26) and fixing $\theta = 1$, we obtain the following dynamics:

$$\frac{dI_t}{I_t} = \mu_{I,t}dt + \begin{bmatrix} \sigma_K & \frac{1}{a_t x - 1}\left\{\sigma_a\left(x + \frac{x_a}{x}\right) + (\bar{a} - a_t)\frac{\nu_{\pi,t}}{\sigma_a}\frac{x_{\widehat{\lambda}}}{x}\right\}\end{bmatrix}\begin{bmatrix} d\widehat{B}_{K,t} \\ d\widehat{B}_{a,t}\end{bmatrix}. \tag{27}$$

Learning about persistence impacts the response of investment to productivity shocks, through the second diffusion term in (27). This term is positive when $\bar{a} = a_t$.[15] Consider now a positive productivity shock. This shock initially increases investment. If the economy is in good times, the impact of this shock on investment is attenuated because the agent perceives now more persistence. In contrast, the impact of this shock on investment is amplified in bad times, because the agent perceives less persistence. Investment, therefore, becomes more

[15] The term $a_t x - 1$ is most of the times positive: consider for instance the productivity to be at its long-term mean $\bar{a} = 0.037$ and the wealth-consumption ratio to be 87, as estimated by Lustig, Van Nieuwerburgh, and Verdelhan (2013); this results in $a_t x - 1 = 2.045$.

volatile in bad times.

Our model has implications for the relationship between investment and consumption, an issue that has received attention in the literature (see, e.g., Croce, 2014). In particular, standard RBC models produce an almost perfect correlation between investment and consumption, whereas in the data this corelation is only mildly positive. In the context of our model, when the agent learns about persistence, the magnitude of the correlation changes with the second diffusion terms in (19) and (27), which, when $\theta = 1$, become:

$$\sigma_{a,t}^C = -\sigma_a \frac{x_a}{x} - \frac{(\bar{a} - a_t)\nu_{\lambda,t}}{\sigma_a}\frac{x_{\widehat{\lambda}}}{x} \tag{28}$$

$$\sigma_{a,t}^I = \frac{1}{a_t x - 1}\left[\sigma_a\left(x + \frac{x_a}{x}\right) + (\bar{a} - a_t)\frac{\nu_{\pi,t}}{\sigma_a}\frac{x_{\widehat{\lambda}}}{x}\right]. \tag{29}$$

The instantaneous covariance between investment and consumption is then

$$\text{Cov}_t\left[\frac{dC_t}{C_t}, \frac{dI_t}{I_t}\right] = \sigma_K^2 + \sigma_{a,t}^C \sigma_{a,t}^I. \tag{30}$$

The terms σ_a^C and σ_a^I depend on the difference $\bar{a} - a_t$ and thus on the state of the economy. In good times, the term σ_a^C is negative but its magnitude is dampened through the extrapolative expectations of the agent (as described in Section 4.3), whereas the term σ_a^I is positive but its magnitude is dampened by the same extrapolative expectations effect. The opposite happens in bad times, when σ_a^C becomes strongly negative and σ_a^I becomes strongly positive. Given this, the second term in (30) becomes strongly negative in bad times and thus decreases the amount of positive correlation generated by the first term. Our model therefore generates a lower covariance between investment and consumption in bad times, a theoretical prediction that can be tested empirically.

5.4 Pricing kernel, risk-free rate, and market prices of risk

As shown in Duffie and Epstein (1992), in an economy with recursive preferences, the state price density, denoted $\{\xi_t\}_{t\geq 0}$, satisfies

$$\frac{d\xi_t}{\xi_t} = \frac{df_C(C_t, J_t)}{f_C(C_t, J_t)} + f_J(C_t, J_t)dt \tag{31}$$

$$= -r_t dt - \lambda_1 d\widehat{B}_{K,t} - \lambda_2 d\widehat{B}_{a,t}. \tag{32}$$

This allows us to compute the risk-free rate and the two market prices of risk. First, we know from the first order condition on consumption that

$$f_C(C_t, J_t) = J_K(K_t, a, \widehat{\pi}, \widehat{\lambda}, \nu_\pi, \nu_\lambda) = K_t^{-\gamma} \beta^{\frac{-\psi(\gamma-1)}{\psi-1}} x(a, \widehat{\pi}, \widehat{\lambda}, \nu_\pi, \nu_\lambda)^{-\frac{\gamma-1}{\psi-1}}, \tag{33}$$

and thus

$$\frac{d\xi_t}{\xi_t} = \frac{dJ_K(K_t, a, \widehat{\pi}, \widehat{\lambda}, \nu_\pi, \nu_\lambda)}{J_K(K_t, a, \widehat{\pi}, \widehat{\lambda}, \nu_\pi, \nu_\lambda)} + f_J(C_t, J_t)dt. \tag{34}$$

The two market prices of risk follow from Itô's Lemma:

$$\lambda_{1,t} = \gamma\sigma_K + (1-\theta)\frac{\gamma-1}{\psi-1}\frac{\nu_{\pi,t}}{\sigma_K}\frac{x_{\widehat{\pi}}}{x} \tag{35}$$

$$\lambda_{2,t} = \frac{\gamma-1}{\psi-1}\left[\sigma_a \frac{x_a}{x} + \theta(\bar{a} - a_t)\frac{\nu_{\lambda_t}}{\sigma_a}\frac{x_{\widehat{\lambda}}}{x}\right]. \tag{36}$$

The first market price of risk is higher when there is uncertainty about the level π, and when $\gamma > 1 > 1/\psi$. The second market price of risk can be decomposed in two parts (see also Cagetti, Hansen, Sargent, and Williams, 2002): the usual price for risk and a price of uncertainty about the persistence parameter λ. Due to the presence of the difference $\bar{a} - a_t$, and knowing that when $\gamma > 1 > 1/\psi$ we have $x_{\widehat{\lambda}}/x > 0$, this second part of the market price of risk becomes countercyclical. Thus, assets that are strongly exposed to productivity shocks $\widehat{B}_a$ have more volatile cash flows when consumption volatility is high and therefore command a higher risk premium. Ostensibly, this theoretical implication is consistent to the findings of Boguth and Kuehn (2013), who conclude that stocks with volatile cash flows in uncertain aggregate times require higher expected returns.

The equilibrium risk-free rate follows:

$$r_t = a_t + (1-\theta)\widehat{\pi}_t - \gamma\sigma_K^2 - (1-\theta)\frac{\gamma-1}{\psi-1}\frac{x_{\widehat{\pi}}}{x}\nu_{\pi,t}. \tag{37}$$

The dynamics of the risk-free rate are not affected by learning about persistence. Uncertainty about the level of productivity decreases the risk-free rate when $\gamma > 1 > 1/\psi$.

The risk premium on the aggregate wealth in the economy can be obtained either by subtracting the risk-free rate from the expected return of capital (which is $a_t + (1-\theta)\widehat{\pi}_t$), or by multiplying σ_K with the market price of risk (35). It equals

$$\mu_t - r_t = \gamma\sigma_K^2 + (1-\theta)\frac{\gamma-1}{\psi-1}\frac{x_{\widehat{\pi}}}{x}\nu_{\pi,t}, \tag{38}$$

and increases with uncertainty about the level of productivity.

6 Conclusion

In this paper, we show that learning about the persistence in productivity impacts consumption's response to long-run productivity shocks and generates countercyclical consumption volatility. In bad economic times, negative news prompt agents to extrapolate that productivity becomes more persistent. The opposite extrapolation occurs in good times, when negative news induce less persistence. This asymmetric response increases consumption volatility in bad times but reduces it in good times, in line with empirical evidence. In contrast, we find that learning about the level of productivity impacts consumption's response to short-run shocks and, under certain calibrations, can potentially yield a procyclical volatility of consumption.

Our analysis essentially compares the equilibrium implications of learning about the level of productivity and learning about the persistence. Notably, the model that we propose is more general and allows the agent to learn simultaneously about these two parameters. Each type of learning would affect consumption differently. Learning about the level of productivity helps reduce the unconditional level of consumption volatility, whereas learning about the mean-reverting speed helps match the conditional properties of consumption volatility. Because these two types of learning have different implications for the behavior of consumption volatility, a useful exercise would be to use this trade-off to structurally estimate the parameter θ and, therefore, determine the type of learning that is actually present in the economy. This analysis is left for future research.

A Learning

We apply the following standard theorem on Bayesian learning (Liptser and Shiryayev, 1977):

Theorem 1 *Consider an unobservable process f and an observable process θ:*

$$df_t = [a_0(t,\theta) + a_1(t,\theta)f_t]\,dt + b_1(t,\theta)dZ_t^f + b_2(t,\theta)dZ_t^\theta \tag{39}$$

$$d\theta_t = [A_0(t,\theta) + A_1(t,\theta)f_t]\,dt + B_1(t,\theta)dZ_t^f + B_2(t,\theta)dZ_t^\theta. \tag{40}$$

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

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

$$\frac{d\nu_t}{dt} = a_1\nu_t + \nu_t a_1^\top + (b\circ b) + [(b\circ B) + \nu_t A_1^\top](B\circ B)^{-1}[(b\circ B) + \nu_t A_1^\top]^\top, \tag{42}$$

where

$$b\circ b = b_1 b_1^\top + b_2 b_2^\top \tag{43}$$

$$B\circ B = B_1 B_1^\top + B_2 B_2^\top \tag{44}$$

$$b\circ B = b_1 B_1^\top + b_2 B_2^\top. \tag{45}$$

Write the dynamics of the observable variables:

$$\begin{bmatrix} dK_t \\ df_t \end{bmatrix} = \left(\underbrace{\begin{bmatrix} K_t a_t - C_t \\ \bar{\lambda}(1-\theta)(\bar{a} - a_t) \end{bmatrix}}_{A_0} + \underbrace{\begin{bmatrix} K_t(1-\theta) & 0 \\ 0 & \theta(\bar{a} - a_t) \end{bmatrix}}_{A_1} \begin{bmatrix} \pi \\ \lambda \end{bmatrix} \right) dt + \underbrace{\begin{bmatrix} \sigma_K K_t & 0 \\ 0 & \sigma_a \end{bmatrix}}_{B_2} \begin{bmatrix} dB_{K,t} \\ dB_{a,t} \end{bmatrix}. \tag{46}$$

All the other matrices in Theorem 1 are equal to zero. Then

$$\begin{bmatrix} d\widehat{\pi} \\ d\widehat{\lambda} \end{bmatrix} = \begin{bmatrix} \frac{(1-\theta)\nu_{\pi,t}}{\sigma_K} & 0 \\ 0 & \frac{\theta(\bar{a}-a_t)\nu_{\lambda,t}}{\sigma_a} \end{bmatrix} \begin{bmatrix} d\widehat{B}_{K,t} \\ d\widehat{B}_{a,t} \end{bmatrix}, \tag{47}$$

where

$$d\widehat{B}_{K,t} \equiv \frac{1}{\sigma_K K_t}\left\{dK_t - [K_t(a_t + (1-\theta)\widehat{\pi}_t) - C_t]\,dt\right\} \tag{48}$$

$$d\widehat{B}_{a,t} \equiv \frac{1}{\sigma_a}\left\{da_t - \left[\theta\widehat{\lambda}_t + (1-\theta)\bar{\lambda}\right](\bar{a} - a_t)dt\right\}. \tag{49}$$

are independent Brownian motions coming from the filtration of the agent. The posterior uncertainties about π and λ evolve according to

$$d\nu_{\pi,t} = -\frac{(1-\theta)^2\nu_{\pi,t}^2}{\sigma_K^2}dt \tag{50}$$

$$d\nu_{\lambda,t} = -\frac{\theta^2(\bar{a} - a_t)^2\nu_{\lambda,t}^2}{\sigma_a^2}dt. \tag{51}$$

B Equilibrium

The differential operator $\mathcal{L}J$ given by

$$\begin{aligned}\mathcal{L}J &= [K(a+(1-\theta)\widehat{\pi}) - C]J_K + [\theta\widehat{\lambda} + (1-\theta)\bar{\lambda}](\bar{a}-a_t) - \frac{(1-\theta)^2\nu_\pi^2}{\sigma_K^2}J_{\nu_\pi} - \frac{\theta^2(\bar{a}-a)^2\nu_\lambda^2}{\sigma_a^2}J_{\nu_\lambda} \\ &+ \frac{1}{2}\left(\sigma_K^2K^2J_{KK} + \sigma_a^2J_{aa} + \frac{(1-\theta)^2\nu_\pi^2}{\sigma_K^2}J_{\widehat{\pi}\widehat{\pi}} + \frac{\theta^2(\bar{a}-a)^2\nu_\lambda^2}{\sigma_a^2}J_{\widehat{\lambda}\widehat{\lambda}}\right) \\ &+ (1-\theta)K\nu_\pi J_{K\widehat{\pi}} + \theta(\bar{a}-a)\nu_\lambda J_{a\widehat{\lambda}}.\end{aligned} \tag{52}$$

The guess of the value function follows from

$$J(K,a,\widehat{\pi},\widehat{\lambda},\nu_\pi,\nu_\lambda) = \frac{C^{1-\gamma}}{1-\gamma}\left[\beta x(a,\widehat{\pi},\widehat{\lambda},\nu_\pi,\nu_\lambda)\right]^\phi \tag{53}$$

$$= \frac{\left(\frac{K}{x(a,\widehat{\pi},\widehat{\lambda},\nu_\pi,\nu_\lambda)}\right)^{1-\gamma}}{1-\gamma}\left[\beta x(a,\widehat{\pi},\widehat{\lambda},\nu_\pi,\nu_\lambda)\right]^\phi \tag{54}$$

$$= \frac{K^{1-\gamma}}{1-\gamma}\beta^\phi x(a,\widehat{\pi},\widehat{\lambda},\nu_\pi,\nu_\lambda)^{-\frac{1-\gamma}{1-\psi}}. \tag{55}$$

The HJB (52) can then be written in terms of the wealth-consumption ratio x to obtain the following partial differential equation:

$$\begin{aligned}0 &= a + (1-\theta)\widehat{\pi} - \frac{\gamma\sigma_K^2}{2} - \frac{\beta\psi}{\psi-1} + \frac{1}{\psi-1}\frac{1}{x} \\ &+ \frac{[\theta\widehat{\lambda}+(1-\theta)\bar{\lambda}](\bar{a}-a)}{\psi-1}\frac{x_a}{x} - \frac{(1-\theta)(\gamma-1)\nu_\pi}{\psi-1}\frac{x_{\widehat{\pi}}}{x} - \frac{(1-\theta)^2\nu_\pi^2}{\sigma_K^2(\psi-1)}\frac{x_{\nu_\pi}}{x} - \frac{\theta^2(\bar{a}-a)^2\nu_\lambda^2}{\sigma_a^2(\psi-1)}\frac{x_{\nu_\lambda}}{x} \\ &+ \frac{1}{\psi-1}\left(\frac{\sigma_a^2}{2}\frac{x_{aa}}{x} + \frac{(1-\theta)^2\nu_\pi^2}{2\sigma_K^2}\frac{x_{\widehat{\pi}\widehat{\pi}}}{x} + \frac{\theta^2(\bar{a}-a)^2\nu_\lambda^2}{2\sigma_a^2}\frac{x_{\widehat{\lambda}\widehat{\lambda}}}{x} + \theta(\bar{a}-a)\nu_\lambda\frac{x_{a\widehat{\lambda}}}{x}\right) \\ &- \frac{\gamma+\psi-2}{(\psi-1)^2}\left(\frac{\sigma_a^2}{2}\frac{x_a^2}{x^2} + \frac{(1-\theta)^2\nu_\pi^2}{2\sigma_K^2}\frac{x_{\widehat{\pi}}^2}{x^2} + \frac{\theta^2(\bar{a}-a)^2\nu_\lambda^2}{2\sigma_a^2}\frac{x_{\widehat{\lambda}}^2}{x^2} + \theta(\bar{a}-a)\nu_\lambda\frac{x_ax_{\widehat{\lambda}}}{x^2}\right).\end{aligned} \tag{56}$$

Let us denote by x_y the partial derivative of the wealth-consumption ratio with respect to the state variable y. From the above Equation, the dynamics of consumption satisfy

$$\frac{dC_t}{C_t} = \mu_t^C dt + \begin{bmatrix}\sigma_{K,t}^C & \sigma_{a,t}^C\end{bmatrix}\begin{bmatrix}d\widehat{B}_{K,t} \\ d\widehat{B}_{a,t}\end{bmatrix}, \tag{57}$$

where the diffusion terms $\sigma_{K,t}^C$ and $\sigma_{a,t}^C$ are given by (20)-(21) in the main text and the drift of consumption is given by

$$\begin{aligned}\mu_t^C &\equiv \psi[a_t + (1-\theta)\widehat{\pi}_t] - \frac{\gamma(\psi-1)}{2}\sigma_K^2 - \beta\psi - \gamma(1-\theta)\nu_{\pi,t}\frac{x_{\widehat{\pi}}}{x} \\ &+ (\gamma+\phi)\left(\frac{\sigma_a^2}{2}\frac{x_a^2}{x^2} + \frac{(1-\theta)^2\nu_{\pi,t}^2}{2\sigma_K^2}\frac{x_{\widehat{\pi}}^2}{x^2} + \frac{\theta^2(\bar{a}-a_t)^2\nu_{\lambda,t}^2}{2\sigma_a^2}\frac{x_{\widehat{\lambda}}^2}{x^2} + \theta(\bar{a}-a_t)\nu_{\lambda,t}\frac{x_ax_{\widehat{\lambda}}}{x^2}\right).\end{aligned} \tag{58}$$

B.1 The case $\theta = 0$ (learning about level)

When $\theta = 0$, the PDE for the wealth-consumption ratio x satisfies:

$$\begin{aligned} 0 &= a + \widehat{\pi} - \frac{\gamma \sigma_K^2}{2} - \frac{\beta \psi}{\psi - 1} + \frac{1}{\psi - 1} \frac{1}{x} \\ &+ \frac{\bar{\lambda}(\bar{a} - a)}{\psi - 1} \frac{x_a}{x} - \frac{(\gamma - 1)\nu_\pi}{\psi - 1} \frac{x_{\widehat{\pi}}}{x} - \frac{\nu_\pi^2}{\sigma_K^2(\psi - 1)} \frac{x_{\nu_\pi}}{x} \\ &+ \frac{1}{\psi - 1} \left(\frac{\sigma_a^2}{2} \frac{x_{aa}}{x} + \frac{\nu_\pi^2}{2\sigma_K^2} \frac{x_{\widehat{\pi}\widehat{\pi}}}{x} \right) - \frac{\gamma + \psi - 2}{(\psi - 1)^2} \left(\frac{\sigma_a^2}{2} \frac{x_a^2}{x^2} + \frac{\nu_\pi^2}{2\sigma_K^2} \frac{x_{\widehat{\pi}}^2}{x^2} \right). \end{aligned} \tag{59}$$

This is an equation in three state variables: a_t, $\widehat{\pi}_t$, and $\nu_{\pi,t}$.

B.2 The case $\theta = 1$ (learning about persistence)

When $\theta = 1$, the PDE for the wealth-consumption ratio x satisfies:

$$\begin{aligned} 0 &= a - \frac{\gamma \sigma_K^2}{2} - \frac{\beta \psi}{\psi - 1} + \frac{1}{\psi - 1} \frac{1}{x} + \frac{\widehat{\lambda}(\bar{a} - a)}{\psi - 1} \frac{x_a}{x} - \frac{(\bar{a} - a)^2 \nu_\lambda^2}{\sigma_a^2(\psi - 1)} \frac{x_{\nu_\lambda}}{x} \\ &+ \frac{1}{\psi - 1} \left(\frac{\sigma_a^2}{2} \frac{x_{aa}}{x} + \frac{(\bar{a} - a)^2 \nu_\lambda^2}{2\sigma_a^2} \frac{x_{\widehat{\lambda}\widehat{\lambda}}}{x} + (\bar{a} - a)\nu_\lambda \frac{x_{a\widehat{\lambda}}}{x} \right) \\ &- \frac{\gamma + \psi - 2}{(\psi - 1)^2} \left(\frac{\sigma_a^2}{2} \frac{x_a^2}{x^2} + \frac{(\bar{a} - a)^2 \nu_\lambda^2}{2\sigma_a^2} \frac{x_{\widehat{\lambda}}^2}{x^2} + (\bar{a} - a)\nu_\lambda \frac{x_a x_{\widehat{\lambda}}}{x^2} \right). \end{aligned} \tag{60}$$

This is an equation in three state variables: a_t, $\widehat{\lambda}_t$, and $\nu_{\lambda,t}$.

References

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

Bansal, R., V. Khatchatrian, and A. Yaron (2005). Interpretable asset markets? *European Economic Review 49*(3), 531–560.

Bansal, R., D. Kiku, I. Shaliastovich, and A. Yaron (2014). Volatility, the macroeconomy, and asset prices. *The Journal of Finance 69*(6), 2471–2511.

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. 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.

Bekaert, G. and E. Engstrom (2015). Asset return dynamics under habits and bad-environment good-environment fundamentals. *Journal of Political Economy*, Forthcoming.

Bhamra, H. S., L.-A. Kuehn, and I. A. Strebulaev (2010). The aggregate dynamics of capital structure and macroeconomic risk. *Review of Financial Studies 23*(12), 4187–4241.

Boguth, O. and L.-A. Kuehn (2013). Consumption volatility risk. *The Journal of Finance 68*(6), 2589–2615.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics 31*(3), 307–327.

Bollerslev, T., M. Gibson, and H. Zhou (2011). Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities. *Journal of Econometrics 160*(1), 235–245.

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.

Cagetti, M., L. P. Hansen, T. Sargent, and N. Williams (2002). Robustness and pricing with uncertain growth. *Review of Financial Studies 15*(2), 363–404.

Campbell, J. Y., G. Chacko, J. Rodriguez, and L. M. Viceira (2004). Strategic asset allocation in a continuous-time var model. *Journal of Economic Dynamics and Control 28*(11), 2195–2214.

Campbell, J. Y. and L. Hentschel (1992). No news is good news: An asymmetric model of changing volatility in stock returns. *Journal of Financial Economics 31*(3), 281–318.

Carr, P. and L. Wu (2009). Variance risk premiums. *Review of Financial Studies 22*(3), 1311–1341.

Chen, H. (2010). Macroeconomic conditions and the puzzles of credit spreads and capital structure. *The Journal of Finance 65*(6), 2171–2212.

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. (2014). Long-run productivity risk: A new hope for production-based asset pricing? *Journal of Monetary Economics 66*, 13–31.

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

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica 50*(4), 987–1007.

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 (1989). Business conditions and expected returns on stocks and bonds. *Journal of Financial Economics 25*(1), 23–49.

Favilukis, J. and X. Lin (2013). Long run productivity risk and aggregate investment. *Journal of Monetary Economics 60*(6), 737–751.

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

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica 57*(2), 357–384.

Hirshleifer, D., J. Li, and J. Yu (2015). Asset pricing in production economies with extrapolative expectations. *Journal of Monetary Economics 76*, 87–106.

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.

Kaltenbrunner, G. and L. A. Lochstoer (2010). Long-run risk through consumption smoothing. *Review of Financial Studies 23*(8), 3190–3224.

Kandel, S. and R. F. Stambaugh (1990). Expectations and volatility of consumption and asset returns. *Review of Financial Studies 3*(2), 207–232.

Kim, C.-J. and C. R. Nelson (1999). Has the US economy become more stable? a bayesian approach based on a markov-switching model of the business cycle. *Review of Economics and Statistics 81*(4), 608–616.

Lettau, M., S. C. Ludvigson, and J. A. Wachter (2008). The declining equity premium: What role does macroeconomic risk play? *Review of Financial Studies 21*(4), 1653–1687.

Liptser, R. and A. Shiryayev (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.

Schwert, G. W. (1989). Why does stock market volatility change over time? *The journal of finance 44*(5), 1115–1153.

Tédongap, R. (2015). Consumption volatility and the cross-section of stock returns. *Review of Finance 19*(1), 367–405.

Yogo, M. (2006). A consumption-based explanation of expected stock returns. *The Journal of Finance 61*(2), 539–580.