Trendy Business Cycles and Asset Prices *

Jesse Davis

Kenan-Flagler Business School, University of North Carolina at Chapel Hill, USA

Gill Segal

Kenan-Flagler Business School, University of North Carolina at Chapel Hill, USA

August 2022

Abstract

The data-generating process of productivity growth includes both trend and business-cycle shocks, generating many counterfactuals for prices under full-information. In practice, agents cannot immediately distinguish between the two shocks, leading to "rational confusion": each shock inherits properties of its counterpart. This confusion magnifies the perceived share of permanent shocks and implies that, in contrast to canonical frameworks, transitory shocks are the main driver of long-run risk through *trendy business-cycles*. With learning, the equity premium turns positive, and both investment and valuation ratios become procyclical, in-line with the data. Consequently, rational confusion is key for bridging disciplined macro-dynamics with equilibrium asset-prices. *(JEL Classification: G10, D8, E32, D5)*

*We would like to thank the helpful comments and suggestions from two anonymous referees, Stefano Giglio (Editor), Ravi Bansal, Ric Colacito, Andrei Goncalves, Lars Lochstoer, Ivan Shaliastovich, and Andreas Stathopoulos. All errors are our own. Send correspondence to Gill Segal, Gill_Segal@kenan-flagler.unc.edu.

1 Introduction

Observed changes in economic growth and financial valuations can be driven by either business-cycle or stochastic-trend shocks. While both are empirically relevant, the common conception, particularly in structural asset-pricing, is that trend shocks are quantitatively more important given their permanence, whereas business-cycle shocks induce only a transitory effect. Many studies highlight the impact of persistent trend shocks on macro dynamics (e.g., Aguiar and Gopinath 2007; Jaimovich and Rebelo 2009a) while others argue that permanent shocks are the main contributors to investors' marginal utility and risk premia (e.g., Alvarez and Jermann 2005). Accordingly, seminal asset-pricing models in endowment economies (e.g., Campbell and Cochrane 1999; Bansal and Yaron 2004; Barro 2006) and production economies (e.g., Berk, Green, and Naik 1999; Tallarini 2000) assume that shocks affect the stochastic trend of either consumption (or productivity).

As a consequence, in the aforementioned models, all fluctuations in consumption's level are permanent in equilibrium. This stands at odds with ample empirical evidence suggesting that most of the variation in output and consumption is purely transitory (see, e.g., King, Plosser, Stock, and Watson 1991; Cochrane 1988; Gali 1999). If business-cycle shocks are incorporated, to be consistent with the estimated data-generating process, would equilibrium asset prices be empirically-consistent? Would permanent shocks remain the primary contributor to risk?

The answer to these questions depends upon the information environment and, in particular, the beliefs of households and firms. In practice, individuals are unable to perfectly distinguish between each type of shock in real-time and must instead learn about their true nature over time. Given an empirically-disciplined mixture of transitory and permanent shocks, we show that introducing this learning process into a general-equilibrium setting not only generates realistic asset prices, but also reveals that transitory, not permanent, shocks can be more important economically.

Specifically, we study a realistic production model that features (1) both trend and business-cycle productivity shocks, (2) imperfect-information about the underlying source of variation in total productivity, (3) Kalman filtering of the underlying shocks, and (4) recursive preferences. We estimate the model using SMM, and find that in-line with extant evidence from the macro literature (e.g., Blanchard, L'Huillier, and Lorenzoni 2013), most productivity fluctuations originate from business-cycle shocks, in contrast to the assumed technological process in most production-based asset-pricing studies.

The estimated data-generating process imposes a significant hurdle for explaining asset prices: under full-information, the estimated model produces several striking counterfactuals. Most strikingly, the equity premium is negative (or positive but very small) while the risk-free rate and valuation ratios are countercyclical.[1] Under imperfect-information, the model is able to overcome these inconsistencies through *"rational confusion"*: each shock (partially) inherits the properties of its counterpart.[2] As a result, business-cycle shocks contribute positively to expected consumption and risk-prices (like permanent shocks under full information), while trend shocks have a positive impact on investment and risk-exposures (like transitory shocks under full information).

In particular, we show that rational confusion generates *"trendy business-cycle shocks"*: realized business-cycle fluctuations that the agent mistakenly (yet rationally) attributes to shifts in the trend. First, for all parameter values, we show that the fraction of uncertainty that is perceived to arise from permanent shocks is larger than under full information, and this effect is uniquely amplified by the presence of transitory shocks. That is, trendy business cycle shocks magnify the endogenous quantity of long-run consumption risk. Second, we find that under our estimated parameters, trendy business-cycle shocks are the main source of fluctuation in expected consumption. In fact, trend shocks never have to actually materialize: as long as the agent believes that permanent shocks might happen, pure business-cycle realizations can endogenously create sizable procyclicality in expected consumption growth. Accordingly, our study is the first, to our knowledge, to quantitatively demonstrate that even when the economy features permanent shocks, the *source* of long-run risk can emanate almost entirely from business-cycle shocks. This stands in contrast with the implications of canonical frameworks, and may help reconcile the challenge faced by empiricists in supporting the long-risk framework using only consumption data.

Likewise, under imperfect-information, a fraction of the fluctuations in the trend is rationally

[1] This failure under full-information is not unique to our specific production setup. Extant asset pricing frameworks can fail to deliver a meaningful premium if sizable business-cycle shocks are introduced, a point we discuss in both Section 2 and Appendix A.

[2] We thank an anonymous referee for suggesting the terminology, "rational confusion", to describe this phenomenon.

attributed to the business-cycle. This additional source of “confusion” bears important implications for the dynamics of investment. While Beaudry and Portier (2004) and Barsky and Sims (2011) show that trend news shocks increase investment in the data, the opposite happens in our model under full-information (see, e.g., Barro and King 1984). Under imperfect-information though, investment turns strictly procyclical, precisely because a sizable fraction of the trend shock is thought to be transitory. Because investment directly affects firms’ risk exposures under production, learning turns the market’s exposure to trend shocks positive.

The combined effect of these forces highlights a novel implication of learning in our setting. While existing studies with learning typically emphasize risk magnification, rational confusion can also generate reversals: it *flips* signs and cyclicality. In our estimated model, for example, the equity premium turns from negative to positive, while the investment, the risk-free rate, and valuation ratios turn from countercyclical to procyclical, consistent with the data.

Rational confusion also bears implications for term-structures and macroeconomic dynamics. First, we show that the slope of the term-structure depends on individuals’ perception of the relative uncertainty they face with respect to each shock. For example, we show that the dividend strip expected returns curve becomes more upward-sloping under learning, qualitatively consistent with attributing a larger fraction of the forecast error variance to permanent shocks, and quantitatively consistent with recent unconditional slope estimates of Bansal, Miller, Song, and Yaron (2021). Second, we illustrate how the interplay between beliefs about business-cycle and trend shocks provide a new perspective on conditional macro moments. For example, after the financial crisis, many deemed the recovery “too slow”. Our model-implied beliefs suggest that there was a persistent decline in beliefs about the trend, an insight which we believe complements existing theories in the literature for the decline in investment growth.[3]

Importantly, while the rational confusion mechanism also arises in an endowment setup, it bears unique implications under production, in which both investment and consumption are endogenous. First, imperfect-information implies that business-cycle shocks are more likely to increase expected consumption growth under production than under endowment, due to the effect of learning on in-

[3]Specifically, we feed into the learning model the observed path of utilization-adjusted productivity growth, as measured by Basu, Fernald, and Kimball (2006) and Fernald (2012), and obtain model-implied paths for real growth rates.

vestment. Second, in the absence of flexible investment, the equity premium under full-information is always positive (albeit small). As described above, this need not be the case under production: under full-information, the equity premium can be negative, and the challenge of reconciling the data is distinct. In particular, in a production setting, the change in the risk premium's sign under learning is not only a reflection of a greater quantity of risk in the SDF – rational confusion also impacts risk exposures via flexible investment. Third, the impact of learning on the slope of term structures can change depending on the endogeneity of consumption.

Before providing more detailed intuition for our pricing results, we first clarify the nature of the shocks. Consider a mean-reverting process, x:

$$x_t = \rho x_{t-1} + \varepsilon_{x,t}, \tag{1}$$

where $\rho \in [0, 1)$. The dynamics of productivity's level, A_t, can relate in two ways to the process x. Consider the specification that A_t features a stochastic trend, $\log(A_t/A_{t-1}) = x_t$, which is widespread in asset-pricing. Whether $\rho = 0$ (i.e, "short-run" trend shocks) or $\rho > 0$ (i.e., "long-run" trend shocks), the shocks to x have a *permanent* effect on the level of A_t, because the process A_t has a unit root. Alternatively, it is possible to specify that $\log(A_t) = x_t$. In this case, which is prevalent in the RBC literature, the shocks $\varepsilon_{x,t+1}$ have a mean-reverting and *transitory* effect on the level of productivity, because $\partial \log(A_{t+k})/\partial \varepsilon_t = \rho^k$. Our model nests both cases, and their relative magnitude is dictated by the data. Because our mechanism exploits the belief dynamics that arise from the interaction of business-cycle and trend shocks, it consequently differs meaningfully from existing studies which analyze the impact of learning on risk premia.[4]

Two factors create a challenge for the full-information model to match the equity premium. First, a trend shock increases expected future consumption, creating a strong wealth effect. Thus, trend shocks, which are akin to news shocks, cause investment to drop even though our setting features the same IES and relative-risk aversion values used in Bansal and Yaron (2004). This implies that the representative firm has a negative exposure to long-run trend productivity risk, as investment

[4]This distinction between permanent and transitory shocks also applies to an economy with exogenous consumption. For instance, several extant papers study imperfect-information economies with i.i.d. shocks to consumption growth (sometimes, labeled as "short-run" shocks). These shocks are not transitory: while their impact on consumption growth lasts one period, they have a permanent impact on the level of consumption. In Appendix A, we solve an endowment economy model which features trend and business-cycle shocks and show that some of our results can arise in this setting as well.

and stock prices comove. As a result, trend shocks (which are traditionally viewed as long-run risk) can drop the equity premium under full-information. Second, while business-cycle shocks increase contemporaneous consumption they also predict future consumption growth negatively due to mean-reversion. These effects offset one another and so the impact on marginal utility is close to zero. As a result, the contribution of business-cycle shocks to the equity premium is negligible and accounting for the impact of the trend, the equity premium is *negative* under full-information.

Imperfect-information allows us to overcome these challenges. Under learning, a fraction of the trend shock is perceived to be transitory, which weakens the perceived wealth effect, allowing the substitution effect to dominate. As a result, the investment response to long-run trend shocks flips its sign, turning the exposure of the firm to these shocks positive. However, while the price of trend productivity risk remains positive, it is significantly attenuated since the firm only learns about its permanent nature over time. Combined, these changes lead the equity premium to flip from negative to slightly positive. Learning's impact on the pricing of business-cycle shocks is more significant. While investment's response to business-cycle shock remains positive under learning and, therefore, so does the firm's exposure, the price of risk for these shocks is amplified considerably. When the agent filters realized productivity shocks, a fraction of the realized transitory shock is perceived as permanent, i.e., a trendy business cycle shock is observed. This endogenously creates a small but highly persistent increase in expected consumption growth in response to a business-cycle shock, creating a large drop in the marginal utility.

The rest of the paper is organized as follows. We cover the related literature in Section 2. We present the production model with imperfect-information in Section 3. We then show the implications of learning: for beliefs about growth and uncertainty in Section 4, for macroeconomic dynamics in Section 5, for asset prices in Section 6, and for fitting empirical paths in Section 7. We provide concluding remarks in Section 8.

2 Related Literature

Our paper connects to three bodies of literature: production based asset pricing under full-information, asset-pricing studies under imperfect-information, and RBC studies with news shocks.

A growing body of research studies asset prices in the context of general equilibrium models (i.e., production-based asset pricing). Within this literature, reconciling the model-implied equity premium with the data is a well-known challenge.[5] Existing papers demonstrate that the ability of the model to match risk premia critically depends on the nature of the underlying productivity shocks. Kaltenbrunner and Lochstoer (2010) show that under Epstein and Zin (1991) preferences featuring early resolution of uncertainty, productivity shocks must have a permanent effect on the level for the model to produce a realistic Sharpe ratio. If productivity shocks are transitory (i.e., if the level reverts to its unconditional mean), then a positive transitory productivity shock is akin to a negative long-run risk shock since expected consumption growth falls through mean-reversion. By contrast, when productivity shocks are permanent, a positive shock to A_t endogenously leads to a small, positive, and predictable component in consumption growth à la Bansal and Yaron (2004) via consumption smoothing. Croce (2014) takes this conclusion a step further, showing not only that permanent shocks are required, but that they should be persistent, i.e., ρ ought to be positive. With a sufficiently high IES, such specification results in a large equity premium. Other studies that follow consider a combination of "short run" permanent shocks ($\rho = 0$) along with "long run" permanent shocks ($\rho > 0$).[6]. In these papers, permanent (short- or long- run) shocks serve as the source of long-run consumption risk as in Bansal and Yaron (2004), while business-cycle risk plays no role.

Our study complements the analysis of Kaltenbrunner and Lochstoer (2010) and Croce (2014) by considering a combination of business cycle and trend productivity shocks jointly. This mixture is consistent with the true data generating process: in-line with the findings of Blanchard, L'Huillier, and Lorenzoni (2013), we estimate the model and find that most productivity fluctuations originate from business cycle shocks. Given these productivity estimates, our contribution vis-a-vis the former papers is two-fold.

First, we show that even when business-cycle shocks dominate in the data, joint beliefs about the

[5]For instance, Jermann (1998) shows that a production model with habit preferences can produce a sizable equity premium, albeit with a counterfactually volatile risk-free rate. In a framework closer to our own, Tallarini (2000) finds that even households with recursive preferences and a very high degree of risk aversion generate a fairly low equity premium.

[6]See, e.g., Ai, Croce, Diercks, and Li 2018; Segal and Shaliastovich 2021, who consider models with "long-term" and "short-term" shocks; In these studies, the short run shocks are i.i.d. growth shocks, or one-time permanent shocks to the level of productivity, resulting in non-stationary dynamics

underlying states evolve in a way that generates empirically-consistent prices and macro dynamics. Under full-information, business cycle shocks barely impact the marginal utility under full-information (as in Kaltenbrunner and Lochstoer (2010)); with learning, the perceived possibility of trend shocks helps business-cycle shocks to have a positive impact on expected consumption and risk-prices. Similarly, under full-information, positive news shocks about the trend lead to lower investment and firm valuation and so, despite having the same preferences as in Bansal and Yaron (2004), long-run productivity shocks à la Croce (2014) can actually decrease the equity premium in our model. With learning, the perceived possibility of business-cycle shocks helps trend shocks to have a positive impact on investment and risk-exposures and, as a result, increase the risk premium.

Second, both papers demonstrate how long-run risk can endogenously arise from consumption smoothing, leading to persistence in expected consumption growth. Our aim is to uncover the true nature of the underlying shocks that drive these expectations. Under full information, as in current studies, long-run risk's underlying shocks are materially permanent (with varying degree of persistence). We show that, under learning, some fraction of mean-reverting shocks are perceived as permanent and, moreover, have a greater quantitative impact than trend shocks themselves, a new outcome in the literature. That is, under our estimates and imperfect-information, transitory business-cycle shocks, not permanent trend shocks, are the *primary* contributor to the equity premium, and constitute the main source of long-run consumption risk. In particular, it is the mere *possibility* of permanent productivity shocks which creates "trendy business cycle shocks" (even if permanent shocks never materialize).

Our paper also relates to a voluminous literature that studies asset prices under imperfect-information. This includes economies with rational state uncertainty (e.g., Cogley and Sargent 2008; Lettau, Ludvigson, and Wachter 2008; Bansal and Shaliastovich 2010; Ai 2010; Boguth and Kuehn 2013), rational parameter uncertainty (Hirshleifer, Li, and Yu 2015; Collin-Dufresne, Johannes, and Lochstoer 2016; Johannes, Lochstoer, and Mou 2016; Andrei, Hasler, and Jeanneret 2019; Andrei, Carlin, and Hasler 2019; Babiak and Kozhan 2020), or subjective beliefs (Cagetti, Hansen, Sargent, and Williams 2002; Jahan-Parvar and Liu 2014; Bidder and Dew-Becker 2016).

A common theme across the aforementioned papers is that they feature a single latent state (typically, about the dynamics of a permanent shock) or unknown parameter (typically, about the evolution of a permanent shock). By contrast, our study features an empirically-representative mixture of *two* latent state variables with *opposing* implications for agents' beliefs. We emphasize that this dual latency is not merely a technical feature: including both business cycle and trend shocks changes the focus, mechanism and outcome of our paper, vis-à-vis the former studies.

First, our focus is distinct from the former papers which study asset-prices under imperfect-information. We focus on the importance of business-cycle relative to trend shocks for expectations and risk premia. We are able to shed new light on this point by studying the *joint* evolution of beliefs about the latent trend and business-cycle states. This focus is particularly relevant, as it helps to bridge the RBC literature with existing asset-pricing theories. Moreover, we demonstrate that contrary to the approach inherited from full-information models, business-cycle shocks should not be "left out" of asset pricing models with learning: their role can exceed that of trend shocks.

Second, in existing studies, learning can enhance long-run risk by creating a martingale in beliefs, for example, or by creating excess comovement between forecast errors and beliefs about the persistent state. Our mechanism adds to existing theoretical insights in two ways. An implication of rational confusion is that the fraction of next period's uncertainty that is perceived to be associated with trend (i.e., permanent) shocks is larger than under full-information. This effect arises because business-cycle and trend shocks have the opposite effect on expectations – a feature which does not appear in models which study these shocks in isolation. Moreover, rational confusion contributes to risk premia beyond magnifying the volatility of the SDF: business-cycle shocks help trend news shocks have a positive impact on investment and risk-exposure, as described above.

Third, the main outcome of existing papers with learning is that it leads to magnification. Rational confusion can go beyond magnification and lead to reversal: learning *flips* signs and cyclicality. For example, the equity premium in our baseline model turns from negative to positive while the risk-free rate turns from countercyclical to procylical.

Lastly, our paper is connected to empirical and theoretical studies in macroeconomics that consider the effect of imperfect-information on real variables (e.g., Moore and Schaller 2002; Edge, Laubach, and Williams 2007; Boz, Daude, Durdu et al. 2008; Blanchard, L'Huillier, and Lorenzoni 2013, among others). For instance, the study of Blanchard, L'Huillier, and Lorenzoni (2013) features both persistent transitory and permanent shocks and estimates the driving forces' parameters using a structural VAR. While our estimation targets separate moments, the parameters governing productivity are strikingly similar: in particular, in both papers, business-cycle shocks are about ten-fold larger than trend shocks. We differ from these papers by considering both the real and asset pricing implications of learning. Moreover, in many RBC frameworks (e.g., Barro and King 1984, among many others), news about the long run growth typically decrease hours worked and investment, for most IES values. The same occurs in our full-information specification. While the literature has proposed several mechanisms to overcome this counterfactual (e.g., flexible utilization as in Jaimovich and Rebelo 2009a), our study suggests a separate channel: rational confusion.

3 Model

We consider an infinite horizon, discrete-time model. The economy is comprised of a representative household who owns a representative firm. The household supplies labor to the firm inelastically. The firm's technology features both a stochastic trend and a business-cycle component. Each component is driven by a persistent state. The household does not observe each state directly, but can learn about these states from output dynamics via Kalman filtering. We provide details on each agent and the learning process below.

3.1 Production

The representative firm produces its output, Y_t, using a constant returns to scale Cobb-Douglas production function, over capital, K_t, and a flow of labor, L_t:

$$Y_t = K_t^{\alpha} \left(A_t L_t\right)^{1-\alpha}, \tag{2}$$

where α is the capital share of output, and A_t is the firm's productivity. The firm owns its capital stock which depreciates at rate δ, and evolves according to the following law of motion:

$$K_{t+1} = (1-\delta)K_t + \phi\left(\frac{I_t}{K_t}\right)K_t, \tag{3}$$

where I_t represents investment, and $\phi(\cdot)$ is a positive, concave function capturing adjustment costs, specified as in Jermann (1998):

$$\phi\left(\frac{I_t}{K_t}\right) = \alpha_1 + \frac{\alpha_2}{1-\frac{1}{\zeta}}\left(\frac{I_t}{K_t}\right)^{1-\frac{1}{\zeta}}. \tag{4}$$

The parameter ζ captures the elasticity of the investment rate. The limiting case $\zeta \to \infty$ ($\zeta \to 0$) represent frictionless (infinitely costly) adjustment. The parameters α_1 and α_2 are set such that there are no adjustment costs in the deterministic steady state.[7]

3.2 Technology

We consider a specification of productivity that combines two components: a real business-cycle and a stochastic trend, similar to Aguiar and Gopinath (2007) and Blanchard, L'Huillier, and Lorenzoni (2013). Specifically, let z_t denote the business-cycle state, while x_t denotes the trend (long-run growth) state. The productivity process, A_t, follows:

$$A_t = \Gamma_t e^{z_t}, \tag{5}$$

$$\Gamma_t = \Gamma_{t-1} e^{\mu + x_t}, \tag{6}$$

where μ is the deterministic drift of productivity, and Γ_t is its stochastic trend. The state of each productivity component is given by a persistent AR(1) process:

$$z_t = \rho_z z_{t-1} + \sigma_z \varepsilon_{z,t} \tag{7}$$

$$x_t = \rho_x x_{t-1} + \sigma_x \varepsilon_{x,t}, \tag{8}$$

where $\varepsilon_{z,t}$ and $\varepsilon_{x,t}$ are i.i.d. standard normal innovations which are independent of one another, i.e., $\mathbb{E}\left[\varepsilon_{z,t}\varepsilon_{x,t}\right] = 0$. We assume that both persistence parameters, ρ_z and ρ_x, are positive but less than one. Thus, a business-cycle shock, $\varepsilon_{z,t}$, has a transitory effect on the level of productivity, A_t. On the other hand, a trend shock, $\varepsilon_{x,t}$, has a permanent effect on A_t because Γ_t is a unit root process. Because the process x_t is persistent, $\varepsilon_{x,t}$ is akin to a long-run news shock about the growth rate of productivity, similarly to Croce (2014) and Bansal and Yaron (2004).

[7]Specifically, $\alpha_1 = (\mu - 1 + \delta)^{\frac{1}{\zeta}}$ and $\alpha_2 = \frac{1}{\zeta - 1}(1 - \delta - \mu)$, where μ is the constant drift of productivity defined in Section 3.2.

3.3 Information and Learning

We assume that the agent (the manager of the representative firm, or equivalently, the household owning the firm) is imperfectly informed about the underlying productivity states. Specifically, the agent does not directly observe the business-cycle state z_t and the trend state x_t separately. Rather, the agent can only observe output, capital, and labor. Because $\log(A_t) = \frac{1}{1-\alpha}(\log(Y_t) - \alpha \log(K_t)) - \log(L_t)$, this is equivalent to assuming that the household can only observe the level of productivity A_t at each date. Thus, the information set as of time t includes the entire history of productivity shocks; $\mathbb{I}_t = \{A_{t-k}\}_{k=0}^{\infty}$. Define $\Delta a_t \equiv \log\left(\frac{A_t}{A_{t-1}}\right)$. It follows from equations (5) and (6) that:

$$\Delta a_t = \mu + x_t + z_t - z_{t-1}, \tag{9}$$

and note that Δa_t can be constructed given $\mathbb{I}_t$. We assume that the agent knows the probability distribution of z, x, and Γ. Hence, we abstract from parameter or model uncertainty, and focus on the implications of imperfect-information about the realization of the underlying shocks.

We assume that the agent forms beliefs about z and x using a Kalman filter. It is possible to cast the dynamics of productivity growth using a stationary state-space representation. Let $\sigma_t = [x_t,\, z_t,\, z_{t-1}]'$ denote the current state, and $\eta_t = [\varepsilon_{x,t},\, \varepsilon_{z,t}]'$ denote the current innovations. We can re-write the dynamics of equation (9) as:

$$\Delta a_t = \mu + \Lambda \sigma_t,$$

$$\sigma_t = T\sigma_{t-1} + R\eta_t.$$

The matrices $\Lambda_{1\times 3}$, $T_{3\times 3}$, and $R_{3\times 2}$ are defined as:

$$\Lambda \equiv [1,\, 1\, -1], \quad T \equiv \begin{bmatrix} \rho_x & 0 & 0 \\ 0 & \rho_z & 0 \\ 0 & 1 & 0 \end{bmatrix}, \quad R \equiv \begin{bmatrix} \sigma_x & 0 \\ 0 & \sigma_z \\ 0 & 0 \end{bmatrix}.$$

Let $\hat{\sigma}_{t+1|t}$ be the prediction of the agent for the underlying state σ at time $t+1$ given time t information set, i.e., $\hat{\sigma}_{t+1|t} = \hat{\mathbb{E}}_t[\sigma_{t+1}]$, where the operator $\hat{\mathbb{E}}$ denotes an expectation taken under the agent's information set, $\mathbb{E}_t[\cdot|\mathbb{I}_{t-1}]$. It follows that the law of motion for $\hat{\sigma}_{t+1|t}$ is given by:

$$\hat{\sigma}_{t+1|t} = T\hat{\sigma}_{t|t-1} + TGv_t, \tag{10}$$

where $v_t = \Delta a_t - \mu - \Lambda\sigma_{t|t-1}$ is the prediction error. The gain vector $G_{3\times 1}$ is given by $P\Lambda' F^{-1}$, where

F is a scalar given by $\Lambda P \Lambda'$, and $P_{3\times 3}$ is the steady-state estimation error covariance given by the solution to the Riccati equation:

$$P = TPT' - TG\Lambda PT' + RR'.$$

Similarly, let $\Delta\hat{a}_{t+1|t}$ denote the agent's prediction for next-period's log-productivity growth, $\hat{\mathbb{E}}_t\left[\Delta a_{t+1}\right]$. We obtain that:

$$\Delta\hat{a}_{t+1|t} \sim N\left(\mu + \Lambda\hat{\sigma}_{t+1|t},\, F\right). \tag{11}$$

3.4 Firm

At each date t, the manager of the representative firm chooses how much to invest I_t and hire L_t in order to maximize firm value given the current stock of capital, the wage rate W_t, the stochastic discount factor of the household $M_{t,t+1}$, and his beliefs about the evolution of productivity. The order of events in the period between time t and $t+1$ is as follows. First, the shocks $\varepsilon_{z,t}$ and $\varepsilon_{x,t}$ realize, and the agent observes A_t. Given this observation, the agent updates his belief about z_t and x_t, and forms a forecast for these state variables and for productivity next period, as in Equations (10) and (11). Next, given these optimal forecasts, the agent selects the factors of production. The dividend of the firm at time t is given by:

$$D_t = Y_t - I_t - W_t L_t. \tag{12}$$

We can write the firm's maximization program recursively as follows:

$$V(K_t, A_t, \hat{\sigma}_{t|t-1}) = \max_{L_t, I_t, K_{t+1}} \quad D_t + \hat{\mathbb{E}}_t\left[M_{t,t+1} V(K_{t+1}, A_{t+1}, \hat{\sigma}_{t+1|t})\right] \tag{13}$$

s.t.

(2), (3), (10), (11), (12).

The realized return of the unlevered firm at time t is given by:

$$R_{d,t}^{\text{UNLEV}} = \frac{V_t}{V_{t-1} - D_{t-1}}.$$

The first-order condition of the firm's problem in (13) implies the following Euler equation:

$$1 = \hat{\mathbb{E}}_t\left[M_{t,t+1} R_{I,t+1}\right], \tag{14}$$

where $M_{t,t+1}$ is the household's stochastic discount factor and $R_{I,t+1}$ is the return on investment. The realized return on investment is given by:

$$R_{I,t} = \frac{\alpha Y_t/K_t + q_t(1-\delta) - \frac{I_t}{K_t} + q_t\phi_t}{q_{t-1}},$$

where $\phi_t = \phi\left(I_t/K_t\right)$, and:

$$q_t = 1/\phi'(I_t/K_t) \tag{15}$$

is the Lagrangian multiplier of Equation (3), i.e., the shadow price of capital.

In our model, firms are all-equity financed, and there is no operating leverage incurred by fixed costs. In the data, there is a substantial amount of financial and operating leverage. Moreover, a considerable component of the dividend growth is due to idiosyncratic payout shocks. To incorporate these features in a transparent fashion, we follow Croce (2014), and utilize the following levered return as a proxy for the market excess return:

$$R^e_{m,t} = \phi_{lev}(R^{\text{UNLEV}}_{d,t} - R_{f,t-1}) + \sigma_d\varepsilon_{d,t},$$

where $\varepsilon_{d,t} \sim N(0,1)$ captures the effect of idiosyncratic dividend shocks. Importantly, the leverage parameter does not affect the sign of the equity premium, and the shocks, $\varepsilon_{d,t}$, do not covary with the SDF, impacting only the volatility of excess returns.

3.5 Household

The representative household exhibits Epstein and Zin (1991) preferences over a consumption stream $\{C_t\}$:

$$U_t = \left[(1-\beta)\, C_t^{1-\frac{1}{\psi}} + \beta\hat{\mathbb{E}}_t\left[U_{t+1}^{1-\gamma}\right]^{\frac{1-\frac{1}{\psi}}{1-\gamma}}\right]^{\frac{1}{1-\frac{1}{\psi}}}, \tag{16}$$

where γ denotes the household's coefficient of relative risk aversion, and ψ denotes its intertemporal elasticity of substitution (IES). When $\psi = \frac{1}{\gamma}$ equation (16) collapses to CRRA preferences. We assume that $\psi > \frac{1}{\gamma}$ so that the household exhibits a preference for early resolution of uncertainty, and dislikes uncertainty about long-run growth.

The household is endowed with one unit of labor. The household maximizes its utility by supplying labor and participating in financial markets. The household can hold a fraction Θ_t of the firm, which pays a dividend D_t as in equation (12). Consequently, the budget constraint of the household is given by:

$$C_t + \Theta_{t+1} V_t = W_t L_t + \Theta_t (V_t + D_t), \tag{17}$$

where L_t is the hours worked, and V_t is the stock price of the representative firm, defined in equation (13). Since the household does not derive disutility from labor, it supplies labor inelastically, and $L_t = 1$ in equilibrium. The first-order condition of the household's maximization program implies that the stochastic discount factor (SDF) is given by:

$$M_{t+1} = \beta \left(\frac{C_{t+1}}{C_t}\right)^{1-\frac{1}{\psi}} \left(\frac{U_{t+1}}{\hat{\mathbb{E}}_t \left[U_{t+1}^{1-\gamma}\right]^{\frac{1}{1-\gamma}}}\right)^{\frac{1}{\psi}-\gamma}. \tag{18}$$

Given this stochastic discount factor, the equilibrium real risk-free rate is pinned by

$$1/R_{f,t} = \hat{\mathbb{E}}_t \left[M_{t,t+1}\right]. \tag{19}$$

3.6 Equilibrium

An equilibrium consists of wage W_t, pricing kernel $M_{t,t+1}$, firm valuation V_t, and allocations for investment, capital, labor, consumption, and equity holding $\{I_t, K_{t+1}, L_t, C_t, \Theta_t\}_{t=0}^{\infty}$ such that: (i) Given W_t and $M_{t,t+1}$, capital and labor allocations maximize program (13), (ii) Given W_t and V_t, consumption, labor and firm holding fraction maximize (16) subject to (17), (iii) good-market clears: $C_t + I_t = Y_t, \quad \forall t$, labor market clears: $L_t = 1, \quad \forall t$, and financial market clears: $\Theta_t = 1, \quad \forall t$.

We solve the model using a perturbation method. To solve the model we first make sure that all variables are covariance-stationary. Let $\tilde{X}_t = X_t / A_{t-1}$. The firm and household problem can be equivalently expressed using the detrended control variables: $\{\tilde{K}_{t+1}, \tilde{C}_t, \tilde{I}_t\}$.

3.7 Estimation

Table 1 shows the benchmark model parameters. Following Christiano, Eichenbaum, and Evans (2005), we classify the parameters into two sets. The first set includes production and preference parameters which we calibrate based on the results of extant studies and are standard to the literature. To match the empirical evidence, capital's share of output, governed by α, is 36%. The quarterly depreciation is 2.1%, suggesting an annual depreciation of about 8%, consistent with the data. Following the strategy of Kaltenbrunner and Lochstoer (2010), we set the adjustment cost parameter ζ to 7.7, to ensure that the ratio between consumption growth volatility and output growth volatility is consistent with the data (approximately 0.7). Consistent with the total degree of leverage (joint operating and financial leverage) estimated in García-Feijóo and Jorgensen (2010), we set ϕ_{lev} to 4.18, which is similar to the leverage parameter used by Bansal and Yaron (2004). We set σ_d such that its annualized volatility $\sigma_d\sqrt{4}$, is 4%, consistent with Croce (2014). We also adopt a standard preference parameter configuration in the production-based asset-pricing literature. Specifically, γ is set to 10, while the IES, ψ, is calibrated to 1.5, implying a preference for early resolution of uncertainty. These values are identical to Bansal and Yaron (2004), among others.

The second set, denoted $\theta = \{\mu, \rho_z, \sigma_z, \rho_x, \sigma_x, \beta\}$ includes the underlying technology parameters and is estimated using SMM. We include the time-discount rate in θ to ensure that the estimated growth and time discount rate parameters satisfy the transversality condition. Our estimate of θ is the solution of the following program:

$$\hat{\theta} = \min_{\theta} \left[\Psi(\theta) - \hat{\Psi}\right]' V^{-1} \left[\Psi(\theta) - \hat{\Psi}\right],$$

where $\hat{\Psi}$ are the points estimates of the empirical moments used in the estimation, V is a diagonal matrix with the empirical variances of each moment along its main diagonal, and $\Psi(\theta)$ are the (imperfect-information) model-implied equivalent of the moments, given the quarterly parameters θ. Given a set of parameters, we compute model-implied moments based on 200 simulations of short sample paths of 200 quarters each. Each model-implied path is time-aggregated to annual observations. The corresponding empirical sample for the moments spans from 1964 to 2018.

The moments utilized include (i) unconditional annual moments: the mean, standard deviation, and autocorrelation of consumption growth, output growth, investment growth and (ii) k-period variance ratios for the time-series of annual consumption and output growth. In all, there are 6 estimated parameters and 14 moments. Table 2 shows these moments in both the model and the data.[8] For each model-implied moment, the table presents the median across all short-sample simulations, as well as the 90% confidence interval. Below, we explain how each parameter is identified.

The mean of annual consumption growth, output growth, and investment growth jointly identify the parameter μ. We estimate μ to be about 0.49%, which implies an annual real consumption growth rate of 1.96%, consistent with its empirical value of 1.93%. The time discount rate β to estimated

[8] The table shows all macro-related moments used in the estimation. See Table 3 for the risk-free moments.

to be 0.991. It is identified using the mean of the real risk-free rate, which is 0.67% in the model compared to 0.69% in the data.

There are four remaining parameters governing the exogenous productivity process: σ_z, ρ_z, σ_x, ρ_x. These parameters are over-identified using our menu of moments. First, the standard deviation (autocorrelation) of consumption, output, and investment growth rates are directly governed by σ_z and σ_x (ρ_z and ρ_x). Second, let the k-variance ratio of variable y be defined as

$$VR(k) = \frac{\mathbb{V}ar \Delta y_{t \to t+1}}{k \mathbb{V}ar \Delta y_{t \to t+k}}.$$

Kaltenbrunner and Lochstoer (2010) show that a production model's endogenous variables feature a unit root if and only if productivity has a stochastic trend. Thus, if the level of productivity is only driven by business-cycle (ε_z) shocks then $VR(k) < 1, \forall k > 1$. However, in the data, the variance ratios of both consumption and output growth rates are greater than one. These ratios can only be larger than one under an amalgamation of transitory and permanent shocks. Consequently, for any fixed k, the variance ratios of consumption and output help to identify the magnitude of σ_x compared to σ_z. All else equal, the magnitude of ρ_x relative to ρ_z dictates how sizable the permanent component is relative to the transitory component in future periods. Consequently, the persistence parameters impact the change in the variance ratios between two consecutive k's (i.e., $VR(3)$ in comparison to $VR(2)$).

The estimation yields that σ_z is 0.017, and σ_x is $8.4 \cdot 10^{-4}$. Both estimates are statistically distinct. We find that ρ_z is 0.88, and that the x process must be highly persistent, similar to Croce (2014), in order to render sufficiently high variance ratios. Interestingly, our estimates are very similar to those of Blanchard, L'Huillier, and Lorenzoni (2013). Due to modeling differences, we cannot employ the exact parametrization of Blanchard, L'Huillier, and Lorenzoni (2013).[9] However, in both studies, business-cycle shocks are estimated to be approximately 20 times larger than trend shocks ($\sigma_z/\sigma_x \approx 20$). For robustness, we apply the Kalman filter to the TFP growth time-series of Basu, Fernald, and Kimball (2006), and re-solve the model in Internet Appendix IA.3.[10]. The filtered ratio σ_z/σ_x is approximately

[9] Specifically, the model of Blanchard, L'Huillier, and Lorenzoni (2013) also includes a noisy signal about the trend shock, x, which is utilized in their estimation.

[10] As we explain in the appendix, we do not employ the TFP time-series of Basu, Fernald, and Kimball (2006) for the benchmark parametrization, because the TFP data can deviate from Δa due to empirical time-variation in markups and utilization.

24, very close to the SMM-implied ratio.

Utilizing the estimated parameters, all empirical moments in Table 2 fall inside the model's confidence intervals under imperfect-information. The empirical autocorrelation of consumption, output and investment growth rates are all very close to the model. While the model-implied volatility of consumption growth is 1.5%, the lower bound of the model's confidence interval is approximately 1.2%, in line with the data, and the implied volatility of output growth precisely matches the data.

Relative to imperfect-information, a full-information environment does not substantially alter most of the model's *unconditional* moments. There are a few exceptions. Under full-information, the autocorrelation of consumption growth rises, which generates a modest overestimate of both VR(2) and VR(3). A similar effect is observed for the autocorrelation of output growth. The effect of learning is more meaningful for *conditional* moments, or the dynamics of macroeconomic variables, as we illustrate next.

4 Learning Implications for Beliefs

We analyze the implications of "rational confusion" for beliefs about the transitory and permanent states of productivity growth. In Section 4.1, we examine learning's impact on growth expectations, which drive investment dynamics (see Section 5) and the sign of the equity premium (see Section 6.2). In Section 4.2, we analyze the effect of learning on uncertainty, which plays a key role in the amplification of the risk premium (see Section 6.3.2). Finally, in Section 4.3 we examine the perceived share of the forecast error variance associated with each shock, which is pivotal for reconciling our results pertaining to long-run risk and term structures (see Sections 6.3.1 and 6.3.3).

4.1 Implications for Perceived Growth

By combining equations (10) and (11) we obtain the agent's expectation of productivity growth:

$$\Delta\hat{a}_{t+1} = \mu + \rho_x \left(\hat{x}_{t|t-1} + g_x v_t\right) + (\rho_z - 1)\left(\hat{z}_{t|t-1} + g_z v_t\right) \tag{20}$$

$$\equiv \mu + \rho_x \hat{x}_{t|t} + (\rho_z - 1)\,\hat{z}_{t|t}, \tag{21}$$

where $g_x \equiv [1,\, 0,\, 0]G$ and $g_z \equiv [0,\, 1,\, 0]G$ are the positive gains with respect to x and z, respectively, and v_t is the prediction error. Equations (20) and (21) have an intuitive interpretation. When productivity growth is higher than expected ($v_t > 0$), both $\hat{z}_{t|t}$ and $\hat{x}_{t|t}$ increase: the agent realizes the positive shock could be due to either a positive business-cycle shock or a positive trend shock. The inability to perfectly distinguish between the two in real time leads to *rational confusion* whereby, under learning, the household partially endows each type of underlying shock (business-cycle or trend) with the characteristics of its counterpart.

Consider the case in which the prediction error, v_t, is driven by a business-cycle shock, $\varepsilon_{z,t}$. Under full-information, such a shock would be fully absorbed into agents' beliefs about z_t and, hence, their expectation of z_{t+1}. In contrast, Equation (20) implies that this shock is only partially absorbed into agents' beliefs about the state of the business-cycle process under learning, $\hat{z}_{t|t}$. Rather, since households cannot perfectly disentangle the shocks, they also rationally update their beliefs about the trend component, $\hat{x}_{t|t}$, i.e., a portion of the business cycle shock is perceived to be permanent. In this case, the innovation to the belief about the trend, $\hat{x}_{t|t} - \hat{x}_{t|t-1}$, is given by $\hat{\varepsilon}_z^{\text{trendy}} = g_x \sigma_z \varepsilon_{z,t}$. We refer to such innovations as **trendy business-cycle shocks**.[11]

Rational confusion distorts the agent's belief about future productivity growth. To demonstrate how, it is constructive to further simplify our expressions by assuming that all beliefs are at their steady-state value. This implies that $v_t = \Delta a_t - \mu$ and so equation (20) simplifies to:

$$\Delta \hat{a}_{t+1} - \mu = (\Delta a_t - \mu)(\rho_x g_x + (\rho_z - 1) g_z) \tag{22}$$

For ease of expression, suppose that productivity growth at date-t is driven only by a business-cycle shock. Let $\mathbb{E}_t [\Delta a_{t+1}]$ denote the expectations under full information. Then, rewriting equation (22), we can compare the full-information forecast to the forecast under imperfect-information:

$$\Delta \hat{a}_{t+1} - \mathbb{E}_t [\Delta a_{t+1}] = \sigma_z \varepsilon_{z,t} \underbrace{(\rho_x g_x + (\rho_z - 1)(g_z - 1))}_{>0}. \tag{23}$$

If $\varepsilon_{z,t} > 0$, a perfectly-informed agent expects lower growth in the future: transitory shocks mean-revert. In contrast, an imperfectly informed agent rationally believes that some portion of the realized forecast error may be persistent. This is because a fraction of the observed Δa_t is absorbed into $\hat{\varepsilon}_z^{\text{Trendy}}$, and positive trend shocks increase productivity growth due to their permanent effect on the productivity *level*. As a result, she expects future productivity growth to be higher than a perfectly-

[11] There is an analogous effect of a trend shock ($\varepsilon_{x,t}$) on beliefs about the business-cycle component since a fraction of it is perceived as transitory. This analogous component plays a non-trivial role for investment behavior, as shown in Section 5.2. However, as explained in Appendix A, the amplification in risk-prices under imperfect-information is driven for the most part by trendy business-cycles.

informed agent. Crucially, as shown in (23), the effect of this trendy business cycle shock on households' growth perception decays at a rate ρ_x, just like a *true* trend shock. These long-run productivity shocks, which are endogenously created via rational confusion, play a key role in explaining asset-prices, as we detail in Sections 4.3 and 6.

The corresponding case in which productivity growth is driven solely by the trend shock is completely symmetric, as can be seen by rewriting equation (22):

$$\Delta \hat{a}_{t+1} - \mathbb{E}_t \left[\Delta a_{t+1}\right] = \sigma_x \varepsilon_{x,t} \underbrace{\left(\rho_x (g_x - 1) + (\rho_z - 1) g_z\right)}_{<0}. \tag{24}$$

Relative to the perfectly-informed household, under imperfect information, the persistence of trend shocks is attenuated. This is because, in updating her beliefs, some of the observed productivity growth is perceived to be driven by the business-cycle component, which is mean-reverting.

Collectively, equations (23) and (24) highlight a novel implication of households' rational confusion. While in standard settings, imperfect-information attenuates the revision in an agent's beliefs, in our setting learning can actually flip the sign of the agent's expected productivity growth. This flip can arise immediately: e.g., a trend shock, which leads to positive growth, may be primarily attributed to the business-cycle process so that households expect negative growth. This flip can also occur over time: e.g., if households initially expect negative growth following a business-cycle shock, their beliefs can become positive before converging to the truth, through the persistence of trendy business-cycle shocks. In Section 5.1, we show that both patterns arise in our benchmark calibration.

4.2 Implications for Uncertainty

Under both full and imperfect information, households face three sources of risk with respect to their beliefs about the underlying productivity process: (i) shocks to realized growth (i.e., prediction errors), (ii) revisions to beliefs about future growth, and (iii) the covariance between these two sources of uncertainty. In what follows, we analytically characterize the impact of learning on each source.[12]

Under imperfect information, agents face more uncertainty about the next period's productivity growth: the forecast error, v_{t+1}, is due not only to the underlying shocks, $\varepsilon_{x,t+1}$ and $\varepsilon_{z,t+1}$, as under full information, but also their imperfect forecasts ($\hat{x}_{t|t}$ and $\hat{z}_{t|t}$) of the true underlying states (x_t and

[12] We prove inequalities (25), (26), and (27) in Appendix A

z_t). This implies that

$$\sigma_v^2 > \sigma_x^2 + \sigma_z^2. \tag{25}$$

In contrast, agents revise their beliefs less aggressively and, as a result, the perceived uncertainty about the state(s) of expected future productivity shrinks under imperfect information. Utilizing the notation from Equation (20), where agents update their belief about $\hat{x}_{t|t}$ and $\hat{z}_{t|t}$ by multiplying the forecast error, v_t, by the gains, g_x and g_z, respectively, this implies that

$$g_x^2\sigma_v^2 < \sigma_x^2 \qquad g_z^2\sigma_v^2 < \sigma_z^2. \tag{26}$$

Thus, the first source of risk is amplified under learning while the second is attenuated.

The impact of learning on the covariance between realized forecast errors and agents' revised beliefs depends upon the type of shock. Equation (20) implies that when productivity growth is higher than expected (i.e., $v_t > 0$), both $\hat{z}_{t|t}$ and $\hat{x}_{t|t}$ increase due to rational confusion. However, a positive forecast error could also be due to either (i) under-estimation of the trend shock (i.e., $x_t > \hat{x}_{t|t}$) or (ii) over-estimation of the business-cycle shock (i.e., $z_t < \hat{z}_{t|t}$). Thus, agents further update their forecast of the trend shock in the same direction as v_t since, on average, a positive forecast error implies that x_t was under-estimated, i.e., the covariance increases. In contrast, the upward revision in agents' beliefs about the business-cycle state is attenuated since, on average, a positive forecast error implies that z_t was over-estimated, i.e., the covariance falls. As a result,

$$g_x\sigma_v^2 > \sigma_x^2 \qquad g_z\sigma_v^2 < \sigma_z^2. \tag{27}$$

As our quantitative analysis in Appendix A suggests, these opposing effects play a significant role in the amplification of the risk premium under imperfect information.

4.3 Implications for Perceived Long-run Productivity Risk

Under imperfect information, it is straightforward to see that

$$\frac{\partial \Delta \hat{a}_{t+1}}{\partial v_t} > 0 \iff \rho_x \frac{g_x}{g_z} + (\rho_z - 1) > 0. \tag{28}$$

Whether agents expect realized productivity shocks to persist ($\frac{\partial \Delta \hat{a}_{t+1}}{\partial v_t} > 0$) or mean-revert ($\frac{\partial \Delta \hat{a}_{t+1}}{\partial v_t} < 0$) depends not only on the persistence of the underlying productivity processes but, critically, on the ratio of the gains, $\frac{g_x}{g_z}$. In equilibrium, this ratio depends on the volatility and persistence parameters via the fixed-point non-linear Riccati equation. While we cannot solve for the ratio in closed-form, we can compare it to the full information counterpart, σ_x/σ_z, utilizing the inequalities in (27):

$$\frac{g_x}{g_z} > \frac{\sigma_x^2}{\sigma_z^2}. \tag{29}$$

Equation (29) implies that, for all parameter values, the share of the forecast error variance associated with the trend (business cycle) shocks is amplified (attenuated) under imperfect information relative to full information. This observation, combined with the dynamics of expected productivity growth in Equation (28), bears two key implications for the magnitude and the persistence of long-run productivity risk.

First, the effect of rational confusion on the perceived growth associated with business-cycle relative to trend shocks do not offset one another. In particular, since $g_x\sigma_z^2 > g_z\sigma_x^2$, the effect of trendy business-cycles dominates the analogous component of trend shocks that is believed to be transitory. Second, this suggests that, on average, households expect more persistence in the realized shocks under learning since a larger fraction of the forecast error variance is believed to be permanent, on average. Intuitively, these implications affect the risk premium the household commands.

5 Learning Implications for Macro Dynamics

We analyze the dynamics of consumption, investment, and output under "rational-confusion." In Section 5.1 we consider the risk-neutral case to establish the importance of beliefs on the results that follow. In Section 5.2 we consider the risk-aversion case, and highlight that, in contrast to full-information, investment is always procyclical under imperfect-information.

5.1 Macroeconomic Dynamics under Risk-Neutrality

Figure 1 plots impulse response function to both business-cycle (ε_z) and trend (ε_x) shocks under risk-neutrality. First, consider the response to a transitory productivity shock (Figure 1a). Under full-information (the solid blue line), $\mathbb{E}_t\left[\Delta a_{t+1}\right] < 0$, and this expectation converges to zero as the impact of the shock decays. Under imperfect-information (the dotted red line), expected growth in the first period is still negative because $\rho_x g_x + (\rho_z - 1)g_z < 0$ in our benchmark estimation, consistent with (22). However, under rational confusion the firm ascribes a portion of the shock to ε_x (a trendy business-cycle shock) and so $\hat{\mathbb{E}}_t\left[\Delta a_{t+1}\right] \approx 0$, consistent with equation (23). Moreover, trendy business-cycle shocks are expected to persist ($\rho_x > 0$) and, in this setting, $\hat{\mathbb{E}}_t\left[\Delta a_{t+1}\right]$ turns positive in the succeeding periods, a reversal of sign driven by the mechanism described in Section 4. In contrast, a permanent productivity shock increases expectations under full-information, i.e., $\mathbb{E}_t[\Delta a_{t+1}] > 0$, while with learning, rational confusion suggests that the firm's expectation decreases precipitously such that $\hat{\mathbb{E}}_t[\Delta a_{t+1}] < 0$ (see Figure 1b). This is consistent with equation (24): imperfect-information causes beliefs to change sign if $\rho_x g_x + (\rho_z - 1)g_z < 0$.

These beliefs drive the firm's investment policy. While a closed form solution to equation (14) is not admissible, we can rewrite the firm's optimality condition for the special case of risk neutrality and no adjustment costs (i.e., $M_{t,t+1} = \beta$ and $\zeta \to \infty$):

$$1 = \hat{\mathbb{E}}_t\left[\beta\left(\alpha A_{t+1}^{1-\alpha} K_{t+1}^{\alpha-1} + (1-\delta)\right)\right] \tag{30}$$

Rearranging, and using Equations (10) and (11), we obtain:

$$K_{t+1} = \left(\frac{1}{\alpha}(\frac{1}{\beta} + \delta - 1)\right)^{\frac{1}{\alpha-1}} A_t^{1-\alpha} \exp\left((1-\alpha)\hat{\mathbb{E}}_t \Delta a_{t+1} + \frac{1}{2}(1-\alpha)^2 F\right). \tag{31}$$

Intuitively, equation (31) suggests that the desired level of capital depends upon both the current productivity, A_t, and the firm's belief about growth, $\hat{E}_t[\Delta a_{t+1}]$. As described in Section 4.1, learning shifts beliefs about expected productivity growth relative to full information and, as a result, alters investment growth in an analogous fashion.

For example, following the realization of a business-cycle chock, the firm increases its investment growth rate under imperfect-information by approximately a factor of five compared to full-information (Figure 1c). Under imperfect-information, the realization of a positive, trendy business-cycle shock implies that $\Delta\hat{a}_{t+1} - \mathbb{E}_t[\Delta a_{t+1}] > 0$ (see equation (23)) and so the agent over-invests. This initial amplification of investment under learning is sufficiently large that the growth rate of consumption falls below zero (Figure 1e), due to market clearing. The opposite arises when a trend shock is realized (Figure 1d) and so the agent underinvests under imperfect-information by equation (24).

These changes in investment under learning shift also the *sign* of expected consumption growth. Under full-information, a transitory shock increases contemporaneous consumption and investment is modest: expected future consumption growth is negative (Figure 1g). Under learning, the realization of a trendy business-cycle shock sufficiently amplifies investment so that expected future consumption growth turns positive. In contrast, a positive trend shock leads to positive investment growth under full-information and, therefore, a sharp increase in expected future consumption (Figure 1h). Under imperfect-information investment is muted which attenuates expected consumption growth.

5.2 Macroeconomic Dynamics under Risk-Aversion

We now turn to the benchmark case in which the household exhibits Epstein and Zin (1991) preferences. The precision of the signal that the firm acquires from realized output does not depend on capital level (as in Van Nieuwerburgh and Veldkamp 2006, for instance), and so the firm's beliefs about productivity are unchanged from Section 5.1.

Under full-information and risk aversion, the firm increases investment in response to a transitory shock relative to risk-neutrality (Figure 2c). With risk aversion (or more precisely, a positive IES), the agents want to smooth consumption which necessitates higher investment. This attenuates the fall in expected consumption growth (Figure 2g). Imperfect-information induces two, countervailing effects on investment: it increases the perceived marginal product of capital, but decreases the incentive to smooth consumption (because trendy business-cycle shocks increase expected consumption growth.) Only the first effect arises under risk-neutrality. With risk aversion, the two effects nearly offset each other, so that the imperfect-information impulse response functions on real quantities are nearly indistinguishable from those derived under full-information (see Figure 2c).

In contrast, following a trend shock, investment growth "flips" signs when comparing learning to full-information (Figure 2d). Trend shocks affect the long-term much more than the short-term (estimated ρ_x is high, while σ_x is relatively small). In essence, the trend shock is akin to a news shock about the future (plus a small short-run innovation). In our estimated model, as well as in many others, these news shocks lead to lower investment.[13] When a positive shock to the long-run trend of productivity materializes, the continuation value increases. Higher expected productivity generates a substitution effect that increases the opportunity cost of consumption, and simultaneously, an income effect, that allows the agent to feel wealthier. The lower (higher) the IES is, the more current consumption and continuation utility are complements (substitutes). If the IES is not too high (in our

[13] Barro and King (1984) shows that a one-sector growth model generates aggregate comovement only in the presence of contemporaneous (short-run) shocks to total factor productivity. News shocks generate a negative correlation between consumption and investment. Jaimovich and Rebelo (2009b) also show that in the absence of flexible utilization, good news about future productivity make agents wealthier, leading to higher consumption, but lower investment and labor supply. More recently, Croce (2014) demonstrates this point, in a model with Epstein-Zin preferences, and when the IES is set close to 1.

case the IES is 1.5, as in Bansal and Yaron 2004), the income effect dominates, and the agent finds it optimal to increase consumption by reducing investment.[14]

The negative response of investment to trend shocks under full-information contradicts extant empirical evidence. Beaudry and Portier (2004) identify TFP news shocks using a vector error correction model and find that investment rises both in response to transitory and trend news shocks shocks. Similarly, Barsky and Sims (2011) identify TFP news shocks using a structural VAR, showing that news shocks have a positive impact on stock prices (which comove with investment).

Rational confusion aligns the model with the data, as each shock partially inherits the attributes of its counterpart. Under full-information, the response of investment to transitory shocks is positive. This is a manifestation of the Barro and King (1984) result where only short-run productivity raises investment. When a trend shock is realized under imperfect-information, a large fraction of the shock is perceived to be transitory ($\sigma_z > \sigma_x$) such that investment increases, in-line with the data. Put differently, with learning, the expected growth rate of consumption is nearly zero following a trend shock. As a result, the income effect largely disappears, and investment growth *increases* (Figure 2d).

6 Learning Implications for Asset Prices

We explore the implications of "rational confusion" for asset pricing. Section 6.1 summarizes key pricing moments: in particular, learning flips the sign of the equity premium and magnifies its absolute value. In Section 6.2 we dissect the key mechanism and connect the change in the equity premium's sign under learning to the firm's investment policy. Section 6.3 examines how learning affects the source of long-run risk, the magnitude of the equity premium, and the term structure of risk. This is done by contrasting the model's predictions to an equivalent endowment economy, allowing us to decouple the role of learning from production. Lastly, in Section 6.4 we consider the impact of learning on the cyclicality of stock prices and the risk free rate.

[14] In our model, trend shocks lower investment as long as the IES is not too high. Given our estimated model parameters, we find that investment declines following a trend shock so long as the IES is lower than 1.8. This is substantially higher compared to recent estimates of the IES, placing it very close to 1 (see, e.g., Calvet, Campbell, Gomes, and Sodini 2021).

6.1 Imperfect-Information and Asset-Pricing Moments

Table 3 reports model-implied asset-pricing moments against their empirical counterparts. The model-implied moments are obtained both from the imperfect-information model detailed in Section 3, as well as from an otherwise identical model in which the agent is fully-informed about z_t and x_t.

Under imperfect-information, the model-implied sample average of the firm's excess return (equity premium) is sizable, at about 2.6% per annum. While the point estimate of the empirical equity premium (4.8%) is larger, the model-implied premium falls within the confidence interval of the data. The model-implied premium is larger than the risk premium on aggregate wealth estimated by Lustig, Van Nieuwerburgh, and Verdelhan (2013). Moreover, the premium under learning is nearly ten times as large as the one obtained in the framework of Kaltenbrunner and Lochstoer (2010) with early resolution of uncertainty. We purposefully abstract from more elaborate adjustment costs (which could further increase the absolute value of the risk premium) for parsimony, since our main interest is its relative magnitude compared to full-information.[15] We find that learning not only *amplifies* the equity premium (the difference between the equity premium under imperfect- and full- information is about 3.2%), but it flips its *sign*: with full-information the equity premium is -0.5%.[16] The amplification effect of learning is even more strongly manifested for the Sharpe ratio. The Sharpe ratio implied by the imperfect-information model is about 50%. This is even more sizable than the empirical counterpart, but the model-implied and data-implied confidence intervals overlap. By contrast, under full-information, the Sharpe ratio (in absolute value) is approximately half of the ratio in the data.

The real annual risk-free rate is about 0.7% in both the imperfect-information model and the data. Interestingly, learning not only amplifies the equity premium, as discussed above, but it also dampens the risk-free rate. Under full-information, utilizing the same parameter values, the risk free rate is about 2.2%. The volatility of the risk free rate under imperfect-information is 0.65% with an upper bound of 1.2%, which is close to the empirical confidence interval. Finally, we note that learning

[15] Notably, the model-implied equity premium under learning can be higher in absolute terms in a model featuring asymmetric adjustment costs, which make it harder to smooth profits in bad states (making dividends less countercyclical; see, e.g., Zhang (2005)). We discuss how additional frictions can impact the equity premium in Section IA.1. For instance, we show that the risk premium of a consumption claim is above 8% p.a. in a production model with learning.

[16] In short sample simulations, the full-information equity premium can be slightly positive, with the presence of idiosyncratic dividend shocks. The population-sample model-implied equity premium in this case is strictly negative, at -0.57%.

slightly attenuates the volatility of the risk free rate.

6.2 Inspecting the Mechanism

The conditional equity premium in the model can be generally written:

$$E_t[R^e_{t+1}] = \beta_{z,t}\lambda_{z,t} + \beta_{x,t}\lambda_{x,t},$$

where

$$\beta_{j,t} = \frac{\partial R^e_{t+1}}{\partial \varepsilon_{j,t+1}},$$
$$\lambda_{j,t} = -\frac{\partial M_{t,t+1}}{\partial \varepsilon_{j,t+1}} \quad j \in \{z, x\}.$$

$\beta_{j,t}$ is the exposure of the firm to shock j, while $\lambda_{j,t}$ is the market price of risk of the shock. While both exposures and prices of risk can be time-varying, to make the mechanism clear we consider their unconditional magnitudes in the (stochastic) steady-state. In what follows, we drop the time subscript from the exposures and risk prices to denote steady-state values. The unconditional prices of risk (exposures) can be computed via model-implied impulse responses from the underlying shocks to the firm's excess return (household's marginal utility). We analyze the exposures and prices of risk under both information environments below to clarify the sign and decomposition of the equity premium.

6.2.1 Full-Information

In our model, the valuation of the representative firm and investment co-move. If the representative firm wants to expand its capital stock, it is required to pay capital adjustment costs, which restrict its ability to fully absorb productivity shocks through investment. Thus, since shocks are not fully absorbed in quantities, they are absorbed in installed capital's price. Because the returns to scale in our model are constant, this shadow price, or Tobin's q from Equation (15), is a sufficient statistic for the ex-dividend firm valuation. In other words, the ex-dividend firm valuation exposure to a productivity shock, $\varepsilon_{j,t}$, is $\partial V_t/\partial \varepsilon_{j,t} = \frac{\partial}{\partial \varepsilon_{j,t}} q(I_t/K_t)K'$. As $q' > 0$, the valuation of the (dis)investing firm rises (falls), and thereby covaries more with aggregate productivity. As discussed in Section 5, a business-cycle shock $\varepsilon_{z,t}$ raises investment, and as a result, $\partial V_t/\partial \varepsilon_{z,t} > 0$ which implies that $\beta_z > 0$. On the other hand, a trend shock $\varepsilon_{x,t}$ initially drops investment under full-information and so $\beta_x < 0$.

The innovation to the stochastic discount factor in Equation (18) can be written as:

$$M_{t-1,t} - \mathbb{E}_{t-1} M_{t-1,t} = -\gamma \left(\Delta c_t - \mathbb{E}_{t-1} \Delta c_t\right) - \left(\gamma - \frac{1}{\psi}\right) \left(\log(\tfrac{U_t}{C_t}) - \mathbb{E}_{t-1} \log(\tfrac{U_t}{C_t})\right). \tag{32}$$

Hence, under early resolution of uncertainty (i.e., $\gamma - \frac{1}{\psi} > 0$), a productivity shock decreases marginal utility (and therefore has a positive price of risk) if it either (i) increases current consumption or (ii) raises the continuation utility. The blue line in Figures 3a and 3b shows how $\log(M_{t-1,t})$ responds under full-information to a business-cycle and a trend shock, respectively. In both cases, a positive shock decreases marginal utility. The magnitudes are extremely asymmetric however: the impact of a trend shock on the marginal utility is approximately 0.4 (in absolute value), while the effect of the business-cycle shock is close to zero.

We briefly outline the intuition for this result, which echoes the finding of Kaltenbrunner and Lochstoer (2010). As previously discussed in Section 5, and shown in Figures 2e and 2f, both the transitory business-cycle shock and the permanent trend shock increase contemporaneous consumption.[17] With early resolution of uncertainty, however, the two shocks have an opposite impact on U_t/C_t. Figure 3c (3d) shows that the business-cycle shock decreases the continuation utility to consumption ratio while a trend shock increases it.

A positive trend shock implies that consumption increases not just contemporaneously, but also in the future, as the shock permanently shifts technology (and output) to a higher level. Note that the innovation of the utility-to-consumption ratio is proportional to $\sum_{s=1}^{\infty} \kappa^j (\mathbb{E}_t - \mathbb{E}_{t-1}) \Delta c_{t+s}$, where κ is a log-linearization parameter that is close to one. Since a positive trend shock increases the expected future consumption growth, as shown in Figure 2h, it raises the continuation utility.

A positive business-cycle shock also implies that consumption increases contemporaneously, but because the technology shift is transitory, it is expected to be lower in the future, eventually converging to its pre-shock level. As demonstrated in Figure 2g, the expected consumption growth is persistently negative and the continuation utility falls. Thus, the impact on the continuation utility amplifies the price of risk for trend shocks, while (almost) completely eliminating the price of risk for business-cycle

[17] Therefore, if the household exhibited CRRA preferences ($\gamma = \frac{1}{\psi}$), both shocks would have a positive price of risk. Furthermore, given that the expected impact of business-cycle shocks on current consumption is larger in our calibration, such shocks would bear a larger risk price.

shocks, so that $\lambda_x >> \lambda_z \approx 0$.

Combined, the contribution of business-cycle shocks to the equity premium is close to zero since $\underbrace{\beta_z}_{(+)}\underbrace{\lambda_z}_{\approx 0}$ while the contribution of trend shocks, $\underbrace{\beta_x}_{(-)}\underbrace{\lambda_x}_{(+)}$, is negative. Taken together, the full-information environment gives rise to a negative equity premium.

Relative Contribution. The relative contribution of business-cycle shocks to the equity premium can be expressed as:

$$\frac{\beta_z \lambda_z}{\beta_z \lambda_z + \beta_x \lambda_x}, \tag{33}$$

while the contribution of permanent shocks is the complement to 100%.[18] Panel A of Table 4 that under full-information, the ratio (33) is approximately -45%, broadly consistent with the findings of Kaltenbrunner and Lochstoer (2010).[19]

6.2.2 Imperfect-Information

The dashed red line in Figure 2c shows that imperfect-information induces only a small attenuation effect on the response of investment to a business-cycle shock. Therefore, β_z is almost identical to the full-information model. For trend shocks, the effect of learning is not only quantitative, but also qualitative (see Figure 2d). As explained in Section 5, trend shocks increase investment under learning so that while β_x is negative under full-information in our estimation, it is positive under learning.

Crucially, learning also affects the stochastic discount factor, as is apparent upon examination of the dashed red lines in Figures 3a and 3b. While both shocks still decrease marginal utility, in contrast to the setting with full-information, the effect of the trend shock is now quantitatively small, while the effect of the business-cycle shock is now sizable. We explain this key result in what follows.

As discussed in Section 6.2.1, an observed increase in output raises the continuation utility if the agent believes that it is driven by a permanent shock. Under imperfect-information, however, the agent does not immediately know if a higher output realization was triggered by a permanent or a transitory shift to productivity. The agent can learn if the shift is permanent only gradually over time, by tracking output's trajectory in future periods. Thus, following any unanticipated increase

[18] The risk exposures β_j, $j \in \{z, x\}$ are computed by the contemporaneous impulse-response of ε_j to R_m^e. Similarly, the market prices of risk λ_j are computed by the impulse-response of ε_j to M.

[19] The contribution of z-shocks in this case is negative because $\beta_z \lambda_z > 0$ while the equity premium is negative, due to permanent shocks.

in output the agent places some likelihood that the shift is permanent and some likelihood that it is transitory where, as discussed in Section 4, the likelihood ascribed to each depends on the relative size of the two shocks.

Consequently, when a trend shock realizes, the agent does not ascribe all of the unexpected increase in output to the permanent component. Rather, the agent initially believes that most of the unexpected variation is caused by the business-cycle shock, since these shocks are much more sizable. Because business-cycle shocks (under full-information) negatively predict consumption growth, this belief dampens the immediate response of the utility-to-consumption ratio to the realization of a trend shock, as seen in the red dashed line of Figure 3d. This, in turn, significantly attenuates the market price of risk of trend shocks, λ_x, in comparison to perfect information.[20]

Similarly, when a business-cycle shock realizes, the agent does not ascribe all of the prediction error to the transitory component. Instead, the agent believes that some of it is due to the permanent component, equivalent to a positive realization of a trendy business-cycle shock ($\hat{\varepsilon}_z^{\text{Trendy}} > 0$). Although the trendy business-cycle shock to productivity is smaller in magnitude compared to the fraction of the prediction error ascribed to the transitory component, the agent believes that a trendy business-cycle shock is highly persistent: through the lens of the Kalman filter, it inherits the same persistence as that of true trend shocks, ρ_x. Consequently, trendy business-cycle shocks induce a persistent increase in expected consumption growth, and the price of risk of business-cycle shocks, λ_z, is positive and larger with imperfect-information.

Combined, the contributions of both business-cycle and trend shocks to the equity premium are positive given $\underbrace{\beta_z}_{(+)}\underbrace{\lambda_x}_{(+)}$ and $\underbrace{\beta_x}_{(+)}\underbrace{\lambda_x}_{(+)}$, respectively. As a result, the equity premium flips its sign compared to full information. We note that under imperfect information the prices of risk of both shocks should share the same sign. Similarly, the risk exposures to both shocks should share the same sign: if the agent only observes the change in total productivity, but not the underlying nature of the change, the sign cannot differently depend on the underlying nature of the shock. However, the relative magnitude, and the contribution to the equity premium from the two shocks is highly

[20]As seen in Figure 3d, the utility-to-consumption ratio gradually but slowly increases over time, as the agent ascribes a growing fraction of the initial shock to the permanent trend component. However, the price of risk is affected by the immediate response of the utility-to-consumption ratio.

asymmetric.

Relative Contribution to the equity premium. Panel A of Table 4 shows the ratio (33) under learning. In contrast to full-information, almost the entire magnitude of the equity premium is explained by transitory business-cycle shocks rather than permanent shocks under learning. When agents must forecast productivity changes from aggregate output, transitory shocks explain as much as 99% of the premium. This arises due to the sharp amplification (attenuation) of λ_z (λ_x) under learning. Next, we link these effects to how transitory and permanent productivity shocks relate to endogenous short- and long-run consumption risk.

6.3 Consumption Risks: The Role of Learning and Production

In what follows, we further dissect the role of imperfect-information and production for endogenous consumption risks. We establish that under learning (i) business-cycle shocks are the primary drivers of expected consumption growth fluctuations; (ii) consumption risk compensation is magnified due to both an amplified quantity of risk as well as altered risk exposures; (iii) the term structure of real yields (expected equity returns) becomes more downward (upward) sloping.

To decouple the role of learning from the impact of flexible investment, we build on the insights which arise from an analogous economy where consumption is *exogenous* but subject to both transitory and permanent shocks as in our benchmark model. Specifically, in Appendix A, we first characterize analytically the risk exposures and prices and then quantify them utilizing the estimated parameters from the benchmark model. We refer to this framework as the "endowment economy", in contrast to our "production economy" from Section 3.

Notably, many of the implications of learning described in Section 6.2.2 arise in the endowment economy. As in the production economy, learning amplifies the risk premium and reduces both the risk-free rate and its volatility. Moreover, business-cycle shocks explain little (most) of the risk premium under full information (learning). Finally, as in production, rational confusion is the primary quantitative driver of the risk premium amplification.

However, not all of the production-based results arise in the endowment economy. Intuitively, just as rational confusion bears implications for beliefs about consumption and the SDF (as can be seen in the endowment economy), rational confusion also bears implication for investment and risk exposures (and those are only captured in production). Furthermore, as we detail next, the endogeneity of consumption also plays a non-trivial role.

Specifically, there are three key features for which production is pivotal: In Section 6.3.1, we map the findings to the Bansal and Yaron (2004) framework, and show that a positive nexus between business-cycle shocks and expected consumption growth shocks depends (partially) on flexible investment. In Section 6.3.2, we decompose the risk premium amplification into distinctive channels, and emphasize how production enhances the impact of learning on the risk premium above and beyond the endowment economy case. In Section 6.3.3, we demonstrate how the term structure of equity risk premia differs between production and endowment under imperfect information.

6.3.1 Economic Origin of Long-Run Risk

Can long-run risk originate from business-cycle shocks? As we show next, the answer hinges on both the information environment and the endogeneity of consumption. To quantify the relation between business-cycle shocks and innovations to long-run risk, we perform a simulation exercise. First, we simulate population paths of the exogenous shocks ε_z and ε_x, from which we obtain model-implied population paths of expected consumption growth. Then, letting $e_t \equiv E_t[\Delta c_{t+1}]$, we fit e_t into an auto-regressive process: $e_{t+1} = const + \rho_e e_t + \eta_t$, where η_t is the expected consumption growth innovation. Finally, we compute the correlation between η_t and the contemporaneous business-cycle shock ($\varepsilon_{z,t}$). We perform this for both the production and endowment economies, under full- and imperfect-information.

Endowment economy. Under full-information, a trend shock induces a positive 'long-run risk' impact on consumption growth, as in Bansal and Yaron (2004). However, due to mean-reversion, business-cycle shocks are unambiguously akin to a negative 'long-run risk' shock. Indeed, Table Appendix-A.1 of the Appendix shows that the correlation between business-cycle shocks and η_t is -0.91 under full-information.

Under learning, the relation between business-cycle shocks and long-run risk is more nuanced. When a business-cycle shock realizes, a part of the observed change in the exogenous process is perceived to be transitory - this lowers expected consumption growth as in full-information. On the other hand, a part of the shock is rationally perceived to be permanent, and a trendy business-cycle shock is realized. Such shocks behave like ordinary 'long-run risk' shocks and render a small but highly persistent increase in future consumption growth. In the endowment economy, realized shocks end-up increasing expected consumption growth under learning if equation (28) holds.

This condition, however, is not satisfied by the benchmark model estimation of Table 1. Indeed, Figure 2a shows that in our parametrization, a positive business cycle shock drops expected productivity growth (which is identical to expected consumption growth in the endowment model), though less so under learning compared to full-information. If the agent consumes the productivity process, as in the endowment case, it would imply that a business cycle shock lowers expected consumption, similar to full-information. Consistently, Table Appendix-A.1 in the Appendix shows that while the correlation between $\varepsilon_{z,t}$ and η_t is higher under learning than full-information, it remains negative.

Production economy. Intuitively, as in the endowment economy, business-cycle shocks lower expected consumption growth under full-information. Indeed, panel B of Table 4 shows that in our benchmark full-information model the correlation between business-cycle shocks and η_t is -0.83.

In contrast to the endowment economy, however, condition (28) does not have to hold in order for business-cycle shocks to increase expected consumption growth. Under production, consumption is endogenous, and its future value depends only partially on expected productivity growth. Under imperfect-information, a business-cycle shock is followed not only by (i) higher expectations for productivity growth (via trendy business-cycle shocks) but also (ii) an increase in investment (see Figure 1c). The combined effect of these two forces implies that perceived future output growth is considerably larger than in the endowment economy, which *flips* the expectation of consumption growth induced by business-cycle shocks: it turns positive under imperfect-information (see Figure 2(g)).

As a result, panel B of Table 4 shows that the correlation between η_t and $\varepsilon_{z,t}$ is both positive and sizable under learning (approximately 0.8). Moreover, the correlation between expected consumption growth innovations and $\varepsilon_{x,t}$ largely disappears (approximately one-tenth as large as under full-information). In short, our estimation suggests that transitory shock realizations are responsible for most of the variation in the beliefs that govern long-run risk. Though both positive trend shock and positive business-cycle shocks contribute positively to the belief about the state of x_t, trendy business-cycle shocks are quantitatively more important because of the relative size of transitory shocks.

The above suggests that production is important for micro-founding *positive* long-run consumption risk via business-cycle shocks. In particular, our estimation implies that rational confusion may not be sufficient to overturn the expectation of productivity growth from negative to positive, but sufficient to overturn the expectation of consumption growth.

6.3.2 The Magnitude of Risk Premia

Through what mechanisms is the risk premium amplified? Learning can affect risk premia through its impact on both the quantity of risk (uncertainty) and risk exposures. While learning amplifies the total quantity of consumption risk in both production and endowment economies, it substantially affects risk exposures only in the production setting.

Endowment economy. Recall that Section 4 highlighted three sources of risk (uncertainty) with respect to beliefs about the underlying stochastic process: (i) uncertainty about realized growth or "prediction errors", (ii) uncertainty over beliefs about future growth states, and (iii) the covariance between these two sources of uncertainty. Under learning, there is more uncertainty about source (i) and less uncertainty about source (ii) (see equations (25) and 26)). Quantitatively, we show in Appendix A that within an endowment economy, the former modestly increases the risk premium while the latter depresses it. Importantly, we find that the learning-induced amplification of the covariance risk (source (iii)) is the crucial channel for understanding the increase in the quantity of risk.

The analysis in Appendix A also leads to two novel observations. First, the change in the covariance risk is *uniquely* magnified in our economy due to the presence of two latent state variables, whose forecast errors are *negatively* correlated. That is, the absolute change in the covariance between the prediction error and the revision of the business-cycle (trend) state becomes even larger, all else equal, due to the presence of the trend (business-cycle) state (see Equations A56) and (A59). Second, due to their size, trendy business cycles are the primary contributor to the amplification of the aforementioned covariance. This implies that total consumption risks are disproportionately driven by business cycle shocks, as we highlighted in Section 6.2.2.

However, Appendix A illustrates that, under the endowment economy, the risk exposures do not materially hinge on the information environment: they are always positive, irrespective of the source of the shock (equivalently, $\beta_z > 0$ and $\beta_x > 0$, with exogenous consumption). Without flexible investment, which creates a wedge between productivity and dividend fluctuations, the consumption claim must positively comove with any shock that raises consumption. As a result, we show that the risk premium under full-information must be positive in the endowment economy – risk exposures and risk prices for each shock share the same sign. Hence, learning only magnifies (an already) positive risk premium.

Production economy. In a production economy, the three sources of productivity risk outlined in the endowment case endogenously propagate into consumption risks. This type of propagation from productivity to consumption is present in other production-based studies (e.g., long-run productivity risk creates long-run consumption risk, as in Croce 2014). This propagation changes the perceived quantity of consumption risks in three ways.

First, learning amplifies uncertainty by inducing a higher correlation between short-term and long-term risk (corresponding to the aforementioned covariance risk from source (iii)). Second, learning implies that the risk price of business-cycle shocks is larger than that of trend shocks (see Figure 3(a) and 3(b)). Third, and perhaps most importantly, the share of the forecast error variance associated with the trend (business cycle) shocks is amplified (attenuated) under learning (see Section 4.3). Since a larger fraction of the forecast error variance is believed to be permanent, the household expects more persistence in future consumption growth. Panel C of Table 4 confirms that expected consumption growth becomes more auto-correlated under learning, contributing to a larger premium.[21]This result is not unique to production: Table Appendix-A.1 shows that under the endowment economy, the autocorrelation of expected consumption growth is similarly amplified with imperfect-information.

In contrast to the endowment economy, however, rational confusion *and* production can *flip* the

[21]The learning-induced change in the auto-correlation of expected consumption, from 0.990 to 0.999, is at first glance, rather small. Nonetheless, a small amplification in the persistence of expectations can have a meaningful impact, similar to the intuition found in the canonical long-run risk model. One way to see this is using Equation (A18) in the Appendix. Because κ is close to 1, the higher ρ_x, the closer the denominator to zero: as a result, even a small change in the degree of persistence can have a large pricing impact, since A_x (in which the risk premium is increasing) is convex in ρ_x.

To illustrate this numerically, we conduct a back-of-the-envelope computation based on calibration #2 of the endowment economy. Under full-information, the implied persistence of expected consumption growth is 0.990, and the associated risk premium is 5.7%. If we increase ρ_x so that the implied persistence of **expected** consumption is the same as under learning (i.e., 0.999), we obtain a risk premium of 9.22% under full information, i.e., the equity premium is amplified considerably (by 3.52%). That is, a small increase in the perceived persistence reflects a large increase in the risk premium demanded.

premium from negative under full information to positive with imperfect-information (as shown in Table 3), since learning also affects risk exposures. In production, risk exposures depend on endogenous investment. A firm that disinvests has a lower growth option value. If the marginal utility contemporaneously drops, the lower valuation creates a hedge that lowers the risk premium. As discussed in Section 5.2, under our benchmark estimation, trend shocks lower investment under full-information, yielding $\beta_x < 0$ and a negative risk premium. Thus, the challenge of reconciling the equity premium under production is substantially distinct from the endowment economy. Because imperfect-information leads to an increase in investment in response to a trend news shock, β_x turns positive under learning, and so does the trend shock's contribution to the risk premium.

In sum, production provides an additional channel through which rational confusion affects the risk premium under learning through its effect on investment (and hence, risk exposures). The risk premium is not merely amplified due to increased uncertainty, as occurs under endowment, but it also flips its sign.

6.3.3 Term Structure of Risk

What is the impact of learning on consumption risks at different horizons? In Appendix B, we study the impact of rational confusion on both the real-yield and equity term structure. In both the endowment and production economies, if there were only trend shocks, the real yield curve would be downward-sloping, whereas the equity curve would be upward-sloping. Intuitively, permanent shocks make long-run cash flows riskier compared to the short-run, and bonds that hedge this long-term uncertainty are more valuable. On the other hand, these slopes would flip if there were only business-cycle shocks because long-run cash flows would be relatively less risky.

The effect of learning on the slope of each term structure is a priori ambiguous and depends on two distinct economic channels. On the one hand, as discussed in Section 6.3.2, agents perceive a larger *share of the forecast error* variance to be associated with permanent shocks. Given the behavior of slopes under full-information, this should decrease (increase) the slope of the real-yield (equity) term structure, all else equal. On the other hand, learning magnifies the price of risk for business-cycle shocks but induces the opposite effect for trend shocks, as can be seen in Figure 3a. These changes in the *relative prices of risk* imply that learning should induce the opposite effect, making the real-yield (equity) term structure less downward (upward) sloping.

In the **endowment economy**, where consumption is exogenous, the relative prices of risk channel dominates: the learning-implied yield (equity) term structure is less downward (upward) sloping, compared to full-information.

In the **production environment**, where consumption is endogenous, the perceived share of the forecast error channel dominates, and the effect of learning is the *opposite* of what we observe in the endowment economy.

The difference between the two settings arises because learning also impacts the riskiness of future investment and output claims (in a manner similar to exogenous consumption). In our parametrization, the quantitative effect of learning on the former is larger. Since endogenous consumption is inversely related to investment (all else equal), the impact of learning on the term structure of endogenous consumption claims flips. Because the prices of consumption claims are tightly connected to the pricing of real yields and dividend claims in equilibrium, the learning-implied yield (equity) term structure is more downward (upward) sloping. As we discuss in the appendix, this helps the model to reconcile the equity term structure's *unconditional* slope as estimated by Bansal et al. (2021).

6.4 Cyclicality under Imperfect-Information: Model vs Data

We demonstrate that the risk free rate and the price-to-consumption ratio are countercyclical (procyclical) under full (imperfect) information.

Simulation-Based Impulse-Responses. Figures 3e and 3f show the impact of business-cycle and trend shocks on the risk free rate. Under full-information, a positive trend (business-cycle) shock increases (decreases) expected consumption growth (see Figures 2g and 2h), and therefore raises (drops) the risk free rate given that the IES is greater than one. Consequently, the cyclicality of the risk free rate is ambiguous when the agent is fully informed. The cyclicality of the price-consumption ratio is also indeterminate. In equilibrium, the ratio is proportional to the utility-to-consumption ratio, where the constant of proportionality is $(1 - \frac{1}{\psi})$. Because it is also driven by expected consumption growth, this valuation ratio increases (decreases) in response to a trend (business-cycle) shock.

Imperfect-information implies that both the risk free rate and the valuation ratio are unambiguously procyclical. The learning-induced trendy business-cycle shock alters how the risk free rate and the price-consumption ratio respond to a business-cycle innovation. In particular, both impulse response functions become positive, as occurs following a true trend shock under perfect information. This procyclicality is in-line with Figure 2g, showing that with learning, consumption growth expectation increases with the business-cycle. On the other hand, learning significantly attenuates the responses of the risk free rate and the price-consumption ratio to a trend shock. These quantities become almost acyclical in relation to the permanent productivity shocks, though the impulse responses remain positive. Given this attenuation, it is evident that prices' procyclicality is driven almost exclusively by business-cycle fluctuations in the imperfect-information framework.

VAR-Based Impulse-Responses. We examine the model-implied cyclicality of prices through the lens of an econometrician who does not observe the underlying productivity components, similar to the agent in the model. The exercise also allows us to determine the cyclicality of prices for an uninformed econometrician, given paths simulated from a model with a perfectly informed agent. This is useful, since the cyclicality of the risk-free rate and the price-consumption ratio from the view of the perfectly informed agent is ambiguous, as suggested by the former discussion. The methodology also admits an empirical equivalent.

We simulate a population path from the imperfect-information model letting the underlying business-cycle and trend shocks fluctuate unrestricted. Let y be the vector of model-implied paths of interest. We set $y = [\Delta a_t, \Delta c_t, r_{f,t}, P/C_t]'$, where Δa_t is log productivity growth, Δc_t is log consumption growth, $r_{f,t}$ is the log risk free rate and P/C_t is the price-to-consumption ratio. We then estimate a vector-auto-regression (VAR) process:

$$Y_t = T_{0,4\times1} + T_{4\times4}Y_{t-1} + \eta_{t,4\times1}.$$

We assume that the econometrician knows that Δa_t is exogenous to the other variables, and thus, restricts the first row of T to [1, 0, 0, 0], suggesting that Δa_t is only affected by its own lag. We then compute the impulse-responses from a one standard deviation Choleskly shock to productivity growth into all variables. We repeat the same exercise using (i) paths obtained from a perfect information model (keeping the underlying population paths of the shocks $\varepsilon_{z,t}$ and $\varepsilon_{x,t}$ unchanged), and (ii) empirical paths. The empirical productivity path is obtained from Basu, Fernald, and Kimball (2006) and Fernald (2012). Real per-capita consumption growth path is obtained from BEA NIPA tables. We proxy for the risk free rate using the yield to maturity on a 3-month T-bill, net of inflation. We proxy for the price-to-consumption ratio using Shiller's cyclically adjusted price-to-earning ratio. The results are shown in Figure 4, along side the empirical impulse-responses. All plotted impulse-responses are cumulative (i.e., period t shows $\sum_{k=0}^{t} IRF_{a,j,k}$, where $IRF_{a,j,k}$ is the response of variable j to a one-standard deviation shock to Δa after k periods). Notably, cumulative impulse responses are similar to the effect of a productivity shock on the future *level* of a variable of interest.

A one standard deviation shock to productivity growth shifts productivity immediately to its new higher steady state level. By construction, this impulse response is identical when based on imperfect- or full-information model paths. Interestingly however, a positive productivity shock increases consumption to a permanently higher level under imperfect-information, whereas the same shock increases consumption in the short-run, but steadily drops consumption afterwards. The latter occurs as business-cycle shocks induce a non-linear mean-reverting effect on equilibrium consumption, and quantitatively dominate trend shocks. The empirical impulse response from productivity to consumption resembles qualitatively (and also quantitatively) the impulse-response implied by the imperfect-information model. The response to consumption in the data also settles at a new permanently higher level.

Figures 4c and 4d show that a productivity shock increases the risk-free rate and the price-consumption ratio when using imperfect-information model paths. This result echoes figures 3e, 3f, 3g, and 3h, showing that with learning both business-cycle shocks and trend shock raise these variables. The imperfect-information impulse-response to the risk-free rate also matches its empirical equivalent. In the data, the price-to-earnings ratio increases following a positive productivity shock. Qualitatively, this is similar to the imperfect-information model, but the function is more modest than the increase in the model-implied P/C ratio. An exact quantitative match between the model and the data is not expected, given that the true P/C ratio is unobserved.

Importantly, using full-information model paths the econometrician infers that the risk-free rate is countercyclical - in contrast to the data: it steadily drops in response to higher productivity. Business-cycle shocks are more sizable, and their negative impact on the risk-free rate dominates the positive impact of trend shocks. Based on full-information paths, the price-consumption ratio is materially acyclical to the econometrician - which is empirically inconsistent. Although transitory shocks are larger, equity prices respond more strongly to permanent shocks, and the two opposite effects of these shocks offset each other.

6.5 Robustness

Section IA.2 of the Internet Appendix shows that the asset pricing implications of the model (and in particular, the importance of business cycle shocks relative to trend shocks for the equity premium) are materially unchanged when the model parameters are perturbed. Importantly, we show that our results are almost identical when there is no asymmetry between the persistence parameter of the business-cycle and trend stats. We also show that the imperfect- (full-information) equity premium is positive (negative), even if the IES is below one. In Internet Appendix IA.3 we obtain similar findings when we use the parameter values obtained from Kalman filtering of the TFP time-series estimated by Basu, Fernald, and Kimball (2006) for the TFP process.

7 Learning Implications for Empirical Paths

In what follows, we examine the model-implied beliefs with respect to both the business-cycle and trend component of productivity, with the aim of providing a new perspective on the empirical evolution of investment and consumption at particular points of interest (e.g., the Great Recession and the COVID-19 pandemic.)

We obtain quarterly utilization-adjusted TFP growth data from the San Francisco Fed, following the methodology of Basu, Fernald, and Kimball (2006) and Fernald (2012), as a proxy for ΔA. We start the simulation when all variables, including beliefs, are at their steady state value. We then simulate model-implied paths, ensuring that the model-implied path for A_t is identical to the empirical counterpart. This exercise does not necessitate simulating the underlying shocks $\varepsilon_{z,t}$ and $\varepsilon_{x,t}$: the agent learns about the state of z_t and x_t exclusively from ΔA_t.

In Figures 5b and 5a we plot standardized model-implied paths for (log) consumption growth Δc and investment growth ΔI against the empirical paths of these variables. Real per-capital consumption growth (nondurables and services) and investment growth (nonresidential) are obtained from BEA NIPA Tables. The correlation between consumption growth in the data and in the model is 0.45, with a similar correlation for investment growth. The fluctuations in the model-implied paths of these variables show a gradual attenuation after the mid-1980s, consistent with Stock and Watson (2002). Figures 5c and 5d show the evolution of beliefs about the trend component $\hat{x}_{t|t}$ and the business-cycle component $\hat{z}_{t|t}$, respectively. The path of $\hat{x}$ shows greater persistence, but smaller volatility in comparison to $\hat{z}$, consistent with the model's estimated parameters.

Several studies have pointed out that the the recovery from the Great Recession seemed slower than expected. One way of illustrating this is found in Figure 6. The black line shows the evolution of (log) real investment before and after the great recession, and the red dashed line shows the constant (log-linear) trend which best fits investment data prior to 2007Q4. The dynamics of investment are puzzling under the belief that the crisis was caused by a (sequence of) negative business-cycle shock(s), $\varepsilon_z < 0$. To see this, note that the dashed blue line depicts the hypothetical path of investment obtained from the full-information model, when productivity is perturbed by a sequence of negative but transitory business-cycle shocks that match the observed drop in investment during the crisis. Under this scenario, and in the absence of further shocks, investment should have converged to the pre-crisis log-linear trend about two years after the trough. In contrast, post 2009Q2 investment rises (almost) in parallel to the constant trend red line, and does not converge to its pre-crisis trend.

While the existing literature has provided compelling theories for the sluggish path of investment (see, e.g., Reinhart and Rogoff 2014; Kozlowski, Veldkamp, and Venkateswaran 2017), we believe that our model can provides a complementary perspective. Specifically, while the recovery dynamics are puzzling *ex-ante* to an agent who believes that most observed fluctuations are driven by business-cycle fluctuations, from the lens of our framework, the recovery dynamics are *ex-post* consistent with the belief that the Great Recession involved primarily negative trend shocks. Consistent with a prior belief that observed fluctuations mostly operate at the business-cycle frequency, we find that as investment fell from 2007Q4 to 2009Q2, the agent ascribed a larger likelihood that the underlying productivity shocks are temporary: in Figure 5d, $\hat{z}$ falls precipitously during this period. However, by the end of 2009, the agent's belief, $\hat{z}$, recovered to a near zero value. In contrast, agent's belief about the trend, $\hat{x}$, shows a small recovery starting at 2009Q1 but it stalls in 2011Q1, remaining firmly lodged in negative territory until the pandemic. The full recovery of $\hat{z}$, alongside the post-2011 dynamics of

$\hat{x}$, are consistent with a posterior belief that the recession was driven (to a large extent) by permanent shifts in the stochastic trend of productivity.[22]

These dynamics stand in contrast to the evolution of $\hat{x}$ in response to other recessions in our sample. In Figure 5c, we denote NBER-dated recessions with the shaded regions. We note that, even though there were secular trends in $\hat{x}$, after each recession beliefs about the trend component had largely rebounded to their pre-recession level. As a result, the accompanying recovery was (relatively) rapid, consistent with the belief that these were largely driven by transitory shocks. A salient example is the latest COVID crisis. Figures 5d and 5c suggest that the pandemic-induced recession was mostly a business-cycle phenomenon. While initially beliefs about both the business-cycle state and the trend state declined, $\hat{z}$ fells substantially more than $\hat{x}$, and both beliefs quickly regained their pre-crisis level, consistent with a V-shaped recovery.

8 Conclusion

We study a framework in which households and firms understand the parameters of the underlying data-generating process, but must form beliefs about two types of shocks, trend (permanent) and business-cycle (transitory), by conditioning only on aggregate productivity. As a result, agents are never fully sure which shock is driving growth, which leads to rational confusion – each shock partially inherits the features of its counterpart. Consequently, the implication of learning depends not only upon the nature of each shock in isolation, but also is crucially determined by their interaction.

We show that rational confusion plays a crucial role in generating a sizable & positive equity premium, a sufficiently low risk-free rate, and pro-cyclical valuation ratios and investment growth, consistent with the data. The interaction between beliefs about business-cycle and trend shocks implies that a larger fraction of agents' uncertainty is perceived to arise from permanent shocks, amplifying the risk premium relative to full information. Moreover, business-cycle shocks play a key

[22]Based on the the dynamics of $\hat{x}$, the model implies that the belief about the underlying shock which drove the Great Recession changed around mid-2011. Interestingly, during this period there is also a notable change in the growth rate of investment. Figure 6 shows the best linear fit of investment pre- (green line) and post- (yellow line) 2012. Investment seems to have grown at a faster rate after the recession but before 2012, while the belief about the permanent component of productivity was recovering, but grows at a slower rate (in comparison to both the green and the red line) post-2012. Therefore, consistent with the implications of our framework, the altered belief about productivity dynamics directly translated into a change in equilibrium investment.

role in explaining asset prices with learning – quantitatively they are the primary source of endogenous long-run consumption risk, via "trendy business-cycles". Our mechanism provides a possible bridge between the importance of (perceived) permanent shocks to the SDF (e.g., Alvarez and Jermann 2005) versus the dominance of transitory shocks in the data (e.g., King et al. 1991; Cochrane 1988).

We believe that further consideration of the learning problem presented in this paper may be a fruitful approach in other areas including models of endogenous growth, capital structure, and the impact of differences in cash-flow duration for cross-sectional asset pricing.

References

Aguiar, M., and G. Gopinath. 2007. Emerging market business cycles: The cycle is the trend. *Journal of political Economy* 115:69–102.

Ai, H. 2010. Information quality and long-run risk: asset pricing implications. *The Journal of Finance* 65:1333–67.

Ai, H., M. M. Croce, A. M. Diercks, and K. Li. 2018. News shocks and the production-based term structure of equity returns. *The Review of Financial Studies* 31:2423–67.

Alvarez, F., and U. J. Jermann. 2005. Using asset prices to measure the persistence of the marginal utility of wealth. *Econometrica* 73:1977–2016.

Andrei, D., B. Carlin, and M. Hasler. 2019. Asset pricing with disagreement and uncertainty about the length of business cycles. *Management Science* 65:2900–23.

Andrei, D., M. Hasler, and A. Jeanneret. 2019. Asset pricing with persistence risk. *The Review of Financial Studies* 32:2809–49.

Babiak, M., and R. Kozhan. 2020. Parameter learning in production economies. *Working Paper. Lancaster University and Warwick Business School* .

Bansal, R., S. Miller, D. Song, and A. Yaron. 2021. The term structure of equity risk premia. *Journal of Financial Economics* 142:1209–28.

Bansal, R., and I. Shaliastovich. 2010. Confidence risk and asset prices. *American Economic Review* 100:537–41.

Bansal, R., and A. Yaron. 2004. Risks for the long run: A potential resolution of asset pricing puzzles. *The journal of Finance* 59:1481–509.

Barro, R. J. 2006. Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics* 121:823–66.

Barro, R. J., and R. G. King. 1984. Time-separable preferences and intertemporal-substitution models of business cycles. *The Quarterly Journal of Economics* 99:817–39.

Barsky, R. B., and E. R. Sims. 2011. News shocks and business cycles. *Journal of monetary Economics* 58:273–89.

Basu, S., J. G. Fernald, and M. S. Kimball. 2006. Are technology improvements contractionary? *American Economic Review* 96:1418–48.

Beaudry, P., and F. Portier. 2004. An exploration into pigou's theory of cycles. *Journal of monetary Economics* 51:1183–216.

Berk, J. B., R. C. Green, and V. Naik. 1999. Optimal investment, growth options, and security returns. *The Journal of finance* 54:1553–607.

Bidder, R., and I. Dew-Becker. 2016. Long-run risk is the worst-case scenario. *American Economic Review* 106:2494–527.

Blanchard, O. J., J.-P. L'Huillier, and G. Lorenzoni. 2013. News, noise, and fluctuations: An empirical exploration. *American Economic Review* 103:3045–70.

Boguth, O., and L.-A. Kuehn. 2013. Consumption volatility risk. *The Journal of Finance* 68:2589–615.

Boz, E., C. Daude, C. B. Durdu, et al. 2008. *Emerging market business cycles revisited: Learning about the trend*. 927. Board of Governors of the Federal Reserve System.

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

Calvet, L. E., J. Y. Campbell, F. Gomes, and P. Sodini. 2021. The cross-section of household preferences. *Working Paper.* .

Campbell, J. Y., and J. H. Cochrane. 1999. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of political Economy* 107:205–51.

Christiano, L. J., M. Eichenbaum, and C. L. Evans. 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of political Economy* 113:1–45.

Cochrane, J. H. 1988. How big is the random walk in gnp? *Journal of political economy* 96:893–920.

Cogley, T., and T. J. Sargent. 2008. Anticipated utility and rational expectations as approximations of bayesian decision making. *International Economic Review* 49:185–221.

Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer. 2016. Parameter learning in general equilibrium: The asset pricing implications. *American Economic Review* 106:664–98.

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

Edge, R. M., T. Laubach, and J. C. Williams. 2007. Learning and shifts in long-run productivity growth. *Journal of Monetary Economics* 54:2421–38.

Epstein, L. G., and S. E. Zin. 1991. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis. *Journal of political Economy* 99:263–86.

Evans, M. D. 1998. Real rates, expected inflation, and inflation risk premia. *The Journal of Finance* 53:187–218.

Fernald, J. 2012. A quarterly, utilization-adjusted series on total factor productivity. *Working paper. Federal Reserve Bank of San Francisco* .

Gali, J. 1999. Technology, employment, and the business cycle: do technology shocks explain aggregate fluctuations? *American economic review* 89:249–71.

García-Feijóo, L., and R. D. Jorgensen. 2010. Can operating leverage be the cause of the value premium? *Financial Management* 39:1127–54.

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

Jahan-Parvar, M. R., and H. Liu. 2014. Ambiguity aversion and asset prices in production economies. *The Review of Financial Studies* 27:3060–97.

Jaimovich, N., and S. Rebelo. 2009a. Can news about the future drive the business cycle? *American Economic Review* 99:1097–118.

———. 2009b. Can news about the future drive the business cycle? *American Economic Review* 99:1097–118.

Jermann, U. J. 1998. Asset pricing in production economies. *Journal of monetary Economics* 41:257–75.

Johannes, M., L. A. Lochstoer, and Y. Mou. 2016. Learning about consumption dynamics. *The Journal of finance* 71:551–600.

Kaltenbrunner, G., and L. A. Lochstoer. 2010. Long-run risk through consumption smoothing. *The Review of Financial Studies* 23:3190–224.

King, R. G., C. I. Plosser, J. H. Stock, and M. W. Watson. 1991. Stochastic trends and economic fluctuations. *American Economic Review* 81:819—-840.

Kozlowski, J., L. Veldkamp, and V. Venkateswaran. 2017. The tail that wags the economy: Beliefs and persistent stagnation .

Lettau, M., S. C. Ludvigson, and J. A. Wachter. 2008. The declining equity premium: What role does macroeconomic risk play? *The Review of Financial Studies* 21:1653–87.

Lustig, H., S. Van Nieuwerburgh, and A. Verdelhan. 2013. The wealth-consumption ratio. *The Review of Asset Pricing Studies* 3:38–94.

Moore, B., and H. Schaller. 2002. Persistent and transitory shocks, learning, and investment dynamics. *Journal of Money, Credit and Banking* 650–77.

Nelson, C. R., and A. F. Siegel. 1987. Parsimonious modeling of yield curves. *Journal of business* 473–89.

Reinhart, C. M., and K. S. Rogoff. 2014. Recovery from financial crises: Evidence from 100 episodes. *American Economic Review* 104:50–5.

Segal, G., and I. Shaliastovich. 2021. Uncertainty, risk, and capital growth. *Working paper. UNC* .

Stock, J. H., and M. W. Watson. 2002. Has the business cycle changed and why? *NBER macroeconomics annual* 17:159–218.

Tallarini, T. D. 2000. Risk-sensitive real business cycles. *Journal of monetary Economics* 45:507–32.

Van Binsbergen, J., W. Hueskes, R. Koijen, and E. Vrugt. 2013. Equity yields. *Journal of Financial Economics* 110:503–19.

Van Binsbergen, J. H., and R. S. Koijen. 2017. The term structure of returns: Facts and theory. *Journal of Financial Economics* 124:1–21.

Van Nieuwerburgh, S., and L. Veldkamp. 2006. Learning asymmetries in real business cycles. *Journal of monetary Economics* 53:753–72.

Verdelhan, A. 2010. A habit-based explanation of the exchange rate risk premium. *The Journal of Finance* 65:123–46.

Zhang, L. 2005. The value premium. *The Journal of Finance* 60:67–103.

Tables and Figures

Table 1: **Model parameters**

Parameter	Value	Source
Production:		
α	0.36	Capital share of output
δ	0.021	Annual depreciation rate from BEA ($0.021 \times 4 = 8.4\%$)
ζ	7.7	Consumption growth's volatility to output growth's volatility, consistent with Kaltenbrunner and Lochstoer (2010)
Preferences:		
γ	10	Bansal and Yaron (2004)
ψ	1.5	Bansal and Yaron (2004)
β	0.991	Estimated
	[0.9894, 0.9920]	
Technology:		
μ	0.0049	Estimated
	[0.0018, 0.0080]	
σ_z	0.0168	Estimated
	[0.0153, 0.0184]	
ρ_z	0.8783	Estimated
	[0.8080, 0.9253]	
σ_x	$8.41 \cdot 10^{-4}$	Estimated
	[0.0005, 0.0012]	
ρ_x	0.9962	Estimated
	[0.9884, 0.9988]	

The table reports the model parameters in the benchmark case. The source indicates the study from which each parameter is obtained or the empirical statistic that it targets. The parameters whose source is "Estimated" are obtained via SMM, and brackets show 90% confidence intervals.

Table 2: **Macroeconomic Moments: Model versus Data**

	Data		Model: Learning		Model: Full Information	
Panel A: Unconditional growth moments.						
Consumption growth:						
mean	1.93	[1.59, 2.28]	1.96	[-0.42, 5.02]	1.96	[-0.52, 5.07]
stdev	1.27	[1.07, 1.57]	1.51	[1.19, 2.24]	1.38	[1.00, 2.16]
AC(1)	0.52	[0.30, 0.74]	0.44	[0.08, 0.75]	0.54	[0.10, 0.83]
Output growth:						
mean	1.94	[1.41, 2.48]	1.96	[-0.66, 5.19]	1.96	[-0.73, 5.21]
stdev	1.98	[1.67, 2.44]	1.98	[1.61, 2.64]	2.00	[1.62, 2.69]
AC(1)	0.30	[0.08, 0.52]	0.33	[0.01, 0.65]	0.34	[0.02, 0.66]
Investment growth:						
mean	3.61	[2.00, 5.22]	1.91	[-2.08, 5.89]	1.91	[-1.82, 5.76]
stdev	5.96	[5.02, 7.34]	4.02	[3.20, 5.09]	4.68	[3.70, 5.64]
AC(1)	0.29	[0.11, 0.48]	0.21	[-0.06, 0.49]	0.17	[-0.07, 0.43]
Panel B: Variance ratios.						
Consumption growth:						
VR(2)	1.52	[1.26, 1.79]	1.43	[1.07, 1.75]	1.53	[1.10, 1.81]
VR(3)	1.80	[1.41, 2.20]	1.74	[1.03, 2.40]	1.94	[1.13, 2.55]
Output growth:						
VR(2)	1.31	[1.04, 1.57]	1.32	[1.01, 1.65]	1.34	[1.01, 1.66]
VR(3)	1.33	[0.94, 1.73]	1.52	[0.96, 2.21]	1.54	[0.97, 2.23]

The table shows model-implied moments for annual consumption growth, output growth and investment growth against empirical counterparts. Panel A shows unconditional growth moments and Panel B shows variance ratio moments. The model-implied moments are computed under learning and under full-information.

Table 3: **Prices Moments: Model versus Data**

	Data		Model: Learning		Model: Full Information	
$E[R_m - R_f](\%)$	4.80	[0.18, 9.42]	2.62	[1.26, 3.70]	-0.54	[-1.61, 0.23]
$SR(R_m - R_f)$	0.31	[0.06, 0.57]	0.55	[0.25, 0.79]	-0.11	[-0.32, 0.05]
$E[R_f](\%)$	0.69	[0.23, 1.15]	0.67	[-0.73, 2.41]	2.20	[0.60, 4.21]
$\sigma[R_f](\%)$	1.71	[1.44, 2.10]	0.65	[0.28, 1.22]	0.72	[0.37, 1.32]

The table shows model-implied annual moments for the market excess return $R_m - R_f$ and the risk free rate R_f, against empirical counterparts. SR denotes Sharpe ratio. σ denotes standard deviation. The model-implied moments are computed under learning and under full-information.

Table 4: **Endogenous Long-Run Consumption Risk**

	Full-Information	Imperfect-Information
Panel A. Contribution to equity premium		
x	144.77%	0.33%
z	-44.77%	99.67%
Panel B. Innovations' correlation		
$corr(\eta_t, \varepsilon_x)$	0.51	0.06
$corr(\eta_t, \varepsilon_z)$	-0.83	0.78
Panel C. AR(1) of expected consumption growth		
ρ	0.990	0.999

The table shows the properties of endogenous long run consumption risk within the full-information and imperfect-information-models. For each model, we simulate a population path of the exogenous shocks ε_z and ε_x and use these to obtain a sample of expected consumption growth path, $e_t = E_t[\Delta c_{t+1}]$, as implied by the model's solution. We then fit e into an AR(1) process: $e_{t+1} = const + \rho e_t + \eta_t$, where η_t is expected consumption growth shock. In Panel A we report the fraction of Eq. (33). In Panel B we report how the shocks η_t correlate with contemporaneous business-cycle shocks and trend shocks. In Panel C we report the population value of ρ.

Figure 1: **Macroeconomic Dynamics Under Learning: Risk Neutrality**

Impulse-Responses: Transitory Shock (ε_z) **Impulse-Responses: Permanent Shock (ε_x)**

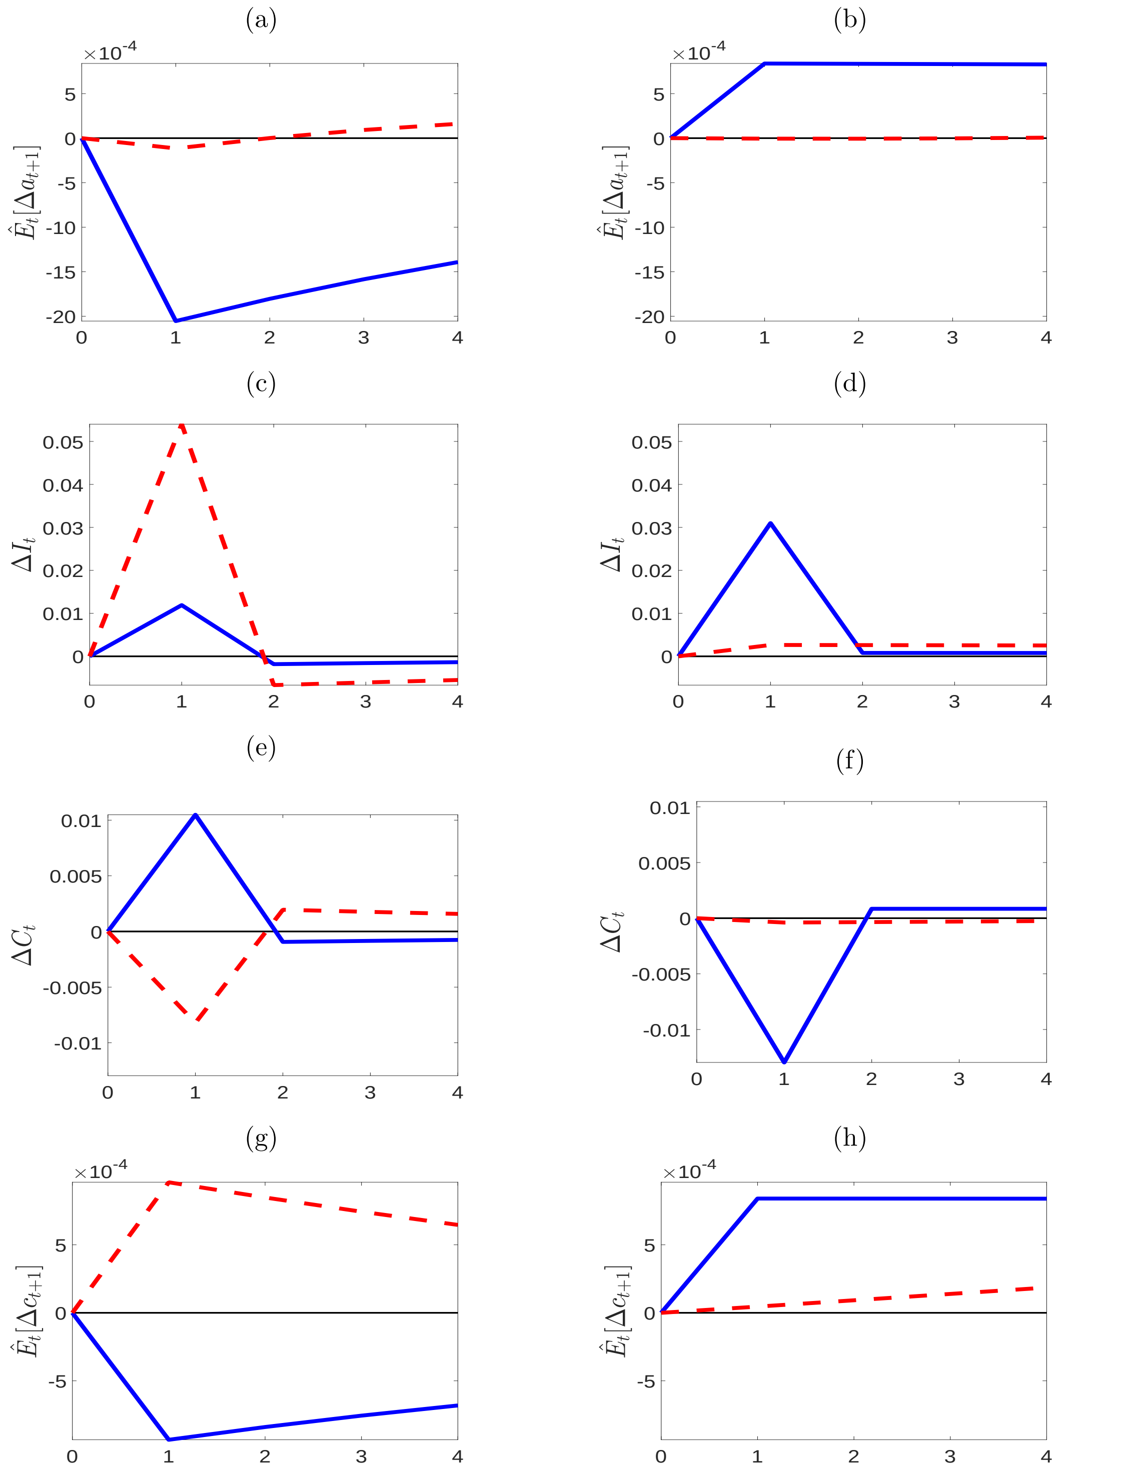

The figure shows impulse responses of transitory and permanent shocks to macroeconomic variables under risk neutrality. The solid blue line shows impulse-responses under full-information. The dashed red line shows impulse-responses under learning. Horizontal axes represent quarters.

Figure 2: **Macroeconomic Dynamics Under Learning: Risk Aversion**

Impulse-Responses: Transitory Shock (ε_z) **Impulse-Responses: Permanent Shock (ε_x)**

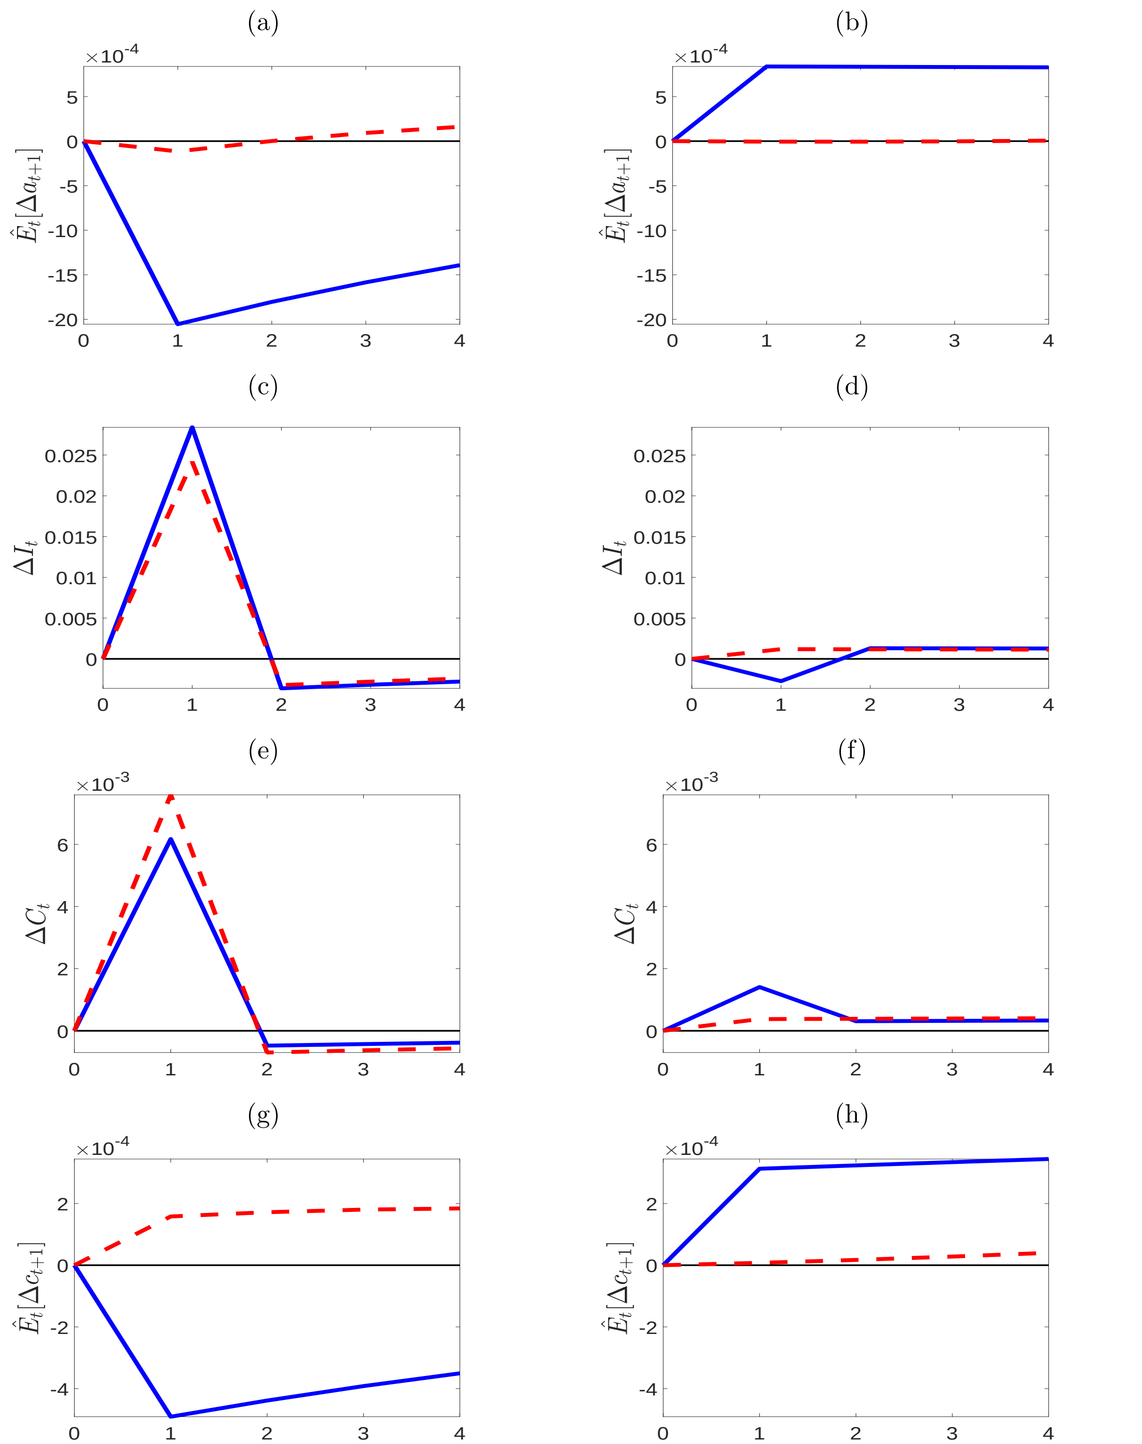

The figure shows impulse responses of transitory and permanent shocks to macroeconomic variables under risk aversion. The solid blue line shows impulse-responses under full-information. The dashed red line shows impulse-responses under learning. Horizontal axes represent quarters.

Figure 3: **Price Dynamics Under Learning**

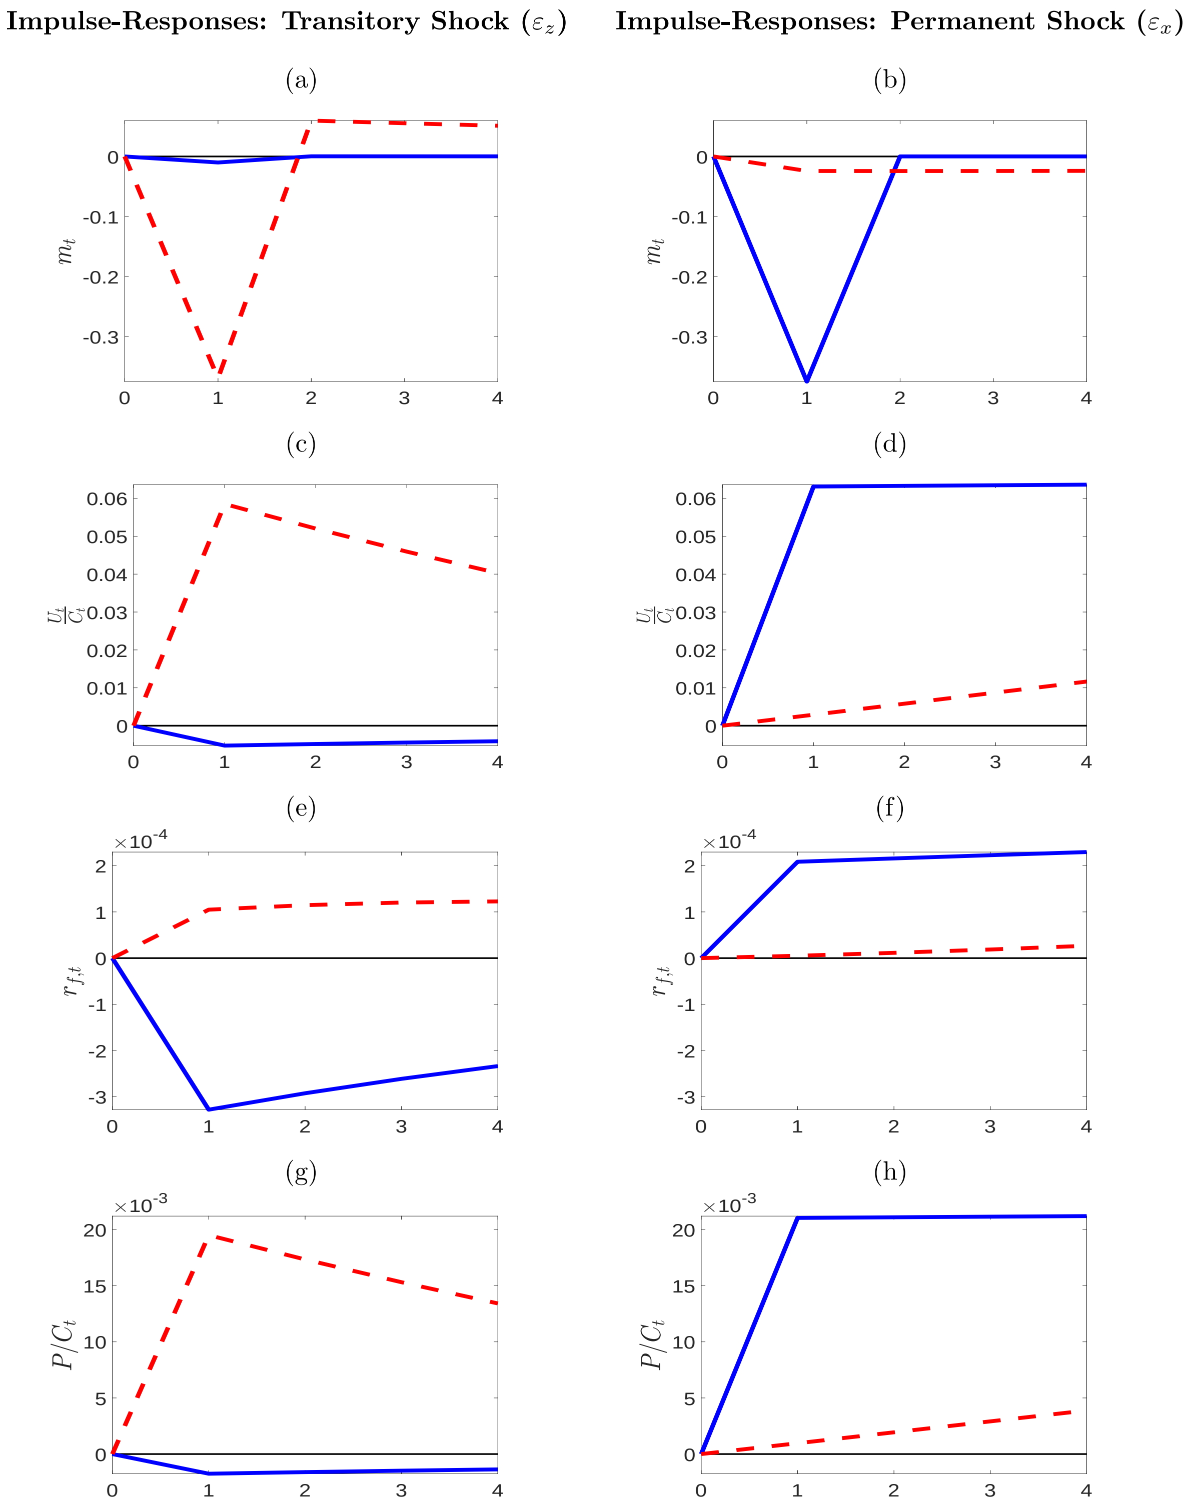

The figure shows impulse responses of transitory and permanent shocks to financial (prices) variables under risk aversion. The solid blue line shows impulse-responses under full-information. The dashed red line shows impulse-responses under learning. Horizontal axes represent quarters.

Figure 4: **Cumulative Impulse-Responses from a VAR(1) specification**

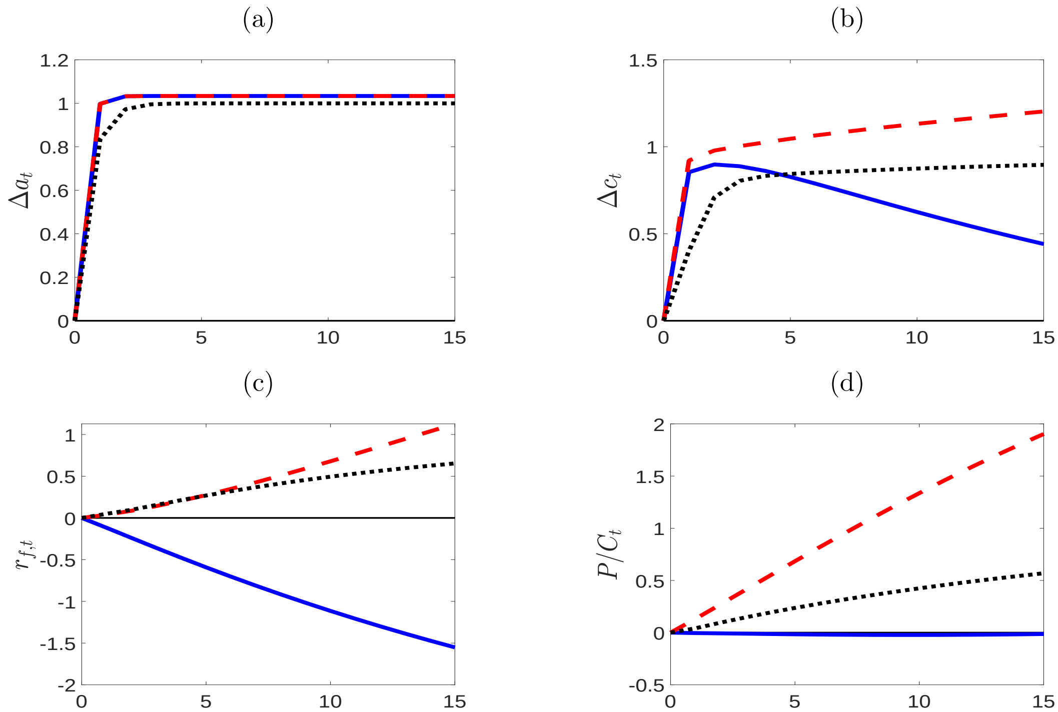

The figure shows impulse responses of one standard deviation Cholesky shock to productivity to consumption growth, the risk free rate, and price-consumption ratio. The impulse-responses are computed by estimating a VAR(1) that includes productivity growth, consumption growth, r_f, and P/C in that order. The plotted impulse response are cumulative (i.e., period t shows $\sum_{k=0}^{t} IRF_{a,j,k}$, where $IRF_{a,j,k}$ is the response of variable j to a one-standard deviation shock to Δa after k periods). The solid blue line shows impulse-responses under full-information. The dashed red line shows impulse-responses under learning. The dotted black line shows the empirical counterpart. All impulse-response functions are standardized by the respective standard deviations of the vertical-axis variables. Horizontal axes represent quarters.

Figure 5: **Model-Implied Paths Given Empirical Productivity Growth Rates**

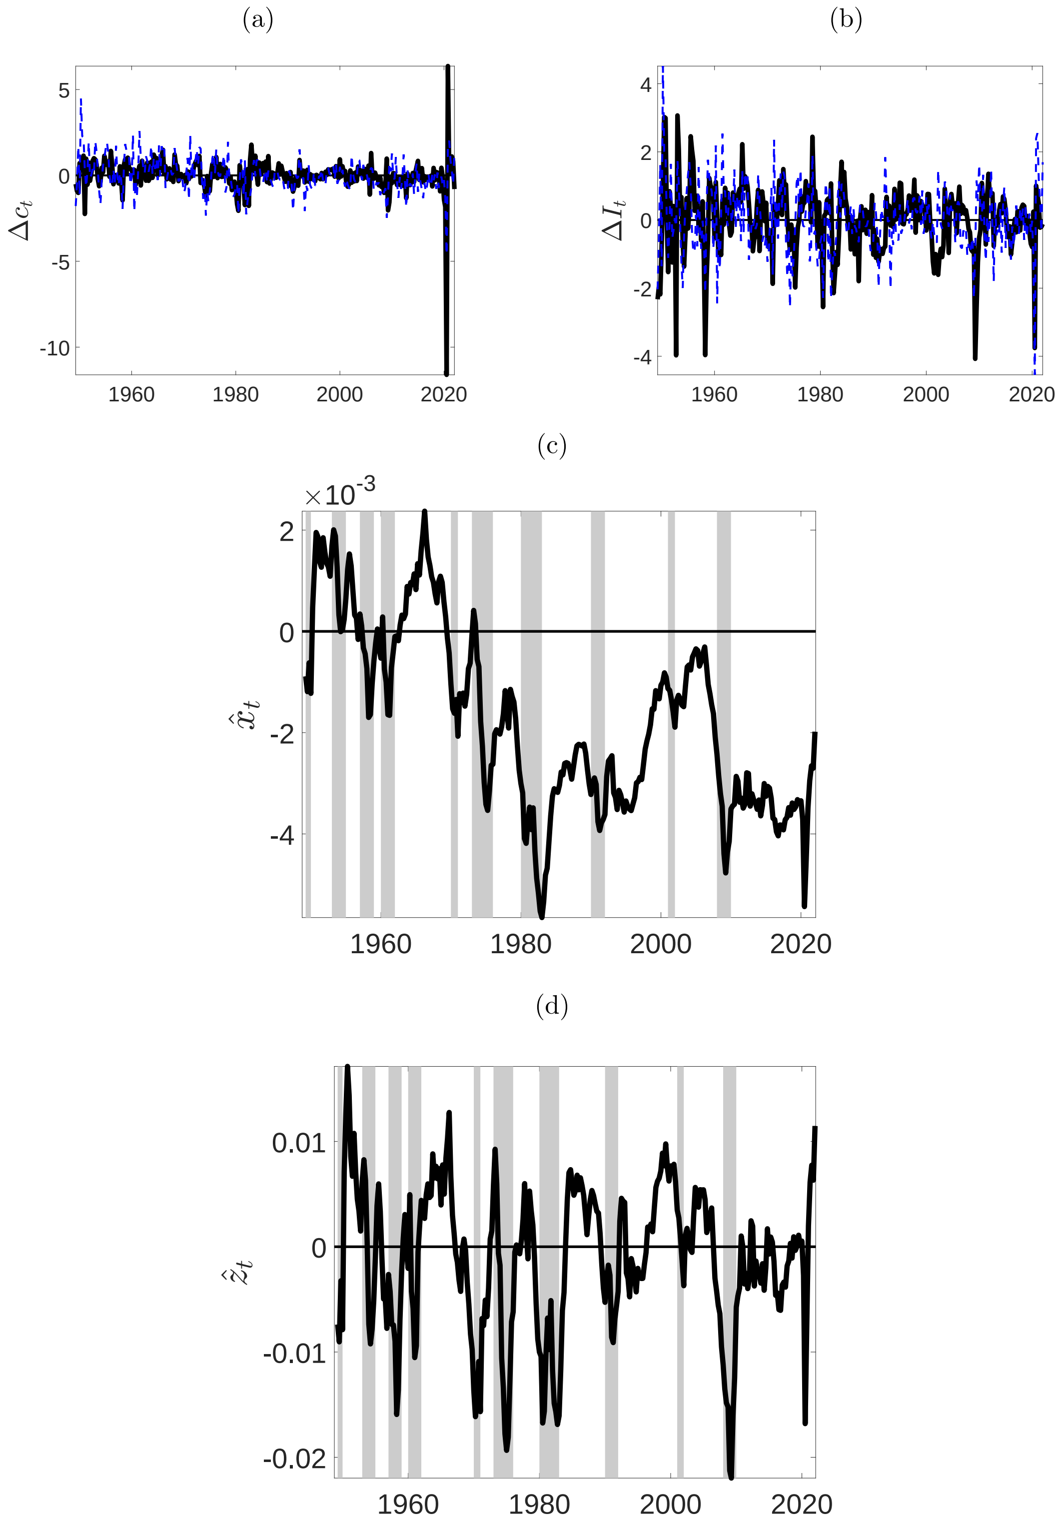

Panels (a) and (b) show the standardized model-implied paths (in blue) of consumption growth Δc and investment growth ΔI given the empirical path of productivity growth rates from 1947-2021, as constructed by Fernald (2012). The corresponding empirical paths are also shown (in black). Panels (c) and (d) shows the model-implied belief for the underlying permanent productivity growth state $\hat{x}$ and the belief for the underlying transitory productivity state $\hat{z}$ given the empirical path of productivity growth rates.

Figure 6: **Investment Around The Great Recession**

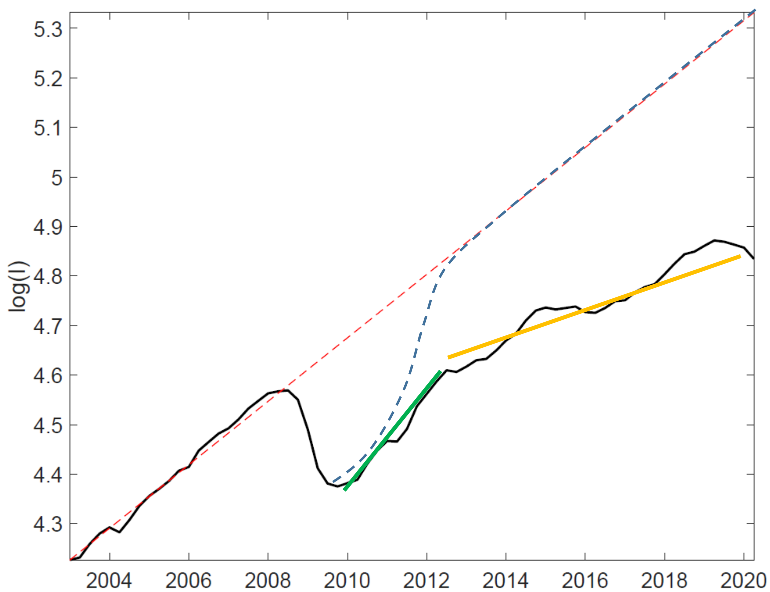

The black line shows the time series of $\log(I)$, where I is the real log level of non-residential investment from NIPA. The red dashed line shows a log-linear trend fitted for the period of 2003Q1-2007Q4. The blue dashed line shows the expected path of I, assuming the drop in investment between 2007Q4-2009Q2 was generated by a sequence of negative business-cycle shocks. The green (yellow) line shows the linear fit for $\log(I)$ between 2009Q2-2012Q1 (2012Q2-2020Q1).

Appendix A Endowment Economy

Appendix-A.1 Model Setup and Solution

Specification. In this section we examine the role of imperfect information in an endowment economy with both trend and business-cycle shocks. Specifically, we assume that

$$\Delta c_{t+1} = \mu + x_{t+1} + z_{t+1} - z_t \tag{A1}$$

$$x_{t+1} = \rho_x x_t + \sigma_x \varepsilon_{x,t+1} \tag{A2}$$

$$z_{t+1} = \rho_z z_t + \sigma_z \varepsilon_{z,t+1} \tag{A3}$$

where $\varepsilon_{x,t+1}, \varepsilon_{z,t+1} \sim N(0,1)$ are independent of one another and over time. This implies that demeaned consumption growth can be written as

$$\Delta c_{t+1} - \mu = \rho_x x_t + (\rho_z - 1) z_t + \sigma_x \varepsilon_{x,t+1} + \sigma_z \varepsilon_{z,t+1}, \tag{A4}$$

so that agents' forecast of consumption growth in the next period is given by

$$\Delta \hat{c}_{t+1|t} = \mu + \rho_x \hat{x}_{t|t} + (\rho_z - 1) \hat{z}_{t|t}. \tag{A5}$$

Under full information, agents observe the realization of each shock independently and update utilizing equations (A2) and (A3), i.e., $\hat{x}_{t|t} = x_t$ and $\hat{z}_{t|t} = z_t$. Under imperfect information, however, agents update their beliefs utilizing the forecast error,

$$v_t = c_t - \hat{c}_{t|t-1} \qquad v_t \sim N\left(0, \sigma_v^2\right). \tag{A6}$$

Specifically, with learning,

$$\hat{x}_{t|t} = \rho_x \hat{x}_{t-1|t-1} + g_x v_t \tag{A7}$$

$$\hat{z}_{t|t} = \rho_z \hat{z}_{t-1|t-1} + g_z v_t. \tag{A8}$$

The variance, σ_v^2, as well as the gains, g_x and g_z are pinned down by the steady state solution of the Ricatti equation, which is the same as in our benchmark analysis, replacing only Δa_t with Δc_t.

Given the assumption of Epstein and Zin (1991) preferences, we can write the log-SDF as

$$m_{t+1} = \theta \log \delta - \tfrac{\theta}{\psi} \Delta c_{t+1} + (\theta - 1) r_{t+1}, \tag{A9}$$

and, letting $r_{c,t+1}$ denote the log return on the consumption claim, it must be the case that

$$\mathbb{E}_t \left[M_{t+1} R_{c,t+1} \right] = \mathbb{E}\left[\exp\left(\theta \log \delta - \tfrac{\theta}{\psi} \Delta c_{t+1} + \theta r_{c,t+1} \right) \right] = 1. \tag{A10}$$

Price-Consumption Ratio. We conjecture that the endogenous log price-consumption follows

$$PC_t = A_0 + A_x x_{t|t} + A_z z_{t|t}, \tag{A11}$$

and then use a Campbell-Shiller approximation of the log return on the consumption claim

$$r_{c,t+1} = \kappa_0 + \kappa_1 PC_{t+1} - PC_t + \Delta c_{t+1}, \tag{A12}$$

where κ_0 and κ_1 are log-linearization parameters[23]. Under full information, this yields

$$r_{c,t+1} = \kappa_0 + A_0 (\kappa_1 - 1) + A_x x_{t|t} (\kappa_1 \rho_x - 1) + A_z z_{t|t} (\kappa_1 \rho_z - 1) \tag{A13}$$

$$+ \Delta c_{t+1} + \kappa_1 A_x \sigma_x \varepsilon_{x,t+1} + \kappa_1 A_z \sigma_z \varepsilon_{z,t+1}, \tag{A14}$$

while under imperfect information this yields

$$r_{c,t+1} = \kappa_0 + A_0 (\kappa_1 - 1) + A_x \hat{x}_{t|t} (\kappa_1 \rho_x - 1) + A_z \hat{z}_{t|t} (\kappa_1 \rho_z - 1) \tag{A15}$$

$$+ \Delta c_{t+1} + \kappa_1 (A_x g_x + A_z g_z) v_{t+1}. \tag{A16}$$

Since $\mathbb{E}_t [M_{t+1} R_{c,t+1}]$ is always equal to one, it must be the case that the coefficients on $x_{t|t}$ and $z_{t|t}$ are zero (since both are known at date-t and are time-varying). Thus, under both full and imperfect information, and given the coefficients on $x_{t|t}$, it must be the case that

$$\left(\theta - \frac{\theta}{\psi}\right) \rho_x + \theta (A_x (\kappa_1 \rho_x - 1)) = 0 \tag{A17}$$

which implies that

$$A_x = \frac{\rho_x \left(1 - \frac{1}{\psi}\right)}{1 - \kappa_1 \rho_x}, \tag{A18}$$

while given the coefficients on $z_{t|t}$, it must be the case that

$$\left(\theta - \frac{\theta}{\psi}\right) (\rho_z - 1) + \theta (A_y (\kappa_1 \rho_z - 1)) = 0 \tag{A19}$$

which implies that

$$A_z = \frac{(\rho_z - 1)\left(1 - \frac{1}{\psi}\right)}{1 - \kappa_1 \rho_z}. \tag{A20}$$

Note that since $\psi > 1$, $1 - \frac{1}{\psi} > 0$, and consequently, $A_x > 0$ while $A_z < 0$, consistent with the notion that ε_x (ε_z) is a positive (negative) long-run risk shock.

Risk Premium. Next, we solve for the prices of risk for each shock in the economy. Under full

[23]The constant κ_1 is positive and close to 1, under both full- and imperfect- information, such that $\kappa_1^F \approx \kappa_1^L$. For notational ease, we omit the superscripts and work with a constant κ_1 in both environments. We confirm in the next subsection that in our first (second) calibration of the model, the difference is just 0.09% (0.04%).

information,

$$m_{t+1} - \mathbb{E}_t\left[m_{t+1}\right] = -\left(\gamma + (1-\theta)\,\kappa_1 A_x\right)\sigma_x\varepsilon_{x,t+1} - \left(\gamma + (1-\theta)\,\kappa_1 A_z\right)\sigma_z\varepsilon_{z,t+1} \tag{A21}$$

$$\equiv -\lambda_x\sigma_x\varepsilon_{x,t+1} - \lambda_z\sigma_z\varepsilon_{z,t+1}. \tag{A22}$$

Under imperfect information, it can be shown that

$$m_{t+1} - \mathbb{E}_t\left[m_{t+1}\right] = -\left(\gamma + (1-\theta)\,\kappa_1\left(A_x g_x + A_z g_z\right)\right)v_{t+1} \tag{A23}$$

$$\equiv -\lambda_v v_{t+1}. \tag{A24}$$

Note that under a preference for early resolution of uncertainty, $\lambda_x, \lambda_z, \lambda_v > 0$.[24] Equipped with the prices of risk, we can solve for the conditional risk premium. Utilizing log-normality, we know that

$$\mathbb{E}_t\left[r_{c,t+1} - r_{f,t}\right] = -\mathrm{Cov}_t\left(m_{t+1} - \mathbb{E}_t\left[m_{t+1}\right], r_{c,t+1} - \mathbb{E}_t\left[r_{c,t+1}\right]\right) - \tfrac{1}{2}\mathrm{Var}_t\left[r_{c,t+1}\right]. \tag{A25}$$

Under full information, the innovations to the return on the consumption claim are

$$r_{c,t+1} - \mathbb{E}_t\left[r_{c,t+1}\right] = \left(1 + \kappa_1 A_x\right)\sigma_x\varepsilon_{x,t+1} + \left(1 + \kappa_1 A_z\right)\sigma_z\varepsilon_{z,t+1} \tag{A26}$$

so that the risk premium is

$$\lambda_x\left(1 + \kappa_1 A_x\right)\sigma_x^2 + \lambda_z\left(1 + \kappa_1 A_z\right)\sigma_z^2 - \tfrac{1}{2}\left[\left(1 + \kappa_1 A_x\right)^2\sigma_x^2 + \left(1 + \kappa_1 A_z\right)^2\sigma_z^2\right] \tag{A27}$$

Under imperfect information, the innovations to the return on consumption are

$$r_{c,t+1} - \mathbb{E}_t\left[r_{c,t+1}\right] = \left(1 + \kappa_1\left(A_x g_x + A_z g_z\right)\right)v_{t+1} \tag{A28}$$

and so we can write the risk premium under learning as

$$\lambda_v\left(1 + \kappa_1\left(A_x g_x + A_z g_z\right)\right)\sigma_v^2 - \tfrac{1}{2}\left[\left(1 + \kappa_1\left(A_x g_x + A_z g_z\right)\right)^2\sigma_v^2\right]. \tag{A29}$$

In both information environments, the risk premium is positive.[25]

Risk-free rate. Under both full and imperfect information,

$$\mathbb{E}_t\left[m_{t+1}\right] = m_o - \tfrac{\rho_x}{\psi}\hat{x}_{t|t} - \tfrac{\rho_z - 1}{\psi}\hat{z}_{t|t}, \tag{A30}$$

where $m_0 = \theta\log\delta - \gamma\mu - (\theta-1)\log\kappa_1$. Then, since $r_{f,t} = \mathbb{E}_t\left[m_{t+1}\right] - \frac{1}{2}\mathrm{Var}_t\left[m_{t+1}\right]$, the unconditional

[24] A preference for early resolution of uncertainty implies that $\gamma > \frac{1}{\psi}$ and so $1-\theta > 0$. Since $A_x > 0$, this implies that $\lambda_x > 0$. Since $\kappa_1 < 1$, and $A_z < 0$

$$\begin{aligned}\lambda_z &> \gamma + (1-\theta)\,A_z \\ &> \gamma + (1-\theta)\,\frac{(\rho_z - 1)\left(1 - \frac{1}{\psi}\right)}{1-\rho_z} = \frac{1}{\psi} > 0\end{aligned}$$

Moreover, since $g_z < 1$, this implies that $\gamma + (1-\theta)g_z A_z > 0$ and so $\lambda_v > 0$ as well.

[25] Since $A_x > 0$, both $1 + \kappa_1 A_x$ and $\kappa_1 A_x g_x > 0$ and since $\kappa_1 < 1$, $1 + \kappa_1 A_z > \frac{1}{\psi} > 0$ which also implies that $1 + \kappa_1 A_z g_z > 0$. If $\gamma > 1$, it is straightforward to show that the Jensen term is sufficiently small.

risk-free rate under full information is

$$\mathbb{E}\left[r_{f,t}\right] = m_0^F - \tfrac{1}{2}\left(\lambda_x^2\sigma_x^2 + \lambda_z^2\sigma_z^2\right), \tag{A31}$$

where m_0^F is the constant under full information while the variance is

$$\mathrm{Var}\left[r_{f,t}\right] = \left(\tfrac{\rho_x}{\psi}\right)^2\left(\frac{\sigma_x^2}{1-\rho_x^2}\right) + \left(\tfrac{\rho_z-1}{\psi}\right)^2\left(\frac{\sigma_z^2}{1-\rho_z^2}\right). \tag{A32}$$

Similarly, with learning, the unconditional risk-free rate can be written as

$$\mathbb{E}\left[r_{f,t}\right] = m_0^L - \tfrac{1}{2}\lambda_v^2\sigma_v^2 \tag{A33}$$

where m_0^L is the constant under learning, while the variance is

$$\mathrm{Var}\left[r_{f,t}\right] = \left(\tfrac{\rho_x}{\psi}\right)^2\left(\frac{g_x^2\sigma_v^2}{1-\rho_x^2}\right) + \left(\tfrac{\rho_z-1}{\psi}\right)^2\left(\frac{g_z^2\sigma_v^2}{1-\rho_z^2}\right). \tag{A34}$$

Appendix-A.2 Risk Premium and Information: Qualitative Analysis

To shed light on the difference in the risk premium under the two information environments in an endowment economy, it is instructive to break the conditional covariance

$$-\mathrm{Cov}_t\left(m_{t+1} - \mathbb{E}_t\left[m_{t+1}\right], r_{c,t+1} - \mathbb{E}_t\left[r_{c,t+1}\right]\right),$$

into three pieces. Each piece captures a unique source of uncertainty driven by households' beliefs. Specifically, we isolate compensation for three sources of risk:

(i) Shocks to realized consumption (i.e., prediction errors);

(ii) Revisions to beliefs about the state variables underpinning future consumption;

(iii) The covariance between these two sources of uncertainty. The first and the third sources amplify the risk premium under learning, while the second attenuates it.

As we describe below, the third component (i.e., the covariance), and its associated compensation, are *uniquely* amplified in our setup due to the negative endogenous association between revisions in trend and business-cycle beliefs. We elaborate on these components in what follows.

First, there is the uncertainty introduced by shocks to realized consumption growth, $\Delta c_{t+1} - \mathbb{E}_t\Delta c_{t+1}$. Under full information, the compensation for this risk is

$$\gamma\left(\sigma_x^2 + \sigma_z^2\right) \tag{A35}$$

while under learning it is given by

$$\gamma\sigma_v^2. \tag{A36}$$

Under imperfect information, households face uncertainty about the true state of the world (i.e., $\mathbb{V}\left[x_t - \hat{x}_{t|t}\right]$ and $\mathbb{V}\left[z_t - \hat{z}_{t|t}\right]$ are non-zero). As a result, the variance of $\Delta c_{t+1} - \mathbb{E}_t[c_{t+1}]$ (i.e., the prediction error) under learning exceeds the equivalent variance under full information, which is driven only by the realization of the underlying shocks. To see this formally, note that we can write

$$v_{t+1} = \rho_x \left(x_t - \hat{x}_{t|t}\right) + (\rho_z - 1)\left(z_t - \hat{z}_{t|t}\right) + \sigma_x \varepsilon_{x,t+1} + \sigma_z \varepsilon_{z,t+1}. \tag{A37}$$

As long as $\rho_x \neq 0$ and $\rho_z \neq 1$, i.e., as long as the shocks to consumption growth are persistent, then it must be that $\sigma_v^2 > \sigma_x^2 + \sigma_z^2$. Consequently, this component of the risk premium is larger under learning.

Second, there is the uncertainty induced by direct shocks to perceived expected future consumption growth, i.e., revisions of beliefs about the state variables. The premium demanded for these shocks under full information is:

$$(1-\theta)\,\kappa_1^2\left[A_x^2\sigma_x^2 + A_z^2\sigma_z^2\right] \tag{A38}$$

while under imperfect information this is

$$(1-\theta)\,\kappa_1^2\left[A_x^2\left(g_x^2\sigma_v^2\right) + A_z^2\left(g_z^2\sigma_v^2\right) + 2A_x g_x A_z g_z \sigma_v^2\right]. \tag{A39}$$

In order to compare the compensation demanded across the two information environments, note that the state variables that determine expected consumption growth are given by the vector $\sigma_t = [\hat{x}_{t|t}, \hat{z}_{t|t}, \hat{z}_{t-1|t}]$. Let $P = Var(\hat{\sigma}_t)$ be the steady-state variance-covariance matrix of $\hat{\sigma}$, obtained from the fixed-point of the Riccati equation. We denote the elements of P by:

$$P \equiv \begin{bmatrix} p_{11} & p_{12} & p_{13} \\ p_{12} & p_{22} & p_{23} \\ p_{13} & p_{23} & p_{33} \end{bmatrix}. \tag{A40}$$

Note that because $g_x = \frac{p_{11}+p_{12}-p_{13}}{\sigma_v^2}$, we obtain

$$g_x^2\sigma_v^2 = \frac{\left(p_{11}+p_{12}-p_{13}\right)^2}{\sigma_v^2}. \tag{A41}$$

From the Ricatti equations, we know that

$$p_{11} = \rho_x^2 p_{11} - \rho_x^2\frac{\left(p_{11}+p_{12}-p_{13}\right)^2}{\sigma_v^2} + \sigma_x^2. \tag{A42}$$

Isolating $\frac{(p_{11}+p_{12}-p_{13})^2}{\sigma_v^2}$ and substituting into (A41) yields

$$g_x^2\sigma_v^2 = \frac{\sigma_x^2 + \left(\rho_x^2 - 1\right)p_{11}}{\rho_x^2} \tag{A43}$$

$$= \sigma_x^2 + \frac{\rho_x^2 - 1}{\rho_x^2} \left(p_{11} - \sigma_x^2 \right) \tag{A44}$$

Finally, since

$$p_{11} = \mathbb{V} \left[x_{t+1} - \hat{x}_{t+1|t} \right] = \rho_x^2 \mathbb{V} \left[x_t - \hat{x}_{t|t} \right] + \sigma_x^2, \tag{A45}$$

it is always the case that $p_{11} > \sigma_x^2$. But this implies that $g_x^2 \sigma_v^2 < \sigma_x^2$ as long as $\rho_x < 1$. Following similar steps with respect to p_{22} yields the analogous result for the business-cycle shock, $g_z^2 \sigma_v^2 < \sigma_z^2$. Since $A_x > 0 > A_z$, these inequalities, in combination with (A38) and (A39), imply that the risk premium required to compensate agents for the variance over their beliefs about the future (long-run) will always be larger under full information. Under imperfect information, agents update less aggressively in response to the prediction error since they recognize that some of this prediction error is due to mistaken beliefs about *earlier* shocks to the underlying state variables.

The third part of the risk premium is attributed to the covariance between state-space belief revisions and the shocks to realized consumption growth. Households demand compensation for this covariance because both contemporaneous shocks to consumption, as well as the revision in their beliefs, affect both the SDF and the return on the consumption claim. Under full information this compensation is

$$(1 - \theta + \gamma) \left(\kappa_1 A_x \sigma_x^2 + \kappa_1 A_z \sigma_z^2 \right) \tag{A46}$$

while under learning the analogous term is

$$(1 - \theta + \gamma) \left(\kappa_1 A_x g_x \sigma_v^2 + \kappa_1 A_z g_z \sigma_v^2 \right). \tag{A47}$$

To facilitate the comparison between (A46) and (A47), we first establish that $p_{12}, p_{13} \leq 0$ in our setting. To understand why, recall that these terms capture the covariance of agents' forecast errors between the underlying state variables. For example, p_{12} is the covariance between agents' forecast errors with respect to x_{t+1} and z_{t+1}. Since the underlying shocks to each process are independent (i.e., $\varepsilon_{x,t+1}$ and $\varepsilon_{z,t+1}$ are orthogonal), the sign of this covariance depends only on the comovement between the time-t prediction errors, $x_t - \hat{x}_{t|t}$ and $z_t - \hat{z}_{t|t}$, which are *negatively* correlated. Ignoring the shocks $\varepsilon_{x,t+1}$ and $\varepsilon_{z,t+1}$ (which affect neither p_{12} nor p_{13}), if

$$\Delta c_{t+1} - \Delta \hat{c}_{t+1} > 0 \implies \tag{A48}$$

$$x_t > \hat{x}_{t|t} \text{ and/or } z_t < \hat{z}_{t|t}. \tag{A49}$$

That is, if consumption growth is higher than expected it could be due to under-estimation of the trend state or over-estimation of the business-cycle state. Similarly, lower than expected consumption growth results from over-estimation of x_t or under-estimation of z_t, i.e.,

$$\Delta c_{t+1} - \Delta \hat{c}_{t+1} < 0 \implies \tag{A50}$$

$$x_t < \hat{x}_{t|t} \text{ and/or } z_t > \hat{z}_{t|t}. \tag{A51}$$

To see the negative covariation formally, note that the Ricatti equations imply that

$$p_{12} = \rho_x \rho_z p_{12} - \rho_x \rho_z g_x g_z \sigma_v^2 \tag{A52}$$

$$= \frac{-\rho_x \rho_z g_x g_z \sigma_v^2}{1 - \rho_x \rho_z} \leq 0, \tag{A53}$$

while

$$p_{13} = \rho_x p_{12} - \rho_x g_x g_z \sigma_v^2 \leq 0 \tag{A54}$$

where the inequalities are strict as long as $\rho_x, \rho_z \neq 0$.

Given $p_{12}, p_{13} \leq 0$, we can establish that $g_x \sigma_v^2 > \sigma_x^2$ while $g_z \sigma_v^2 < \sigma_z^2$. The covariance between the prediction error to current consumption and the revision in beliefs regarding the trend state is given by

$$g_x \sigma_v^2 = p_{11} + p_{12} - p_{13} \tag{A55}$$

$$= p_{11} + (\rho_z - 1) p_{13}. \tag{A56}$$

Since $\rho_z < 1$ and $p_{13} \leq 0$, this implies that $g_x \sigma_v^2 > p_{11} > \sigma_x^2$. In contrast,

$$g_z \sigma_v^2 = p_{22} + p_{12} - p_{23} \tag{A57}$$

$$= \rho_z^2 p_{33} + \sigma_z^2 + p_{12} - \rho_z p_{33} \tag{A58}$$

$$= \sigma_z^2 + \rho_z p_{33} (\rho_z - 1) + p_{12} \tag{A59}$$

Since $p_{12} \leq 0$ and the variance of any prediction error is positive, i.e., $p_{33} > 0$, this implies that $g_z \sigma_v^2 < \sigma_z^2$.

Intuitively, with respect to the trend shock, the covariance is higher than σ_x^2 for two reasons. The first reason depends on rational learning. A positive (negative) prediction error implies that the belief about the trend state, $\hat{x}_{t|t}$, was too low (high) which leads to a larger upward (downward) revision in beliefs about the trend state than under full information. This channel would arise even if the model featured only trend shocks (with, e.g., i.i.d. noise.) The second reason depends on the *interaction* between beliefs about the trend versus the business cycle state, and hence is *unique* to our setup. A positive (negative) prediction error implies that past beliefs about the business-cycle state were too high (low) as well, which given the negative comovement, suggests the beliefs about the trend were deflated (inflated) above and beyond what would have been perceived if only the one state variable existed. Thus, the covariance is amplified, all else equal, through the $(\rho_z - 1)\, p_{13}$ term found in (A56).

In contrast, with respect to the business-cycle shock, the covariance is lower than σ_z^2 since a positive prediction (negative) prediction error implies that the belief about the business-cycle state, $\hat{z}_{t|t}$, was too high (low) and so there is a smaller upward revision in beliefs about the business-cycle relative to full information. This revision is further deflated through the interaction with the trend shock, and there is further attenuation in the covariance term ($p_{12} < 0$ in (A59)).

These observations, in combination with the fact that $A_x > 0 > A_z$ and equations (A46) and (A47), imply that households demand more compensation for this type of "learning" risk under imperfect information than under full information. More precisely, with respect to the trend (business-cycle) shock, the covariance is higher (lower) under learning; given that the return on consumption's exposure to long-run consumption shocks induced by trend (business-cycle) shocks is positive (negative), the compensation for this source of risk is more positive (less negative), which operates to raise the total risk premium.

Appendix-A.3 Comparing Full to Imperfect Information: Quantitative Analysis

Table Appendix-A.1 summarizes the key model-implied asset-pricing moments from two calibrations of the exogenous consumption process.[26] In the first, we adopt the estimated parameters underlying the productivity process found in our baseline analysis, i.e., our estimated process for productivity, A_t becomes our estimated process for consumption, C_t. In the second, we simulate a population path of the endogenous consumption process from the production model and then estimate the parameters of the Kalman filter using MLE.

There are three key observations. (1) Learning amplifies the risk premium.[27] (2) Learning reduces

[26] The parameters underlying household preferences are the same as in the main text.

[27] While the risk premium of the consumption claim is substantially higher than the estimated *equity* risk premium, both calibrations' results are in line with the risk premium of a *consumption* claim in our production economy, as shown in Section IA.1.

Table Appendix-A.1: **Endowment economy: full information versus learning**

	Calibration 1 $\Delta c = \Delta a$		Calibration 2 $\Delta c =$ Kalman(endogenous consumption)	
Parameters.				
ρ_z	0.8783		0.8804	
σ_z	0.0168		0.0074	
ρ_x	0.9962		0.9978	
σ_x	0.00084		0.00054	
Consumption moments.				
$E[\Delta c]$ (%)	2.11		2.11	
$\sigma[\Delta c]$ (%)	2.99		1.54	
$AC1[\Delta c]$	0.27		0.46	
$VR2[\Delta c]$	1.28		1.44	
$VR3[\Delta c]$	1.40		1.80	
Model solution.				
	Full information	*Learning*	*Full information*	*Learning*
$E[R_c^e]$ (%)	8.50	13.86	5.70	7.69
Contribution z (%)	1.85	99.66	0.47	99.32
$Corr(\eta_t, \varepsilon_z)$	-0.91	-0.45	-0.84	0.87
ρ_η	0.971	0.999	0.990	0.999
$E[R_f]$	1.56	-2.50	2.27	0.73
$\sigma[R_f]$	1.44	0.85	0.82	0.62

The table shows model-implied moments from an endowment economy, in which the business-cycle and trend parameters of consumption are: (1) identical to those of aggregate productivity in the production model, or (2) obtained by Kalman filtering the endogenous consumption process from the simulated production economy. The model-implied moments are computed under learning and under full-information.

both the risk-free rate and its volatility, relative to the full-information benchmark. (3) Business-cycle shocks explain little of the risk premium under full information but are the primary driver under imperfect information.

Risk-premium and relative contribution. To understand the quantitative importance of sources (i)–(iii) (as described in subsection Appendix-A.2) for this amplification, we consider the fraction of the risk premium amplification (i.e., $\mathbb{E}\left[R_c^e|\text{Learning}\right] - \mathbb{E}\left[R_c^e|\text{Full Info}\right]$) which is attributable to each channel. Under the second calibration, the fraction of the risk premium amplification explained by shocks to realized consumption growth (i.e., channel (i)) is only 3%.[28] In contrast, the reduced variation in beliefs about the underlying state variables (i.e., channel (ii)) reduces the relative risk premium. Consequently, its relative contribution to the risk premium amplification is negative, accounting for approximately -47%. Thus, the key amplification of the risk premium under learning is due to channel (iii), the covariance between the current shock and households' beliefs about the future. This channel quantitatively dominates channel (ii), and in our calibration, explains just over 144% of the risk premium difference. Notably, channel (iii) is amplified by the interplay between trend and business-cycle beliefs, and consequently, is economically distinct from other amplification channels discussed in the literature.

Moreover, this breakdown explains why business-cycle shocks explain so little of the risk premium under full information while driving most of the risk under learning. Positive consumption growth shocks driven by $\varepsilon_{z,t}$ ($\varepsilon_{x,t}$) reduce (increase) future growth which offsets (amplifies) their contemporaneous impact on the risk premium. Consequently, $|A_z| << |A_x|$, and under full information the relative contribution of business-cycle shocks is less than 1%.[29] Under imperfect information, the forecast error variance is driven primarily by business-cycle shocks and their associated state uncertainty: $\frac{\sigma_z^2+(\rho_z-1)^2 p22}{\sigma_v^2}$ is greater than 90%.

We note that, in particular, *trendy business cycle shocks*, $g_x\sigma_z^2$, play the *largest* role in determining the risk premium given the outsized importance of the covariance term in explaining the relative risk premium under learning. To see this, note that under learning, the covariance with respect to the trend shock can be written as $g_x\sigma_v^2 = g_x(\sigma_z^2 + \sigma_x^2 + \rho_x^2 p_{11} + (\rho_z - 1)^2 p_{22})$. Since $\sigma_x < p_{11} << \sigma_z$, and since ρ_z is sufficiently high, the trendy business cycle component $g_x\sigma_z^2$ dominates quantitatively. In summary, trendy business-cycles amplify the covariance risk, which in turn, amplifies the total compensation for perceived expected consumption growth fluctuations (that is, long-run risk).[30]

[28]The relative contribution of each component is similar under both calibrations and so we focus on the process which more closely matches the baseline production economy here.

[29]To be more precise, under either calibration, the price of risk for each shock ($\lambda_x\sigma_x$ and $\lambda_z\sigma_z$ in the notation of equation (A22)) is positive under full information. However, because transitory shocks are expected to mean-revert and given households' preference for early-resolution of uncertainty, $A_z < 0$. Thus, even though $\sigma_z >> \sigma_x$, the price of risk for trend shocks is an order of magnitude larger.

[30]Just as a fraction of business-cycle shocks are perceived to be trend shocks (i.e., trendy business-cycles), a fraction of the trend shocks are perceived to be part of the business-cycle (i.e., transitory trends). Recall that the covariance between the prediction error and the beliefs about the business-cycle is attenuated under learning, in part due to the

Risk-free rate. While we cannot show analytically that the risk-free rate is always lower under learning in the endowment economy (similar to the risk premium, there are opposite forces at work), our numerical analysis suggests that this is the case under both calibrations. In contrast, it is always the case that the variance of the risk-free rate is lower under learning. This stems immediately from Equation (A34) in conjunction with $g_x^2\sigma_v^2 < \sigma_x^2$ while $g_z^2\sigma_v^2 < \sigma_z^2$, as previously shown. Intuitively, because agents revise their beliefs about the current state variables less aggressively, the variance of the equilibrium risk-free rate is lower.

Source of long-run risk. We conclude by discussing the contribution of business cycle shocks to long-run consumption risk. As in the corresponding analysis of the production economy (Section 6.3), we estimate an AR(1) process to a simulated path of expected consumption. The residuals, denoted by η_t, are the shocks to expected consumption growth. As shown in Table Appendix-A.1, business cycle shocks are negatively correlated with shocks to expected consumption growth under full information: positive (negative) realizations of $\varepsilon_{z,t}$ are associated with lower (higher) growth in the future.

The impact under learning, however, differs across the two calibrations due to the differential response of households to forecast errors. Specifically, the change in household beliefs is given by

$$\mathbb{E}_t\left[\Delta c_{t+1}\right] - \mathbb{E}_{t-1}\left[\Delta c_{t+1}\right] = \left[\rho_x g_x + \left(\rho_z - 1\right) g_z\right] v_t. \tag{A60}$$

This suggests that the sign of $corr(\eta_t, \varepsilon_{z,t})$ should be driven by the sign of $\rho_x g_x + (\rho_z - 1)\, g_z$. In the second calibration, $\rho_x g_x + (\rho_z - 1)\, g_z > 0$. Positive realizations of the business-cycle shock are therefore associated with an increase in expected consumption growth and the estimated correlation is positive. In the first calibration, this is not the case, despite featuring the same parameters as the productivity process in our benchmark model. While the persistence parameters (ρ_z and ρ_x) are nearly identical, the gains are not. g_x is smaller (0.045 versus 0.067) while g_z is larger (0.425 versus 0.347). This is driven by the relative increase in the variance of the business-cycle shocks in the first calibration: more of the forecast error is associated with $\varepsilon_{z,t}$ relative to the second calibration. As a result, positive forecast errors decrease expectations of consumption growth in this calibration, and business-cycle shocks are negatively-correlated with shocks to expected consumption growth, even under learning - in contrast to the equivalent result under production.

interaction with the trend shock. Since this covariance bears a *negative* risk premium, one could argue that transitory trends also play a role in the risk premium amplification, albeit a modest one.

Appendix B Term Structure

In this section we compare the term structure of real interest rates and of risk premia on dividend strips under a model with perfect information versus a model with learning. We repeat this exercise twice: once for the endowment economy described in Appendix A, and another for our benchmark production economy described in Section 3. Learning tilts the slope of the term structure compared to full information, but its effect is ex-ante ambiguous and depends upon two opposing factors: the share of the forecast error variance attributed to each shock and the relative prices of risk. Moreover, the relative strength of these factors depends upon whether the consumption process is endogenous or exogenous (i.e., the flexibility of investment).

In the endowment setting (a setting with implicitly inflexible investment), imperfect-information generates term structures that look closer to a model that is dominated by business-cycle shocks. We show that this is driven by the amplification (attenuation) of the price of business-cycle (trend) risk under learning. When consumption is endogenous (as in our production economy with flexible investment), learning yields term structure slopes that appear to be more impacted by trend shocks. We show that this arises because observers attribute a larger share of the forecast error variance to trend shocks under learning.

Appendix-B.1 Model Solution

The n-period bond log-price, $B_{n,t}$, is given by:

$$\exp(B_{n,t}) = \mathbb{E}_t[\Pi_{j=1}^{n} M_{t+j-1,t+j}],$$

and the n-period dividend-strip price is given by:

$$\exp(D_{n,t}) = \mathbb{E}_t[\Pi_{j=1}^{n} M_{t+j-1,t+j} CF_{t+n}].$$

While the log-prices $B_{n,t}$ and $D_{n,t}$ can only be solved for numerically under production, they can be expressed in closed form in the endowment economy. This analytical solution allows us to more precisely identify the different forces that drive the term structure under both full information and under learning. [31]

[31] Under production, $CF_t = D_t$, given by equation (12). Under an endowment setup, $D_t = C_t$. Since the only wedge between D_t and C_t under production is the labor cost, and given that labor is inelastic, this wedge does not materially affect the cyclicality of dividends compared to endogenous consumption. Consequently, all results below are qualitatively

Appendix-B.1.1 Real Yields Curve

We begin by analyzing the term structure of real yields. Conjecture that the n-period log-bond price is given by:

$$B_{n,t} = -B_{0,n} - B_{x,n}\hat{x}_{t|t} - B_{z,n}\hat{z}_{t|t}. \tag{B1}$$

We solve for the bond prices recursively utilizing the stochastic discount factor and the observation that the n-period log-bond price satisfies

$$\exp(B_{n,t}) = \mathbb{E}\left[\exp(m_{t+1} + B_{n-1,t+1})\right]. \tag{B2}$$

Applying log to both sides yields

$$B_{n,t} = \mathbb{E}_t\left[m_{t+1} + B_{n-1,t+1}\right] + \tfrac{1}{2}\mathrm{Var}_t\left[m_{t+1} + B_{n-1,t+1}\right], \tag{B3}$$

given the properties of the log-normal distribution.

Going forward, we will distinguish between full information and learning using the superscripts FI and L, respectively. Under full information, agents observe the underlying shocks perfectly and so

$$B^{FI}_{n-1,t+1} = -B^{FI}_{0,n-1} - B^{FI}_{x,n-1}\left(\rho_x x_t + \sigma_x\varepsilon_{x,t+1}\right) - B^{FI}_{z,n-1}\left(\rho_z z_t + \sigma_z\varepsilon_{z,t+1}\right), \tag{B4}$$

while the full-information SDF is given by

$$m^{FI}_{t+1} = m_0 - \tfrac{\rho_x}{\psi}x_t - \tfrac{\rho_z - 1}{\psi}z_t - \lambda_x\sigma_x\varepsilon_{x,t+1} - \lambda_z\sigma_z\varepsilon_{z,t+1}. \tag{B5}$$

Substituting Equations (B4) and (B5) into Equation (B3) yields

$$B^{FI}_{n,t} = m_0 - B^{FI}_{0,n-1} - \left(\tfrac{\rho_x}{\psi} + B^{FI}_{x,n-1}\rho_x\right)x_t - \left(\tfrac{\rho_z-1}{\psi} + B_{z,n-1}\rho_z\right)z_t + \tfrac{1}{2}\left[\left(\lambda_x + B^{FI}_{x,n-1}\right)\sigma_x^2 + \left(\lambda_z + B^{FI}_{z,n-1}\right)\sigma_z^2\right]. \tag{B6}$$

Under learning however, agents update their beliefs utilizing the forecast error, v_{t+1}, so that

$$B^{L}_{n-1,t+1} = -B^{L}_{0,n-1} - B^{L}_{x,n-1}\left(\rho_x\hat{x}_{t|t} + g_x v_{t+1}\right) - B^{L}_{z,n-1}\left(\rho_z\hat{z}_{t|t} + g_z v_{t+1}\right), \tag{B7}$$

while the SDF in this case can be expresses as

$$m^{L}_{t+1} = m_0 - \tfrac{\rho_x}{\psi}\hat{x}_{t|t} - \tfrac{\rho_z-1}{\psi}\hat{z}_{t|t} - \lambda_v v_{t+1}. \tag{B8}$$

Substituting Equations (B7) and (B8) into Equation (B3) yields that

$$B^{L}_{n,t} = m_0 - \left(\tfrac{\rho_x}{\psi} + B^{L}_{x,n-1}\rho_x\right)\hat{x}_{t|t} - \left(\tfrac{\rho_z-1}{\psi} + B^{L}_{z,n-1}\rho_z\right)\hat{z}_{t|t} + \tfrac{1}{2}\left(\lambda_v + B^{L}_{x,n-1}g_x + B^{L}_{z,n-1}g_z\right)\sigma_v^2. \tag{B9}$$

Utilizing these expressions, we determine the coefficients recursively. First, note that $B^J_{0,0} = B^J_{x,0} = B^J_{z,0} = 0$, for $J \in \{FI, L\}$. Second, in both cases $B^{FI}_{x,n} = B^{L}_{x,n} \equiv B_{x,n}$, and $B^{FI}_{z,n} = B^{L}_{z,n} \equiv B_{z,n}$, where

identical when we set $CF_t = C_t$ in the production environment.

$B_{x,n}$ and $B_{z,n}$ are given by

$$B_{x,n} = \rho_x B_{x,n-1} + \tfrac{\rho_x}{\psi}, \tag{B10}$$

$$B_{z,n} = \rho_z B_{z,n-1} + \tfrac{\rho_z - 1}{\psi}. \tag{B11}$$

This implies that for all $n \geq 1$

$$B_{x,n} = \left(\sum_{j=0}^{n-1} \rho_x^j\right) \tfrac{\rho_x}{\psi} = \left(\tfrac{1-\rho_x^n}{1-\rho_x}\right) \tfrac{\rho_x}{\psi} > 0 \tag{B12}$$

$$B_{z,n} = \left(\sum_{j=0}^{n-1} \rho_z^j\right) \tfrac{\rho_z - 1}{\psi} = \left(\tfrac{1-\rho_z^n}{1-\rho_z}\right) \tfrac{\rho_z - 1}{\psi} < 0 \tag{B13}$$

However, the unconditional log-bond price varies across the two settings with

$$B_{0,n}^{FI} = B_{0,n-1}^{FI} - m_0 - \tfrac{1}{2}\left[\left(\lambda_x + B_{x,n-1}\right)^2 \sigma_x^2 + \left(\lambda_z + B_{z,n-1}\right)^2 \sigma_z^2\right], \tag{B14}$$

$$B_{0,n}^{L} = B_{0,n-1}^{L} - m_0 - \tfrac{1}{2}\left(\lambda_v + B_{x,n-1} g_x + B_{z,n-1} g_z\right)^2 \sigma_v^2. \tag{B15}$$

Slope of the Yield Curve

The conditional n-period yield is

$$y_{n,t}^J = -\tfrac{1}{n} B_{n,t}^J = \tfrac{1}{n}\left[B_{0,n}^J + B_{x,n}\hat{x}_{t|t} + B_{z,n}\hat{z}_{t|t}\right], \quad j \in \{FI, L\}, \tag{B16}$$

which implies that the unconditional yield y_n^J, is given by

$$y_n^J \equiv \mathbb{E}\left[y_{n,t}^J\right] = \tfrac{1}{n} B_{0,n}^J. \tag{B17}$$

For ease of expression, we will focus on $y_2 - y_1$ in our analysis of the determinants of the slope of the real yield curve.[32] We start by deriving the yields under full information:

$$y_1^{FI} = -m_0 - \tfrac{1}{2}\left[\lambda_x \sigma_x^2 + \lambda_z \sigma_z^2\right], \tag{B18}$$

$$y_2^{FI} = \tfrac{y_1^{FI}}{2} - \tfrac{m_0}{2} - \tfrac{1}{4}\left[\left(\lambda_x + B_{x,1}\right)\sigma_x^2 + \left(\lambda_z + B_{z,1}\right)^2 \sigma_z^2\right]. \tag{B19}$$

Substituting in for $y_{0,1}^{FI}$ and $B_{x,1}, B_{z,1}$ allows us to write the full-information slope as

$$slope^{FI} \equiv y_2^{FI} - y_1^{FI} = -\tfrac{1}{4}\left[\left(\tfrac{\rho_x}{\psi}\right)^2 \sigma_x^2 + \left(\tfrac{\rho_z - 1}{\psi}\right)^2 \sigma_z^2\right] - \tfrac{1}{2}\left[\lambda_x \tfrac{\rho_x}{\psi}\sigma_x^2 + \lambda_z \tfrac{\rho_z - 1}{\psi}\sigma_z^2\right]. \tag{B20}$$

Similarly, the equivalent yields under learning are given by

$$y_1^L = -m_0 - \tfrac{1}{2}\lambda_v \sigma_v^2, \tag{B21}$$

$$y_2^L = \tfrac{y_1^L}{2} - \tfrac{m_0}{2} - \tfrac{1}{4}\left(\lambda_v + B_{x,1} g_x + B_{z,1} g_z\right)^2 \sigma_v^2. \tag{B22}$$

[32] We can show analytically that the economic channels we highlight arise in any $y_n - y_{n-1}$ and hence apply to any slope, $y_{n+k} - y_n$.

Substituting in for $y_{0,1}^L$ allows us to write the learning slope as

$$slope^L \equiv y_2^L - y_1^L = -\frac{1}{4}\left[\left(\frac{\rho_x}{\psi}\right)^2 g_x^2\sigma_v^2 + \left(\frac{\rho_z-1}{\psi}\right)^2 g_z^2\sigma_v^2\right] - \frac{1}{2}\left[\lambda_v\frac{\rho_x}{\psi}g_x\sigma_v^2 + \lambda_v\frac{\rho_z-1}{\psi}g_z\sigma_z^2\right]. \tag{B23}$$

Comparing $slope^{FI}$ to $slope^L$ yields several observations.

First, note that if one of the shocks is shut down, the slopes under full-information and learning are identical. For instance, if $\sigma_z = 0$, then $g_x = 1$, $\sigma_v^2 = \sigma_x^2$, $\lambda_x = \lambda_v$ and, therefore, $slope^{FI} = slope^L$. Thus, it is straightforward to see that the slope is negative if there are only trend shocks (i.e., if $\sigma_z = 0$) while the slope can be positive if there are only business-cycle shocks (i.e., if $\sigma_x = 0$)[33].

Second, we obtain that

$$\text{Sign}(slope^L - slope^{FI}) \approx \text{Sign}\left(\rho_x\left[\lambda_x\sigma_x^2 - \lambda_v g_x\sigma_v^2\right] + (1-\rho_z)\left[\lambda_v g_z\sigma_v^2 - \lambda_z\sigma_z^2\right]\right) \tag{B24}$$

where the suppressed terms are quantitatively negligible.[34] This expression yields the two aforementioned opposing forces which determine the effect of learning on the slope of the real yield curve.

The first force is related to the perceived quantity of risk of each shock. As we showed in Appendix-A.2, under learning a larger share of the forecast error variance is associated with the trend shocks (i.e., $\frac{g_x}{g_z} > \frac{\sigma_x^2}{\sigma_z^2}$), because the following inequalities always hold:

$$g_x\sigma_v^2 > \sigma_x^2, \text{ and } g_z\sigma_v^2 < \sigma_z^2. \tag{B25}$$

As shown in Equation (B24), this implies that all else equal (in particular, keeping the *price* of risks for each shock equal), $\text{Sign}(slope^L - slope^{FI}) < 0$. Henceforth, we refer to this force as **the share of the forecast error** channel.

The second force is related to the change in the magnitudes of risk prices. As can be seen in Figure 3a the price of risk for business-cycle (trend) shocks is attenuated (amplified). Intuitively, a portion of the business-cycle shock is perceived as permanent (raising the price of risk of such shocks) while a portion of the trend shock is perceived as transitory, lowering its price. Formally, it is straightforward

[33] In this case, a necessary condition for the upward sloping curve is that $\lambda_z > \frac{1-\rho_z}{2\psi}$, which is satisfied under our calibration

[34] Note that since $g_x^2\sigma_v^2 < \sigma_x^2$ and $g_z^2\sigma_v^2 < \sigma_z^2$, the term $-\frac{1}{4}\left[\left(\frac{\rho_x}{\psi}\right)^2 g_x^2\sigma_v^2 + \left(\frac{\rho_z-1}{\psi}\right)^2 g_z^2\sigma_v^2\right] - \frac{1}{4}\left[\left(\frac{\rho_x}{\psi}\right)^2 g_x^2\sigma_v^2 + \left(\frac{\rho_z-1}{\psi}\right)^2 g_z^2\sigma_v^2\right]$ would be positive, suggesting that all else equal $Slope^L > Slope^{FI}$. However, unlike the remainder of the slope differential found in (B24), this term is not scaled by the respective prices of risk. As a result, we confirm that in our calibration this term has a negligible quantitative effect.

to show that

$$\lambda_x > \lambda_v > \lambda_z. \tag{B26}$$

All else equal (in particular, keeping the perceived *quantities* of risk constant), this implies that $\text{Sign}(slope^L - slope^{FI}) > 0$. Below, we refer to this force as the **relative price of risk** channel.

In Appendix-B.2, we show that which of the two channels dominates depends not only on the calibration, but also on the environment (i.e., the flexibility of investment).

Appendix-B.1.2 Dividend Strips Expected Returns Curve

We turn now to the term structure of dividend strips (equivalently, consumption strips in the endowment setup, since $CF_t = C_t = D_t$, in the absence of leverage of idiosyncratic shocks). Let $\hat{p}_{n,t}$ denote the log-price of an n-period dividend claim and $p_{n,t} \equiv \hat{p}_{n,t} - c_t$ denote the log-price to consumption ratio (to ensure stationarity). We conjecture that

$$p_{n,t} = D_{0,n} + D_{x,n}\hat{x}_{t|t} + D_{z,n}\hat{z}_{t|t}. \tag{B27}$$

Note that

$$exp\left(p_{n,t}\right) = \mathbb{E}\left[exp\left(m_{t+1} + \Delta c_{t+1} + p_{n-1,t+1}\right)\right] \tag{B28}$$

and so by the properties of the log-normal distrbution,

$$p_{n,t} = \mathbb{E}_t\left[m_{t+1} + \Delta c_{t+1} + p_{n-1,t+1}\right] + \tfrac{1}{2}Var_t\left[m_{t+1} + \Delta c_{t+1} + p_{n-1,t+1}\right]. \tag{B29}$$

Under full information, the stochastic discount factor is given by (B5) while

$$p^{FI}_{n-1,t+1} = D^{FI}_{0,n-1} + D^{FI}_{x,n-1}\left(\rho_x x_t + \sigma_x \varepsilon_{x,t+1}\right) + D^{FI}_{z,n-1}\left(\rho_z z_t + \sigma_z \varepsilon_{z,t+1}\right). \tag{B30}$$

Substituting Equations (B30) and (B5) into Equation (B29) yields that

$$\begin{aligned} p^{FI}_{n,t} = {} & m_0 + \mu + D^{FI}_{0,n-1} + \left(-\tfrac{\rho_x}{\psi} + \rho_x + D^{FI}_{x,n-1}\rho_x\right)x_t + \left(-\tfrac{\rho_z - 1}{\psi} + \rho_z - 1 + D^{FI}_{z,n-1}\rho_z\right)z_t \\ & + \tfrac{1}{2}\left[\left(1 - \lambda_x + D^{FI}_{x,n-1}\right)^2\sigma_x^2 + \left(1 - \lambda_z + D^{FI}_{z,n-1}\right)^2\sigma_z^2\right]. \end{aligned}$$

Under learning, the SDF is given by (B8) while

$$p^{L}_{n-1,t+1} = D^{L}_{0,n-1} + D^{L}_{x,n-1}\left(\rho_x \hat{x}_{t|t} + g_x v_{t+1}\right) + D_{z,n-1}\left(\rho_z \hat{z}_{t|t} + g_z v_{t+1}\right). \tag{B31}$$

Substituting Equations (B31) and (B8) into Equation (B29) yields that

$$\begin{aligned} p^{L}_{n,t} = {} & m_0 + \mu + D^{L}_{0,n-1} + \left(-\tfrac{\rho_x}{\psi} + \rho_x + D^{L}_{x,n-1}\rho_x\right)x_t + \left(-\tfrac{\rho_z - 1}{\psi} + \rho_z - 1 + D^{L}_{z,n-1}\rho_z\right)z_t \\ & + \tfrac{1}{2}\left(1 - \lambda_v + D^{L}_{x,n-1}g_x + D^{L}_{z,n-1}g_z\right)^2\sigma_v^2. \end{aligned}$$

This implies that $D^{L}_{x,n} = D^{FI}_{x,n} \equiv D_{x,n}$ and $D^{L}_{z,n} = D^{FI}_{z,n} \equiv D_{z,n}$ where

$$D_{x,n} = \rho_x \left(1 - \frac{1}{\psi}\right) + \rho_x D_{x,n-1}, \tag{B32}$$

$$D_{z,n} = (\rho_z - 1)\left(1 - \frac{1}{\psi}\right) + \rho_z D_{z,n-1}, \tag{B33}$$

for all $n \geq 1$ since $D_{0,0} = D_{x,0} = D_{z,0} = 0$.

Suppressing superscripts for convenience, the log-return on the $n-$period ahead dividend (consumption) strip, $r_{n,t+1}$, can be written as

$$r_{n,t+1} = \hat{p}_{n-1,t+1} - \hat{p}_{n,t} = p_{n-1,t+1} - p_{n,t} + \Delta c_{t+1} \tag{B34}$$

which implies that its risk-premium is[35]

$$RP_{n,t} \equiv \mathbb{E}_t\left[r_{n,t+1} - r_{f,t}\right] \approx -\mathrm{Cov}_t\left(m_{t+1} - \mathbb{E}_t\left[m_{t+1}\right], r_{n,t+1} - \mathbb{E}_t\left[r_{n,t+1}\right]\right). \tag{B35}$$

Under full information,

$$r^{FI}_{n,t+1} - \mathbb{E}_t\left[r^{FI}_{n,t+1}\right] = \left(D_{x,n-1} + 1\right)\sigma_x \varepsilon_{x,t+1} + \left(D_{z,n-1} + 1\right)\sigma_z \varepsilon_{z,t+1} \tag{B36}$$

which implies that unconditionally

$$RP^{FI}_{n} = \lambda_x \left(D_{x,n-1} + 1\right)\sigma_x^2 + \lambda_z \left(D_{z,n-1} + 1\right)\sigma_z^2. \tag{B37}$$

Under learning,

$$r^{L}_{n,t+1} - \mathbb{E}_t\left[r^{L}_{n,t+1}\right] = D_{x,n-1} g_x v_{t+1} + D_{z,n-1} g_z v_{t+1} + v_{t+1} \tag{B38}$$

which implies that unconditionally

$$RP^{L}_{n} = \lambda_v \left(D_{x,n-1} g_x + D_{z,n-1} g_z + 1\right)\sigma_v^2 \tag{B39}$$

Slope of the Strips Curve

As with the real yield curve, we focus on $RP_{2,t} - RP_{1,t}$ for parsimony. Under full information,

$$RP^{FI}_{1} = \lambda_x \sigma_x^2 + \lambda_z \sigma_z^2$$
$$RP^{FI}_{2} = \lambda_x \left(\rho_x \left(1 - \frac{1}{\psi}\right) + 1\right)\sigma_x^2 + \lambda_z \left(\left(\rho_z - 1\right)\left(1 - \frac{1}{\psi}\right) + 1\right)\sigma_z^2$$

and so the slope is given by

$$slope^{FI} \equiv RP^{FI}_{2} - RP^{FI}_{1} = \lambda_x \rho_x \left(1 - \frac{1}{\psi}\right)\sigma_x^2 + \lambda_z \left(\rho_z - 1\right)\left(1 - \frac{1}{\psi}\right)\sigma_z^2 \tag{B40}$$

Under learning,

$$RP^{L}_{1} = \lambda_v \sigma_v^2 \tag{B41}$$

[35] It can be shown that the Jensen's term, $\frac{1}{2}\mathrm{Var}_t\left[r_{n,t+1}\right]$, is quantitatively unimportant and so for parsimony we exclude it from the qualitative discussion which follows.

$$RP_2^L = \lambda_v \left(1 + \rho_x \left(1 - \frac{1}{\psi}\right) g_x + (\rho_z - 1)\left(1 - \frac{1}{\psi}\right) g_z\right) \sigma_v^2 \tag{B42}$$

and so the slope is given by

$$slope^L \equiv RP_2^L - RP_1^L = \lambda_v \left(\rho_x \left(1 - \frac{1}{\psi}\right) g_x + (\rho_z - 1)\left(1 - \frac{1}{\psi}\right) g_z\right) \sigma_v^2$$

As with the real yield curve, if either shock is shut down, $slope^L = slope^{FI}$. This implies that for the term structure of dividend strips, trend shocks tend to increase the slope while business cycle shocks decrease it. Note that this implies that the impact of each shock is flipped relative to its impact on the real yield curve.

We can see this more generally by signing the relative slope utilizing the following expression:

$$\text{Sign}\left(slope^L - slope^{FI}\right) = \text{Sign}\left(\rho_x \left(\lambda_v g_x \sigma_v^2 - \lambda_x \sigma_x^2\right) + (1 - \rho_z)\left(\lambda_z \sigma_z^2 - \lambda_v g_z \sigma_v^2\right)\right). \tag{B43}$$

This expression is exactly the opposite of Equation (B24), which approximates the impact of learning for the real yield curve. Thus, in this case, the **share of the forecast error** channel suggests that $slope^L > slope^{FI}$, since agents assign a larger fraction of the forecast error variance to the trend shock, making the slope more positive (or less negative). In contrast, the **relative price of risk channel** suggests that $slope^L < slope^{FI}$ since the price of risk is larger (smaller) for the business-cycle (trend) shock.

Appendix-B.2 Quantitative Results

It is clear from the preceding analysis that the effect of learning on the slope of the term structure is a priori ambiguous and, importantly, depends upon the relative strength of two distinct economic channels. To resolve this ambiguity, we numerically solve for the term structure of real yields and equity risk premia under the endowment and the production economies. To so so, we utilize the parameters of the production economy which were estimated using SMM and reported in Table 1. For the endowment economy, the parameters governing the household preferences are the same as in Table 1. To ensure that the only difference between the two economies is the production environment (i.e., flexible investment), the parameters that govern the evolution of exogenous consumption, C_t, in the endowment setup are identical to those that govern the evolution of productivity, A_t, in the production setup.[36]

[36] This is the first calibration utilized in Appendix A. The results are qualitatively similar if we utilize the second.

Panel A of Table Appendix-B.1 shows model-implied slopes for the real yield curve, where

$$slope(ny - 2y) = E[y_{n,t}] - E[y_{2,t}],$$

and $y_{n,t}$ ($y_{2,t}$) is the n- (2-) year real yield. We focus on annual frequency yields and on the two-year yield for the short end to facilitate a comparison with the empirical counterparts for the US and the UK. Panel B of Table Appendix-B.1 shows model simulated slopes for the term structure of dividend strip expected returns, where

$$slope(ny - 1y) = E[R_{n,t+1}] - E[R_{1,t+1}],$$

and $R_{n,t+1}$ ($R_{1,t+1}$) is the n- (1-) year dividend strip return. Due to the availability of the corresponding data, we focus on the one-year dividend strip for the short end.

This analysis generates three main observations. First, under full-information, the model's yield (equity) term structure is downward (upward) sloping in a population sample, as implied by both the endowment or the production model. Recall that under full-information, trend (business-cycle) shocks induce a downward (upward) sloping real yield term structure, and an upward (downward) sloping equity term structure. Thus, while trend shocks are smaller than business-cycle shocks, the implications of the former for term structures dominates the effect of the latter. This result mirrors the dominance of trend shocks for the equity premium under full-information. We note, however, that the production-economy equity term structure slopes of horizons longer than two years are statistically indistinguishable from zero, which indicates that in *short samples* empirical estimates may suggest a larger role for business-cycle shocks.

Second, in the endowment economy, both term structures become flatter under learning (relative to the full-information benchmark). The learning-implied yield (equity) term structure is less downward (upward) sloping. This result suggests a larger role for the relative price of risk channel in our calibration; that is, the amplification (attenuation) of business-cycle (trend) risk-prices under learning leads to a flatter term structure.

Third, in our production setting, learning pushes the slopes in the *opposite* direction than what we observe in the endowment economy. Specifically, the learning-implied yield (equity) term structure is more downward (upward) sloping. This is particularly pronounced for the dividend strip term structure. While the five year slope is 2.2% under full-information, it is about 11% under learning.

These results indicate that, under production, the amplification of the trend shock's perceived share of the forecast error variance becomes the dominant channel. Consequently, term structures behave as if the data generating process exhibits more sizable permanent shocks than they actually do. This outcome is consistent with trendy business cycle shocks driving a large fraction of the equity premium as we discussed in Section 6.

Importantly, the qualitative change in learning's impact on the slope of the term structures emphasizes the importance of production for studying the pricing implications of imperfect information. By assuming that consumption is exogenous, one may spuriously conclude that learning flattens term structure slopes (or equivalently, enhances the contribution of the business cycle component). Yet, in the data, consumption is truly endogenous, and as we show in Table Appendix-B.1, the effect of learning is reversed in this setting (or equivalently, learning ends up enhancing the contribution of trend shocks) relative to an endowment economy.

This reversal in the impact on the slope bears an intuitive explanation. Untabulated quantitative results confirm that the impact of learning on the term structure of *output* strips (i.e., $CF_t = Y_t$) is similar to the impact of learning on *equity* strips in the endowment economy. That is, all else equal, the term structure of output strips is less upward-sloping under learning. A similar result applies to the term structure of investment strips, which becomes less upward-sloping as well. The key observation is that investment is more volatile than output, so the effect of learning on investment strips is more pronounced than output strips. Because consumption is output *minus* investment in the production economy, this implies that the effect of learning on the term structure of consumption strips should be the opposite: learning turns the term structure of consumption strips' expected returns to be *more* upward sloping. Put differently, the endogeneity of consumption in the model implies that the dominant channel in determining the impact of imperfect information on the term structure is the fact that a larger fraction of the forecast error variance is attributable to trend shocks under learning. Because dividends are almost identical to consumption (net of wages), the impact of learning on dividend strips is similar to that on consumption strips. In both cases, the slope turns more positive.

Whether learning helps to quantitatively reconcile the empirical term structures depends both on the the endogeneity of consumption as well as on one's prior on the data, given the relatively short samples available. For the US, data on real yields encompass a short period starting with the introduction of TIPs in 1997. Panel A of Table Appendix-B.1 shows that within this sample, the five-year yield slope is positive. However, for the UK, data on inflation-indexed bonds is available for a longer sample, starting in 1983. Evans (1998) and Verdelhan (2010) use a Nelson and Siegel (1987) model along with Bank of England zero-coupon real yields to obtain yields for fixed annual maturities. The five year real-yield slope is -0.91%. Under production, the five year yield slope turns from -0.47% under full information to -0.62% under learning. This is closer to the data, so long as one maintains a belief that the UK sample better reflects the population dynamics of real yields. Notably, our model is estimated using data from the 1960s onward, which is closer to the timeline of the UK sample than the US sample.

The empirical evidence about the equity term structure is also somewhat inconclusive. Papers by Van Binsbergen, Hueskes, Koijen, and Vrugt (2013) and Van Binsbergen and Koijen (2017) argue that the term structure of equity risk premia is downward sloping. More recent evidence by Bansal et al. (2021) suggests otherwise. Updated data shows that the term structure of dividend strip expected returns is downward sloping only in recessions. Unconditionally, the term structure has a positive slope, as shown in Panel B of Table Appendix-B.1. The five year slope using Bansal et al. (2021) data amounts to 7%. Under production, the five equity term structure slope is 11% (2.2%) under learning (full-information). Slopes are amplified at other maturities as well, which helps to reconcile the model-implied curve with the *unconditional* empirical curve.

Table Appendix-B.1: **Term Structures: full information versus learning**

Panel A: Real yields term structure

	Endowment Economy					
	Full Information		*Learning*			
slope(3y-2y)	-0.28	[-0.36,-0.28]	-0.13	[-0.14,-0.11]		
slope(4y-2y)	-0.57	[-0.71,-0.55]	-0.30	[-0.31,-0.28]		
slope(5y-2y)	-0.86	[-1.04,-0.81]	-0.51	[-0.51,-0.49]		
	Production Economy					
	Full Information		*Learning*		*UK Data*	*US Data*
slope(3y-2y)	-0.16	[-0.18,-0.12]	-0.20	[-0.27,-0.16]	-0.37	0.17
slope(4y-2y)	-0.31	[-0.37,-0.25]	-0.41	[-0.52,-0.33]	-0.75	0.32
slope(5y-2y)	-0.47	[-0.56,-0.38]	-0.62	[-0.76,-0.51]	-0.91	0.44

Panel B: Dividend strips term structure

	Endowment Economy					
	Full Information		*Learning*			
slope(2y-1y)	0.29	[0.27,0.33]	0.11	[0.10,0.12]		
slope(3y-1y)	0.60	[0.53,0.66]	0.29	[0.26,0.34]		
slope(4y-1y)	0.89	[0.79,1.00]	0.52	[0.47,0.61]		
slope(5y-1y)	1.19	[1.04,1.35]	0.78	[0.78,0.91]		
	Production Economy					
	Full Information		*Learning*		*UK Data*	*US Data*
slope(2y-1y)	1.06	[0.04,3.02]	4.57	[2.93,8.34]	-	3.85
slope(3y-1y)	1.69	[-0.33,5.16]	7.63	[4.67,14.30]	-	5.54
slope(4y-1y)	2.04	[-0.85,6.90]	9.69	[5.67,18.63]	-	6.43
slope(5y-1y)	2.21	[-1.33,8.14]	11.09	[6.18,21.81]	-	7.03

The table shows model-implied moments from an endowment economy, in which the business-cycle and trend parameters of consumption are identical to those of aggregate productivity in the production model, and from a production economy that is identical to the one described in Section 3, in which the parameters are described in Table 1. Panel A shows the slopes of the real yield curve, where slope(ny-$2y$) = $E[y_{n,t}] - E[y_{2,t}]$. Panel B shows the slopes of the real yield curve, where slope(ny-$1y$) = $E[R_{n,t+1}] - E[R_{1,t+1}]$, with $R_{n,t+1} = P_{n-1,t+1}/P_{n,t}$. The model-implied moments are computed under learning and under full-information. UK inflation-indexed yields cover the period 1983 to 2006, where the 1983-1995 yields are obtained from Evans (1998), and the 1995-2006 yields are from Verdelhan (2010). US inflation-indexed (TIPS) yields are from 1997-2006. US dividend strip returns are from Bansal et al. (2021), covering the period 2005-2017. All yields and returns are in annual terms expressed in percents.