Structural Breaks in an Endogenous Growth Model

Timothy Cogley* and Boyan Jovanovic†

May 5, 2021

Abstract

We study the effects of parameter uncertainty prompted by structural breaks. In our model, agents respond differently to uncertainty prompted by regime shifts in shock processes than they react to comparable perceived increases in shock volatility. The magnitude of the response to an increase in uncertainty about TFP associated with a structural break is greater than that of a response to a comparable perceived rise in volatility. This is because lifetime utility varies more when shocks shift beliefs and perceived wealth.

1 Introduction

Are learning episodes similar to times when shock volatilities rise? We study this question in a general-equilibrium model for an economy that is subject to recurrent structural breaks, breaks that create parameter uncertainty. Our model is based on an Ak growth model with two aggregate shocks, a neutral *TFP* shock in the final-goods sector and an investment-specific technology shock that affects production of capital goods. We deal only with aggregate shocks and abstract from idiosyncratic risk.

Structural breaks are introduced as occasional shifts in the parameters governing technology shocks. A break alters medium- and long-run forecasts of the level of technology and initiates a period of higher parameter uncertainty. We examine how this uncertainty affects savings, investment, and growth. Under uncertainty, beliefs

*Department of Economics, New York University, 19 W. 4th St 6FL, New York, New York 10012. Tel. 212-992-8679. Email: tim.cogley@nyu.edu.

†Department of Economics, New York University, 19 W. 4th St 6FL, New York, New York 10012. Email: boyan.jovanovic@nyu.edu. Tel. 212-998-8953. Fax: 212-995-4186.

over the distributions of future shocks are an aggregate state that responds to current shock realizations. Shocks thus have wealth effects through the beliefs channel – a channel that is absent in models where the change in the shock distributions is known, and one that creates both a quantitative and a qualitative difference between the two classes of models.

To motivate interest in structural breaks, figure 1 plots empirical counterparts of the shock processes. In our model, TFP equals the ratio of real output to real capital. We constructed a measure of y/k using data from Kendrick (1961) along with the National Income and Product Accounts and Fixed Asset Accounts from the US Bureau of Economic Analysis.[1] Appendix A explains in detail how our measure is constructed. Data for aggregate q come from Wright (2004) and the Federal Reserve Board's Financial Accounts of the United States. Wright's "equity q" measure covers the period 1900-2002, and we extend this through 2018 by ratio splicing the Federal Reserve Board's measure of equity q for the nonfinancial corporate business sector.[2]

The top row of figure 1 plots the level of each series along with *NBER* recession dates (shown as vertically shaded areas).[3] The bottom row depicts ten-year rolling standard deviations.

TFP is shown in the upper left panel. As highlighted by Gordon (2016, pp. 545-548), productivity surges in the 1940s and never returns to its pre-war level. Although *TFP* also varies procylically, the mid-century surge is the dominant source of volatility and presumably also of uncertainty. For instance, the rolling 10-year standard deviation is about 5 times higher around the years of the rise than at 'normal' times.[4]

The case for a structural break in the investment-specific technology shock – shown in the upper right panel – is less obvious, but an argument could be made that the

[1] See Bureau of Economic Analysis (2017a-e), available at https://apps.bea.gov/iTable/

[2] This is defined as corporate equities as a percentage of net worth (Board of Governors of the Federal Reserve System 2021, Financial Accounts of the United States - Z.1, table B.103, line 45, available at https://www.federalreserve.gov/releases/Z1/.

[3] Recession dates can be found at National Bureau of Economic Research (2017), available at https://www.nber.org/research/data/us-business-cycle-expansions-and-contractions

[4] The same pattern holds if multifactor *TFP* is used instead of the output-capital ratio: Fig. 16-5 of Gordon (2016) shows that the growth of *TFP* over the 1940-50 decade was about twice as high as it was over the adjacent decades, and about four times as high as it was over the other decades since 1900.

Data in Piketty (2014) suggest that structural breaks in Y/K also occurred in Britain and France, although perhaps at different times (compare Piketty's tables 3.1 and 3.2 for Britain and France with his table 4.6 for the US).

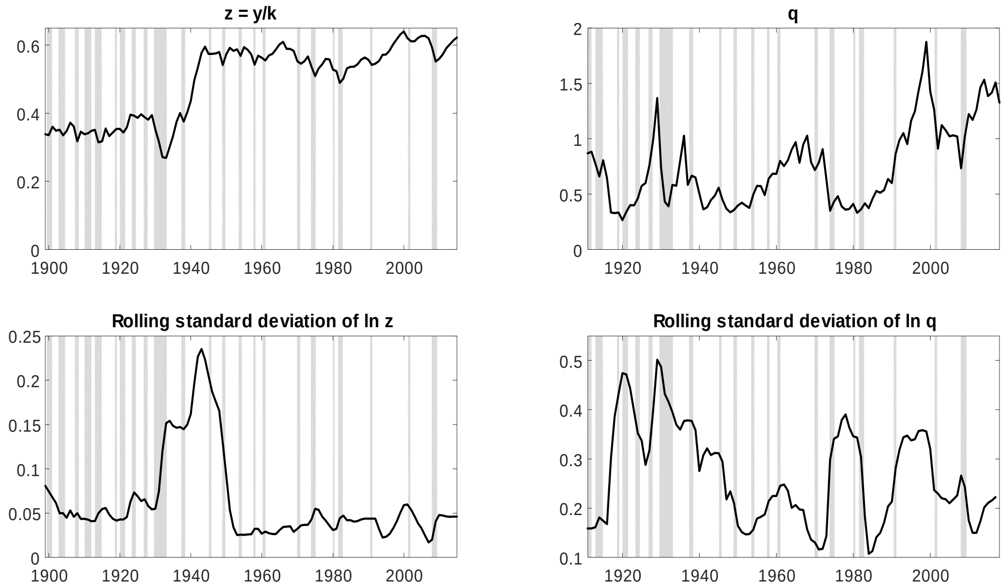

Figure 1: TFP (z) and investment cost (q) over a century. Levels are depicted in the top row, and 10 year rolling standard deviations are shown in the bottom. Shaded areas mark the dates of NBER recessions.

run-up in the 1990s and early 2000s represented a shift in its mean. For instance, the peak and trough after 1995 are two and three times higher, respectively, than those of the previous 90 years, and the trough in 2002 is almost as high as the peaks in the 1920s and 1960s. The volatility associated with this shift is comparable to that experienced in earlier years, neither dominating nor being dominated by ordinary cyclical variation.[5]

Our goal is to understand how shifts such as these affect savings, investment, and growth in general equilibrium. Our analysis is based on three assumptions. The first is that if a structural break has occurred in the past, others are possible in the future. Furthermore, agents are aware of this possibility and account for it when formulating their plans. Structural shifts are not complete surprises and as Sims (1982) argued in a related context, regime shifts can be viewed as outcomes of a rule that agents understand. This contrasts with anticipated-utility models in which agents update beliefs after a structural break but fail to account for the possibility of future breaks

[5] We acknowledge that this is not the only respectable interpretation of the data. Alas, because statistical tests for structural breaks have low power, a decisive characterization is unlikely.

when formulating plans (e.g., Kreps 1998 or Evans and Honkapohja 2001).

Second, history never repeats itself exactly. Regime realizations do not necessarily cycle between a small number of possible states, as in the literature on Markov-switching models. Entirely new regimes can and do emerge. The timing of the shift follows a Poisson process and conditional on a structural break new parameter values are drawn from a distribution conditioned on the previous parameter realizations; the exact nature of regimes is first-order Markov process on the space of shock-distribution parameters.

Third, because each realization is distinct, a structural break initiates a period of uncertainty. The new regime is not suddenly revealed after a break. On the contrary, agents must learn the new parameters governing the shocks. It follows that a structural break also raises uncertainty. A structural break combines a persistent change in the level of technology with a transitory increase in uncertainty. Both affect behavior and macroeconomic outcomes. Using our model, we can quantify and disentangle them.

The model has no heterogeneity among agents, and our study is therefore limited to the effects of macro uncertainty and macro risk. The model also has no financial sector and would be inappropriate for analyzing financial crises. On the other hand, the model is solved in closed form, without any linearization. Moreover the Ak structure allows a direct mapping between the observables and the shocks, and the endogeneity of growth amplifies the transitional dynamics – for instance, precautionary savings will have a stronger effect on future output.

For plausible calibrations, we find that uncertainty episodes entail lower consumption and higher growth compared with an economy that is identical except that agents know the parameters of the distributions governing the shocks. Uncertainty prompts precautionary savings that insures agents against an unfavorable change. Consumption falls and the higher savings then raise growth above its no-uncertainty level. Each type of learning, be it about TFP or about investment efficiency, has an expansionary effect on growth, although the effect is smaller in the second case. Uncertainty about TFP leads to more saving because in our closed economy low-TFP realizations lead to consumption disasters whereas high-investment-cost realizations do not because agents can reduce investment accordingly.[6]

[6]Even a bad TFP realization would not cause a consumption disaster in an open economy because agents could then borrow abroad. Baxter and Jermann (1997) note, however, that investors do not

To isolate the effects of a rise in uncertainty, however, we then ask whether these effects arise because agents live through episodes during which they are learning the parameters of shock distributions, or whether they would occur even if agents simply faced higher shock variances. The effect on growth of a mean-preserving increase in variance (MPS) that agents immediately understand can go in either direction. An MPS activates two forces: higher investment as a consequence of precautionary savings versus lower investment because of the option to wait (e.g., Dixit and Pindyck 1994). Shock persistence matters, however, because waiting is more attractive when shocks are weakly persistent. Since Bayesian learning makes one of the state variables a martingale (Doob 1949), a strongly persistent state is present in all our learning models. The presence of a persistent state variable weakens incentives to wait and brings the precautionary-savings motive to the fore.[7]

The next step is to compare the magnitude of responses associated with subjective and objective uncertainty. Toward that end, we equate the sequences of one-step-ahead marginal distributions of shocks in the two cases. Under learning, subjective uncertainty falls over time, and the no-learning economy is put on the same footing by assuming a corresponding decline in the conditional variance of shocks. For plausibly calibrated models, consumption and growth respond more strongly under learning: after ten years, output is permanently higher by two percent. In contrast, a mean-preserving increase in variance matching the profile of subjective uncertainty has very little effect on savings or growth.

Models featuring TFP and investment-efficiency shocks but without learning include Greenwood, Hercowitz, and Krusell (2000), Fisher (2006), Justiniano, Primiceri and Tambalotti (2010) and Jovanovic and Rousseau (2014). Gilchrist and Williams (2005) find that a rise in dispersion of shocks over vintages of capital is expansionary because labor reallocates to the more productive vintages, but their MPS is cross-sectional. Bloom (2014) summarizes results on time-varying risk when shifts in the distributions of shocks are known.

Endogenous growth models have studied how growth responds to volatility in aggregate TFP shocks (Jones, Manuelli, Siu and Stacchetti 2005) and to policy shocks (Hopenhayn and Muniagurria 1996). We find that the qualitative effects of shocks to

diversify internationally to any significant extent.

[7] Collin-Dufresne, et al. (2016) also emphasize the importance of this martingale state. They study implications for asset pricing in economies in which consumption growth is exogenous, whereas we study production economies in which uncertainty influences consumption and investment.

the efficiency of investment are similar to those of investment taxes or subsidies.

Among related models with learning, Huffman and Kiefer (1994) and Koulovantianos, Mirman and Santugini (2009) studied learning the TFP process in the production of final goods, but not that of capital goods. Bernanke (1983) and Stokey (2016) study partial equilibrium models with delayed information arrival about investment profitability. As in our model, learning occurs exogenously, as a result of the passage of time alone. It does not depend on actions that agents take – on how much to invest, for instance. This is in contrast to Chamley and Gale (1994), Veldkamp (2005), and Pastor and Veronesi (2009) where actions do affect the speed of learning.

Parameter learning also implies momentum effects on stock prices – serially correlated price changes caused by learning. As Lewellen and Schanken (2002) have shown, parameter uncertainty drives a wedge between the distribution perceived by investors and the distribution estimated by empirical tests. Following a regime shift, an econometrician should expect to find predictability of returns. As agents learn the parameters stock prices in general start to reflect this. For instance, if a regime shift lowered the distribution of TFP, returns would initially be low, reflecting the fall in dividends and would then gradually rise as stock prices fall to reflect the lower mean of the dividend process.

In our model momentum effects arise for the risk-free rate, but not for stock prices. If investment is positive, Tobin's Q (the market value of capital relative to its replacement cost) is unity. Only when investment hits zero can Q fall below unity. As Sargent (1980, p. 111) puts it, roughly speaking, Q drops farther below unity as the constraint that investment be irreversible becomes more binding.

In that case, parameter uncertainty also leads to a temporary rise in the volatility of the stock market, and some regime changes may raise the possibility of disasters as in Veronesi (2004), Gourio (2008), and Wachter (2013). Regime shifts also raise uncertainty and if we had ambiguity-averse agents regime shifts would depress stock prices and raise returns that would gradually decline as learning takes place, as in Epstein and Schneider (2008).

The paper that is closest to ours is Bianchi and Melosi (2016). They study a family of DSGE models with Markov switching and Bayesian learning about latent states. Importantly, unlike anticipated-utility models of learning (Kreps 1998, Evans and Honkapohja 2001), their agents take regime shifts and learning fully into account when forming plans. For a real-business-cycle model in which TFP growth switches between high or low values, Bianchi and Melosi demonstrate that the anticipated-utility assumption is not innocuous because past mistakes about TFP growth are built into the capital stock and thus have persistent effects on output, consumption, and investment.[8]

We extend their analysis by breaking free of the assumption that there are finite number of potential states. In our model, there is a continuum of possible future states, and entirely new regimes can emerge. Because there is no worst possible state, the influence of uncertainty is magnified. When there is a worst possible state and agents are in it, they have nowhere to go but up. In addition, when in some other state, the extent to which they could fall is bounded. The bounding of downside risk lessens the need for precautionary wealth. On the other hand, with a continuum of possible regimes, conditions can always worsen, and the magnitude of a potential decline is unbounded. Hence, precautionary motives are stronger when downside risk is unbounded.[9,10]

Last but not least, what we call "parameter uncertainty" differs from Knightian uncertainty because agents can still form predictive distributions over variables – they are Bayesians. Knight's definition would more appropriately apply to boundedly rational agents. For example, Pintus and Suda (2019) feature adaptive learning and find, as we do for Bayes learning, that the responses of output, investment, and other aggregates under adaptive learning are significantly larger than when the agents know the parameters governing the shocks.[11] One could also analyze the effects of shifts when agents have ambiguous beliefs (Hansen and Sargent 2001; Bianchi, Ilut, and Schneider 2018).

The paper is organized as follows. Section 2 describes an Ak growth model with these features and characterizes decision rules and the equilibrium law of motion. Section 3 presents a number of simple examples to build intuition. Section 4 solves and simulates calibrated versions in order to assess the strength of various forces, especially those governing uncertainty and precautionary behavior.

[8] Bianchi and Melosi (2018) use similar methods to examine the role of policy uncertainty in a new Keynesian model with recurrent shifts in the monetary-policy rule.

[9] That period utility is also unbounded below is also important.

[10] For endowment economies, Geweke (2001) and Cogley (2009) examine the consequences for expected utility and the market price of risk, respectively.

[11] However, their adaptive-learning framework shuts down precautionary effects associated with uncertainty.

2 Model and Planner's Solution

Technology.—The model has two sectors: a final goods sector and a capital-goods sector. The production function for the final good is

$$Y = zk, \tag{1}$$

where Y is the output of final goods, k the capital-good input, and z a shock to the final goods technology. Capital depreciates at the rate δ; the law of motion for capital is

$$k' = (1-\delta)k + \frac{1}{q}X, \tag{2}$$

and the aggregate resource constraint is

$$zk = C + X, \tag{3}$$

where X is final goods devoted to the production of capital, and q is the shock to the capital-goods technology.

Preferences.—A representative agent has recursive preferences as in Epstein and Zin (1989) and Weil (1989),

$$U_t = \left[(1-\beta)C_t^{1-\rho} + \beta\left(\left(E_t[U_{t+1}^{1-\gamma}]\right)^{\frac{1-\rho}{1-\gamma}}\right)\right]^{\frac{1}{1-\rho}}, \tag{4}$$

where β is the subjective discount factor, γ is the coefficient of relative risk aversion, and ρ is the inverse of the elasticity of intertemporal substitution.[12] Expectations are taken with respect to subjective beliefs.

Aggregate resource constraint.—The Ak production structure in Eqs. (1) and (3) along with the homothetic preferences in Eq. (4) allows us to drop k from the set of states and scale variables by k as follows

$$c = C/k, \quad x = X/k, \quad y = Y/k \quad \text{and} \quad g = k'/k, \tag{5}$$

and write the law of motion for k in (2) as

$$g = 1 - \delta + \frac{1}{q}x, \tag{6}$$

and the aggregate resource constraint in (3) as

$$z = c + x. \tag{7}$$

[12] This simplifies to time-separable isoelastic preferences when $\rho = \gamma$.

Law of motion of the shocks.—Let $S = (q, z)$ follow a Markov process with distribution $\psi(S' \mid S, \theta)$ where θ represents the parameters of the transition law. Some components of θ will be unknown. We will be interested in comparing the behavior of an economy in which agents learn about θ over time to an economy in which θ is known.

Updating of beliefs when there is no regime shift.—A "regime" is denoted by θ. Let $\mu(\theta)$ denote beliefs over θ. Each period the agents observe S_t and update μ using Bayes law. After observing the value S',

$$\mu'_{\text{no shift}}(\theta) = \frac{\psi(S' \mid S, \theta)\,\mu(\theta)}{\int \psi(S' \mid S, \theta)\,d\mu(\theta)} \equiv b(\theta \mid S', S, \mu). \tag{8}$$

Regime shifts.—We assume that each period a regime shift occurs with probability λ. We expect λ to be close to zero given the above discussion of the probable infrequency of regime shifts. The simplest treatment is to posit a transition law for regimes upon a shift. So we assume that the distribution of the new regime θ' conditional on being in regime θ and conditional on there being a regime shift is $\pi(\theta' \mid \theta)$ with $\theta \in \Theta$ and π a measure on Θ. In some examples we shall assume that Θ is a finite set so that θ follows a discrete-state Markov process as in Hamilton (1989), Bianchi and Melosi (2016), and Foerster *et al.* (2016).

Timing of a regime shift.—Let $\chi = 1$ denote a regime shift and $\chi = 0$ denote no shift. When a shift occurs, it does so at the start of the production period so that agents update by observing S and that $\chi = 1$, i.e., with the knowledge that a break has occurred.[13]

Updating of beliefs after a regime shift.—We shall assume that agents know the law $\pi(\cdot)$ and whether a regime shift has just occurred, but that they know neither θ (over which they hold beliefs μ) nor θ'. In that case

$$\mu'_{\text{shift}}(\theta') = \int \pi(\theta' \mid \theta)\, b(\theta \mid S', S, \mu)\, d\theta. \tag{9}$$

I.e., S' is generated by the pre-regime-shift θ via the likelihood ψ.

[13] The assumption that break dates are known is for tractability. Some of our examples are simple enough to handle unobserved regime shifts (e.g. the finite-state Markov model in section 3.2.2), but others are not. For instance, in the models of section 4, time since the last break is a state variable. If break dates were unobserved, the *distribution* over time since the last break would be a state variable, and the resulting curse of dimensionality would make the calculations intractable. We are still thinking about ways to address this issue.

General updating.—Beliefs update as follows:

$$\mu'(\theta) = (1 - I_{\chi=1})\,\mu'_{\text{no shift}}(\theta) + I_{\chi=1}\mu'_{\text{shift}}(\theta)\,. \tag{10}$$

Plan for the analysis.—Because the model features no monopoly power and has no externalities, the (recursive) competitive equilibrium will coincide with the planner's solution. Of course the planner too has to learn θ by observing s_t in the same way as private agents do. We shall therefore solve the planner's problem first, and then discuss the markets for goods, for capital and for assets that decentralize the optimum.

2.1 Planner's solution

Let μ denote the planner's belief and S the current shock. Denote the augmented state by $s \equiv (S, \mu)$.[14] The next period's state is $s' = (S', \mu')$. The distribution of S' is $\int \Psi(S' \mid S, \theta) d\mu(\theta)$ and it is the same whether a regime occurs or not. The only thing that is affected by a regime shift is μ'. Therefore,

$$F(s' \mid s, \chi') = \begin{cases} F\left(S', \mu'_{\text{shift}} \mid s\right) & \text{if } \chi' = 1 \\ F\left(S', \mu'_{\text{no shift}} \mid s\right) & \text{if } \chi' = 0. \end{cases} \tag{11}$$

so the planner's predictive distribution is

$$F(s' \mid s) = \lambda F(S', \mu'_{\text{shift}} \mid s) + (1 - \lambda)\, F(S', \mu'_{\text{no shift}} \mid s), \tag{12}$$

and it underlies the expectations operator in the next two equations and in Appendix B. The planner's state vector is (s, k). He chooses C and X to maximize his value function V

$$V(s_t, k_t) = \max_{C_t, X_t} \left[(1 - \beta) C_t^{1-\rho} + \beta \left(\left(E_t[V(s_{t+1}, k_{t+1})^{1-\gamma}] \right)^{\frac{1-\rho}{1-\gamma}} \right) \right]^{\frac{1}{1-\rho}},$$

subject to (2), (7), (10) and (12).

Because production has constant returns and preferences are homogeneous of degree $1 - \rho$, we can eliminate scale, k, from V and deal in terms of s alone. We then have the following:

Proposition 1 • *The value function takes the form*

$$V(s, k) = v(s)k, \tag{13}$$

[14]Since χ is i.i.d. and since beliefs μ update right after χ is observed, including μ as a part of the state means that χ drops out of the state vector.

where

$$v(s) = (z + q(1-\delta))\left[(1-\beta)(1+q^{1-\frac{1}{\rho}}w(s))^{\rho}\right]^{\frac{1}{1-\rho}}. \tag{14}$$

- *Consumption c takes the form*

$$c = \frac{z + q(1-\delta)}{1 + q^{1-\frac{1}{\rho}}w(s)}. \tag{15}$$

- *Growth g takes the form*

$$g = 1 - \delta + \frac{zq^{-\frac{1}{\rho}}w(s) - (1-\delta)}{1 + q^{1-\frac{1}{\rho}}w(s)}, \tag{16}$$

and $w(s)$ satisfies the recursion

$$w(s) = \left(\frac{\beta}{1-\beta}\right)^{\frac{1}{\rho}}\left[E\left((z' + q'(1-\delta))\left[(1-\beta)(1+q'^{1-\frac{1}{\rho}}w(s'))^{\rho}\right]^{\frac{1}{1-\rho}}\right)^{1-\gamma} \mid s\right]^{\frac{1-\rho}{\rho(1-\gamma)}}. \tag{17}$$

The proof is in Appendix B. The relative simplicity of the value in (14) and the policy rules in (15) and (16) stems from the model's Ak structure. The proposition helps because the right-hand side of (17) contains no max operator; only the level of w and not its derivatives enter the solutions in (14), (15) and (16), and this will help characterize them.

Proposition 1 encompasses a number of applications. Each uses a different predictive distribution for (z', q') for the expectations operator in (17), but no other change to equations (13)-(17). The applications include

1. conventional finite-state Markov-switching models in which the dates of breaks and parameters of regimes are known (e.g., see Foerster, *et al.* 2016 and the references therein);[15,16]

2. finite-state Markov-switching models in which the dates of breaks are known but parameters of some regimes are unknown (e.g., Schorfheide 2005, Bianchi and Melosi 2016);

[15] The posterior μ would drop out of the state vector in this case.

[16] Because proposition 1 provides a semi-analytic solution, it could be used to assess the accuracy of numerical approximations for models in this class.

3. finite-state hidden Markov models in which both the dates of breaks and regime-specific parameters are unknown (e.g., Bianchi and Melosi 2016);

4. models with nonrecurrent structural breaks across a continuum of regimes in which dates of breaks are known but regime-specific parameters are unknown (section 4);

5. models with nonrecurrent structural breaks across a continuum of regimes in which dates of breaks and regime-specific parameters are unknown (future research).

Section 3 provides examples of type-2 and 3 models, while section 4 examines a type-4 model. Type-5 models are left for future research, primarily for computational reasons.[17]

3 Examples of type-2 and 3 models

This section specializes the model so as to focus on two specific questions. We assume that

(*i*) $\rho = \gamma$ so that utility is the time additive CRRA form $E\sum_0^\infty \beta^t \frac{1}{1-\gamma} c_t^{1-\gamma}$, and

(*ii*) S is i.i.d. conditional on θ, i.e., $\psi(S' \mid S, \theta)$ is independent of S.

We shall analyze parameter learning about the distribution of z separately from that of parameter learning about the distribution of q. I.e., we shall analyze them one at a time. When learning about the distribution of z we shall assume that

$$z = a_z + \frac{A_z}{1 + e^{-x_z}}, \tag{18}$$

where $x_z \sim N(\theta_z, \sigma_z^2)$. The parameters a_z, A_z and σ_z^2 will be known, and θ_z will be unknown. And when learning about the distribution of q we shall assume that

$$q = a_q + \frac{A_q}{1 + e^{-x_q}}, \tag{19}$$

where $x_q \sim N\left(\theta_q, \sigma_q^2\right)$, where $\left(a_q, A_q, \sigma_q^2\right)$ are known, and θ_q is unknown. We shall henceforth drop the z and q subscripts from x and from the parameters (a, A, θ). We now ask two questions:

[17] This case is computationally very hard to handle because it involves keeping track of a posterior distribution over dates of structural breaks and parameters of past regimes. This complication is more than we can currently manage.

1. How does a learning-about-θ episode compare to an episode in which θ is known but in which the variance of the shocks x temporarily rises? Sec. 3.1 will show that uncertainty induced by structural breaks has substantially larger and qualitatively different effects than uncertainty triggered by volatility of the exogenous shocks. This is because uncertainty about the parameters of the data generating processes has larger effects on agents' lifetime utility.

2. Second, how does the case where the occurrences of regime shifts, χ, are observed differ from the case where they are not observed? Sec. 3.2 will show that the value of being informed about regime shifts is three to four times as large for the case of z as it is for the case of q.

3.1 Parameter learning vs. higher shock volatility

Suppose that a structural break has just occurred, and that agents know this. They also know that thereafter there will be no more breaks, so that $\lambda = 0$ in (12). Following the structural break, the prior variance over θ rises, implying a temporary rise in the variance of the predictive distribution over s in (17). This will be an example of a type 2 model.

To put the two cases on a comparable footing, we equate the one-step ahead marginal distributions of the shocks at each t. We wish to compare mean preserving spreads in the distributions of shocks to learning, while keeping the means the same for the two cases. We shall make the comparison for a single learning episode that is not interrupted by a new regime shift ($\lambda = 0$).

The learning economy.—Agents (or the social planner) know the exact values of the parameters $\{\theta_L, \theta_M, \theta_H\}$, but they do not know the regime in place, they have to learn that. Let $h^t = (x_s)_{s=0}^{t-1}$ denote the history of realizations of x,[18] and let $L(h^t \mid \theta_i)$ denote the likelihood of h^t. Assuming that the prior over $\{\theta_L, \theta_M, \theta_H\}$ is 1/3,1/3,1/3, Bayes rule yields the posterior probability that $\theta = \theta_i$ as

$$\mu_i\left(h^t\right) = \frac{L\left(h^t \mid \theta_i\right)}{\sum_{j\in\{L,M,H\}} L\left(h^t \mid \theta_j\right)},$$

[18] Agents see z_t or q_t and can recover x_t by inverting (18) or (19) to get

$$x_t = \ln\left(\frac{z_t - a}{A - z_t + a}\right).$$

and the posterior mean and variance as

$$\begin{aligned} \bar{\theta}\left(h^{t}\right) &= \sum \theta_{j}\mu_{j}, \quad \text{and} \\ \sigma_{\theta}^{2}\left(\mu\right) &= \sum \left(\theta_{j}-\bar{\theta}\left[h^{t}\right]\right)^{2}\mu_{j}. \end{aligned} \tag{20}$$

Suppose that the draw following the regime shift is $\theta = \theta_M$, although this will only gradually become apparent to agents as they learn. We condition the distribution of the history on $\theta = \theta_M$, so that $x_t = \theta_M + \varepsilon_t$. The total variance of x is

$$\sigma_{x,t}^{2} = \sigma_{\theta,t}^{2} + \sigma_{\varepsilon}^{2}, \tag{21}$$

where $\sigma_{\theta,t}^{2}$ is the expected variance conditional on θ_M,

$$\sigma_{\theta,t}^{2} = \int \sigma_{\theta}^{2}\left(\mu\left[t,u\right]\right)\frac{\sqrt{t}}{\sigma_{\varepsilon}\sqrt{2\pi}}\exp\left(-\frac{\left(u-\theta_{M}\right)^{2}}{2\sigma_{\varepsilon}^{2}t^{-1}}\right)du. \tag{22}$$

The no-learning economy.—In this economy agents *know* that $\theta = \theta_M$, and they face a time-varying variance of ε given by the right-hand side of Eq. (21), $\sigma_{\varepsilon,t}^{2} = \sigma_{x,t}^{2}$. We shall use the symbol F_t to denote the resulting sequence of time-varying distributions. We treat z learning and q learning one at a time. The treatment of the two is parallel, and we illustrate the treatment of the varying risk value function for the case where z is fixed and only q varies and only q is subject to a possible regime shift. Since the shock variances now depend on t, so does the value:

$$V_t(q,k) = \max_{C,X}\left\{\frac{C^{1-\gamma}}{1-\gamma} + \beta\textstyle\int V_{t+1}(q',k')dF_t(q')\right\} = v_t(q)k_t^{1-\gamma},$$

where, using the aggregate resource constraint (7) and law of motion for capital (2),

$$\begin{aligned} v_t(q) &= \max_{x}\left\{\frac{\left(z-x\right)^{1-\gamma}}{1-\gamma} + \beta\left(1-\delta+\frac{1}{q}x\right)^{1-\gamma}\textstyle\int v_{t+1}(q')dF_t(q')\right\}, \\ &= \max_{x}\left\{\frac{\left(z-x\right)^{1-\gamma}}{1-\gamma} + \beta\frac{\left(1-\delta+\frac{1}{q}x\right)^{1-\gamma}}{\sigma_{\varepsilon,t}\sqrt{2\pi}}\textstyle\int v_{t+1}(q')\exp\left(-\frac{\left(q'-\theta_{M}\right)^{2}}{2\sigma_{\varepsilon,t}^{2}}\right)dq'\right\}. \end{aligned} \tag{23}$$

For the known-θ case, c_t and g_t are computed directly from equations (15), (16), and (17).[19]

[19] The proof of Proposition 1 does not impose the constraint that $g \geq 1-\delta \iff x \geq 0$ (see Eq. (6)) but the inequality always holds in our simulations.

Parameters are calibrated so that the time series means of (c, g) in the learning economy match $c = C/K = 0.3$ and $g = 1.02$ per annum. Results are shown in table 1. The same parameters are used for the known case at $\theta = \theta_M$.

TABLE 1: Parameters used for figure 2

	γ	β	δ	q	z	A	a	θ_L	θ_M	θ_H
Vary z	4	0.95	0.05	1.160	random	0.113	0.164	-0.780	-0.428	0.213
Vary q	4	0.95	0.05	random	0.340	0.932	1.355	-0.523	-0.124	0.521

The top row of figure 2 compares the two economies when the injected uncertainty is over the distribution of z. After 10 years, cumulative growth is 18.49% in the learning economy and 16.61% in the MPS case. Since the two economies have the same C/K ratios after a decade and K is 1.88 percent higher in the learning economy, it follows that consumption is permanently higher by 1.88 percent.

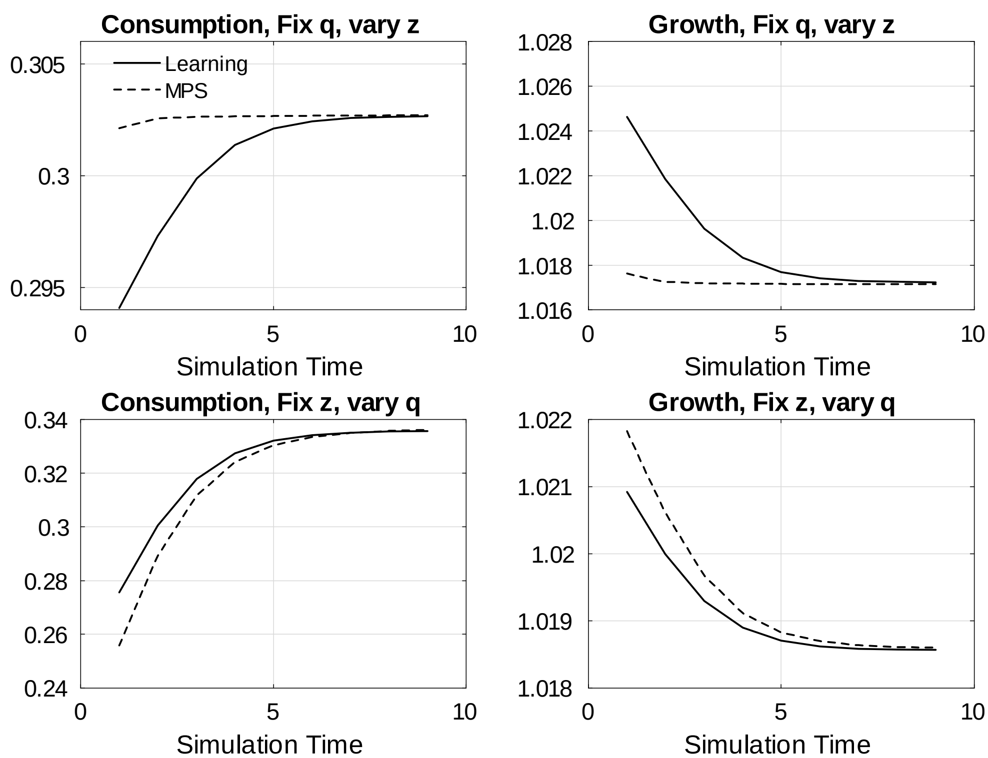

Figure 2: Comparison of learning (solid line) and MPS (dashed line) when the shock distribution is unknown. Top row, z. bottom row, q.

The bottom row does the same thing for learning about the distribution of q. In this case, the results go in the opposite direction. Cumulative 10-year growth is

22.22% for the MPS and 21.45% for the learning economy. The difference amounts to a permanent consumption dividend of 0.77 percent per annum for agents in the MPS economy.

Precautionary saving and investment versus the option to wait.—The marginal distributions of the shocks are equated so that $\sigma^2_{\varepsilon,t} = \sigma^2_{x,t}$, but since beliefs are a Martingale in the learning economy, there is a permanent component to shocks that is absent in the MPS economy. Thus in the top two panels a shock to z has a longer expected duration in the learning case, and to prepare for this agents save more than in the MPS case. In the bottom two panels it is the shock to q that now has a longer expected duration in the learning case, and in Sec. 4 we will find that under learning, precautionary savings dominates the option to wait, but by much less for learning about q than it does for learning about z.[20] In line with that finding, learning now generates more consumption and less growth than does MPS.

3.2 Observed vs. unobserved χ_t when breaks recur

We now assume breaks are recurrent ($\lambda > 0$), and that there are only two regimes $\theta \in \{\theta_L, \theta_H\}$. Agents again know the parameters $\{\theta_L, \theta_H\}$, but they do not know the regime in place. We aim to compare a type-2 version of the model in section 3.2.1 with a type-3 version of its equivalent in section 3.2.2.

With just two possible values of θ, beliefs can be summarized by the real number $\mu = \Pr(\theta = \theta_H)$ or, equivalently, by the expectation $E(\theta) = \mu\theta_H + (1-\mu)\theta_L$. Moreover, we assume $\theta_L = -1$ and $\theta_H = 1$, so that

As in the previous subsection, we shall study regime shifts as shifts in the distribution of one of the shocks alone, keeping the value of the other shock fixed. In light of (18) and (19) for z, the favorable state is θ_H whereas for q, the favorable state is θ_L.

3.2.1 Observed break dates

When a break occurs, each regime is equally likely; in other words, $\pi(\theta' \mid \theta) = 1/2$ for all θ and θ'. Following a regime shift, $E(\theta) = 0$, and then $E(\theta)$ drifts towards the realized $\theta \in \{-1, 1\}$ unless the next regime shift occurs in which case $E(\theta)$ reverts to zero.

[20] Compare the top-left panels of figures 7 and 10.

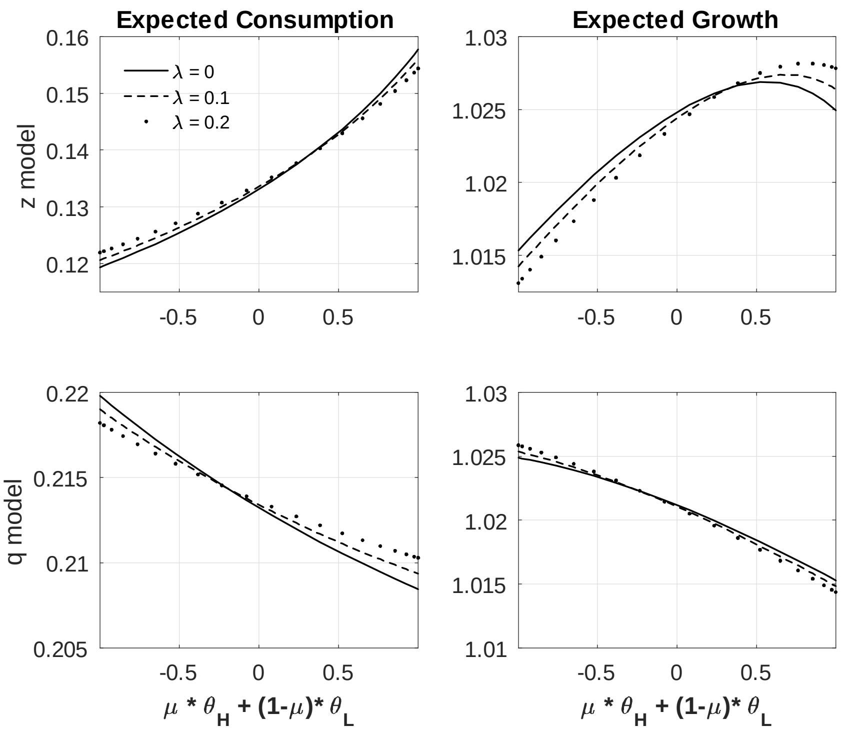

Figure 3: Expected consumption and growth when χ is observable. Top row: only θ^z varies; bottom row: only θ^q does.

Fig. 3 portrays expected consumption and growth for various values of λ. On the horizontal axis is $E(\theta)$. At higher λ the belief μ matters less because the regime is more likely to soon change. As λ rises, regime persistence falls and the curves portrayed in the left-hand panels become flatter. A higher λ lowers consumption and raises growth in 'good' regimes and does the opposite in 'bad" regimes. For z, the 'good' θ state is θ_H and consumption and savings (and growth) both rise as beliefs shift towards the regime θ_H. For q the good state is θ_L and so the curves slope down, but the consumption curve again flattens out as λ rises.

3.2.2 Unobserved break dates

We now assume the Markov transition probabilities,

$$\theta_z \quad \begin{array}{c} \\ \theta_L \\ \theta_H \end{array} \overset{\theta'_z}{\begin{array}{cc} \theta_L & \theta_H \\ \left[\begin{array}{cc} \pi^z_{LL} & \pi^z_{LH} \\ \pi^z_{HL} & \pi^z_{HH} \end{array}\right] \end{array}}, \qquad \theta_q \quad \begin{array}{c} \\ \theta_L \\ \theta_H \end{array} \overset{\theta'_q}{\begin{array}{cc} \theta_L & \theta_H \\ \left[\begin{array}{cc} \pi^q_{LL} & \pi^q_{LH} \\ \pi^q_{HL} & \pi^q_{HH} \end{array}\right] \end{array}},$$

where $\left(\theta_z', \theta_q'\right)$ are mutually independent.[21] We also allow the structural break indicator χ_t to be observed or unobserved. If the χ_t are not observed, situations may arise in which agents think that a break occurred, while in fact it did not.[22] For this example, we are mainly interested in building intuition about the consequences of assuming that agents know when a break has occurred.

Observability of χ adds nothing if $\lambda \in \{0, 1\}$. When $\lambda = 0$ the regime never changes and when $\lambda = 1$ it changes every period. In these extreme cases, however, there are *de facto* no regime shifts since even if $\lambda = 1$, θ is absorbed into an i.i.d. component of the residual. We therefore derive the consequences for $\lambda \in (0, 1)$.

If χ_t is observable, we have

$$\begin{aligned} \mu'_{\text{shift}}\left(\theta'\right) &= \sum_{\theta \in \{\theta_L, \theta_H\}} \mu'_{\text{no shift}}\left(\theta\right) \pi_{\theta, \theta'} \qquad \text{when } \chi = 1, \\ &= \mu_{\text{no shift}}\left(\theta'\right) \qquad \text{when } \chi = 0. \end{aligned}$$

Defining $\mu \equiv \mu\left(\theta_H\right)$, it follows that

$$\mu' = (1 - \lambda)\, b\left(\theta_H \mid S', \mu\right) + \lambda\left(b\left(\theta_L \mid S', \mu\right) \pi_{LH} + b\left(\theta_H \mid S', \mu\right) \pi_{HH}\right). \tag{24}$$

Then $w\left(s, \mu\right)$ is defined via (8), (9), (12) and (17).

On the other hand, if χ_t is unobserved, beliefs are updated as follows. Agents start the period with a prior μ over $\theta = \theta_H$. An observation S leads to the posterior belief

$$\mu\left(\theta_H \mid S, \mu\right) = \frac{L\left(S \mid \theta_H\right) \mu}{L\left(S \mid \theta_H\right) \mu + L\left(S \mid \theta_L\right)(1 - \mu)} \equiv b\left(S, \mu\right).$$

Then next-period belief is

$$\mu' = \lambda\left(\pi_{HH} b\left(S, \mu\right) + \pi_{LH}\left[1 - b\left(S, \mu\right)\right]\right) + (1 - \lambda)\, b\left(S, \mu\right) \tag{25}$$

The joint distribution of (μ', s') in Eq. (11) now changes as follows. In $F(s', \mu' \mid s, \mu)$, χ is no longer a state so that $s = S$. Then $\mu' = b\left(S, \mu\right)$ is not random once we condition on $\left(S, \mu\right)$ because we no longer have μ'_{shift} and $\mu'_{\text{no shift}}$ induced by an observation of χ', and in Eq. (17) $F(s', \mu' \mid s, \mu)$ is replaced by

$$F(S' \mid S, \mu) = \begin{cases} L\left(S' \mid \theta_H\right) & \text{w. prob. } \mu' \\ L\left(S' \mid \theta_L\right) & \text{w. prob. } 1 - \mu' \end{cases}.$$

[21] In Appendix C, we study the alternating-regime case (in which $\pi_{\text{LL}} = \pi_{\text{HH}} = 0$) analogous to that of Hopenhayn and Muniagurria (1996).

[22] Bianchi and Melosi (2016) study such a scenario, but with a different learning mechanism. Andolfatto and Gomme (2003) study a monetary economy with two money-growth regimes in which agents may or may not observe the break dates.

where μ' is given in (25). The decision rules remain as in Eqs. (15) and (16). The unconditional long-run expectation of θ is

$$\hat{E}(\theta) = \frac{\pi_{LH}}{\pi_{LH} + \pi_{HL}} \theta_H + \frac{\pi_{HL}}{\pi_{LH} + \pi_{HL}} \theta_L. \tag{26}$$

Beliefs remain non-degenerate.—Unless θ is forever fixed and that agents know that it is fixed, beliefs will remain non-degenerate. When χ is unobserved the agents never learn which of the two values of θ is operative: Beliefs remain a non-degenerate stochastic process even if (unbeknownst to the agents) θ remained forever fixed at θ_H, say. If agents' beliefs placed all weight on θ_H, say, at $t-1$, they would become non-degenerate at t, because of the positive probability $\lambda\pi_{HL}$ that θ will have shifted due to a possible regime shift. In other words, although they know the set $\{\theta_L, \theta_H\}$, agents will typically not know which regime θ is in place.

For example, suppose parameters are fixed as in table 2. These parameters are based on the following accounting: The NBER data[23] indicate 215 expansions, of which 204 (94.88%) are followed by an expansion, and 11 (5.1%) are followed by a recession. On the other hand, there are 41 recessions, of which 11 (26.8%) are followed by an expansion, and 30 (73.1%) are followed by another recession. We use the notation π_{LH} and π_{HL} for both cases – learning about z and learning about q since here they are obtained by the same accounting method.

TABLE 2: Parameters chosen for figure 4

θ_L	θ_H	π_{LH}	π_{HL}	λ
-1	1	0.27	0.05	0.29

With these parameters, $\hat{E}(\theta) = 0.68$, which is indicated by the vertical lines in figure 4. Not observing χ means that the planner acts on the basis of less information than in the observable-χ case. This appears to imply more uncertainty over the future realizations of (z, q), and we expected to see a rise in precautionary savings in response. For most values of μ, this occurs when learning about shifts in q (see the bottom row) but not for shifts in z (shown in the top row). The learning-about-z model has the surprising outcome that for most values of μ, agents consume more when regime switches are not observed.

[23]See National Bureau of Economic Research (2017, https://www.nber.org/research/data/us-business-cycle-expansions-and-contractions.

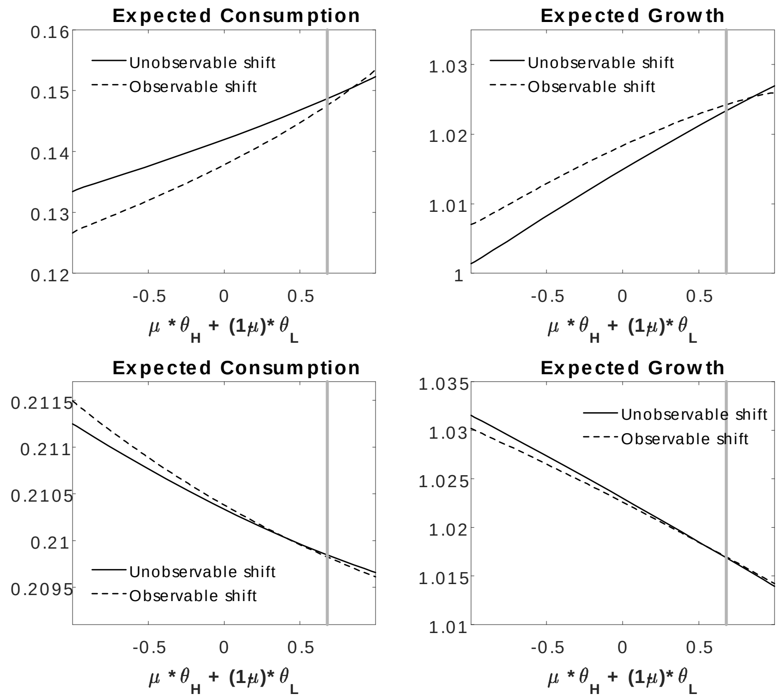

Figure 4: Expected consumption and growth when χ is unobservable. Top row: only θ^z varies; bottom row: only θ^q does.

3.2.3 The value of being able to observe χ

As noted above, the ability to observe χ_t is of no value if $\lambda \in \{0, 1\}$, because then one would know that the regime never changes (when $\lambda = 0$) or that it changes every period (when $\lambda = 1$).

Figure 5 portrays the relative value of information $v^I(s, \mu)/v^U(S, \mu)$ as λ varies. The experiment fixes $\mu = 0.5$ and sets $z = 0.34$ and $q = 1$. With $\gamma > 1$, $v^I(s, \mu)$ and $v^U(s, \mu)$ are both negative. Then $v^I(s, \mu) > v^U(s, \mu) \Rightarrow v^I(s, \mu)/v^U(S, \mu) < 1$. At $\lambda = 0$ and 1, the ratio is equal to 1. For intermediate values of λ, being informed about regime shifts raises $v(s, \mu)$, and the value of information rises with the degree of risk aversion.

One may have expected the value of information to be highest for values of λ around 1/2. Figure 5 shows, however, that the value of information is consistently maximized at values of λ well below 1/2, and closer to a value of 1/3 or even 1/4, depending on the value of γ. Knowing the current regime has little value when regimes are not persistent. Similarly, the benefits of knowing when to start learning is lower

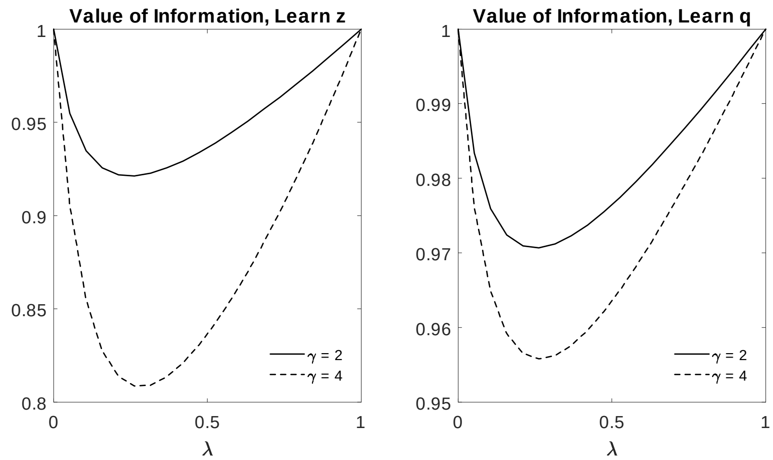

Figure 5: Value of information about regime shifts

when regimes are not persistent.

The value of information is larger when γ is large, and larger for the case of z than it is for the case of q. Overall, the relative change $\left(v^I - v^U\right)/v^U$ is three to four times as large for z as it is for q.

4 A Type-4 example: non-recurrent structural breaks and autoregressive shocks

Next we turn to examples in which shocks are autoregressive processes,

$$\begin{aligned} \ln z_t &= \mu_{zt} + \rho_z \ln z_{t-1} + \sigma_z \varepsilon_{zt}, \\ \ln q_t &= \mu_{qt} + \rho_q \ln q_{t-1} + \sigma_q \varepsilon_{qt}. \end{aligned} \tag{27}$$

To keep things simple, we activate one at a time. The first example features neutral technology shocks with a constant level of q, while the second fixes z and activates investment shocks. Letting a_t represent the active shock, the autoregressive parameter ρ_a and conditional standard deviation σ_a are assumed to be known, and the intercept μ_{at} is unknown. Agents must learn about μ_{at}.[24]

[24]Proposition 1 could also handle breaks in the autoregressive parameter ρ_a and/or volatility parameter σ_a. We focus on breaks in μ_a because that is the most parsimonious way to fit breaks in data on $\ln z$ and $\ln q$ (see figure 1).

The intercept is subject to occasional structural breaks,

$$\begin{aligned} \mu_{at} &= \mu_{at-1} \quad w/\ pr\ 1-\lambda, \\ &= m_t \sim N(m, \sigma_m^2) \quad w/\ pr\ \lambda. \end{aligned} \tag{28}$$

A Bernoulli random variable χ_t governs whether a structural break occurs. With probability $1-\lambda$, no break occurs ($\chi_t = 0$), and the intercept remains unchanged. With probability λ, a break occurs ($\chi_t = 1$), and a new intercept is drawn from a normal distribution. The random variables $\varepsilon_{at}, \chi_t, m_t$ are mutually independent, and the parameters $\theta = [\lambda, m, \sigma_m, \rho_a, \sigma_a]$ are known.

Agents observe χ_t and a_t, but not μ_{at} or m_t. Because χ_t is observable, agents know when a break occurs, but they don't know whether μ_{at} has increased or decreased or by how much. They update beliefs about μ_{at} by applying Bayes theorem.

To fit this example into our general framework, let $s = [a, \chi]$ represent the observable states, and write the Markov transition equation as $\Psi(s_{t+1}, \mu_{t+1}|s_t, \mu_t, \theta)$. The predictive density is

$$F(s_{t+1}|s^t, \theta) = \iint \Psi(s_{t+1}, \mu_{t+1}|s_t, \mu_t, \theta) p(\mu_{t+1}, \mu_t|s^t, \theta) d\mu_{t+1} d\mu_t, \tag{29}$$

where $p(\mu_{t+1}, \mu_t|s^t, \theta)$ is the posterior for the unobserved intercept. Proposition 1 goes through with this specialization of notation. All that remains is to derive the posterior $p(\mu_{t+1}, \mu_t|s^t, \theta)$ and solve the integral equation 17. Appendix D addresses these problems and also explains how the solution for $w(s)$ is computed.

4.1 Neutral technology shocks

Table 3 calibrates parameters for an economy in which the neutral technology shock z is active and q is constant. The discount factor β is set *a priori* to a standard RBC value for annual data, and the coefficient of relative risk aversion γ is set to 4. Consumers are myopic if $\gamma = 1$, and a higher value is needed for learning to matter. However, Pratt-style thought experiments suggest that γ should not be much higher than 1. A value of 4 seems like a reasonable compromise.

TABLE 3: Calibration for structural breaks in *TFP*

β	γ	ρ	δ	ρ_z	σ_z	λ	m	σ_m	$\bar{z}$	$\bar{q}$
1.01^{-4}	4	12.95	0.119	0.8867	0.0422	0.0091	−0.0893	0.0234	0.464	1

Parameters governing $\ln z$ are estimated by combining the data shown in figure 1 with informative priors over $(\rho_z, \sigma_z, \lambda, m, \sigma_m)$. In a nutshell, the process for $\ln z$ can be expressed as a non-Gaussian state-space model whose log likelihood function can be evaluated via a particle filter. Because structural breaks are infrequent, λ and σ_m are weakly identified in a frequentist sense. We therefore add an informative prior and maximize the log posterior. Details can be found in appendix E. Table 3 reports the posterior mode for $(\rho_z, \sigma_z, \lambda, m, \sigma_m)$, and $\bar{z}$ is the implied unconditional geometric mean. Since no model for $\ln q$ is on the table at this point, $\bar{q}$ is set to 1.

Even with structural breaks, $\ln z$ is a stationary random process, although it would be hard to distinguish from an integrated process in samples of 100 or 200 years. That breaks are rare means that agents face long-run risks analogous to those of Bansal and Yaron (2004). The left hand column of figure 6 illustrates the long-run risk by comparing unconditional distributions for processes with and without structural breaks.[25] When structural breaks are absent (solid lines), $\ln z$ is unconditionally normal with mean $m/(1-\rho_z)$ and variance $\sigma_z^2/(1-\rho_z^2)$. With rare structural breaks ($\lambda = 0.0091$, dashed lines), the distribution is still centered on $m/(1-\rho_z)$, but the tails are fatter. In fact, in the tails, the log density for the rare-break process is 10 times greater than that of the no-break process.[26]

The remaining parameters – ρ and δ – are calibrated by matching aspects of the deterministic steady state.[27] When all shocks are deactivated ($\sigma_z = \lambda = 0$), steady-state growth and the consumption-income ratio are

$$
\begin{aligned}
g_d &= \beta^{1/\rho}\left(\frac{\bar{z}}{\bar{q}} + 1 - \delta\right)^{1/\rho}, \\
\left(\frac{C}{Y}\right)_d &= 1 + (1 - \delta - g_d)\frac{\bar{q}}{\bar{z}}.
\end{aligned}
\tag{30}
$$

[25] The $\ln q$ calibration and stochastic volatility model are discussed below.

[26] Because structural breaks are so rare, an extremely long time series would be needed to consistently estimate moments of the unconditional distribution. E.g., for $\lambda \approx 0.01$, 1000 years of data would contain 10 breaks, on average. For samples like ours, however, a comparison between the model's unconditional distribution and empirical moments would be extremely noisy. For other applications in which λ is larger, showcasing the model's implications for the entire stationary distribution would be interesting and more promising.

[27] Results are similar if ρ and δ are calibrated to the mean of a stochastic sequilibrium in which shocks to $\ln z$ are active but structural breaks are not ($\sigma_z \neq 0, \lambda = 0$).

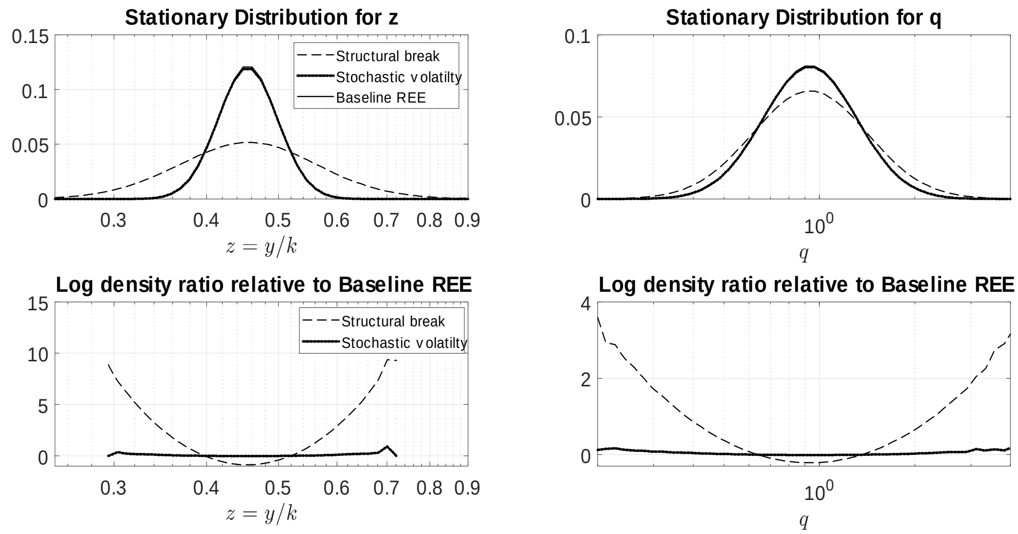

Figure 6: Unconditional shock distributions

The parameters ρ and δ are chosen to match $g_d = 1.02$ and $(C/Y)_d = 0.7$.[28,29] The solution for $\delta = 0.119$ is close to the standard RBC calibration for annual data. The solution for $\rho = 12.95$ implies that the elasticity of intertemporal substitution (EIS) is 0.077 which is broadly consistent with estimates based on micro data (e.g., Hall 1988, Vissing-Jorgensen 2002, Yogo 2004, and Havranak 2015). For instance, in a meta-analysis of 169 publications, Havranak (2015, p.1196) reports a 'best practice' estimate of $1/\rho = 0.33$ with a confidence interval of (-0.2,0.8). It follows that our consumers have a strong preference for smooth consumption streams.

The *EIS* is lower than typical calibrations in macroeconomics, which are often based on log utility ($\rho = \gamma = 1$). In this case, the unique solution to equation (17) is $w(s) = \beta/(1-\beta)$, the posterior μ drops out of the decision rules, and the model reduces to a conventional AK structure. We depart from log utility in order to activate a role for subjective beliefs. In addition, our model could not hit the macro

[28]The target for the consumption-income ratio is based on US data for the period 1950-2015. Consumption is measured by real expenditures on nondurable goods and services, and investment is real gross private domestic investment plus expenditures on consumer durables. Because the model has no government or net exports, output is measured by $C + I$. If government spending were included and split between consumption and investment, the mean C/Y ratio would be around 0.75.

[29]In section 3.1, we targeted the ratio $C/k = 0.3$. Since $Y = zk$, $C/Y = (C/k)/z$. Based on the values assumed in Tables 4.1 and 4.2, where $\bar{z} = 0.46$, the implied C/Y in section 3.1 would be $(C/k)/z = 0.3/0.46 = 0.65$, which is not far from the value 0.7 imposed here.

targets $g_d = 1.02$ and $(C/Y)_d = 0.7$ if the *EIS* were equal to 1.[30]

The *EIS* is also lower than typical calibrations in the finance literature; e.g. Bansal and Yaron (2004) assume a value of $1/\rho = 1.5$. As they emphasize, a model with *EZW* preferences requires both long-run risk and an *EIS* greater than 1 to resolve asset pricing puzzles. Structural breaks generate long run risk, but hitting the macro targets would be even more difficult if the *EIS* were greater than 1. Furthermore, attempts to calibrate to the risk-free rate plus one of the macro targets were unsuccessful, with moment conditions either failing to solve altogether or resulting in implausible values for other parameters. We chose the calibration in table 4.1 so that the model generates plausible macro outcomes, but we acknowledge that the model will not match asset prices with an *EIS* close to zero.

Incidentally, there is also a close relation between our model and the finance literature on savings and portfolio choice problems with uncertain returns (e.g. Campbell and Viceira 1999). Define the state variable a as savings in the risky asset last period, and write the budget constraint as

$$a' = Ra - c, \tag{33}$$

where the return on savings is

$$R = \frac{z' + (1-\delta)q'}{q}. \tag{34}$$

Then the value function is linear in a, and the model can be solved by adapting the equations in proposition 1.

Of course, the growth model is different from the finance setup because the way the return process is set up is different, and shock variances enter in particular way that have a technological interpretation. Another important difference concerns the calibration of the elasticity of intertemporal substitution. As explained in section 4, an EIS around 0.075 is needed for our model to target a mean consumption-GDP

[30] Given the macro targets $g_d = 1.02$ and $(C/Y)_d = 0.7$ plus $\bar{z}/\bar{q} = 0.464$, the second row of equation (30) pins down $\delta = 0.1192$. Taking logs of of the first row yields

$$\ln g_d = \rho^{-1} \ln \beta + \rho^{-1} \ln(\bar{z}/\bar{q} + 1 - \delta). \tag{31}$$

With $\beta = 0.961$, we have

$$0.0198 = \rho^{-1}(-0.0398 + 0.296), \tag{32}$$

implying $\rho^{-1} = 0.077$. One or more of the other parameters would have to be grossly implausible to raise ρ^{-1} to 1 or above.

ratio of 0.7 and a mean growth rate of 0.02. Rare structural breaks do generate a form of long-risk, but the finance literature emphasizes that resolving asset pricing puzzles also requires an EIS greater than 1. Our model could not hit plausible targets for C/Y or g with an EIS above 1

As a point of departure, we deactivate structural breaks and parameter uncertainty and compute the *REE* for a conventional AK model. The solid lines in the top row of figure 7 depict policy functions for C/Y and g when μ_z is constant and known with certainty. As expected, the consumption-output ratio is decreasing in $\ln z$, and growth is increasing. The former reflects consumption smoothing, because a low current value of $\ln z$ signals higher future productivity. The latter just reflects that an increase in the investment share raises growth in an AK model.

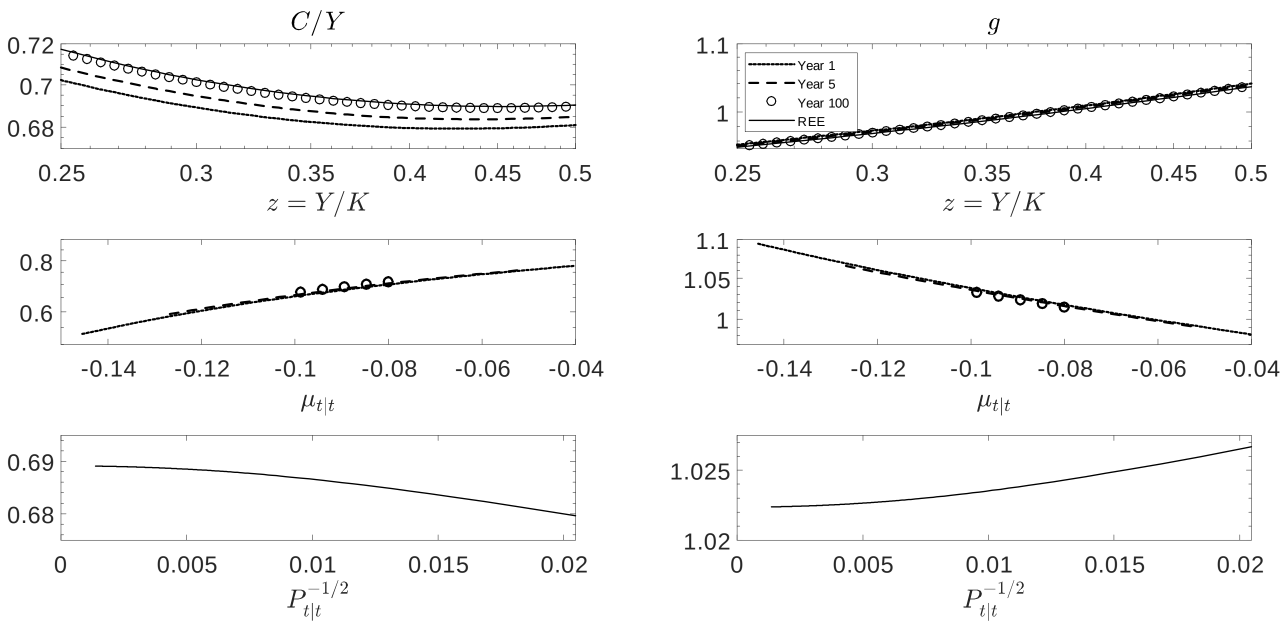

Figure 7: Policy functions when z is active and q is inactive.

The other curves in the top row portray slices of the policy functions when structural breaks and learning are active. This model has three state variables, $\ln z$, the posterior mean $\mu_{t|t}$, and the posterior precision $P_{t|t}$. In the top row, the posterior mean $\mu_{t|t}$ is held constant at the true μ_z, and precision $P_{t|t}$ is indexed by time since the last break, with uncertainty being highest in year 1 and diminishing with t. Uncertainty promotes precautionary saving and shifts the policy function for C/Y downward. Because capital is the only aggregate store of wealth, the additional savings are channeled into investment, and g increases. As parameter uncertainty diminishes, precautionary savings decline, and the policy functions shift toward those for the benchmark no-break model.

Since the model has complete markets, a safe asset is implicit. One might wonder why agents don't park their precautionary savings there to await more information. The answer is that the safe asset is in zero net supply and hence is not a store of aggregate wealth. The risk-free rate and return on capital adjust so that agents do not park precautionary savings in the safe asset. In principle, inventories or other durable commodities could serve as an aggregate store of precautionary wealth, but they are not riskless.

The second row illustrates how variation in point estimates $\mu_{t|t}$ influence C/Y and g with $\ln z$ held constant at its unconditional mean. We limit attention to plausible ranges by plotting $\mu_{t|t}$ over intervals of ± 3 posterior standard deviations around the true value. Variation in $\mu_{t|t}$ matters a lot, especially shortly after a break when point estimates can bounce all over the place. Conditional on no future break, $\mu_{t|t}$ is a martingale under the subjective probability law (Doob 1949), so a revision of the point estimate shifts the agent's long run forecast for productivity. Hence, consumption responds more strongly than to a transitory innovation ε_{zt}. Although the estimates eventually settle down as t increases, this channel continues to amplify volatility for a long time.

The magnitude of precautionary effects can be seen more clearly in the third row. Here, C/Y and g are depicted as functions of the posterior standard deviation $P_{t|t}^{-1/2}$ with the posterior mean $\mu_{t|t}$ and technology shock $\ln z$ held constant at their unconditional means. The peak effect occurs in the year of a break (shown at the far right of each panel). The consumption-income ratio is about 1 percentage point lower than in the benchmark no-break model, and growth is about 40 basis points higher. As uncertainty is resolved (moving from right to left), the investment share falls and growth slows. It takes about 5 years to move halfway across the graph and 20 years to move two-thirds of the way. Beyond that point, the policy functions become flatter, and the effects of parameter uncertainty diminish. Most of the precautionary effects occur in the first two decades. The effects of variation in $\mu_{t|t}$ persist beyond that point, however.

Further insight can be obtained by simulating the model. Figure 8 compares outcomes for economics with identical sequences of (scaled) *TFP* shocks. The simulations

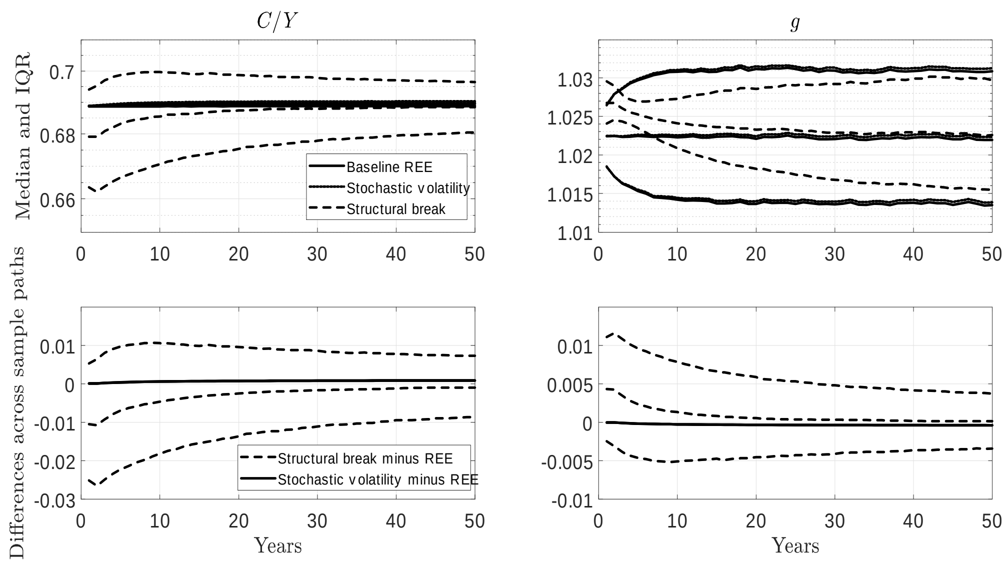

Figure 8: Fan charts when z is active and q is inactive. In the first row, dashed lines portray the median and interquatile range for a pure uncertainty shock, while solid and dotted lines depict those for the baseline *REE* and a *REE* with stochastic volatility, respectively. The second row depicts the mean and interquartile range for differences between the learning and stochastic volatility models relative to the benchmark *REE* model.

are initialized by setting $\ln z_0$ at its unconditional mean, drawing 10,000 sample paths for $\ln z_t$, and then calculating outcomes for c, g by plugging the simulated shocks into the policy functions. Dashed lines in the top row portray the median and interquartile range at each date for the structural-break model, while solid lines depict those for the baseline *REE*. The second row depicts differences across sample paths on which standardized realizations of ε_{zt} are held constant.

In addition to conventional innovations in *TFP* (ε_{zt}), the learning economy is also subjected to a 'pure uncertainty' shock at date zero. To create a pure uncertainty shock, we posit a structural break in which the newly drawn value of μ_z happens to coincide with the old value and the prior mean; i.e., $\chi_t = 1$, $m_t = \mu_{t-1} = m$. In other words, a pure uncertainty shock is fake news: agents believe a structural break has occurred when in fact it has not.

A structural break activates two belief states, $\mu_{t|t}$, and $P_{t|t}$. A fake news shock creates uncertainty about μ_z, raising the posterior variance $P_{t|t}^{-1}$ and increasing the sensitivity of the conditional mean $\mu_{t|t}$ to incoming data. Their combined effect is to depress the average consumption share, raise average growth, and amplify consumption volatility.

The precautionary effects can be seen most clearly in the second row. The middle dashed line portrays cross-sectional average differences between the structural-break model and the no-break economy. Because no actual break has occurred, averaging across sample paths removes the effects of $\mu_{t|t}$ and isolates the influence of $P_{t|t}$.[31] Furthermore, because the posterior variance is a deterministic function of time since the last break, it is the same on all sample paths. The cross-sectional average therefore illustrates how higher uncertainty about μ_z influences C/Y and g with the mean of $\mu_{t|t}$ held constant at the true μ_z. The peak differences occur in the first 5 years, when the mean consumption share is 0.8 to 1.1 percentage points lower and mean growth is 25 to 45 basis points higher. By the end of the second decade, the effects of μ_z uncertainty are weaker, but the risk of a future break remains. Consequently, the mean consumption share remains about 25 basis points lower than in the baseline *REE*, and mean growth is about 5 basis points higher. These small growth effects continue to add up, and after 50 years the mean capital stock is 4.75 percent higher than in the *REE*.[32]

The structural-break model also features greater consumption volatility, due largely to variation in the conditional mean $\mu_{t|t}$. Because $\mu_{t|t}$ evolves as a martingale (conditional on no future break), it has a strong influence on permanent income. Hence variability in $\mu_{t|t}$ strongly amplifies consumption volatility. For a model like ours, Jones, *et al.* (2005) report that consumption would be too smooth under *REE* for an *EIS* as low as ours. This is also true in our model (the solid fan chart in the upper left panel is so narrow that it looks like a single line). When $\mu_{t|t}$ is active, however, the model exhibits plenty of consumption volatility despite the low *EIS* (see the dashed fan chart in the upper left panel). Indeed, on many sample paths, the pure precautionary effect is swamped by movements in permanent income. In other words, while average C/Y is lower in the learning model, there are many sample paths on which it is higher than in the *REE* (compare the means and interquartile ranges in the second row). Of course, variation in $\mu_{t|t}$ is itself a consequence of the higher uncertainty

[31] I.e., $\mu_{t|t}$ averages to μ_{zt} across sample paths. Since there are no actual breaks, μ_{zt} remains constant at m.

[32] A higher EIS would amplify the response of precautionary savings because consumers would be more willing to trade lower present consumption for higher future consumption. For instance, for $1/\rho = 0.33$ (Havranak's point estimate), the mean consumption share would be 140 basis points lower than in the *REE* at impact, mean growth would be 55 basis points higher, and the capital stock would be 8 percent higher after 50 years.

because $\mu_{t|t}$ would be less sensitive to incoming data if $P_{t|t}^{-1}$ did not rise.

Since a pure uncertainty shock is somewhat contrived in this environment, we also look at the combined effects a structural break along with the associated increase in uncertainty. Figure 9 illustrates the differences between a fake news shock and breaks of magnitude $\pm 1\sigma_m$. As in the second row of figure 8, the solid lines represent mean differences across sample paths, with realizations of ε_{zt} held constant, while the dashed lines represent the interquartile range.

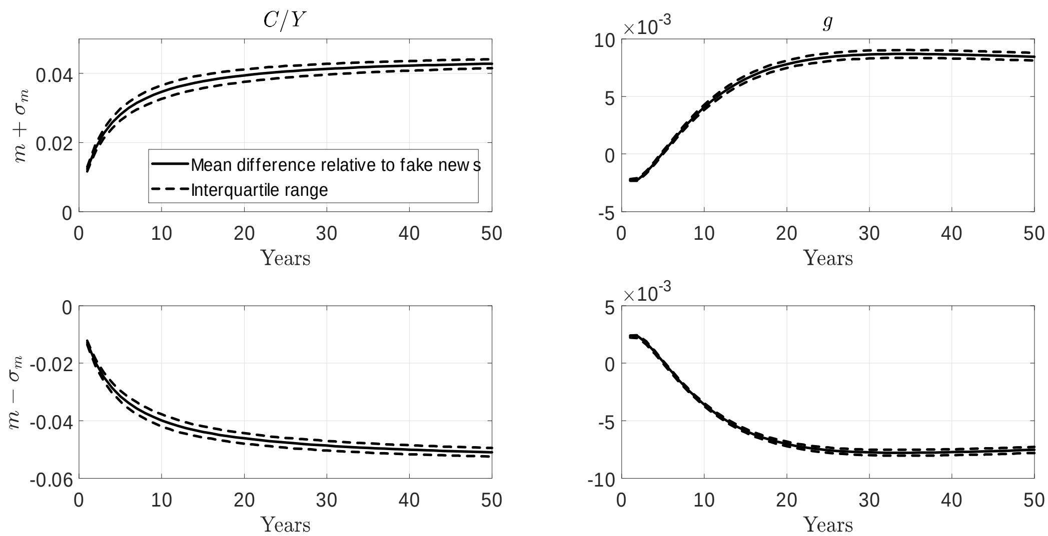

Figure 9: Fake news about μ_{zt} versus actual breaks

After a signal that $\chi_{t+1} = 1$, the posterior variance rises by the same amount for both shocks, and the precautionary effects net roughly to zero when differencing across sample paths. What remains are the effects of $\mu_{t|t}$. When μ_z rises (top row), agents attribute the increase in productivity partly to an increase in μ_z and partly to a positive innovation in ε_z. The former raises the consumption share, while the latter depresses it. Relative to a fake news shock, more of the increase in $\ln z_t$ is attributed to μ_z. As a consequence, the net impact effect on the mean C/Y ratio is positive. As agents gradually verify that μ_z is actually higher, both C/Y and g rise. After 50 years, the mean consumption share is 4 percentage points higher than for a fake news shock, average growth is about 90 basis points higher, and the mean capital stock is about 30 percent higher. The effects of a downward shift in μ_z are similar but have the opposite sign (bottom row).

Last but not least, we contrast uncertainty about μ_z with uncertainty about future innovations ε_{zt}. We do this by introducing stochastic volatility in a model in which μ_z is constant and known with certainty. In the structural-break model, the prediction error variance for $\ln z_t$ is

$$\begin{aligned} var(\ln z_{t+1} | \ln z^t, \chi_{t+1} &= 1) = \sigma_z^2 + \sigma_m^2, \\ var(\ln z_{t+1} | \ln z^t, \chi_{t+1} &= 0) = \sigma_z^2 + P_{t|t}^{-1}, \end{aligned} \tag{35}$$

where t indexes time since the last volatility break. To put a stochastic volatility model on the same footing, we assume that the conditional variance for ε_{zt} jumps to $\sigma_z^2 + \sigma_m^2$ with probability λ and then declines with probability $1 - \lambda$ to $\sigma_z^2 + P_{t|t}^{-1}$. The red lines in figure 8 portray outcomes for the stochastic volatility model.

Adding stochastic volatility to the *REE* model matters slightly, but does not change the big picture. This is especially clear in the second row, where red lines depicts differences across sample paths relative to the *REE* economy. When shown on the same scale as sample-path differences for the structural-break economy, those for the stochastic-volatility model are too small to see. The reason why structural breaks are more important is that they generate more persistent variation in the conditional mean. As a consequence, they have a greater impact on permanent income.

Furthermore, stochastic volatility creates less long-run risk. As shown in figure 6, structural breaks fatten the tails of the unconditional distribution of lnz more. Indeed, the long-run risk depicted there depends mainly on variation in the conditional mean. Matching the conditional variance of ε_{zt} to the prediction-error variance of the structural-break process fattens the tails on the unconditional distribution only slightly.

To summarize, structural breaks have big effects in this model, as do movements in the conditional mean $\mu_{t|t}$. Movements in the conditional variance $P_{t|t}^{-1}$ matter, but their effects are smaller in magnitude. Smaller still are the effects of stochastic volatility in ε_{zt}.

4.2 Investment-specific technology shocks

Results for a model in which q is active and z is inactive are qualitatively similar. An uncertainty shock promotes savings and growth, and the responses are larger in magnitude than those associated with a comparable increase in conditional variances.

Parameters for this model are shown in table 4. The subjective discount factor β, coefficient of relative risk aversion γ, and steady-state value of z are the same as in the previous section. The parameters governing $\ln q_t$ are estimated by maximizing the log posterior for the data shown in figure 1, with priors as described in appendix E.3. Given these parameters, the inverse-EIS ρ and depreciation rate δ are chosen so that the deterministic steady states for C/Y and g are 0.7 and 2 percent per annum, respectively. Once again, the implied EIS is close to zero.

TABLE 4: Calibration for structural breaks in q

β	γ	ρ	δ	ρ_q	σ_q	λ	m	σ_m	$\bar{z}$	$\bar{q}$
1.01^{-4}	4	13.8045	0.129	0.81	0.198	0.0111	-0.013	0.048	0.464	0.934

The solid lines in the top row of figure 10 portray policy functions for a version that abstracts from breaks and learning. As q increases, the capital-goods technology becomes less efficient, and each unit of investment is transformed into fewer units of new capital. To compensate, agents feed more investment goods into the capital-goods sector. Thus, the consumption share decreases with q. The rise in the investment share only partly offsets the decline in the efficiency of the capital-goods technology, however, implying that the capital stock also grows more slowly as q increases.

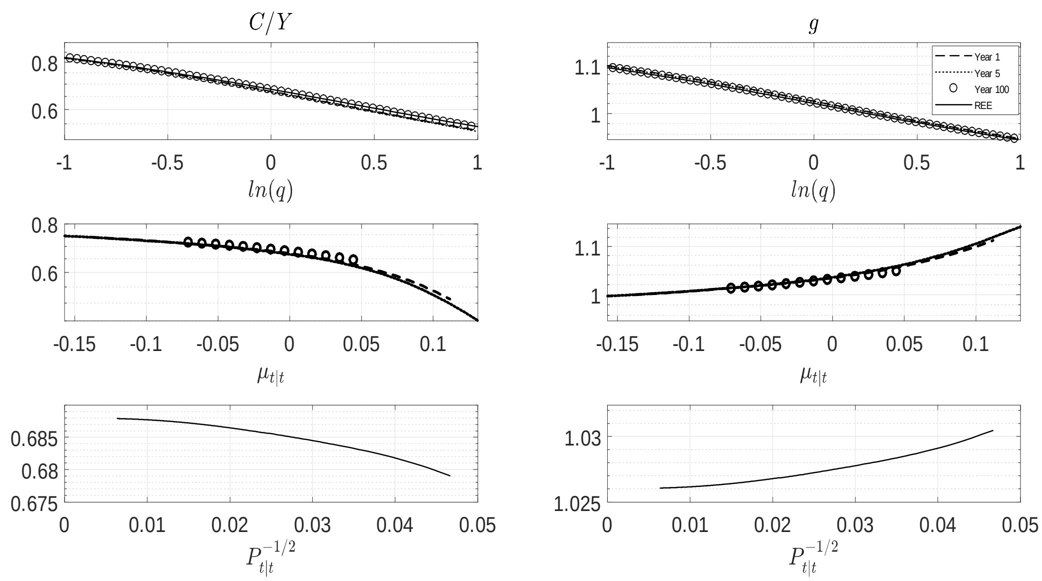

Figure 10: Policy functions when q is active and z is inactive

The circles and dashed and dotted lines in the top row illustrate the effects of uncertainty. In this panel, the posterior mean $\mu_{t|t}$ is held constant at the true μ_q, and precision $P_{t|t}$ is indexed by time since the last break, with uncertainty being highest immediately after a structural break. The qualitative effects are the same as for a neutral *TFP* shock: uncertainty promotes precautionary saving, increases the investment share, and shifts the policy function for g upward. As uncertainty is resolved, precautionary savings decline, and the policy functions shift toward those for the no-break benchmark.[33]

The middle row illustrates how variation in $\mu_{t|t}$ influences C/Y and g with $\ln q$ held constant at its unconditional mean. Attention is limited to plausible values by plotting $\mu_{t|t}$ over intervals of ± 2 posterior standard deviations around the true value. Variation in $\mu_{t|t}$ matters a lot, especially shortly after a break when point estimates are most sensitive to realizations of $\ln q_t$. Recall once more that $\mu_{t|t}$ is a martingale conditional on no break, so the agents's (conditional) long run forecast for $\ln q_t$ reacts strongly to this state variable. Hence, consumption and growth also respond strongly. This effect is nonlinear, being largest for high values of $\mu_{t|t}$.

The third row isolates the influence of $P_{t|t}$. Here C/Y and g are plotted as functions of the posterior standard deviation $P_{t|t}^{-1/2}$ with $\ln q$ and $\mu_{t|t}$ held constant at their respective unconditional means. Because $P_{t|t}^{-1/2}$ declines deterministically with time since the last break, time runs from right to left in this row. The points closest to the origin illustrate the effects of background uncertainty associated with conventional shocks to lnq and potential future breaks. Since $C/Y = 0.7$ and $g = 1.02$ in the deterministic steady state, background uncertainty reduces the consumption share and raises growth by roughly 120 and 60 basis points, respectively. The points on the on the far right depict the impact effects of an uncertainty shock. In the year of a structural break, the consumption share falls by another 90 basis points and growth increases by a further 45 basis points. As uncertainty is resolved, the investment share declines back toward its stochastic steady state, and growth declines.

The next figure compares outcomes for economies with and without structural breaks. There are two of the latter, one in which the conditional variance of ε_{qt} is constant and another in which it is calibrated to match the prediction-error variance of the structural-break model (cf. the discussion surrounding equation 35). To create a pure uncertainty shock, we again imagine that agents believe a structural break

[33] The shifts hard to see on this scale because they are small relative to variation induced by $\ln q$.

has occurred when in fact it has not. As before, the three simulations used the same scaled ε_{qt} shocks, so shocks are held constant on all sample paths.

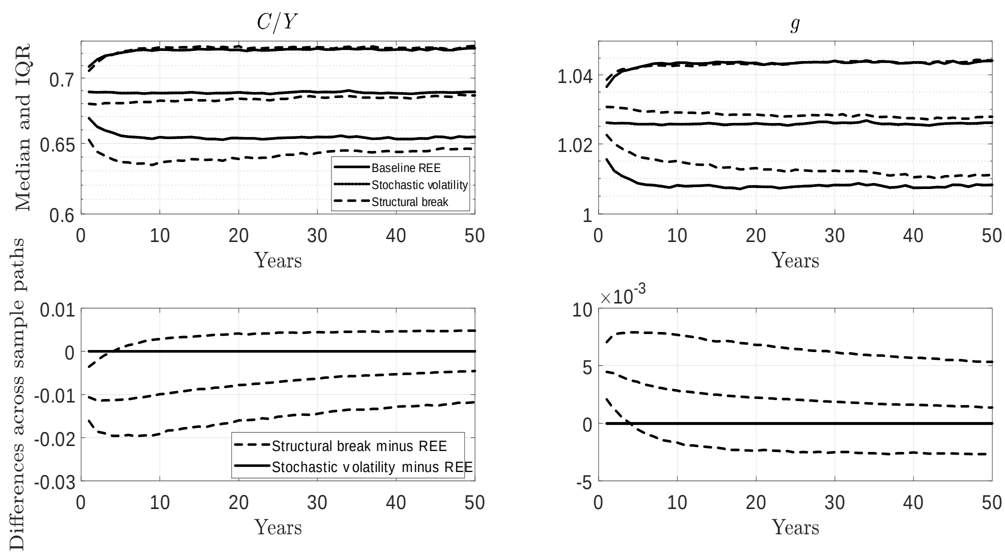

Figure 11: Fan charts when q is active and z is inactive. In the first row, dashed lines portray the median and interquatile range for a pure uncertainty shock, while solid and dotted lines depict those for the baseline *REE* and a *REE* with stochastic volatility, respectively. The second row depicts the mean and interquartile range for differences between the learning and stochastic volatility models relative to the benchmark *REE* model.

In the top row, dashed lines portray the median and interquartile range for the structural-break model, while solid and dotted lines depict those for the benchmark *REE* and stochastic-volatility models. The fact that the dotted and solid lines lie on top of one another implies that stochastic volatility matters very little when μ_q constant and known with certainty. In contrast, in response to a pure uncertainty shock, the median consumption share falls by about 100 basis points and then rises gradually as uncertainty is resolved. The median growth rate rises by about 45 basis points at impact, and declines slowly.

The second row highlights differences between the economies. Solid lines portray differences between the stochastic-volatility and *REE* economies, while dashed lines depict differences between the strucutral break model and the *REE*.

As before, the cross sectional average illustrates how higher uncertainty influences C/Y and g with the mean of $\mu_{t|t}$ held constant at the true μ_q. For the structural-break model, the biggest differences occur in the first 5 years, when the mean consumption share is about 100 basis points percentage points lower and mean growth is 30 to 50 basis points higher. By the end of the second decade, the effects uncertainty are weaker, but the risk of a future break remains. Consequently, the mean consumption share remains about 80 basis points lower than in the *REE*, and mean growth is about 25 basis points higher. Although this might not seem like much, the cumulative effect over 50 years is an 11 percent increase in the level of the capital stock.

As before, when breaks are absent, outcomes for a stochastic-volatility economy are essentially the same as for the benchmark *REE* model with constant conditional variances. Differences across sample paths are too small to see when graphed on the same scale as the structural-break economy. Structural breaks matter more because they activate persistent variation in the conditional mean of $\ln q$.

Finally, figure 12 illustrates the differences between structural breaks of magnitude $\pm 1\sigma_m$ and a fake news shock. Solid lines again represent mean differences across sample paths, with realizations of ε_{qt} held constant, while the dashed lines represent the interquartile range.

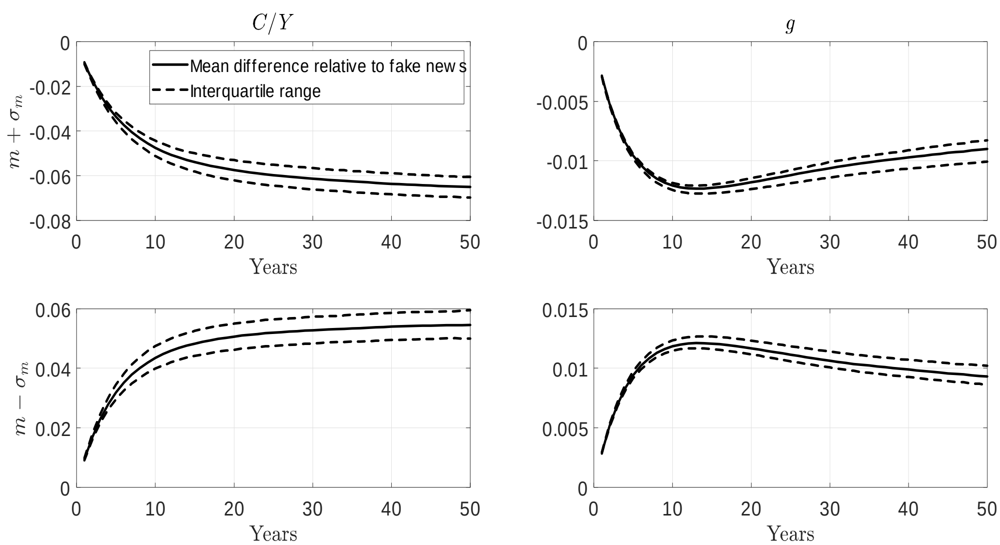

Figure 12: Structural breaks in μ_q versus fake news

Since the effects of positive and negative structural breaks are approximately symmetric, it suffices to discuss the effects of a structural break that increases μ_q (top row). Recall that the posterior variance rises by the same amount for actual breaks and fake news, implying that the precautionary effects net roughly to zero when differencing across sample paths. Thus, the differences primarily reflect the influence of $\mu_{t|t}$. After both shocks, agents attribute the decrease in capital-goods productivity partly to an increase in μ_{qt} and partly to a positive innovation in ε_{qt}, but more to μ_{qt} when there is an actual break. After an actual break, as agents gradually verify that μ_{qt} is higher, the decline in C/Y is magnified. After 50 years, the mean consumption share for an actual break is 6 percentage points lower than for a fake-news shock. The effects on g are similar, but with some over-shooting in the medium run.

5 Conclusion

This paper analyzed learning by agents of two aggregate-shock distributions. The economy is competitive and has no external effects. We solved the planner's problem and then explained how decentralization would effect the same outcome.

We contrasted the learning economy to the standard no-learning economy with two shocks. Under no learning, a mean-preserving increase in the variance of TFP shocks increases long-run growth by raising precautionary savings. By contrast, a mean-preserving increase in the variance of investment shocks lowers growth because it raises the option value of waiting to invest, thereby reducing savings. Such a contrast does not arise in the learning economy: For the case of TFP shocks, the qualitative effects of learning work in the same direction as a permanent rise in variance. For shocks to the efficiency of investment, the qualitative effects can go in the opposite direction. In the learning economy, growth rises whereas in the shock-variance-spreads economy the growth rate falls.

Endogeneity of growth reverses the implications for the risk-free rate: As in standard models, a rise in uncertainty raises precautionary savings, but because growth is endogenous, the rise in the growth rate is accompanied by a rise in the real interest rate, in contrast to what happens in an endowment economy.

When we compare the learning economy to one in which shock variances increase and equate the marginal distributions of the shocks at each date, the qualitative effects in the two cases do go the same way for TFP shocks, but quantitatively the effects are larger in the case of learning because expected lifetime utility varies more as beliefs respond to the shocks.

The data and programs underlying this article are available on Zenodo, at

https://dx.doi.org/10.5281/zenodo.4749595.

Acknowledgement 2 *We thank the NSF and the C.V. Starr Center for support. For comments and suggestions, we are grateful to the editor and three referees plus F. Bianchi, R. Lucas, S. Ludvigson, V. Midrigan, P. Rousseau, L. Veldkamp and G. Violante. Gaston Navarro, Sai Ma, Angelo Orane, and Xi Xiong provide excellent research assistance.*

References

[1] ANDOLFATTO, D. AND GOMME P. (2003), "Monetary Policy Regimes and Beliefs," *International Economic Review* 44(1), 1-30.

[2] BANSAL, R. and YARON, A. (2004), "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance* 59(4), 1481-1509.

[3] BAXTER, M. and JERMAN, U.J. (1997), "The International Diversification Puzzle Is Worse Than You Think," *American Economic Review* 87(1), 170-180.

[4] BERNANKE, B. (1983), "Irreversibility, uncertainty, and cyclical investment," *Quarterly Journal of Economics* 98(1), 85-106.

[5] BIANCHI, F., ILUT, C. AND SCHNEIDER, M. (2018), "Uncertainty Shocks, Asset Supply and Pricing over the Business Cycle", *Review of Economic Studies* 85(2), 810-854.

[6] BIANCHI, F. and MELOSI, L. (2016), "Modeling the Evolution of Expectations and Uncertainty in General Equilibrium," *International Economic Review* 57(2), 717-756.

[7] BIANCHI, F. and MELOSI, L. (2018), "Constrained Discretion and Central Bank Transparency," *Review of Economics and Statistics* 100(1), 187-202.

[8] BLOOM, N. (2014), "Fluctuations in Uncertainty," *J. Econ. Perspectives* 28(2), 153-176.

[9] BOARD OF GOVERNORS OF THE FEDERAL RESERVE SYSTEM (2021). Financial Accounts of the United States, Nonfinancial corporate business;

corporate equities as a percentage of net worth, Table Z.1. Available at https://www.federalreserve.gov/datadownload/Choose.aspx?rel=z1. Last accessed April 7, 2021.

[10] BUREAU OF ECONOMIC ANALYSIS (2017a). Fixed Asset Accounts, Current-Cost Net Stock of Private Fixed Assets, Equipment, Structures, and Intellectual Property Products by Type, Table 2.1. Available at https://apps.bea.gov/iTable/iTable.cfm?ReqID=10&step=2. Last accessed May 9, 2017.

[11] BUREAU OF ECONOMIC ANALYSIS (2017b). Fixed Asset Accounts, Chain-Type Quantity Indexes for Net Stock of Private Fixed Assets, Equipment, Structures, and Intellectual Property Products by Type, Table 2.2. Available at https://apps.bea.gov/iTable/iTable.cfm?ReqID=10&step=2. Last accessed May 9, 2017.

[12] BUREAU OF ECONOMIC ANALYSIS (2017c). National Income and Product Accounts, Real Gross Domestic Product (Chained Dollars), Table 1.1.6. Available at https://apps.bea.gov/iTable/iTable.cfm?reqid=19&step=2&isuri=1&1921=survey#reqid=19&step=2&isuri=1&1921=survey. Last accessed May 9, 2017.

[13] BUREAU OF ECONOMIC ANALYSIS (2017d). National Income and Product Accounts, Nominal Gross Domestic Product, Table 1.1.5. Available at https://apps.bea.gov/iTable/iTable.cfm?reqid=19&step=2&isuri=1&1921=survey#reqid=19&step=2&isuri=1&1921=survey. Last accessed May 10, 2017.

[14] BUREAU OF ECONOMIC ANALYSIS (2017e). National Income and Product Accounts, Nominal Government Expenditures, Table 3.1. Available at https://apps.bea.gov/iTable/iTable.cfm?reqid=19&step=2&isuri=1&1921=survey#reqid=19&step=2&isuri=1&1921=survey. Last accessed June 5, 2017.

[15] CAMPBELL, J.Y. and VICEIRA, L.M. (1999), "Consumption and Portfolio Decisions When Expected Returns Are Time Varying," *Quarterly Journal of Economics* 114(2), 433-495.

[16] CHAMLEY, C. and GALE, D. (1994), "Information Revelation and Strategic Delay in a Model of Investment,". *Econometrica* 62(5),1065-1085.

[17] COGLEY, T. (2009), "Is the Market Price of Risk Infinite?" *Economic Letters* 102(1), 13-16.

[18] COLLIN-DUFRESEN, P., JOHANNES, M. and LOCHSTOER, L. (2016), "Parameter Learning in General Equilibrium: The Asset Pricing Implications," *American Economic Review* 106(3), 664-698.

[19] DIXIT, A.K. and PINDYCK, R.S., (1994), *Investment Under Uncertainty* (Princeton: Princeton University Press).

[20] DOOB, J.L. (1949), "Application of the Theory of Martingales", *Colloques Internationaux du Centre National de la Recherché Scientifique* 36, 23–27.

[21] EPSTEIN, L., and SCHNEIDER, M. (2008), "Ambiguity, Information Quality, and Asset Pricing," *Journal of Finance*, 63(1), 197-228.

[22] EPSTEIN, L.G. and ZIN, S.E. (1989), "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica* 57(4), 937-969.

[23] EVANS, G.W. and HONKAPOHJA, S. (2001), *Learning and Expectations in Macroeconomics* (Princeton: Princeton University Press).

[24] FISHER, J. D. (2006), "The dynamic effects of neutral and investment-specific technology shocks," *Journal of Political Economy*, 114(3), 413-451.

[25] FOERSTER, A., RUBIO-RAMIREZ, J., WAGGONER, D.F. and ZHA, T (2016), "Perturbation Methods for Markov-Switching DSGE Models," *Quantitative Economics* 7(2), 637-669.

[26] GEWEKE, J. (2001), "A Note on Some Limitations of CRRA Utility," *Economics Letters* 71(3), 341-345.

[27] GILCHRIST, S. and WILLIAMS, J. (2005), "Investment, Capacity and Uncertainty: A Putty-Clay Approach," *Review of Economic Dynamics*, 8(1), 1-27.

[28] GORDON, R.J. (2016), *The Rise and Fall of American Growth* (Princeton: Princeton University Press).

[29] GOURIO, F. (2008), “Time-Series Predictability in the Disaster Model,” *Finance Research Letters*, 5(4), 191-203.

[30] GREENWOOD, J., HERCOWITZ, Z., and KRUSELL, P. (2000), “The Role of Investment-Specific Technological Change in the Business Cycle, *European Economic Review* 44(1), 91-115.

[31] HALL, R.E. (1988), “Intertemporal Substitution in Consumption,” *Journal of Political Economy* 96(2), 339-57

[32] HAMILTON, J. D. (1989), “A New Approach to the Economic Analysis of Non-stationary Time Series and the Business Cycle,” *Econometrica* 57(2), 357-384.

[33] HANSEN, L.P. and SARGENT, T. J. (2001), “Robust Control and Model Uncertainty,” *American Economic Review*, 91(2), 60-66.

[34] HAVRANEK, T. (2015), “Measuring Intertemporal Substitution: The Importance of Method Choices and Selective Reporting,” *Journal of the European Economic Association* 13(6), 1180-1204.

[35] HOPENHAYN, H. and MUNIAGURRIA, M.E. (1996), “Policy Variability and Economic Growth,” *Review of Economic Studies* 63(4), 611-626.

[36] HUFFMAN, G. AND KIEFER, N. (1994), “Optimal Learning in a One-Sector Model,” unpublished manuscript.

[37] JONES, L. E., MANUELLI, R., SIU, H. and STACCHETTI, E. (2005), “Fluctuations in Convex Models of Endogenous Growth, I: Growth effects,” *Review of Economic Dynamics* 8(4), 780-804.

[38] JOVANOVIC, B. and ROUSSEAU, P. L. (2005), “General Purpose Technologies.” In P. Aghion and S.N. Durlauf (eds.) *Handbook of Economic Growth*, vol. 1B (Amsterdam: Elsevier B.V), 1181-1224.

[39] JOVANOVIC, B. and ROUSSEAU, P. L. (2014), “Extensive and Intensive Investment over the Business Cycle.” *Journal of Political Economy*, 122(4), 863-908.

[40] JUSTINIANO, A., PRIMICERI, G. E., and TAMBALOTTI, A. (2010), "Investment Shocks and Business Cycles," *Journal of Monetary Economics* 57(2), 132-145.

[41] KENDRICK, J.W. (1961), *Productivity Trends in the United States*. National Bureau of Economic Research General Series no. 71. (Princeton: Princeton University Press).

[42] KOULOVATIANOS, C., MIRMAN, L. and SANTUGINI, M. (2009), "Optimal Growth and Uncertainty: Learning," *Journal of Economic Theory* 144(1), 280–295.

[43] KREPS, D. (1998), "Anticipated Utility and Dynamic Choice," in D. Jacobs, E. Kalai, M. Kamien, and N. Schwartz (Eds.) *Frontiers of Research in Economic Theory*, (Cambridge, UK: Cambridge University Press), 242-274.

[44] KRUSELL, P., OHANIAN, L.E., RIOS-RULL, J.V. and VIOLANTE, G.L. (2000), "Capital-Skill Complementarity: A Macroeconomic Analysis," *Econometrica* 68(5), 1029-1053.

[45] LEWELLEN, J., & SHANKEN, J. (2002), "Learning, Asset-Pricing Tests, and Market Efficiency," *Journal of Finance* 57(3), 1113-1145.

[46] NATIONAL BUREAU OF ECONOMIC RESEARCH (2017). US Business Cycle Expansions and Contractions. Available at https://www.nber.org/research/data/us-business-cycle-expansions-and-contractions. Last accessed May 9, 2017.

[47] PASTOR, L., & VERONESI, P. (2009), "Technological Revolutions and Stock Prices," *American Economic Review* 99(4), 1451-1483.

[48] PIKETTY, T. (2014), *Capital in the Twenty First Century* (Cambridge, MA: Harvard University Press).

[49] PINTUS, P. and SUDA, J. (2019), "Learning Financial Shocks and the Great Recession," *Review of Economic Dynamics* 31, 123-146.

[50] SARGENT, T. J. (1980), "'Tobin's q' and the Rate of Investment in General Equilibrium," *Carnegie-Rochester Conference Series on Public Policy* 12, 107-154.

[51] SCHORFHEIDE, F. (2005), "Learning and Monetary Policy Shifts," *Review of Economic Dynamics* 8(2), 392-419.

[52] SIMS, C.A. (1982), "Policy Analysis With Econometric Models," *Brookings Papers on Economic Activity*, 107-152.

[53] STOKEY, N. (2016), "Wait-and-See: Investment Options under Policy Uncertainty," *Review of Economic Dynamics* 21, 246-65.

[54] VELDKAMP, L. L. (2005), "Slow Boom, Sudden Crash," *Journal of Economic Theory* 124(2), 230-257.

[55] VERONESI, P. (2004), "The Peso Problem Hypothesis and Stock Market Returns," *Journal of Economic Dynamics and Control* 28(4), 707-725.

[56] VISSING-JORGONSEN, A. (2002), "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution," *Journal of Political Economy* 110(4), 825-853.

[57] WACHTER, J. A. (2013), "Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?" *Journal of Finance* 68(3), 987-1035.

[58] WEIL, P. (1989), "Non-Expected Utility in Macroeconomics," *Quarterly Journal of Economics* 105(1), 29-42..

[59] WRIGHT, S. (2004), "Measures Of Stock Market Value And Returns For The U.S. Nonfinancial Corporate Sector, 1900-2002," *Review of Income and Wealth* 50(4), 561-584. Available at http://piketty.pse.ens.fr/files/capitalisback/CountryData/USA/Other/?C=N;O=D. Last accessed November 18, 2019.

[60] YOGO, M. (2004), "Estimating the Elasticity of Intertemporal Substitution When Instruments Are Weak," *Review of Economics and Statistics* 86(3), 797-810.

[61] WHELAN, K. (2002), "A Guide to US Chain Aggregated NIPA Data," Review of Income and Wealth 48(2), 217-233.