

FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

Learning about Regime Change

Andrew Foerster
Federal Reserve Bank of San Francisco

Christian Matthes
Indiana University

December 2021

Working Paper 2020-15

<https://www.frbsf.org/economic-research/publications/working-papers/2020/15/>

Suggested citation:

Foerster, Andrew, Christian Matthes. 2021. “Learning about Regime Change,” Federal Reserve Bank of San Francisco Working Paper 2020-15.
<https://doi.org/10.24148/wp2020-15>

The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco or the Board of Governors of the Federal Reserve System.

Learning about Regime Change^{*}

Andrew Foerster[†] Christian Matthes[‡]

December 8, 2021

Abstract

Total factor productivity (TFP) and investment specific technology (IST) growth both exhibit regime-switching behavior, but the regime at any given time is difficult to infer. We build a rational expectations real business cycle model where the underlying TFP and IST regimes are unobserved. We develop a general perturbation solution algorithm for a wide class of models with unobserved regime-switching. Using our method, we show learning about regime-switching fits the data, affects the responses to regime shifts and intra-regime shocks, increases asymmetries in the responses, generates forecast error bias even with rational agents, and raises the welfare cost of fluctuations.

Keywords: Bayesian learning; regime switching; technology growth

JEL Codes: E13, E32, C63

^{*}We thank the Editor Jesus Fernandez-Villaverde and two anonymous referees, plus David Andolfatto, Francois Gourio, Nils Gornemann, and conference participants at the NBER Summer Institute, the Society for Computational Economics, Midwest Macro, the Norges Bank's Nonlinear Models Conferences as well as seminar participants at the University of Hong Kong and William & Mary for helpful comments. Lily Seitelman provided excellent research assistance. The views expressed are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

[†]Federal Reserve Bank of San Francisco, 101 Market St., San Francisco, CA 94105 , andrew.foerster@sf.frb.org.

[‡]Indiana University, 100 S Woodlawn Ave., Bloomington, IN 47405, matthes@iu.edu.

1 Introduction

Understanding technological progress is crucial for the study of the macroeconomy, both in determining the long run rate of growth but also fluctuations at higher frequencies. Growth in technology tends to move through prolonged periods of relatively high growth, such as during the information technology boom in the 1990s, and periods of relatively low growth such as the years since the Great Recession. Both total factor productivity (TFP) and investment specific technology (IST) exhibit this behavior of switching between prolonged periods of higher or lower growth.

At the same time, technology growth is a noisy process around these longer-term trends; for example, there were quarters of relatively lackluster growth numbers during the 1990s, and relatively stronger growth since the Great Recession. So while both TFP and IST show shifts in longer-run trends, identifying the lower frequency trend in real-time can be difficult. In this paper, we document this behavior of technological growth and study its economic implications.

To address the issue of the economic effects of learning about technology growth regimes, we take four steps in this paper. First, we document empirically that technology goes through distinct periods characterizing its growth. This finding is not new, but reaffirms the analysis of TFP growth by Fernald (2012) and of IST growth by Cummins and Violante (2002), among others. Our empirical conclusion is based upon estimates from a Markov-switching model, which, as a by-product, produces filtered estimates about which regime is in place given information at that time. These filtered estimates show that, despite the distinct regime-switching behavior, inferring the underlying regime based upon incoming data produces uncertainty about which regime is actually in place.

Second, with evidence that technology regimes are not perfectly observed but instead must be inferred, we formulate an otherwise canonical real business cycle model with TFP and IST processes that switch regimes, but with the underlying regime being learned or inferred rather than observed.

Next, we develop a perturbation-based solution algorithm for a general class of rational expectations Markov-switching models with Bayesian learning about the regime. This solution method is necessary because, due to the complexity of the learning process we consider, existing solution methods would be unable to handle our objective of merging the learning process with a structural model of economic behavior. Our solution method is more generally applicable, making it possible to efficiently solve a large wide variety of regime-switching models with learning. As by-product, we explore the properties of our solution method, as well as document

its accuracy in a simple framework. In particular, we note that a check of Euler equation errors alone can be misleading in imperfect information environments because the probability measure that is used to compute the conditional expectations in the Euler equation might be approximated poorly. Hence, we demonstrate our methodology shows good accuracy properties for both the probability measure and Euler equation errors.

Finally, building upon our empirical estimates and using our developed solution methodology, we study the economic implications of learning about technology growth regimes. We compare model fit for several different estimated specifications of the TFP and IST growth processes and find that the model with multiple regimes and learning rather than full information has superior fit. Using impulse responses, we show that a key feature of the learning mechanism is muting responses to regime changes: whereas the full information version shows quick adjustments when the growth regime changes, the learning version slows these responses due to the inability to quickly infer a regime change. Impulse responses to intra-regime shocks are affected by learning, since a temporary shock induces a partial belief that the regime has changed. Learning also produces biased forecast errors even in the presence of rational agents, as expectations of future growth are distorted by an incomplete knowledge of the current regime in place. Lastly, we study the welfare implications of learning and show that the cost of fluctuations is substantially increased with learning relative to full information.

Beyond the analysis of learning and technology growth regimes, our solution method represents an innovation in how to handle learning in a broad class of economic models. Solving full information Markov-switching rational expectations models is more involved than solving corresponding models without Markov-switching. In this paper we face further difficulties because we introduce learning via Bayes' rule, which naturally introduces a nonlinearity in the equilibrium conditions. Furthermore, we find evidence of switching in the volatilities of both technology growth processes, so a linear approximation will not suffice for our purposes. We build on Foerster et al. (2016) to show how to construct higher order perturbation solutions in Markov-switching models where agents have to infer the current economic environment. Using a perturbation-based method allows us to solve the model much faster than global solution methods would allow. The key insight to our approach is that we jointly approximate the decision rules and the learning process via Bayes' rule. The joint approximation allows us to rely on extensions to the methods developed by Foerster et al. (2016). Previous papers that have introduced partial information into rational expectations Markov-switching models either had to resort to time-intensive global solution methods (Andolfatto and Gomme, 2003; Davig, 2004; Richter and Throckmorton, 2015; Alpanda, 2019), restrict the types of regime-switching

they model (Schorfheide, 2005; Bullard and Singh, 2012), or have used a different information structure such as agents observing the current economic environment, but not knowing how persistent it is (Bianchi and Melosi, 2016).¹ Markov switching environments when agents are not fully rational, but instead learn adaptively, have been studied by Branch et al. (2013). In terms of model solution, their assumption of adaptive learning makes solving the model easier because the way expectations are formed is predetermined—it is part of the model description in adaptive learning models, rather than jointly determined with the rest of the decision rules as in our approach.²

The question of what impact imperfect information about the trends in productivity has on economic outcomes is not new. In fact, the seminal paper on real business cycle (RBC) models, Kydland and Prescott (1982), features a signal extraction problem about the trend in TFP that is broadly similar to the signal extraction problem that our agents are confronted with when observing TFP data; however, the stochastic process for TFP is different in our case as we emphasize the regime-switching nature of changes in productivity.³ Other related papers to ours are Saito (2006), Bullard and Singh (2012), Bocola and Gornemann (2013), and Alpanda (2019). All these papers study environments with TFP only where TFP grows according to a regime-switching process, whereas we emphasize that taking into account both TFP and IST fluctuations is crucial to understanding the evolution of the US economy (Greenwood et al., 1997). Taking into account both types of productivity changes is especially important in imperfect information environments, because the stochastic processes for these two types of productivity differ so substantially that the associated regimes are not uncovered by agents at the same speed. This fact in turn leads to substantial differences when it comes to impulse response functions.

Our methodology enables the analysis of a wide class of regime-switching models with learning. Bullard and Singh (2012) use a perturbation-based approach, but one that is specifically

¹In particular, the algorithm in Schorfheide (2005) is more restrictive than ours along two dimensions: (i) it focuses on models that are linearized except for the Markov-switching part, and (ii) it only allows for certain types of switches in intercepts and volatilities. The benefit of that approach is that, within the restricted class of models that it applies to, it is faster than ours and eliminates the need to approximate the updating of beliefs. Recent work on projection-based solution methods for Bayesian learning problems (Gust et al., 2018) allows for more general environments than ours, but comes with higher computational costs.

²To be very clear, we are not making any statement on whether adaptive learning or fully rational learning as in our approach is preferable. Which of these is preferred by the data most likely depends on the application. We are interested in developing a fast and reliable algorithm to solve Markov-switching rational expectations models with partial information, which would allow such a comparison to be undertaken.

³Similar specifications to Kydland and Prescott (1982) appear in Edge et al. (2007) and Gilchrist and Saito (2006).

geared towards their two-state TFP process and does not generalize. Instead, we present a general algorithm that can approximate models driven by various Markov chains with any number of states to any order of approximation. Saito (2006), Bocola and Gornemann (2013), and Alpanda (2019) use global solution methods, but this method comes with computational costs that eliminate the ability to consider a richer model. In contrast, our solution method is general enough to be able to handle regime switching behavior in both average growth rates and volatilities as well as computationally fast enough to handle both TFP and IST processes for drivers of growth. We demonstrate that both of these features are key for understanding the effects of technology growth regimes. Similar to Alpanda (2019), we emphasize the importance of nonlinearities to fully understand the effects of learning on equilibrium outcomes.


The remainder of the paper is as follows. In Section 2, we provide empirical evidence for technology growth regimes, and highlight how the underlying regime may be hard to learn or infer from the data. Section 3 lays out a canonical real business cycle model with TFP and IST processes to match the estimated ones, but includes a learning environment where the regime is not observed. In Section 4 we develop a general, perturbation-based solution method for handling a wide class of Markov-switching models where the underlying regime must be inferred via Bayesian learning. Section 5 presents our results, showing the effects of learning. Finally, Section 6 concludes.

2 Empirical Evidence for Technology Growth Regimes

In this section, we first look at measures of total factor productivity and investment specific technology, provide evidence that growth in each progresses through distinct regimes, and highlight that there is often considerable uncertainty about the regime in place.

Figure 1 shows the measures of TFP and IST on a quarterly basis from 1947 to 2019. The data on TFP are the utilization-adjusted series by Fernald (2012). The plot of the data suggests periods of different growth, with a period of higher growth early in the sample and in the late 1990s and early 2000s, and lower growth in the 1980s and last part of the sample. Fernald (2012) breaks the series into specific sub-samples; from 1947-1973 and 1996-2004, TFP grew at an average of 2 percent per year, while from 1974-1995 and 2005-2019 the growth rates were lower, at 0.61 percent and 0.22 percent, respectively. Less immediately obvious from the Figure is the fact that the quarterly growth series exhibits significant volatility, with a standard deviation of 3.38. Thus, the quarter-to-quarter volatility tends to be large relative to the average growth over the full sample and across sub-samples. This fact suggests that isolating shifts in

Figure 1: Total Factor Productivity and Investment Specific Technology



Notes: Figure shows the the log-level of utilization-adjusted TFP from Fernald (2012) in blue, and the log-level of IST based upon NIPA deflators in red; both series are indexed to be 100 in 1947:Q1. Gray bars indicate NBER recessions.

the average growth rate will be difficult to detect.

Turning to investment-specific technology, the measure of IST is based on the relative price of capital goods from the National Income and Product Accounts. This measure also shows periods of distinct behavior, with growth averaging 0.19 percent from 1947 to 1981, and 2.20 percent growth from 1982 to 2019.⁴ Similar to the case with TFP, the volatility of the quarterly series is large relative to the full sample and within sub-sample averages; the standard deviation of 3.0 indicates that detecting shifts in the average growth rate will be difficult.

To formalize the notion of growth regimes suggested by Figure 1, we estimate separate regime-switching models for the growth rate of TFP,

$$\Delta \log z_t = \mu_z(s_t^{\mu z}) + \sigma_z(s_t^{\sigma z}) \varepsilon_{z,t}, \quad (1)$$

⁴Specifically, we construct a nondurable goods plus services deflator, and a durables plus private investment deflator, and take the ratio. Justiniano et al. (2011) highlight the 1982 breakpoint in the growth of the relative price of investment goods, both in the NIPA accounts and using the series from Cummins and Violante (2002). We use the NIPA accounts for our analysis since in recent years there have been improvements in quality measurement that have been extended back to the 1950s, and because we want our analysis to include recent data.

where z_t denotes TFP, and a similar one for IST growth,

$$\Delta \log u_t = (1 - \rho_u) \mu_u (s_t^{\mu u}) + \rho_u \Delta \log u_{t-1} + \sigma_u (s_t^{\sigma u}) \varepsilon_{u,t} \quad (2)$$

where u_t denotes IST. The parameters μ_j and σ_j for $j \in \{z, u\}$ denote the average growth rate and the volatility of growth, respectively. Each of these parameters is governed by an independent Markov process indexed by s_t^{ij} , for $i \in \{\mu, \sigma\}$ and $j \in \{z, u\}$, where s_t^{ij} takes either a “low” or “high” value $s_t^{ij} \in \{L, H\}$, and has transition matrix

$$P^{ij} = \begin{bmatrix} p_{LL}^{ij} & p_{LH}^{ij} \\ p_{HL}^{ij} & p_{HH}^{ij} \end{bmatrix} = \begin{bmatrix} p_{LL}^{ij} & 1 - p_{LL}^{ij} \\ 1 - p_{HH}^{ij} & p_{HH}^{ij} \end{bmatrix}. \quad (3)$$

The parameter ρ_j indicates the persistence of the growth process. Based upon our testing of different model specifications, we found this parameter to be statistically zero in the TFP case, so we impose that result from the outset. In addition, we estimate four different specifications for the number of regimes, allowing either one or two regimes for each of the μ_j and σ_j processes. We estimate equations (1) and (2) via maximum likelihood as in Hamilton (1989).

Estimates of our regime-switching model for TFP and IST are shown in Table 1. The estimates for the TFP process show some support for a specification of two regimes for μ_z and two regimes for σ_z . In particular, the point estimates for the parameters show statistically distinct estimates across regimes. The log likelihood also increases modestly, by a few log points, with the introduction of additional regimes. For IST, the improvement in the log likelihood with additional regimes is more substantial, particularly in the volatility. In Section 5, we return to this model evaluation and conclude that the specification with two regimes each for μ_j and σ_j for both processes fits the data best, and thus we focus the remaining discussion on that specification.

The estimates for TFP of the low and high average growth regimes indicate growth of 0.50 and 1.95 percent annually, which corresponds closely to the sub-sample averages from Fernald (2012). Each regime shows extremely high persistence, with diagonal entries of the transition matrices that imply an expected durations of almost 20 years for the low growth regime, and over 15 years for the high growth regime. For the volatilities, both regimes show relatively large standard deviations of growth when compared to the average growth rates.

The IST estimates likewise show distinct average growth and volatility regimes. The average growth rates in the low and high regimes is 0.58 and 2.18 percent, respectively; the diagonal elements of the transition matrices indicate an even longer expected duration than TFP growth,

Table 1: Parameter Estimates

(a) TFP Growth Process

Model	$\mu_z(L)$	$\mu_z(H)$	$\sigma_z(L)$	$\sigma_z(H)$	$P_{LL}^{\mu z}$	$P_{HH}^{\mu z}$	$P_{LL}^{\sigma z}$	$P_{LL}^{\sigma z}$	$\log L$
One μ , One σ	1.22 (0.21)		3.38 (0.14)						-767
One μ , Two σ	1.10 (0.21)		2.78 (0.26)	3.86 (0.28)			0.995 (0.01)	0.995 (0.01)	-763
Two μ , One σ	0.77 (0.47)	2.03 (0.47)	3.32 (0.14)		0.995 (0.01)	0.992 (0.01)			-765
Two μ , Two σ	0.50 (0.56)	1.95 (0.55)	2.67 (0.27)	3.77 (0.27)	0.985 (0.02)	0.984 (0.02)	0.994 (0.01)	0.995 (0.01)	-761

(b) IST Growth Process

Model	$\mu_u(L)$	$\mu_u(H)$	$\sigma_u(L)$	$\sigma_u(H)$	ρ_u	$P_{LL}^{\mu u}$	$P_{HH}^{\mu u}$	$P_{LL}^{\sigma u}$	$P_{LL}^{\sigma u}$	$\log L$
One μ , One σ	1.24 (0.25)		2.90 (0.09)		0.24 (0.04)					-723
One μ , Two σ	1.62 (0.17)		1.82 (0.10)	4.70 (0.50)	0.22 (0.06)			0.984 (0.01)	0.962 (0.02)	-680
Two μ , One σ	0.25 (0.28)	2.17 (0.47)	2.80 (0.08)		0.16 (0.05)	0.996 (0.01)	0.996 (0.01)			-717
Two μ , Two σ	0.58 (0.24)	2.18 (0.24)	1.71 (0.10)	4.46 (0.46)	0.12 (0.06)	0.996 (0.01)	0.996 (0.01)	0.983 (0.01)	0.963 (0.02)	-672


Note: Table shows the maximum likelihood estimates of the regime switching models in equations (1) and (2). Numbers in parentheses are standard errors.

which is in line with Figure 1. The volatility regimes have standard deviations which are again quite high compared to the average growth rates.

Both processes therefore show two key features. First, they show evidence of regime-switching in both the means and the volatilities, where switches are low-probability events. Second, the estimates of the standard deviations are large relative to the estimates of the means, which suggests that, while each technology may grow at different average rates in various periods, quarter-to-quarter fluctuations in the growth rate tend to be relatively large.

Figure 2 presents a main conclusion from this empirical exercise, which is that, despite the estimates of distinct regimes, inference about which regime is in place at any given point in time is not necessarily clear. The Figure shows, for each of the four independent Markov-processes

Figure 2: Filtered Estimates of TFP and IST Regimes



Notes: Sub-panels show the filtered probabilities of the relevant regime implied by the estimated regime-switching models in equations (1) and (2), at the estimated values in Table 1. The filtered probability is the probability of being in the regime at time t , conditional on all data up to time t .

and conditional on the parameter estimates presented in Table 1, the filtered probability that the economy is in the low regime at any given quarter. These filtered estimates about time t are based on all realized data until time t . For each plot, probabilities near zero or unity reflect relative certainty that the high or low regime, respectively, is in place; on the other hand, intermediate values suggest that incoming observations about TFP or IST growth generate some uncertainty about which regime is generating them. More specifically, with both TFP and IST growth being noisy processes, high realizations of growth could signal that the high mean regime is in place, but they could be generated by the low mean regime with just a very positive ε shock. The fact that volatility changes as well further complicates the filtering, since the standard deviation of the shocks is also not known with certainty.

Figure 2 shows that the average TFP growth regimes move in line with the narrative discussed earlier, and are broadly consistent with the sub-samples Fernald (2012) identifies. High growth dominates in the early part of the sample until the mid 1970s, followed by low growth until the 1990s, with a brief period of high growth before a return to lower growth towards the end of the sample. Of note is that these regimes tend to lag the breakpoints in Fernald (2012), as for example, the high growth regime isn't considered to be in place until around 2000. The low IST growth regime plot also is consistent with the narrative discussed earlier, with low growth holding until around 1980, and then high growth in place until the end of the sample. Both low volatility plots show some movement, with lower volatilities in both towards the latter half of the sample, consistent with the Great Moderation period.

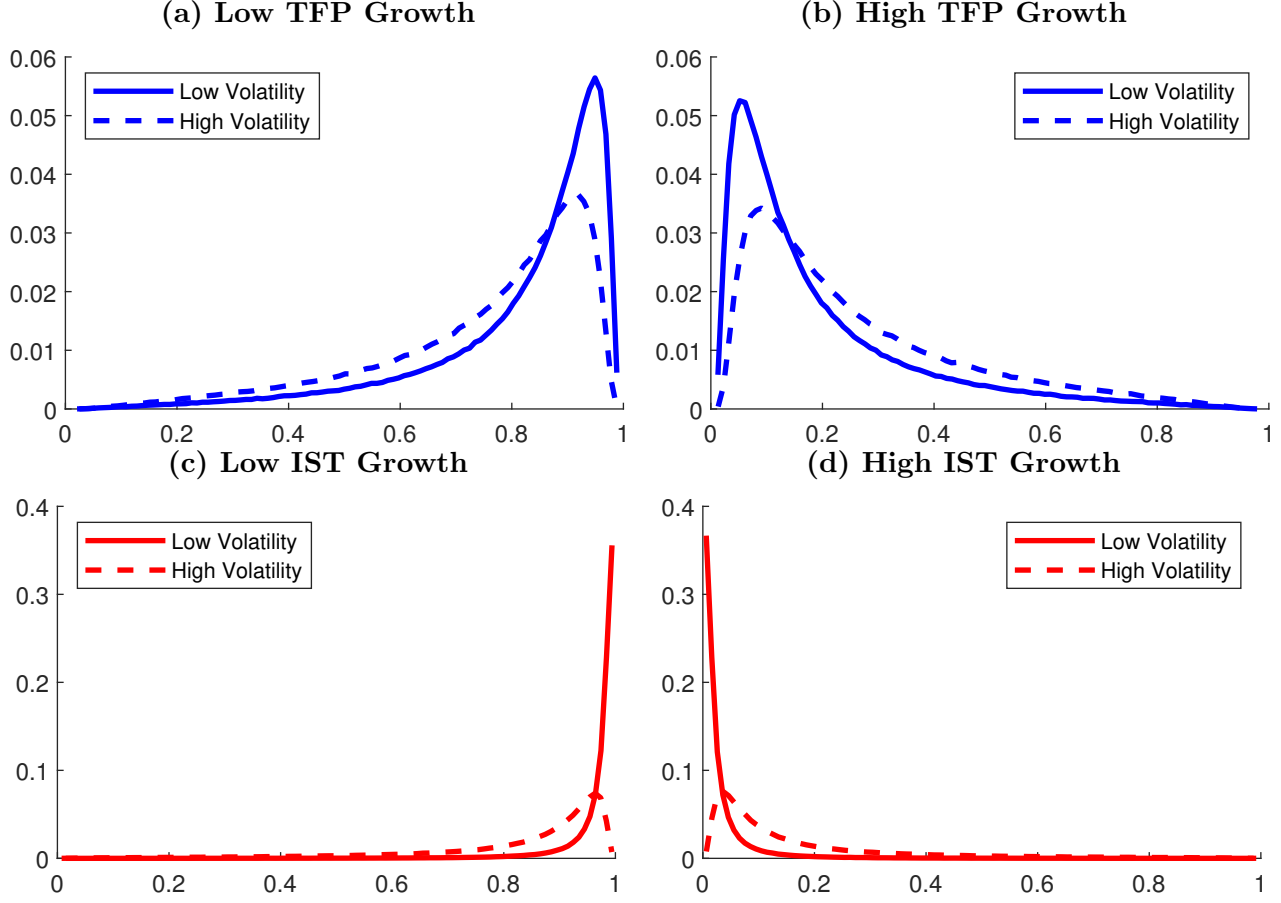
All of these plots show uncertainty about the regime in place at any given quarter, although the magnitude of this uncertainty varies by the regime. In particular, the TFP low growth regime appears to be relatively more difficult to learn than the IST low growth regime, since the IST process has filtered probabilities that tend to be closer to zero or one.

Lastly, we highlight the fact that the probabilities in Figure 2 are not simply artifacts of a distinct sample for the US over this period, but instead features of the estimated processes in Table 1. To illustrate this fact, Figure 3 plots the distributions of beliefs that the TFP or IST process is in the low growth regime for a number of simulated datasets. Specifically, in the left two panels, we plot the beliefs that TFP (top left) and IST (bottom left) are in the low growth regime, when the simulated dataset of 100,000 periods comes solely from the low growth regime; the two lines show results for data coming from the low volatility (solid line) or high volatility (dashed line) regimes. The right two panels likewise have the simulated data coming solely from the high growth regime.

If the regimes were perfectly known, the left panels in Figure 3 would show a spike at unity, and the right panels would spike at zero. The fact that they do not indicates the regimes are in fact hard to learn, even without the presence of regime changes in a long dataset. The distribution of TFP beliefs more dispersed, indicating that process is more difficult to learn than the IST process. In addition, higher volatility makes each of the belief distributions less peaked, as the wider variance of realizations makes ascertaining the regime in place more difficult. We therefore conclude that learning the true regime that is generating the TFP and IST data for the US is difficult, and there is often considerable uncertainty about which regime is in place based upon incoming data.

This empirical exercise therefore suggests that incoming data about technology growth regimes makes it difficult to infer the combination of regimes that is generating that data.

Figure 3: Distributions of Beliefs for the Low Growth Regimes



Notes: Panels show simulation-based distribution of beliefs of being in the low growth regime for TFP (top panels) and IST (bottom panels), conditional on data only being drawn from the low growth regimes (left panels) or the high growth regimes (right panels). The data is simulated using the point estimates for the Markov-switching processes obtained via Maximum Likelihood and reported in Table 1.

As a result, a structural model in which economic agents have full information about the regime in place may give these agents too much information about the economic environment, and that learning the regime from incoming data may be a more plausible assumption. In order to assess the economic implications of learning about technology growth regimes, we will embed the estimates presented in Table 1 in an otherwise standard real business cycle model, and consider how filtering incoming technology growth data to learn about the regime—rather than observing the regime directly—affects the consumption, investment, and labor decisions in the economy.

3 RBC Model with Learning and Growth Regimes

In this Section, we discuss a canonical rational expectations real business cycle model (Cooley, 1995) where growth and fluctuations are driven by changes in total factor productivity and investment-specific technology. We first lay out a planner's problem of choosing consumption, labor, and investment, and then discuss the learning problem the planner faces. We assume that the planner has the same information set as agents in the economy. Hence the solution to the planner's problem coincides with the limited information rational expectations equilibrium in this otherwise standard RBC model. Even though this framework is standard, we present it in detail to facilitate the discussion of our solution method in Section 4.

3.1 Planner's Problem

We consider an otherwise standard real business cycle model where growth in total factor productivity and investment specific technology are subject to both regime changes and idiosyncratic shocks, but the composition of these factors is not observed. The planner maximizes preferences of the form

$$\tilde{\mathbb{E}}_0 \sum_{t=0}^{\infty} \beta^t [\log c_t + \xi \log (1 - l_t)] \quad (4)$$

where c_t denotes consumption, l_t is labor, β is the discount factor, and ξ controls the utility from leisure. The expectations operator $\tilde{\mathbb{E}}$ indicates that the planner has a limited information set to be discussed. The planner is subject to a budget constraint

$$c_t + x_t = y_t \quad (5)$$

where x_t denotes investment and y_t denotes output. The latter is produced according to a Cobb-Douglas production function with capital share α ,

$$y_t = z_t^{1-\alpha} k_{t-1}^{\alpha} l_t^{1-\alpha}, \quad (6)$$

where z_t is the TFP process.⁵ Capital is accumulated according to

$$k_t = (1 - \delta) k_{t-1} + u_t x_t, \quad (7)$$

where u_t is the IST process scaling the effectiveness of turning new investment into capital. In a decentralized version of this economy, the IST shock is the inverse of the relative price of capital goods, which supports the use of the data on relative deflators presented in Figure 1.

The two exogenous processes are identical to those estimated in Section 2. The TFP process has independent regimes for the mean and standard deviation

$$\Delta \log z_t = \mu_z (s_t^{\mu z}) + \sigma_z (s_t^{\sigma z}) \varepsilon_{z,t}, \quad (8)$$

and similarly for the IST process

$$\Delta \log u_t = (1 - \rho_u) \mu_u (s_t^{\mu u}) + \rho_u \Delta \log u_{t-1} + \sigma_u (s_t^{\sigma u}) \varepsilon_{u,t}. \quad (9)$$

Again, the regimes are governed by an independent Markov process indexed by s_t^{ij} , for $i \in \{\mu, \sigma\}$ and $j \in \{z, u\}$, where $s_t^{ij} \in \{L, H\}$ denoting “low” or “high.” Each regime variable has the transition matrix

$$P^{ij} = \begin{bmatrix} p_{LL}^{ij} & p_{LH}^{ij} \\ p_{HL}^{ij} & p_{HH}^{ij} \end{bmatrix} = \begin{bmatrix} p_{LL}^{ij} & 1 - p_{LL}^{ij} \\ 1 - p_{HH}^{ij} & p_{HH}^{ij} \end{bmatrix}. \quad (10)$$

3.2 Information and Learning

The planner’s information set \mathcal{I}_t consists of the complete histories of TFP z^t and IST u^t , but not the regimes or the shocks.⁶ This assumption makes the planner equivalent to the econometrician that, conditional on estimates of the model parameters, filters incoming data to perform inference about the current regime, as in Figure 2. The planner updates beliefs about which combination of regimes is in place by Bayesian updating.

⁵The term z_t is raised to the power $1 - \alpha$ to account for the fact that we have a second source of growth in the model; see the accounting in Fernald (2012). Quantitatively, we can obtain very similar results to ours in a model with z_t instead of $z_t^{1-\alpha}$ in the production function. In that scenario, we would have to recalibrate α to match the same set of target moments that we describe later.

⁶The introduction of classical measurement error is in theory possible: the planner in our model would then have to use the Kim (1994) filter to obtain estimates of the state probabilities and the conditional distribution of the unobserved variables. Instead of approximating Bayes’ rule, as we do in our approach, we would then need to approximate the filtering equations.

For example: with TFP, the planner combines the observed growth rate $\Delta \log z_t$ with prior beliefs about the regime to produce an estimated probability of being in the low-growth regime ($s_t^{\mu z} = L$) according to

$$\psi_{L,t}^{\mu z} = \frac{\sum_{j=L,H} \frac{1}{\sigma_z(j)} \phi\left(\frac{\Delta \log z_t - \mu_z(L)}{\sigma_z(j)}\right) (p_{Lj}^{\sigma z} \psi_{L,t-1}^{\sigma z} + p_{Hj}^{\sigma z} \psi_{H,t-1}^{\sigma z}) (p_{LL}^{\mu z} \psi_{L,t-1}^{\mu z} + p_{HL}^{\mu z} \psi_{H,t-1}^{\mu z})}{\sum_{i,j=L,H} \frac{1}{\sigma_z(j)} \phi\left(\frac{\Delta \log z_t - \mu_z(i)}{\sigma_z(j)}\right) (p_{Lj}^{\sigma z} \psi_{L,t-1}^{\sigma z} + p_{Hj}^{\sigma z} \psi_{H,t-1}^{\sigma z}) (p_{Li}^{\mu z} \psi_{L,t-1}^{\mu z} + p_{Hi}^{\mu z} \psi_{H,t-1}^{\mu z})}, \quad (11)$$

Similar equations hold for TFP volatility as well as for IST volatility and growth regimes; these equations can be found in Appendix A. A key takeaway is that increases in the volatility of the innovations (σ_z in this case) flatten the likelihood function of the observed TFP data.

3.3 Equilibrium Conditions

Standard optimization implies a labor-leisure choice trade-off characterized by

$$\xi \frac{c_t}{1 - l_t} = (1 - \alpha) \frac{y_t}{l_t} \quad (12)$$

and an Euler equation given by

$$1 = \beta \tilde{\mathbb{E}}_t \frac{c_t}{c_{t+1}} \left[\alpha u_t \frac{y_{t+1}}{k_t} + \frac{u_t}{u_{t+1}} (1 - \delta) \right]. \quad (13)$$

Since the economy has two unit roots, we define detrended variables $\tilde{y}_t = y_t / z_t u_t^{\frac{\alpha}{1-\alpha}}$, $\tilde{c}_t = c_t / z_t u_t^{\frac{\alpha}{1-\alpha}}$, $\tilde{x}_t = x_t / z_t u_t^{\frac{\alpha}{1-\alpha}}$, $\tilde{k}_t = k_t / z_t u_t^{\frac{1}{1-\alpha}}$, $\tilde{z}_t = z_t / z_{t-1}$, and $\tilde{u}_t = u_t / u_{t-1}$. The normalized equilibrium conditions are therefore

$$\frac{\xi \tilde{c}_t}{1 - l_t} = \frac{1}{\tilde{z}_t^\alpha \tilde{u}_t^{\frac{\alpha}{1-\alpha}}} (1 - \alpha) \tilde{k}_{t-1}^\alpha l_t^{-\alpha}, \quad (14)$$

$$1 = \beta \tilde{\mathbb{E}}_t \frac{\tilde{c}_t}{\tilde{c}_{t+1}} \left(\alpha \frac{\tilde{k}_t^{\alpha-1} l_{t+1}^{1-\alpha}}{\tilde{z}_{t+1}^\alpha \tilde{u}_{t+1}^{\frac{\alpha}{1-\alpha}}} + \frac{1 - \delta}{\tilde{z}_{t+1} \tilde{u}_{t+1}^{\frac{1}{1-\alpha}}} \right), \quad (15)$$

$$\tilde{k}_t = (1 - \delta) \frac{\tilde{k}_{t-1}}{\tilde{z}_t \tilde{u}_t^{\frac{1}{1-\alpha}}} + \tilde{x}_t, \quad (16)$$

$$\tilde{c}_t + \tilde{x}_t = \tilde{y}_t, \quad (17)$$

$$\tilde{y}_t = \frac{1}{\tilde{z}_t^\alpha \tilde{u}_t^{\frac{\alpha}{1-\alpha}}} \tilde{k}_{t-1}^\alpha l_t^{1-\alpha}, \quad (18)$$

$$\log \tilde{z}_t = \mu_z(s_t^{\mu z}) + \sigma_z(s_t^{\sigma z}) \varepsilon_{z,t}, \quad (19)$$

and

$$\log \tilde{u}_t = (1 - \rho_u) \mu_u(s_t^{\mu u}) + \rho_u \log \tilde{u}_{t-1} + \sigma_u(s_t^{\sigma u}) \varepsilon_{u,t}. \quad (20)$$

Equations (14)-(20), along with the learning processes exemplified in equation (11), characterize the full set of equilibrium conditions for the planner's problem. Given this definition of an equilibrium, we turn to our solution method.

4 Solving Markov-Switching Models with Learning

This Section presents a general perturbation-based solution method for Markov-switching models with learning. Our ultimate goal in Section 5 is to use the model developed in Section 3 to characterize how learning about TFP and IST regimes affects economic dynamics. However, a typical solution method for this class of model, policy function iteration (Davig, 2004, for example), runs into the curse of dimensionality very quickly, which makes solving a model with learning about multiple independent Markov processes an arduous task. In contrast, our perturbation method solves very quickly, which allows for likelihood based estimation even in models with rich dynamics.

To describe our solution method, we lay out the general framework, develop a refinement to the Foerster et al. (2016) partition principle that helps us define a steady state around which to approximate, and characterize first- and higher-order approximations. We relegate some of the derivations to Appendix B, while Appendix C includes measures of accuracy in a simplified RBC framework.

4.1 The General Framework

This section lays out the general framework that we consider, showing how a general class of models is combined with a Bayesian learning process, and we refer to our model derived in the previous section as one specific case.

We consider a general class of dynamic, stochastic general equilibrium models where some of the parameters follow a discrete Markov process that is indexed by $s_t \in \{1, \dots, n_s\}$. The regime variable s_t is not directly observed, but has known transition matrix $P = [p_{i,j}]$, where

$p_{i,j} = \Pr(s_{t+1} = j | s_t = i)$. The set of parameters that follow a Markov process is given by $\theta_t = \theta(s_t)$.

The equilibrium conditions for the class of models that we study—which includes first-order conditions, budget constraints, and exogenous processes—can be written as

$$\tilde{\mathbb{E}}_t \mathbf{f}(\mathbf{y}_{t+1}, \mathbf{y}_t, \mathbf{x}_t, \mathbf{x}_{t-1}, \varepsilon_{t+1}, \varepsilon_t, \theta_{t+1}, \theta_t) = 0, \quad (21)$$

where \mathbf{y}_t denotes non-predetermined variables at time t , \mathbf{x}_t denotes predetermined variables, and ε_t denote the innovations which are serially uncorrelated and jointly distributed according with density function ϕ^ε .

In the RBC model discussed in the previous Section, equations (14)-(20) constitute the function in equation (21), and $\mathbf{y}_t = [\tilde{y}_t, \tilde{c}_t, \tilde{x}_t, l_t, \tilde{z}_t]'$, $\mathbf{x}_t = [\tilde{k}_t, \tilde{u}_t]'$, $\varepsilon_t = [\varepsilon_{z,t}, \varepsilon_{u,t}]'$, and $\theta_t = [\mu_{z,t}, \mu_{u,t}, \sigma_{z,t}, \sigma_{u,t}]'$.

The expectations operator $\tilde{\mathbb{E}}_t$ denotes rational expectations based on an information set given by $\mathcal{I}_t = \{\mathbf{y}^t, \mathbf{x}^t\}$. The history of innovations ε^t , parameters θ^t , and regimes s^t is not part of the information set.⁷ The information set produces subjective probabilities of being in each regime $\{1, \dots, n_s\}$, denoted by a vector ψ_t , where $\psi_{i,t} = \Pr(s_t = i | \mathcal{I}_t)$. The subjective probabilities are updated via Bayesian learning.

Since the regime s_t is not observed directly, the equilibrium dynamics depend on the subjective probabilities of being in each regime at time t , conditional on all past observables, which is given by ψ_t , which is a n_s -vector with i -th element given by $\psi_{i,t} = \Pr(s_t = i | \mathcal{I}_t)$. The subjective probabilities are updated via Bayesian learning, which involves combining prior beliefs in the form of last period's subjective probabilities ψ_{t-1} , with information about a known signal $\tilde{\mathbf{y}}_t \subseteq \mathbf{y}_t$.⁸ The signal is generated by a combination of the predetermined variables \mathbf{x}_{t-1} and the shocks ε_t and depends on the regime s_t :

$$\tilde{\mathbf{y}}_t = \tilde{\lambda}_{s_t}(\mathbf{x}_{t-1}, \varepsilon_t). \quad (22)$$

We assume that, given the regime s_t and the predetermined variables \mathbf{x}_{t-1} , there is a one-to-one

⁷Variations on this information set where some shocks are observed but not others can be handled by using identities in the signal process discussed below.

⁸The restriction that signals are part of the non-predetermined variables is without loss of generality, since auxiliary variables can be used to link elements of the set of predetermined and non-predetermined variables.

mapping between shocks and signals, and hence can write

$$\varepsilon_t = \lambda_{s_t}(\tilde{\mathbf{y}}_t, \mathbf{x}_{t-1}). \quad (23)$$

The Jacobian of this mapping is given by

$$J_{s_t}(\tilde{\mathbf{y}}_t, \mathbf{x}_{t-1}) = \left| \frac{\partial \lambda_{s_t}(\tilde{\mathbf{y}}_t, \mathbf{x}_{t-1})}{\partial \tilde{\mathbf{y}}_t} \right|, \quad (24)$$

where $|\cdot|$ denotes the determinant.

In the RBC model of Section 3, the signals are \tilde{z}_t and \tilde{u}_t . The function $\tilde{\lambda}_{s_t}$ is the exogenous processes, and the inverse mappings λ_{s_t} are given by

$$\varepsilon_{z,t} = \frac{\log \tilde{z}_t - \mu_z(s_t^{\mu z})}{\sigma_z(s_t^{\sigma z})} \quad (25)$$

and

$$\varepsilon_{u,t} = \frac{\log \tilde{u}_t - (1 - \rho_u) \mu_u(s_t^{\mu u}) - \rho_u \log \tilde{u}_{t-1}}{\sigma_u(s_t^{\sigma u})} \quad (26)$$

Given known signals $\tilde{\mathbf{y}}_t \subseteq \mathbf{y}_t$, predetermined variables \mathbf{x}_{t-1} , prior probabilities ψ_{t-1} , transition probabilities $P = [p_{i,j}]$, and a joint density function of the errors ϕ^ε , Bayesian updating implies

$$\psi_{i,t} = \frac{J_{s_t=i}(\tilde{\mathbf{y}}_t, \mathbf{x}_{t-1}) \phi^\varepsilon(\lambda_{s_t=i}(\tilde{\mathbf{y}}_t, \mathbf{x}_{t-1})) \sum_{s=1}^{n_s} p_{s,i} \psi_{s,t-1}}{\sum_{j=1}^{n_s} J_{s_t=j}(\tilde{\mathbf{y}}_t, \mathbf{x}_{t-1}) \phi^\varepsilon(\lambda_{s_t=j}(\tilde{\mathbf{y}}_t, \mathbf{x}_{t-1})) \sum_{s=1}^{n_s} p_{s,j} \psi_{s,t-1}}. \quad (27)$$

For ease in constructing approximations that appropriately bound the probabilities between zero and one, define the logit of the probabilities $\eta_{i,t} = \log \left(\frac{\psi_{i,t}}{1-\psi_{i,t}} \right)$, which in turn implies $\psi_{i,t} = \frac{1}{1+\exp(-\eta_{i,t})}$.⁹ These logits are expressed as

$$\exp \eta_{i,t} = \frac{J_{s_t=i}(\tilde{\mathbf{y}}_t, \mathbf{x}_{t-1}) \phi^\varepsilon(\lambda_{s_t=i}(\tilde{\mathbf{y}}_t, \mathbf{x}_{t-1})) \sum_{s=1}^{n_s} p_{s,i} \frac{1}{1+\exp(-\eta_{s,t-1})}}{\sum_{j=1, j \neq i}^{n_s} J_{s_t=j}(\tilde{\mathbf{y}}_t, \mathbf{x}_{t-1}) \phi^\varepsilon(\lambda_{s_t=j}(\tilde{\mathbf{y}}_t, \mathbf{x}_{t-1})) \sum_{s=1}^{n_s} p_{s,j} \frac{1}{1+\exp(-\eta_{s,t-1})}}. \quad (28)$$

We denote the vector of logits of the probabilities as η_t .¹⁰

⁹In principle, any function mapping $(0, 1)$ to $(-\infty, \infty)$ in a one-to-one fashion could be used in this step, but the logit function is a natural choice since it uses exponential and log functions consistent with log-approximations, has a readily computed inverse function as noted, and derivatives that are easily calculated symbolically.

¹⁰In practice, we can use the fact that $\sum_{i=1}^{n_s} \psi_{i,t} = 1$ and only generate approximations to at most $n_s - 1$ elements of the vector η_t . In our specific RBC case, we can use independence of the processes to only generate approximations for $\eta_t^{\mu z}$, $\eta_t^{\sigma z}$, $\eta_t^{\mu u}$, and $\eta_t^{\sigma u}$.

As a result, we write the equations characterizing the learning process as

$$\Phi(\mathbf{y}_t, \mathbf{x}_{t-1}, \eta_t, \eta_{t-1}, \Theta) = 0 \quad (29)$$

where the i -th equation of Φ is given by rearranging (28) into implicit form. In our RBC example, the logistic versions of equations shown in Appendix A constitute the equations Φ . Whereas θ_t and θ_{t+1} in equation (21) refer to specific realizations of the Markov-switching parameters in a period, $\Theta = [\theta(1), \dots, \theta(n_s)]'$ denotes the complete set of regime-switching parameters, which naturally appear in the definitions of the probabilities. The full vector of regime-switching parameters Θ , and not just the current regime's values θ_t , matter for the learning process because the Bayesian updating weighs the relative likelihood of the observables being generated by each possible regime.

Note that, since the learning process is independent of future variables, we trivially have

$$\tilde{\mathbb{E}}_t \Phi(\mathbf{y}_t, \mathbf{x}_{t-1}, \eta_t, \eta_{t-1}, \Theta) = \Phi(\mathbf{y}_t, \mathbf{x}_{t-1}, \eta_t, \eta_{t-1}, \Theta). \quad (30)$$

4.2 Equilibrium Conditions with Learning and Solutions

To characterize the full equilibrium conditions with the Bayesian updating of subjective probabilities, we can simply append the equations (29) to the original equilibrium conditions in equation (21). This produces a system of the form

$$\begin{aligned} \tilde{\mathbb{E}}_t \tilde{\mathbf{f}}(\mathbf{y}_{t+1}, \mathbf{y}_t, \mathbf{x}_t, \mathbf{x}_{t-1}, \eta_t, \eta_{t-1}, \varepsilon_{t+1}, \varepsilon_t, \theta_{t+1}, \theta_t, \Theta) &= \\ \tilde{\mathbb{E}}_t \begin{bmatrix} \mathbf{f}(\mathbf{y}_{t+1}, \mathbf{y}_t, \mathbf{x}_t, \mathbf{x}_{t-1}, \varepsilon_{t+1}, \varepsilon_t, \theta_{t+1}, \theta_t) \\ \Phi(\mathbf{y}_t, \mathbf{x}_t, \mathbf{x}_{t-1}, \eta_t, \eta_{t-1}, \Theta) \end{bmatrix} &= 0. \end{aligned} \quad (31)$$

It is important to point out that the presence of all full regime-switching parameters Θ is due to the introduction of learning—it is absent from the corresponding equation in the corresponding full information system as described in Foerster et al. (2016). This seemingly small difference will force us to modify the solution algorithm from Foerster et al. (2016), as described in the next section.

The expectation can be decomposed into the subjective probabilities, transitions between

regimes, and expectations over future shocks:

$$\sum_{s=1}^{n_s} \frac{1}{1 + \exp(-\eta_{s,t})} \sum_{s'=1}^{n_s} p_{s,s'} \int \tilde{\mathbf{f}}(\mathbf{y}_{t+1}, \mathbf{y}_t, \mathbf{x}_t, \mathbf{x}_{t-1}, \eta_t, \eta_{t-1}, \varepsilon', \varepsilon_t, \theta(s'), \theta(s_t), \boldsymbol{\Theta}) \phi^\varepsilon(\varepsilon') = 0. \quad (32)$$

Extending Foerster et al. (2016) to the case where subjective probabilities are now a state variable, minimum state variable (MSV) solutions to the model in equation (32) have the form

$$\mathbf{y}_t = g_{s_t}(\mathbf{x}_{t-1}, \eta_{t-1}, \varepsilon_t, \chi), \quad (33)$$

$$\mathbf{x}_t = h_{s_t}^x(\mathbf{x}_{t-1}, \eta_{t-1}, \varepsilon_t, \chi), \quad (34)$$

and

$$\eta_t = h_{s_t}^\eta(\mathbf{x}_{t-1}, \eta_{t-1}, \varepsilon_t, \chi), \quad (35)$$

The form of these solutions show that the evolution of non-predetermined variables \mathbf{y}_t , predetermined variables \mathbf{x}_t , and beliefs η_t , depends upon the actual realized regime s_t , the previous values of \mathbf{x}_{t-1} and η_{t-1} , the realization of shocks ε_t , and a perturbation parameter χ .

Perturbation seeks to generate Taylor series expansions to the functions g_{s_t} , $h_{s_t}^x$, and $h_{s_t}^\eta$ around a given point. The following turns to how to define this point.

4.3 The Refined Partition Principle

A key feature of the perturbation solution of full-information rational expectation models with Markov-switching as in Foerster et al. (2016) is that the steady state does not depend on the regime. The same argument holds in our setup. In order to increase numerical efficiency, the Partition Principle in Foerster et al. (2016) dictates separating the switching parameters $\theta(k)$ into blocks denoted $\theta_1(k)$ and $\theta_2(k)$, for $k \in \{1, \dots, n_s\}$, where the first set are perturbed and the second set are not. This partition of the parameters allows for finding a steady state and preserving the maximum information at lower orders of approximation. In particular, the perturbation function is

$$\theta(k, \chi) = \chi \begin{bmatrix} \theta_1(k) \\ \theta_2(k) \end{bmatrix} + (1 - \chi) \begin{bmatrix} \bar{\theta}_1 \\ \theta_2(k) \end{bmatrix} \quad (36)$$

for for $k \in \{1, \dots, n_s\}$. The set of parameters included in $\theta_2(k)$ is chosen to be the maximal set such that a steady state is defined. In our RBC example, $\theta_1(k) = [\mu_z(k), \mu_u(k)]'$ and

$$\theta_2(k) = [\sigma_z(k), \sigma_u(k)]'.$$

The definition of a steady state is when $\varepsilon_t = 0$, $\chi = 0$, and hence for all s_t

$$\mathbf{y}_{ss} = g_{s_t}(\mathbf{x}_{ss}, \eta_{ss}, \mathbf{0}, 0), \mathbf{x}_{ss} = h_{s_t}^x(\mathbf{x}_{ss}, \eta_{ss}, \mathbf{0}, 0), \text{ and } \eta_{ss} = h_{s_t}^\eta(\mathbf{x}_{ss}, \eta_{ss}, \mathbf{0}, 0). \quad (37)$$

In the case with learning as shown in the set of equilibrium equations (32), the presence of all the regime-switching parameters Θ in the learning process poses a challenge. A natural extension of the Partition Principle would suggest perturbing the same sets of parameters in Θ so that the steady state is independent of any switching parameters. That is, we could write

$$\Theta(\chi) = \chi \begin{bmatrix} \theta_1(1) \\ \theta_2(1) \\ \vdots \\ \theta_1(n_s) \\ \theta_2(n_s) \end{bmatrix} + (1 - \chi) \begin{bmatrix} \bar{\theta}_1 \\ \theta_2(1) \\ \vdots \\ \bar{\theta}_1 \\ \theta_2(n_s) \end{bmatrix}. \quad (38)$$

However, this assumption would lead to a loss of information in the steady state and at low orders of approximation. To see this point, note that when all parameters are perturbed, $J_{s_t=i}(\tilde{\mathbf{y}}_{ss}, \mathbf{x}_{ss}) \phi^\varepsilon(\lambda_{s_t=i}(\tilde{\mathbf{y}}_{ss}, \mathbf{x}_{ss}))$ is constant for all i , and hence the steady state satisfies

$$\psi_{i,ss} = \frac{\sum_{s=1}^{n_s} p_{s,i} \psi_{s,ss}}{\sum_{j=1}^{n_s} \sum_{s=1}^{n_s} p_{s,j} \psi_{s,ss}}, \quad (39)$$

which implies that ψ_{ss} is the ergodic vector associated with the transition matrix P .

The fact that perturbing Θ in a symmetric fashion to $\theta(k)$ generates steady state beliefs equal to the ergodic probability across regimes is a reasonable choice, but in practice leads to a loss of information due to the fact that the implied probabilities from each regime are different at the steady state of $(\tilde{\mathbf{y}}_{ss}, \mathbf{x}_{ss})$. In other words, a straightforward application of the Partition Principle would eliminate too much information about the distribution of the signal from the steady state, and this loss of information would in turn lower approximation quality.

To resolve this issue, one option is to not perturb any part of Θ , treating it differently than θ_t entirely. In this *Partition Principle Refinement*, we would leave Θ unchanged, and hence the steady state would be defined by

$$\tilde{\mathbf{f}}(\mathbf{y}_{ss}, \mathbf{y}_{ss}, \mathbf{x}_{ss}, \mathbf{x}_{ss}, \eta_{ss}, \eta_{ss}, \mathbf{0}, \mathbf{0}, \bar{\theta}_1, \bar{\theta}_2(s'), \bar{\theta}_1, \bar{\theta}_2(s), \Theta) = 0 \quad (40)$$

for all s and s' . In Appendix C we show that accuracy is improved with the Refinement.

Having discussed the Partition Principle Refinement, we return to a full definition of the equilibrium, which is

$$\sum_{s=1}^{n_s} \sum_{s'=1}^{n_s} \frac{p_{s,s'}}{1 + \exp(-\eta_{s,t})} \int \tilde{\mathbf{f}}(\mathbf{y}_{t+1}, \mathbf{y}_t, \mathbf{x}_t, \mathbf{x}_{t-1}, \eta_t, \eta_{t-1}, \varepsilon', \varepsilon_t, \theta_1(s'), \theta_2(s'), \theta_1(s_t), \theta_2(s_t), \boldsymbol{\Theta}) \phi_\varepsilon(\varepsilon') = 0. \quad (41)$$

Using the functional forms (33), (34), (35), and (36) produces an equation of the form

$$\mathbf{F}_{s_t}(\mathbf{x}_{t-1}, \eta_{t-1}, \varepsilon_t, \chi) = 0. \quad (42)$$

We will take derivatives of this function, evaluated at steady state, to find approximations to the policy functions (33), (34), (35).

The steady state is given by the set of equations

$$\mathbf{F}_{s_t}(\mathbf{x}_{ss}, \eta_{ss}, \mathbf{0}, 0) = 0,$$

for all s_t . Since the first n equations of $\tilde{\mathbf{f}}$ are the original equilibrium conditions in equation (21), then the steady state satisfies

$$f(\mathbf{y}_{ss}, \mathbf{y}_{ss}, \mathbf{x}_{ss}, \mathbf{x}_{ss}, \mathbf{0}, \mathbf{0}, \bar{\theta}_1, \theta_2(s_{t+1}), \bar{\theta}_1, \theta_2(s_t)) = 0, \quad (43)$$

or in other words, is identical to the steady state to a version of the model with full information and can be used to solve for the n unknowns $\{\mathbf{y}_{ss}, \mathbf{x}_{ss}\}$. The second set of n_s equations of \tilde{f} are the Bayesian updating equations, and so

$$\Phi(\mathbf{y}_{ss}, \mathbf{x}_{ss}, \mathbf{x}_{ss}, \eta_{ss}, \eta_{ss}, \boldsymbol{\Theta}) = 0$$

pins down the n_s unknowns η_{ss} .

4.4 Approximations and Their Properties

We can take derivatives of $\mathbf{F}_{s_t}(\mathbf{x}_{t-1}, \eta_{t-1}, \varepsilon_t, \chi) = 0$ with respect to its arguments to get equations that allow us to solve for the coefficients in the expansions of (33), (34), and (35). We relegate details of the derivations to Appendix B, but summarize the main results below.

Proposition 1 (Irrelevance of Beliefs in First-Order) *In a first-order approximation, the decision rules for \mathbf{y}_t and \mathbf{x}_t are invariant to η_{t-1} . Mathematically,*

$$g_{\eta,s_t} = \left. \frac{\partial g_{s_t}(\mathbf{x}_{t-1}, \eta_{t-1}, \varepsilon_t, \chi)}{\partial \eta_{t-1}} \right|_{\mathbf{x}_{t-1}=\mathbf{x}_{ss}, \eta_{t-1}=\eta_{ss}, \varepsilon_t=0, \chi=0} = 0$$

and

$$h_{\eta,s_t} = \left. \frac{\partial h_{s_t}(\mathbf{x}_{t-1}, \eta_{t-1}, \varepsilon_t, \chi)}{\partial \eta_{t-1}} \right|_{\mathbf{x}_{t-1}=\mathbf{x}_{ss}, \eta_{t-1}=\eta_{ss}, \varepsilon_t=0, \chi=0} = 0.$$

Proof. See Appendix B. ■

The implications of Proposition 1 are striking, in that it immediately says that we must go to at least a second-order approximation using our perturbation method if we want learning to play a role in dynamics. In first-order, the decision rules for the economic variables—such as consumption, investment, and labor in our RBC model—do not respond to the evolution of beliefs, and hence learning does not appear in the dynamics of the economy. This result is analogous to the result that second-order terms are needed to capture precautionary behavior in models without regime switching (Schmitt-Grohe and Uribe, 2004), or that second-order terms are needed to capture the effects of endogenous regime switching (Benigno et al., 2020).

Proposition 2 (Limited Impact of Beliefs in Second-Order) *If the signal is independent of predetermined variables, so that $\varepsilon_t = \lambda_{s_t}(\tilde{\mathbf{y}}_t)$, then a second-order approximation has no state-dependent effects of beliefs. Mathematically,*

$$g_{x\eta,s_t} = \left. \frac{\partial g_{s_t}(\mathbf{x}_{t-1}, \eta_{t-1}, \varepsilon_t, \chi)}{\partial \mathbf{x}_{t-1} \partial \eta_{t-1}} \right|_{\mathbf{x}_{t-1}=\mathbf{x}_{ss}, \eta_{t-1}=\eta_{ss}, \varepsilon_t=0, \chi=0} = 0$$

and

$$h_{x\eta,s_t} = \left. \frac{\partial h_{s_t}(\mathbf{x}_{t-1}, \eta_{t-1}, \varepsilon_t, \chi)}{\partial \mathbf{x}_{t-1} \partial \eta_{t-1}} \right|_{\mathbf{x}_{t-1}=\mathbf{x}_{ss}, \eta_{t-1}=\eta_{ss}, \varepsilon_t=0, \chi=0} = 0.$$

Proof. See Appendix B. ■

Proposition 2 matters for our RBC example in that it states that when the signal is independent of the endogenous predetermined variables—TFP and IST are independent of the capital stock in our case—then second-order approximations do not capture interactions between beliefs and economic states. To the extent that we are interested in capturing the effects of learning as much as possible, this Proposition then implies that we need to use a third-order approximation in order to have dynamics that reflect interactions between the learning and the capital stock.

In other words, we want to capture the state-dependence implied by different consumption or labor decisions when beliefs are close to 0 or 1, given some capital stock. As we will show, these dynamics are important for generating asymmetries in the impulse responses to shocks. This result is similar to the fact that second-order approximations are not sufficient to capture the effects of time-varying volatility in models without regime switching (Fernandez-Villaverde et al., 2015).

Lastly, our methodology can be used to characterize existence and uniqueness of MSV solutions using the concepts of Gröbner Bases and Mean Square Stability (MSS) as in Foerster et al. (2016). Gröbner bases are used to solve a generalized quadratic form that appears when solving the first-order approximation. Foerster et al. (2016) find all possible solutions to the quadratic equation, and check each for stability according to the MSS concept, which requires finite first and second-moments in expectation:

$$\lim_{j \rightarrow \infty} \tilde{\mathbb{E}}_t [\mathbf{x}_{t+j}] = \bar{\mathbf{x}}, \text{ and } \lim_{j \rightarrow \infty} \tilde{\mathbb{E}}_t [\mathbf{x}'_{t+j} \mathbf{x}_{t+1}] = \Sigma.$$

After checking stability of each possible solution, we are able to characterize whether the approximation has a unique MSV solution or not.¹¹ In our RBC application we find a unique equilibrium, but more broadly our solution algorithm allows researchers to study equilibrium non-uniqueness in non-linear models of Bayesian learning.

5 The Economic Effects of Learning about Technology Growth Regimes

In this Section we present our main analysis of how learning about technology growth regimes affects macroeconomic dynamics. We use the canonical real business cycle model from Section 3, coupled with the technology processes according to our estimates in Section 2, and solve the model using the method developed in Section 4. First we compute the fit of the model to data on consumption and investment growth, assessing whether learning or full information about regime switches, as well as what specification of the TFP and IST growth processes, lead to a better fit. Then we focus on our preferred specification and investigate the implications of learning about regime change by studying impulse responses, followed by an analysis of forecast

¹¹In our context, MSS as opposed to bounded stability (for example, Davig and Leeper, 2007; Barthlemy and Marx, 2017), as a stability concept is an important choice since we need to be able to construct a realization of regimes that generates arbitrarily high confidence in a given regime, with $\eta_{i,t} \rightarrow \pm\infty$ ($\psi_{i,t} \rightarrow 0$ or 1).

errors, and finally discuss the welfare implications of learning information.

5.1 Model Fit

We first study the effects of learning by examining the ability of the real business cycle model to fit data, and how this fit changes with learning or with full information and with different assumptions on the nature of regime switches. In particular, we use standard parameter values describing preferences and production, along with our estimates of TFP and IST processes from Table 1, along with a non-linear filter to find the likelihood associated with different combinations of information structure and driving processes.¹² Given our relatively parsimonious model, we focus on matching the growth rates of per capita real consumption and investment along with our data on TFP and IST.

We start by picking parameters for preferences and production. We choose the discount factor β to be 0.997, which, given the average growth in TFP and IST, implies a steady state annual real interest rate of 3%. The capital share in production α is chosen to be 0.298, which likewise given the average growth in TFP and IST yields a annual per capita consumption growth equal to the data. The steady state of hours is picked to be 1/3, and the annual depreciation rate of capital is 10%.

Given this parameterization, we solve the model to a third-order, in order to include the state-dependent effects of learning highlighted in Propositions 1 and 2, as well as due to the accuracy checks in Appendix C. Given this solution, we use the unscented Kalman filter (UKF) with sigma points to compute the likelihood associated with each specification of the model. The UKF has the advantage of being well-suited for our setup with multiple regimes, and decision rules that are solved to the third-order conditional on the regime. The UKF has been shown to work well in similar situations with nonlinearities in regimes and higher-order approximations (Binning and Maih, 2015; Benigno et al., 2020). We use the TFP and IST growth values studied in Section 2, plus the growth rates of per capita real consumption and investment. Since we are explaining four observables with two shocks, we include measurement errors equal to 0.1% of the variance of the observed series.

Recall from the empirical estimates presented in Table 1 that while the estimates seemed

¹²Our empirical strategy of first estimating nonlinear exogenous forcing processes and then using those in an otherwise calibrated equilibrium model mirrors the approach in Fernandez-Villaverde et al. (2015), who first estimate non-linear tax and spending processes and then use simulated method of moments to obtain the structural parameters. In our case, our parameters are readily calibrated using steady state relationships and we do not have to rely on a method of moments estimator.

to indicate the existence of distinct regimes for the means and volatilities of both processes, the improvement in the log likelihoods was possibly minor. In particular, the log likelihoods indicated relatively weak support for regime switching in the TFP process and means of the IST process, and relatively strong evidence for switching in volatilities of the IST process.

Table 2 shows the log likelihoods for different model specifications. We compare five different specifications for regime switching behavior, and full information versus learning for both. In the first specification, shown in row (1) of the Table, we use a fixed coefficient model with our estimates with no regime switching in either TFP or IST; in this case learning is irrelevant. In the next three specifications, shown in rows (2)-(4) of the Table, we use different combinations of regime switching for each process: either TFP or IST have two mean and two volatility regimes, or both do.¹³ In this case, we can compare the version of the model where the planner has full information and observes the regime with the version where the planner does not observe the regime and must learn it. Finally, the last specification that we study is one in which there are shifts in the TFP and IST processes at given break dates discussed in Section 2, but these come as surprises; in the full information case for this specification, the planner observes that a shift has occurred and assumes the new regime will be in place forever, while in the learning case the planner learns that a shift may have occurred and must update the beliefs about whether the regime actually changed. This last version has the feature of shutting down expectation effects since in-between switches the planner does not expect future switches in the regime.

The log likelihoods in Table 2 show strong evidence for both the presence of regime switching behavior in the processes and for learning about the regime. The fixed coefficient model (1) sets a baseline log-likelihood. For models (2)-(4) that introduce regime switching, there are very sizable gains in the log likelihood regardless of the source of the switches. For example, model (3) introduces switching only in TFP; the statistical model in Section 2 showed weak evidence supporting this switching. By contrast, making the real business cycle model fit consumption and investment growth data as well demonstrates that there are major gains to model fit by introducing TFP switches. The most favored model for both full information and learning is model (4) with switches in means and volatilities for both processes. Focusing on this preferred model, there is a modest but notable increase in fit for the learning version, which is the best fit among all versions by a sizable margin. Finally, model (5) shows that expectation effects about the possibility of future regimes are important, since assuming the planner thinks the current regime will last forever leads to less fit relative to the preferred model.

¹³We focus on the case with two μ and two σ for simplicity, cases with one of either that use our estimates from Table 1 do not change the conclusions from this exercise.

Table 2: Model Fit

Model	Full Info	Learning
(1) Fixed Coefficient Model	-51023	-
Regime Switching Models		
(2) TFP: One μ , One σ , IST: Two μ , Two σ	-41369	-41124
(3) TFP: Two μ , Two σ , IST: One μ , One σ	-43765	-46934
(4) TFP: Two μ , Two σ , IST: Two μ , Two σ	-33816	-33391
(5) Break Dates that are Surprises	-38264	-47566

Note: Table shows the value of the likelihood function at the Maximum Likelihood estimates displayed in Table 1. The data used to compute these estimates are quarterly growth rates of TFP, IST, real per-capita consumption and real per-capita investment.

Figure 4 shows the data for consumption and investment growth that we use as additional observables, along with the fitted values from the full information and learning versions of the preferred model (4) with regime switches in both means and volatilities for both processes. Both models generally track the data most of the time, but the full information model generates a number of large positive or negative spikes in the growth rates of both consumption and investment growth that are not present in the data. These large spikes are not present in the learning model, and account for the superior fit of that version.


Comparing the timing of the large spikes with the filtered probabilities of each regime in Figure 2, the large moves produced in the full information case tend to be around dates in which one of the regimes changes. Intuitively, one of the reasons for the superior fit of the learning model is that it tends to smooth the responses of the endogenous variables seen around regime changes, whereas in the full information case a switch in the regime leads to a counterfactually large response by the planner. To develop this intuition and better understand the differences between the learning and full information models, we next turn to analyzing the impulse responses.

5.2 Impulse Responses


The model has multiple sources of impulses: regime changes in the means or volatilities of each process, and the responses to intra-regime shocks captured by fluctuations in the $\epsilon_{z,t}$ and $\epsilon_{u,t}$ terms. We focus our analysis on regime changes in the mean and intra-regime shocks rather than regime changes in volatility, since those are the relatively more important drivers of economic

Figure 4: Model Fit

(a) Consumption Growth (Percent, Annual Rate)



(b) Investment Growth (Percent, Annual Rate)

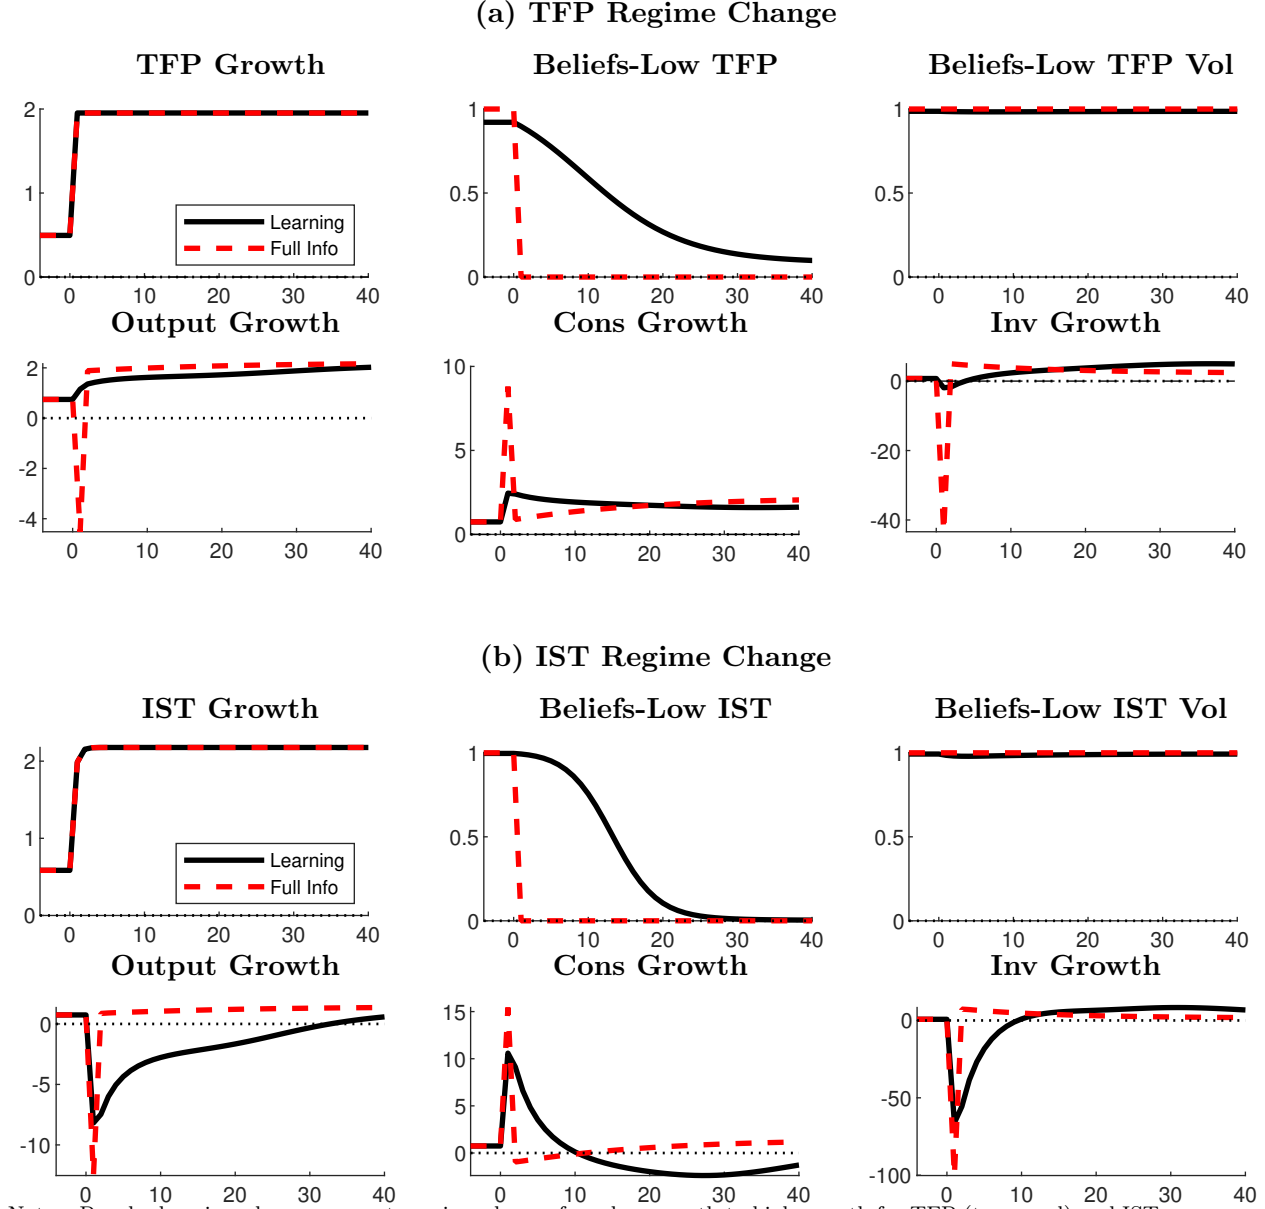


Notes: Plots show the growth rates of per capita consumption and investment from the data and the smoothed estimates from the learning and full information models with regime switches in means and volatilities in both TFP and IST growth (see row (4) of Table 2).

dynamics.

First, examining the effects of regime switches, a major factor in generating the differences between the full information model and the learning model is the endogenous response to regime

Figure 5: Impulse Response to Regime Changes



Notes: Panels show impulse responses to regime change from low growth to high growth for TFP (top panel) and IST (bottom panel), conditional on all other regimes being in the low regime. Growth rates are in annualized percents.

changes. Figure 5a shows the response to a switch from the low TFP growth regime to the high growth regime, conditional on all other regimes staying in the low regime. The change boosts realized TFP growth in period 0, and in the full information case the subjective probability of being in the low growth regime drops from 1 to 0. Upon realization of this regime, there is an immediate spike in consumption growth due to a wealth effect; likewise there is a boost in

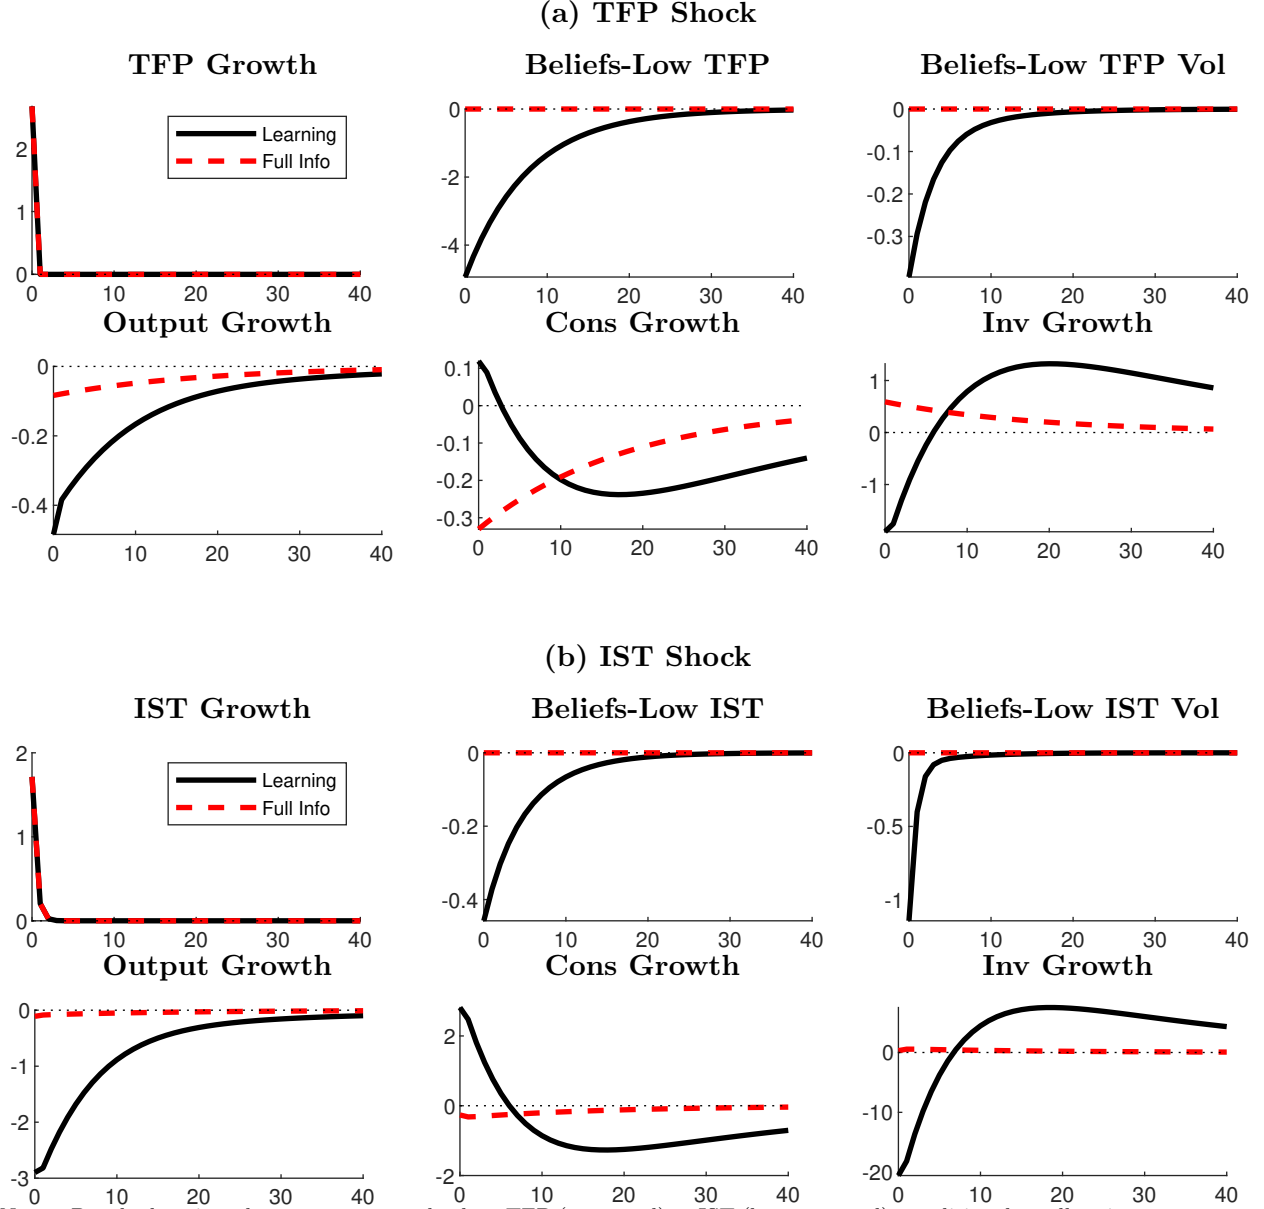
leisure that lowers labor, and hence output and investment. The response of these variables is immediate and large, but they quickly converge to the new level in the period after a shock.

With learning, on the other hand, it takes the planner quite some time to realize the regime has shifted, as beliefs of being in the low TFP growth regime fall very slowly—even 10 periods after the fact the planner puts about 50 percent weight on each regime, even though prior to the shift the probability of being in the low regime was less than 1. This uncertainty in the regime negates the wealth effects that spurred much of the movement in decisions in the full information case. Since the planner is not sure if the high realizations of TFP were from a switch or just positive idiosyncratic shocks, consumption, investment, and labor all increase slightly before leveling off. This mechanism is similar to that studied by Milani (2007), in that learning generates endogenous persistence rather than relying on additional frictions.

A switch in the IST regime from low to high, conditional on all other regimes being low, is shown in Figure 5b, and has similar results to the TFP switch, but the learning dynamics show slight differences that are informative about the importance of the filtering. The economic behavior for the full information case show sharp responses to the regime change, with surges in consumption and leisure. The learning model shows more modest, gradual reductions. However, relative to the TFP switch where the planner learned the new regime very gradually, the IST switch is learned relatively quickly, with beliefs of being in the low growth regime converging to zero within about 10 quarters. This dynamic is understandable given the parameter estimates in Table 1 and especially given the filtered probabilities in Figure 2, which showed the IST regimes as relatively far apart and hence easier distinguished than the TFP regimes. The fact that the filtered probabilities for IST tended to be closer to zero or one than the TFP regimes suggests that IST regimes are more easily identified in real time, and Figure 5b demonstrates that intuition carries over to the economic model.

The fact that the planner in the learning case is not certain when the regime change occurs has implications for the response of an intra-regime shock. Figure 6a shows the response to a TFP shock, conditional on staying in the low-growth regime for the duration of the time plotted, with all other switching parameters likewise in the low regime. For intra-regime shocks, for ease of comparison, we plot the detrended variables rather than growth, and relative to their pre-shock period level. With a positive shock, the planner is uncertain about whether there has been a regime change or simply a transitory shock. Beliefs of being in the low TFP regime decline by about 4 percentage points in this case. Relative to the full information case, the output response is larger, as the learning planner internalizes the possibility that a longer-duration regime shift has occurred. The dynamics of investment flip between the two cases, as the full-information

Figure 6: Impulse Response to Shocks



Notes: Panels show impulse responses to a shock to TFP (top panel) or IST (bottom panel), conditional on all regimes being in the low regime. All panels show annualized differences relative to pre-shock period, except beliefs which are percentage point differences.

agent realizes a temporary shock is a good time to invest immediately, whereas the learning agent who believes that a more permanent shift in growth has occurred lowers investment on impact, but then makes up for it in the future. The fact that dynamics of investment flip is notable in that it reverses the sign of the conditional correlation between consumption and

investment growth on impact from positive to negative.

Figure 6b shows the impulse response to an intra-regime IST shock, conditional on all switching processes in the low regime. Similar to the TFP case, this shock triggers a decline in beliefs that the low IST growth is in place, although the magnitude is much smaller. The temporary increase in IST leads to a slight increase in investment that crowds out consumption, as the planner takes advantage of the opportunity with low relative cost of investment. With learning, the planner responds with a large decrease in investment and an increase in consumption as wealth effects interact with the beliefs of the shock being a regime change.

Because the learning planner who receives a shock thinks there is the possibility that a regime change occurred, there is scope for non-linearities and asymmetries in the responses. An implication of Proposition 2 is that these features only appear once we go beyond a second-order approximation in our solution method. Figure 7a shows the extent of these non-linearities by plotting shocks to TFP of different magnitudes and signs, conditional on the economy being in the low TFP growth regime. Figure 7b likewise shows the asymmetries generated by IST shocks. More precisely, each Figure panel shows the differences between a impulse responses to a two standard deviation shock, and a negative one standard deviation shock. We re-scale both of these sets of impulse responses—we divide the responses to a two standard deviation shock by 2 and multiply the responses to the negative one standard deviation shock by -1 . We then plot both sets of impulse responses relative to a positive one standard deviation shock.¹⁴ In a purely linear model, the impulse responses would be identical and thus the lines in Figures 7a and 7b would be identically equal to 0 for all horizons; the fact that there are differences highlights the non-linearities.

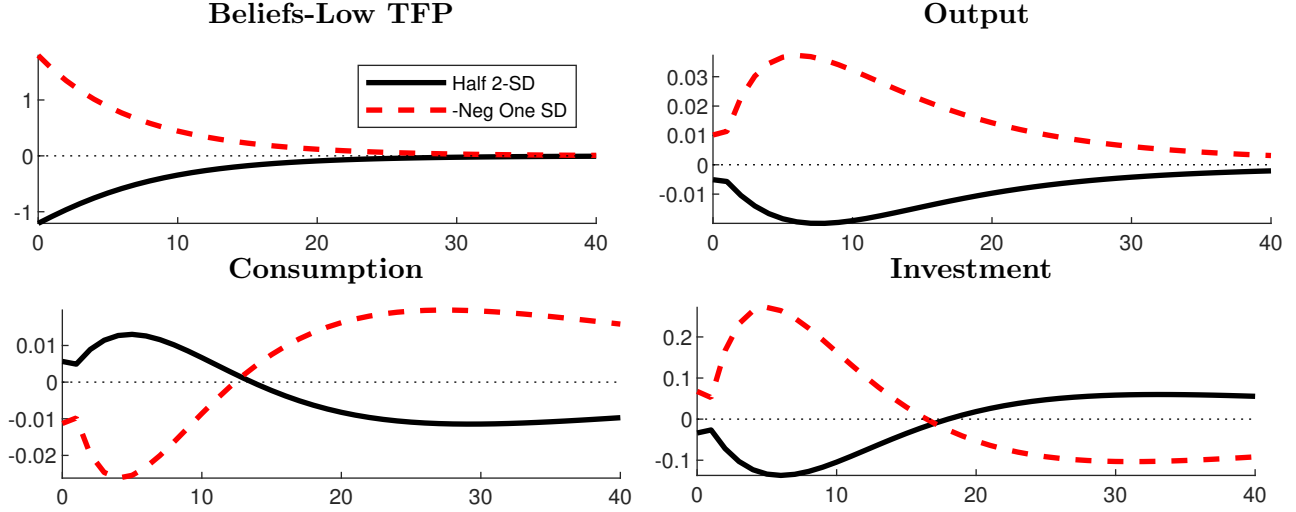
As noted in Figures 6a and 6b, a key feature of the response to shocks is to what extent the planner views the shock as permanent, and hence can wait to invest, as well as take more leisure, which ultimately lowers output. Figures 7a and 7b indicates that this effect increases the shock standard deviation shock, since investment and output increase less than proportionally. Indeed, a negative shock mutes this effect, as investment and output move by a larger amount than from a positive shock.

The results of these impulse responses therefore point to several conclusions about the economic effects of learning. First, learning tends to smooth out the responses to regime switches, since the planner takes time to realize the switch has actually occurred. Second, learning distorts the responses to shocks, because the planner thinks a regime change has possibly occurred. And

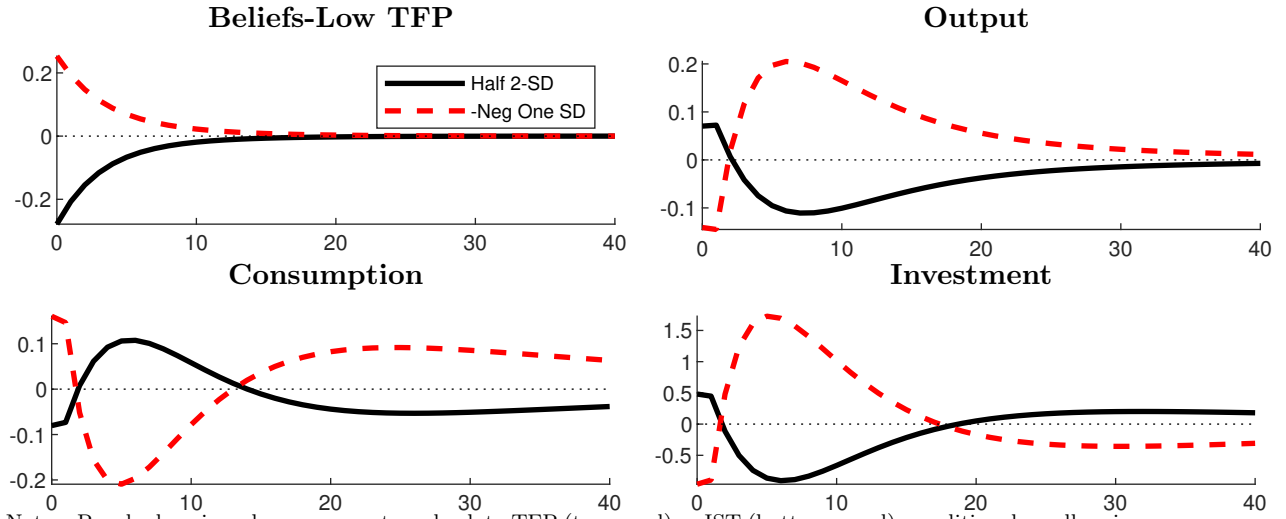
¹⁴Mathematically, if $IRF_{i,t}(\varepsilon_{z,0})$ denotes the response of variable i at time t to a shock $\varepsilon_{z,0}$, we are plotting $\frac{1}{2}IRF_{i,t}(\varepsilon_{z,0} = 2) - IRF_{i,t}(\varepsilon_{z,0} = 1)$ and $-IRF_{i,t}(\varepsilon_{z,0} = -1) - IRF_{i,t}(\varepsilon_{z,0} = 1)$ for TFP, and similarly for IST.

Figure 7: Asymmetries in Impulse Responses to Shocks with Learning

(a) TFP Shock



(b) IST Shock



Notes: Panels show impulse responses to a shock to TFP (top panel) or IST (bottom panel), conditional on all regimes being in the low regime. Shocks are relative to the impulse response to a one standard deviation shock. All plots show annualized percent deviations relative to a one standard deviation shock.

third, these responses to shocks are non-linear and asymmetric, as different magnitude or signs of shocks have different implications for the evolution of beliefs.

Table 3: The Economic Effects of Learning

	Full Info			Learning		
Panel A: Mean Forecast Errors (pp)	TFP	IST	GDP	TFP	IST	GDP
$s^{\mu z} = L, s^{\mu u} = L$	-0.02	-0.01	-0.03	-0.13	-0.02	-0.13
$s^{\mu z} = L, s^{\mu u} = H$	-0.02	0.01	-0.02	-0.13	0.02	-0.12
$s^{\mu z} = H, s^{\mu u} = L$	0.02	-0.01	0.02	0.14	-0.02	0.13
$s^{\mu z} = H, s^{\mu u} = H$	0.02	0.01	0.03	0.14	0.02	0.14
Panel B: Welfare (% Cons Equiv Units)	-0.98			-3.23		

Notes: Panel A shows model-implied forecast errors conditional on a given regime. Welfare in Panel B is conditional on starting from the steady state, prior to the regime being drawn.

5.3 Forecast Errors

We next turn to the possibility of learning about growth regimes to generate systematic forecast errors, even though the planner is behaving optimally and has rational expectations. Survey evidence of professional forecasters indicates that expectations of future growth are often biased (Conlisk, 1996; Laster et al., 1999).

Panel A of Table 3 shows that learning can capture biased expectations about growth. The fact that *rational* Bayesian learning can generate biases and predictability is not new (for example, Lewis, 1989; Lewellen and Shanken, 2002). Note that biases and predictability can be found in long samples, but they can *not* be exploited in real time by Bayesian decision makers because from their point of view they incorporate information optimally.¹⁵ In the presence of full information, the average forecast error is essentially zero. With full information, the regime is known and so within-regime shocks are the main source of variability at a one-quarter-ahead frequency. Since these are mean zero, the planner's expectations about future growth are essentially unbiased.

By contrast, with learning about regimes, there is a major source of uncertainty in that the regime is unknown, which can create biased forecasts. For example, if the economy is truly in a low TFP and low IST growth regime, but the planner places some probability mass of the subjective belief distribution on being in one of the high growth regimes, then the mean expectation will tend to be higher than realized future growth values; this result means that forecast errors in this situation will tend to be negative. Further, upon realizing a draw from the low TFP and low IST growth regime, the planner will only slightly update the subjective belief distribution, so when the subsequent period also has low growth draws, the forecast error

¹⁵In Appendix D we show an example highlighting how biases and autocorrelation in forecast errors can arise.

will again tend to be negative. Likewise, when the economy is in the high growth regimes, the planner will tend to underpredict output growth and the forecast error will tend to be positive. We see in Table 3 that these patterns broadly hold in our model.

5.4 Welfare Cost of Information

Our last set of results examine how much worse off the planner is with imperfect information about trends in productivity. We can use our perturbation solution method to address this question of welfare, since the third-order approximation that we use captures higher-order effects in the decision rules for consumption and labor. To compute welfare, we consider the recursive formulation of the planner's preferences in equation (4):

$$W_t = \log c_t + \xi \log(1 - l_t) + \beta \tilde{E}_t W_{t+1}. \quad (44)$$

For each of the different information environments, the planner is asked to assess the percentage of steady state consumption that would make them indifferent between the steady state and stochastic economy with regime switches and shocks. Since we have multiple regimes, we consider a conditional welfare measure that makes this comparison before the regime is drawn, and with capital at its steady state value. That is, if W_0 denotes the expected welfare at this initial state, then the welfare cost of fluctuations is Υ , given by:

$$W_0 = \sum_{t=0}^{\infty} \beta^t [\log [c_{ss} (1 + \Upsilon)] + \xi \log (1 - l_{ss})]. \quad (45)$$

The fact that the steady state is identical across information assumptions makes this a comparable benchmark. Note that $\Upsilon < 0$ implies that the agents would be willing to give up consumption to move to a deterministic economy.

Panel B of Table 3 shows the measure of welfare across information sets. The full information case shows a relatively small cost of fluctuations, with the planner being willing to give up about 1% of steady state consumption. The small cost under full information is in line with the celebrated results of Lucas (1987) and Lucas (2003). While we are analyzing a different type of lower-frequency shifts in technology than standard business cycle analysis, this result mimics the finding that business cycles have low welfare cost in representative agent models. In the learning case, the welfare cost of fluctuations is more than three times as large as under full information, at around 3.2% of steady state consumption. The fact that the planner cannot distinguish the

sources of TFP and IST growth are costly from a welfare perspective. Our analysis is in line with Barillas et al. (2009), who also highlight the large cost of imperfect information, albeit in a non-Bayesian context.

It is important to remember here that our agents are not only fully rational, but the amount of uncertainty they have about the trends in productivity are informed actual estimates coming from US data. So agents confronted with a *realistic* amount of uncertainty are much worse off than agents who are fully aware of movements in trends, even though they still face uncertainty.

6 Conclusion

In this paper we have studied the economic implications of learning about technology growth regimes. Data from the US shows distinct regime-switching behavior for both total factor productivity and investment specific technology, and that incoming data is usually insufficient for perfectly identifying the regime. We developed a general perturbation-based solution algorithm that allows for accurate second- or higher-order approximations. Using this solution method, our analysis of a RBC model with learning about TFP and IST regimes shows that learning alters the response to regime changes and intra-regime shocks, generates forecast error bias, and increases the welfare cost of fluctuations.

A key advantage of our framework is that, while identifying key implications of the filtering of incoming technology growth data, we can also easily build upon and extend the model to consider additional channels. For example, with our perturbation-based methodology we can solve larger models that have nominal and real rigidities in order to study the implications for inflation dynamics, wages, and monetary or fiscal policy. Further, the ability to consider second- or higher-order approximations implies that studying asset price implications of learning about regimes is possible. We leave these issues for future work.

References

- Alpanda, S. (2019). Regime-Switching Productivity Growth and Bayesian Learning in Real Business Cycles. Forthcoming, Macroeconomic Dynamics.
- Andolfatto, D. and P. Gomme (2003). Monetary Policy Regimes and Beliefs. *International Economic Review* 44(1), 1–30.
- Barillas, F., L. P. Hansen, and T. Sargent (2009). Doubts or Variability? *Journal of Economic Theory* 144(6), 2388–2418.
- Barthlemy, J. and M. Marx (2017). Solving Endogenous Regime Switching Models. *Journal of Economic Dynamics and Control* 77(C), 1–25.
- Benigno, G., A. Foerster, C. Otrok, and A. Rebucci (2020). Estimating Macroeconomic Models of Financial Crises: An Endogenous Regime-Switching Approach. Working Paper 2020-10, Federal Reserve Bank of San Francisco.
- Bianchi, F. and L. Melosi (2016). Modeling The Evolution Of Expectations And Uncertainty In General Equilibrium. *International Economic Review* 57, 717–756.
- Binning, A. and J. Maih (2015). Sigma Point Filters for Dynamic Nonlinear Regime Switching Models. Working Paper 2015/10, Norges Bank.
- Bocola, L. and N. Gornemann (2013). Risk, economic growth and the value of U.S. corporations. Working Papers 13-10, Federal Reserve Bank of Philadelphia.
- Branch, W. A., T. Davig, and B. McGough (2013). Adaptive Learning In Regime-Switching Models. *Macroeconomic Dynamics* 17(05), 998–1022.
- Bullard, J. and A. Singh (2012). Learning And The Great Moderation. *International Economic Review* 53(2), 375–397.
- Conlisk, J. (1996). Why Bounded Rationality? *Journal of Economic Literature* 34(2), 669–700.
- Cooley, T. F. (1995). *Frontiers of Business Cycle Research*. Princeton University Press.
- Costa, O., M. Fragoso, and R. Marques (2005). *Discrete-Time Markov Jump Linear Systems*. Springer.
- Cummins, J. G. and G. L. Violante (2002). Investment-Specific Technical Change in the US (1947-2000): Measurement and Macroeconomic Consequences. *Review of Economic Dynamics* 5(2), 243–284.
- Davig, T. (2004). Regime-Switching Debt and Taxation. *Journal of Monetary Economics* 51(4), 837–859.

- Davig, T. and E. Leeper (2007). Generalizing the Taylor Principle. *American Economic Review* 97(3), 607–635.
- Edge, R. M., T. Laubach, and J. C. Williams (2007). Learning and shifts in long-run productivity growth. *Journal of Monetary Economics* 54(8), 2421 – 2438.
- Farmer, R., D. Waggoner, and T. Zha (2009). Understanding Markov-Switching Rational Expectations Models. *Journal of Economic Theory* 144(5), 1849–1867.
- Fernald, J. (2012). A Quarterly, Utilization-Adjusted Series on Total Factor Productivity. Working Paper Series 2012-19, Federal Reserve Bank of San Francisco.
- Fernandez-Villaverde, J., P. Guerron-Quintana, K. Kuester, and J. Rubio-Ramrez (2015). Fiscal volatility shocks and economic activity. *American Economic Review* 105(11), 3352–84.
- Fernandez-Villaverde, J., P. Guerron-Quintana, and J. Rubio-Ramirez (2015). Estimating Dynamic Equilibrium Models with Stochastic Volatility. *Journal of Econometrics* 185(1), 216–229.
- Foerster, A., J. Rubio-Ramirez, D. Waggoner, and T. Zha (2016). Perturbation Methods for Markov-Switching Dynamic Stochastic General Equilibrium Models. *Quantitative Economics* 7(2), 637–669.
- Gilchrist, S. and M. Saito (2006, August). Expectations, asset prices, and monetary policy: The role of learning. Working Paper 12442, National Bureau of Economic Research.
- Greenwood, J., Z. Hercowitz, and P. Krusell (1997, June). Long-Run Implications of Investment-Specific Technological Change. *American Economic Review* 87(3), 342–362.
- Gust, C. J., E. P. Herbst, and J. D. Lopez-Salido (2018, October). Forward Guidance with Bayesian Learning and Estimation. Finance and Economics Discussion Series 2018-072, Board of Governors of the Federal Reserve System (U.S.).
- Hamilton, J. (1989). A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica* 57(2), 357–384.
- Hamilton, J. (2016). Macroeconomic Regimes and Regime Shifts. In J. B. Taylor and H. Uhlig (Eds.), *Handbook of Macroeconomics*, Volume 2 of *Handbook of Macroeconomics*, pp. 163–201. Elsevier.
- Justiniano, A., G. Primiceri, and A. Tambalotti (2011). Investment Shocks and the Relative Price of Investment. *Review of Economic Dynamics* 14(1), 101–121.
- Kim, C.-J. (1994). Dynamic linear models with markov-switching. *Journal of Econometrics* 60(1), 1–22.

- Kydland, F. E. and E. C. Prescott (1982). Time to Build and Aggregate Fluctuations. *Econometrica* 50(6), 1345–1370.
- Laster, D., P. Bennett, and I. S. Geoum (1999). Rational Bias in Macroeconomic Forecasts. *The Quarterly Journal of Economics* 114(1), 293–318.
- Lewellen, J. and J. Shanken (2002). Learning, Asset-Pricing Tests, and Market Efficiency. *The Journal of Finance* 57(3), 1113–1145.
- Lewis, K. K. (1989). Changing Beliefs and Systematic Rational Forecast Errors with Evidence from Foreign Exchange. *American Economic Review* 79(4), 621–636.
- Lucas, R. E. (1987). *Models of Business Cycles*. Basil Blackwell.
- Lucas, R. E. (2003). Macroeconomic Priorities. *American Economic Review* 93(1), 1–14.
- Milani, F. (2007). Expectations, Learning and Macroeconomic Persistence. *Journal of Monetary Economics* 54(7), 2065–2082.
- Richter, A. W. and N. A. Throckmorton (2015). The Consequences of an Unknown Debt Target. *European Economic Review* 78(C), 76–96.
- Saito, M. (2006). Changes in Perceptions of the Future: Macroeconomic Implications . Working Paper Series 2006-E-11, FBank of Japan.
- Schmitt-Grohe, S. and M. Uribe (2004). Solving Dynamic General Equilibrium Models using a Second-order Approximation to the Policy Function. *Journal of Economic Dynamics and Control* 28(4), 755–775.
- Schorfheide, F. (2005). Learning and Monetary Policy Shifts. *Review of Economic Dynamics* 8(2), 392–419.

Appendix A Bayesian Updating in the RBC Model

For TFP, the planner combines the observed growth rate $\Delta \log z_t$ with prior beliefs about the regime to produce an estimated probability of being in the low-growth regime ($s_t^{\mu z} = L$) according to

$$\psi_{L,t}^{\mu z} = \frac{\sum_{j=L,H} \frac{1}{\sigma_z(j)} \phi\left(\frac{\Delta \log z_t - \mu_z(L)}{\sigma_z(j)}\right) (p_{Lj}^{\sigma z} \psi_{L,t-1}^{\sigma z} + p_{Hj}^{\sigma z} \psi_{H,t-1}^{\sigma z}) (p_{LL}^{\mu z} \psi_{L,t-1}^{\mu z} + p_{HL}^{\mu z} \psi_{H,t-1}^{\mu z})}{\sum_{i,j=L,H} \frac{1}{\sigma_z(j)} \phi\left(\frac{\Delta \log z_t - \mu_z(i)}{\sigma_z(j)}\right) (p_{Lj}^{\sigma z} \psi_{L,t-1}^{\sigma z} + p_{Hj}^{\sigma z} \psi_{H,t-1}^{\sigma z}) (p_{Li}^{\mu z} \psi_{L,t-1}^{\mu z} + p_{Hi}^{\mu z} \psi_{H,t-1}^{\mu z})}, \quad (\text{A.1})$$

and the low-volatility regime ($s_t^{\sigma z} = L$) by

$$\psi_{L,t}^{\sigma z} = \frac{\sum_{j=L,H} \frac{1}{\sigma_z(L)} \phi\left(\frac{\Delta \log z_t - \mu_z(j)}{\sigma_z(L)}\right) (p_{Lj}^{\mu z} \psi_{L,t-1}^{\mu z} + p_{Hj}^{\mu z} \psi_{H,t-1}^{\mu z}) (p_{LL}^{\sigma z} \psi_{L,t-1}^{\sigma z} + p_{HL}^{\sigma z} \psi_{H,t-1}^{\sigma z})}{\sum_{i,j=L,H} \frac{1}{\sigma_z(i)} \phi\left(\frac{\Delta \log z_t - \mu_z(j)}{\sigma_z(i)}\right) (p_{Lj}^{\mu z} \psi_{L,t-1}^{\mu z} + p_{Hj}^{\mu z} \psi_{H,t-1}^{\mu z}) (p_{Li}^{\sigma z} \psi_{L,t-1}^{\sigma z} + p_{Hi}^{\sigma z} \psi_{H,t-1}^{\sigma z})}, \quad (\text{A.2})$$

where $\phi(\cdot)$ denotes the standard normal density function, $\psi_{L,t}^{\mu z} = \Pr(s_t^{\mu z} = L|z^t)$, $\psi_{H,t}^{\mu z} = 1 - \psi_{L,t}^{\mu z}$, $\psi_{L,t}^{\sigma z} = \Pr(s_t^{\sigma z} = L|z^t)$, and $\psi_{H,t}^{\sigma z} = 1 - \psi_{L,t}^{\sigma z}$.

Similarly, for IST, the planner combines the observed growth rate $\Delta \log u_t$ with prior beliefs to produce

$$\psi_{L,t}^{\mu u} = \frac{\sum_{j=L,H} \frac{1}{\sigma_u(j)} \phi\left(\frac{\Delta \log u_t - (1-\rho_u)\mu_u(L) - \rho_u \Delta \log u_{t-1}}{\sigma_u(j)}\right) (p_{Lj}^{\sigma u} \psi_{L,t-1}^{\sigma u} + p_{Hj}^{\sigma u} \psi_{H,t-1}^{\sigma u}) (p_{LL}^{\mu u} \psi_{L,t-1}^{\mu u} + p_{HL}^{\mu u} \psi_{H,t-1}^{\mu u})}{\sum_{i,j=L,H} \frac{1}{\sigma_u(j)} \phi\left(\frac{\Delta \log u_t - (1-\rho_u)\mu_u(i) - \rho_u \Delta \log u_{t-1}}{\sigma_u(j)}\right) (p_{Lj}^{\sigma u} \psi_{L,t-1}^{\sigma u} + p_{Hj}^{\sigma u} \psi_{H,t-1}^{\sigma u}) (p_{Li}^{\mu u} \psi_{L,t-1}^{\mu u} + p_{Hi}^{\mu u} \psi_{H,t-1}^{\mu u})}, \quad (\text{A.3})$$

and

$$\psi_{L,t}^{\sigma u} = \frac{\sum_{j=L,H} \frac{1}{\sigma_u(L)} \phi\left(\frac{\Delta \log u_t - (1-\rho_u)\mu_u(j) - \rho_u \Delta \log u_{t-1}}{\sigma_u(L)}\right) (p_{Lj}^{\mu u} \psi_{L,t-1}^{\mu u} + p_{Hj}^{\mu u} \psi_{H,t-1}^{\mu u}) (p_{LL}^{\sigma u} \psi_{L,t-1}^{\sigma u} + p_{HL}^{\sigma u} \psi_{H,t-1}^{\sigma u})}{\sum_{i,j=L,H} \frac{1}{\sigma_u(i)} \phi\left(\frac{\Delta \log u_t - (1-\rho_u)\mu_u(j) - \rho_u \Delta \log u_{t-1}}{\sigma_u(i)}\right) (p_{Lj}^{\mu u} \psi_{L,t-1}^{\mu u} + p_{Hj}^{\mu u} \psi_{H,t-1}^{\mu u}) (p_{Li}^{\sigma u} \psi_{L,t-1}^{\sigma u} + p_{Hi}^{\sigma u} \psi_{H,t-1}^{\sigma u})}, \quad (\text{A.4})$$

where $\psi_{L,t}^{\mu u} = \Pr(s_t^{\mu u} = L|u^t)$, $\psi_{H,t}^{\mu u} = 1 - \psi_{L,t}^{\mu u}$, $\psi_{L,t}^{\sigma u} = \Pr(s_t^{\sigma u} = L|u^t)$, and $\psi_{H,t}^{\sigma u} = 1 - \psi_{L,t}^{\sigma u}$.

In this learning environment, the filtering for TFP and IST occur independently without inter-dependence. However, for each of the processes, the existence of changes in the mean and the volatility interact within the learning process.

Appendix B Derivations of Approximations

This Appendix shows the derivations for the solution method, and gives the proofs of Propositions 1 and 2.

First, the derivative with respect to \mathbf{x}_{t-1} evaluated at steady state is given by

$$\mathbf{F}_{\mathbf{x}_{t-1}, s_t} = \sum_{s=1}^{n_s} \sum_{s'=1}^{n_s} \frac{p_{s,s'}}{1 + \exp(-\eta_{s,ss})} \begin{bmatrix} \tilde{\mathbf{f}}_{\mathbf{y}_{t+1}}(s', s_t) (g_{\mathbf{x}_{t-1}, s'} h_{\mathbf{x}_{t-1}, s_t}^x + g_{\eta_{t-1}, s'} h_{\mathbf{x}_{t-1}, s_t}^\eta) \\ + \tilde{\mathbf{f}}_{\mathbf{y}_t}(s', s_t) g_{\mathbf{x}_{t-1}, s_t} + \tilde{\mathbf{f}}_{\mathbf{x}_t}(s', s_t) h_{\mathbf{x}_{t-1}, s_t}^x \\ + \tilde{\mathbf{f}}_{\mathbf{x}_{t-1}}(s', s_t) + \tilde{\mathbf{f}}_{\eta_t}(s', s_t) h_{\mathbf{x}_{t-1}, s_t}^\eta \end{bmatrix} = 0 \quad (\text{B.1})$$

and with respect to η_{t-1} is given by

$$\mathbf{F}_{\eta_{t-1}, s_t} = \sum_{s=1}^{n_s} \sum_{s'=1}^{n_s} \frac{p_{s,s'}}{1 + \exp(-\eta_{s,ss})} \begin{bmatrix} \tilde{\mathbf{f}}_{\mathbf{y}_{t+1}}(s', s_t) (g_{\mathbf{x}_{t-1}, s'} h_{\eta_{t-1}, s_t}^x + g_{\eta_{t-1}, s'} h_{\eta_{t-1}, s_t}^\eta) \\ + \tilde{\mathbf{f}}_{\mathbf{y}_t}(s', s_t) g_{\eta_{t-1}, s_t} + \tilde{\mathbf{f}}_{\mathbf{x}_t}(s', s_t) h_{\eta_{t-1}, s_t}^x \\ + \tilde{\mathbf{f}}_{\eta_t}(s', s_t) h_{\eta_{t-1}, s_t}^\eta + \tilde{\mathbf{f}}_{\eta_{t-1}}(s', s_t) \end{bmatrix} = 0 \quad (\text{B.2})$$

These two expressions can be concatenated together and across regimes to produce a system of the form

$$\mathbf{F}_{[\mathbf{x}_{t-1} \eta_{t-1}]} = \begin{bmatrix} \mathbf{F}_{\mathbf{x}_{t-1}, s_t=1} & \mathbf{F}_{\eta_{t-1}, s_t=1} \\ \vdots & \vdots \\ \mathbf{F}_{\mathbf{x}_{t-1}, s_t=n_s} & \mathbf{F}_{\eta_{t-1}, s_t=n_s} \end{bmatrix} = 0, \quad (\text{B.3})$$

which is a quadratic form in the unknowns $\left\{ g_{\mathbf{x}_{t-1}, s_t}, g_{\eta_{t-1}, s_t}, h_{\mathbf{x}_{t-1}, s_t}^x, h_{\eta_{t-1}, s_t}^x, h_{\mathbf{x}_{t-1}, s_t}^\eta, h_{\eta_{t-1}, s_t}^\eta \right\}_{s_t=1}^{n_s}$. Further, the expression is of the form shown by Foerster et al. (2016) to be a general quadratic form that is solvable by using Gröbner bases rather than the standard Eigenvalue problem found in constant parameter models. Gröbner bases have the advantage that they will find all possible solutions to the quadratic form (B.3), each of which can be checked for stability. The concept of mean square stability (MSS), defined in Costa et al. (2005), and advocated by Farmer et al. (2009) and Foerster et al. (2016) allows for unbounded paths provided that first and second moments of the solutions are finite. In the context of learning, MSS has the advantage over an alternative concept of bounded stability—which requires all possible solutions to be bounded—since it allows for $\eta_{i,t} \rightarrow \pm\infty$ and hence the subjective probabilities to become arbitrarily close to either 0 or 1. In other words, it includes possible realizations of shocks and regimes such that the current regime is learned with near-perfect precision.

To show Proposition 1, focus on the first $n_x + n_y$ equations of (B.2), which from equation

(31) in the main text, are the equilibrium conditions denoted by \mathbf{f} :

$$\mathbf{F}_{\eta_{t-1}, s_t}^{[1:n_x+n_y]} = \sum_{s=1}^{n_s} \sum_{s'=1}^{n_s} \frac{p_{s,s'}}{1 + \exp(-\eta_{s,ss})} \begin{bmatrix} \mathbf{f}_{\mathbf{y}_{t+1}}(s', s_t) \left(g_{\mathbf{x}_{t-1}, s'} h_{\eta_{t-1}, s_t}^x + g_{\eta_{t-1}, s'} h_{\eta_{t-1}, s_t}^\eta \right) \\ + \mathbf{f}_{\mathbf{y}_t}(s', s_t) g_{\eta_{t-1}, s_t} + \mathbf{f}_{\mathbf{x}_t}(s', s_t) h_{\eta_{t-1}, s_t}^x \end{bmatrix} = 0. \quad (\text{B.4})$$

Note that this expression, collected for s_t , is a linear homogeneous system for the unknowns $\{g_{\eta_{t-1}, s_t}, h_{\eta_{t-1}, s_t}^x\}_{s_t=1}^{n_s}$, and hence its solution is zero as shown in the Proposition.

After solving equation (B.3), the other terms that make up the first order approximations can be solved via standard linear methods. First, the derivative with respect to ε_t is given by

$$\mathbf{F}_{\varepsilon_t, s_t} = \sum_{s=1}^{n_s} \sum_{s'=1}^{n_s} \frac{p_{s,s'}}{1 + \exp(-\eta_{s,ss})} \begin{bmatrix} \tilde{\mathbf{f}}_{\mathbf{y}_{t+1}}(s', s_t) (g_{\mathbf{x}_{t-1}, s'} h_{\varepsilon_t, s_t}^x + g_{\eta_{t-1}, s'} h_{\varepsilon_t, s_t}^\eta) \\ + \tilde{\mathbf{f}}_{\mathbf{y}_t}(s', s_t) g_{\varepsilon_t, s_t} + \tilde{\mathbf{f}}_{\mathbf{x}_t}(s', s_t) h_{\varepsilon_t, s_t}^x \\ + \tilde{\mathbf{f}}_{\eta_t}(s', s_t) h_{\varepsilon_t, s_t}^\eta + \tilde{\mathbf{f}}_{\varepsilon_t}(s', s_t) \end{bmatrix} \quad (\text{B.5})$$

which can be concatenated to generate

$$\mathbf{F}_{\varepsilon_t} = \begin{bmatrix} \mathbf{F}_{\varepsilon_t, s_t=1} \\ \vdots \\ \mathbf{F}_{\varepsilon_t, s_t=n_s} \end{bmatrix} = 0, \quad (\text{B.6})$$

which is a linear system in the unknowns $\{g_{\varepsilon_t, s_t}, h_{\varepsilon_t, s_t}^x, h_{\varepsilon_t, s_t}^\eta\}_{s_t=1}^{n_s}$.

Lastly the derivative with respect to χ is given by

$$\mathbf{F}_{\chi, s_t} = \sum_{s=1}^{n_s} \sum_{s'=1}^{n_s} \frac{p_{s,s'}}{1 + \exp(-\eta_{s,ss})} \begin{bmatrix} \tilde{\mathbf{f}}_{\mathbf{y}_{t+1}}(s', s_t) (g_{\mathbf{x}_{t-1}, s'} h_{\chi, s_t}^x + g_{\eta_{t-1}, s'} h_{\chi, s_t}^\eta + g_{\chi, s'}) \\ + \tilde{\mathbf{f}}_{\mathbf{y}_t}(s', s_t) g_{\chi, s_t} + \tilde{\mathbf{f}}_{\mathbf{x}_t}(s', s_t) h_{\chi, s_t}^x + \tilde{\mathbf{f}}_{\eta_t}(s', s_t) h_{\chi, s_t}^\eta \\ \tilde{\mathbf{f}}_{\theta_{1,t+1}}(\theta_1(s') - \bar{\theta}_1) + \tilde{\mathbf{f}}_{\theta_{1,t}}(\theta_1(s_t) - \bar{\theta}_1) \end{bmatrix} \quad (\text{B.7})$$

which produces a system given by

$$\mathbf{F}_{\chi} = \begin{bmatrix} \mathbf{F}_{\chi, s_t=1} \\ \vdots \\ \mathbf{F}_{\chi, s_t=n_s} \end{bmatrix} = 0, \quad (\text{B.8})$$

which is linear in the unknowns $\{g_{\chi, s_t}, h_{\chi, s_t}^x, h_{\chi, s_t}^\eta\}_{s_t=1}^{n_s}$.

These expressions thus can be used to solve for the coefficients of the first-order expansion. We can take derivatives of $(\mathbf{x}_{t-1}, \eta_{t-1}, \varepsilon_t, \chi)$ multiple times to solve for second- or higher-order

approximations, which are simply progressively larger linear systems.

To show Proposition 2 we first note that the hypothesis of the Proposition implies that $h_{x,s}^\eta = 0$ since the evolution of beliefs will be independent of predetermined variables. We use this fact, plus the result from Proposition 1, while taking derivatives $\mathbb{F}_{x\eta}^{[1:n_x+n_y]}$ to produce

$$\begin{aligned} \sum_{s=1}^{n_s} \sum_{s'=1}^{n_s} p_{s,s'} \{ & \frac{\exp(-\eta_{s,ss})}{1 + \exp(\eta_{s,ss})} \eta_{\eta,s} [\mathbf{f}_{x_{t-1}}(s', s_t) + \mathbf{f}_{y_t}(s', s_t) g_{\mathbf{x},s} + (\mathbf{f}_{\mathbf{x}_t}(s', s_t) + \mathbf{f}_{\mathbf{y}_{t+1}} g_{\mathbf{x},s'}) h_{\mathbf{x},s}] \\ & + [(\mathbf{f}_{\mathbf{x}_t}(s', s_t) + \mathbf{f}_{\mathbf{y}_{t+1}}(s', s_t) g_{\mathbf{x},s'}) h_{\mathbf{x}\eta,s} + \mathbf{f}_{\mathbf{y}\eta}(s', s_t) g_{\mathbf{x}\eta,s} + \mathbf{f}_{\mathbf{y}_{t+1}}(s', s_t) h_{\eta,s}^\eta h_{\mathbf{x},s}^x g_{\mathbf{x}\eta,s}] \} = 0 \end{aligned}$$

Note that the bracketed term in the first line equals zero from equation (B.1). This implies that the above equation is a linear homogeneous system for $\{h_{\mathbf{x}\eta,s}, g_{\mathbf{x}\eta,s}\}_{s=1}^{n_s}$, and hence its solution is zero as shown in the Proposition.

Appendix C Accuracy in a Simplified Model

To illustrate the impact of approximating the subjective probabilities, we consider a simplified version of the RBC model that can facilitate accuracy comparisons. Specifically, we eliminate IST growth from the model and assume an inelastic labor supply. Since we have eliminated IST as a driver of long-run growth, we also eliminate the exponent $1 - \alpha$ on TFP in the production function and write the production function as $y_t = z_t k_{t-1}^\alpha l_t^{1-\alpha}$. Table C.1 shows the parameters. For the growth processes, we calibrate the means and standard deviations as in the estimates presented in Table 1, with the caveat that we've now assumed synchronized regimes for simplicity. Likewise, broadly in line with the estimated values, we set the diagonal of the transition matrix to 0.98. Conditional on these growth rates, the discount factor implies a steady state real interest rate of 3%.

For our accuracy checks, we focus on two measures of accuracy and compare them across alternative solution methods. In particular, we solve the model to first, second, and third orders, first using the Partition Principle where we perturb both μ and σ in the economic equations and the learning set of equations, and second using the refinement where we only perturb μ in the economic equations. Finally, we solve via nonlinear policy function iteration, as in Davig (2004). In each case, the accuracy checks are based on a simulation of 10,000 periods, with the first 1,000 of that simulation discarded. Table C.2 shows these accuracy measures across models.

Our first measure of accuracy is the mean square error of beliefs relative to the truth. In our

Table C.1: Calibration of Simple RBC Model

Parameter	Value
Discount Factor	$\beta = 0.9983$
Capital Share	$\alpha = 0.33$
Depreciation Rate	$\delta = 0.025$
Regime Probability	$P_{j,j} = 0.98$
Low Growth Rate (ann, pct)	$\mu(1) = 0.6892$
High Growth Rate (ann, pct)	$\mu(2) = 1.9221$
Low Volatility Rate (ann, pct)	$\sigma(1) = 2.5614$
High Volatility Rate (ann, pct)	$\sigma(2) = 3.9002$

case, this measure is given by

$$MSE^{order} = T^{-1} \sum_{t=1}^T \left(\hat{\psi}_{1,t}^{order} - \psi_{1,t} \right)^2 \quad (C.1)$$

where $\hat{\psi}_{1,t}^{order}$ is the approximated beliefs for $order \in \{1, 2, 3\}$, and $\psi_{1,t}$ is the true beliefs based upon the fully nonlinear updating. The results show that, when we perturb the parameters in the learning equations according to the Partition Principle, accuracy is low and does not improve with the order of approximation. With the Refinement, accuracy is low at first order, but improves markedly at second-order; moving to third-order shows no improvement in this specific example. Policy function iteration, by using the fully nonlinear updating of beliefs, achieves exact accuracy by definition.

Figure C.1 shows some intuition as to why the refinement produces much higher accuracy than the Partition Principle. In steady state, the TFP process repeatedly generates signals $z_{ss} = \bar{\mu}$. Under the true distributions that are left intact via the refinement, the likelihood of generating a signal of z_{ss} is much higher in regime 2 (red line) than in regime 1 (blue line). The difference between these two probabilities affects the likelihood component of the updating of beliefs. As a result, combining the likelihood with the symmetric transition probabilities generates a steady state vector of beliefs of $\psi_{ss} = [0.0393, 0.9607]'$. On the other hand, the Partition Principle, by perturbing both μ and σ , effectively collapses the signal distributions at steady state to an identical distribution across regimes (the black line). This collapsing removes any relative differences between the regimes, making the steady state beliefs solely a function of the transition matrix, namely the ergodic distribution across regimes, $\psi_{ss} = [0.5, 0.5]'$ in this example. By preserving information about the relative likelihoods at the steady state value

Table C.2: Accuracy Check of Simple RBC Model

Method	MSE of Beliefs	Euler Eqn Error
Partition Principle		
First Order	0.4033	-3.0253
Second Order	0.3811	-2.4216
Third Order	0.3794	-2.2628
Refinement		
First Order	0.3224	-4.3419
Second Order	0.0505	-3.5078
Third Order	0.0505	-3.6729
Policy Function Iteration	0	-4.2904

Note: Euler equation errors in base-10 log points; -4 implies \$1 error for \$10,000 in consumption.

of the signal, the refinement incorporates more information into the steady state beliefs, and produces a higher degree of accuracy of beliefs, especially at higher orders.


Our second measure of accuracy is Euler equation errors, which measure the accuracy of the economic decision-making. We compute the average Euler equation error in our simulation by computing

$$EEE_t = 1 - \beta \tilde{\mathbb{E}}_t \left[\exp(\tilde{z}_{t+1})^{\frac{1}{\alpha-1}} \frac{\tilde{c}_t}{\tilde{c}_{t+1}} \left(\alpha \exp(\tilde{z}_{t+1}) \tilde{k}_t^{\alpha-1} + 1 - \delta \right) \right],$$

where we numerically approximate the expectation using 10,000 draws of the shocks. Table C.2 shows that the refinement improves accuracy relative to the Partition Principle, and performs not that much worse than policy function iteration. Note that, in contrast to full information Euler equation errors where probability measures for the computation of expectations are identical across solution methods, our method generates different expectations jointly with different policy functions for capital and consumption. In other words, these errors must be viewed in conjunction with the accuracy of beliefs as measured by the MSE. For example, simply examining the Euler equation errors would suggest that the first-order approximation using the refinement was the most accurate; however, this high level of accuracy of the consumption and capital policy functions is conditioned on inaccurate beliefs. We conclude that the refinement with second- or third-order approximations does the best in terms of accurately capturing both beliefs and economic decisions.

The trade-off between solution time and accuracy is important in our context. While policy function iteration is a global solution that exactly matches the updating of beliefs and has

Figure C.1: Distributions of TFP Growth



a high degree of accuracy as measured by Euler equation errors, it is quite costly in terms of computational time. In this policy function iteration example, we solved for the decision rules on a relatively sparse grid of 25,500 points distributed across the state space of capital, beliefs, and the technology shock. Even in this case, a single iteration of updating the policy function approximation across all grid points took 93 seconds in Matlab using 16 processors in parallel. In contrast, our perturbation solution solved in around a second, making it possible to do likelihood-based estimation. In addition, policy function iteration suffers from the curse of dimensionality, so more involved learning processes, such as our case with TFP and IST growth, and two regimes each for the mean and standard deviation, would become increasingly difficult, while perturbation scales much more easily.

Appendix D An Illustration of the Sources of Bias in Rational Forecast Errors

In this section we highlight how a rational (Bayesian) agent can have forecast errors that, to an econometrician observing them ex-post, appear biased. Consider the following simple example:

$$x_t = \mu(s_t) + \varepsilon_t \quad (\text{D.1})$$

where, as in our RBC model, $\mu(s_t)$ follows a two state Markov chain indexed by $s_t = 1$ or $s_t = 2$ with realizations μ_1 and μ_2 , and $\varepsilon_t \sim_{iid} N(0, 1)$. We now consider three versions of this model.

First, consider a fixed coefficient model where $\mu(s_t) = \mu$ for all s_t . In that case we have $\mathbb{E}_t x_{t+1} = \mu$ and the one-step ahead forecast error $x_{t+1} - \mathbb{E}_t x_{t+1} = \varepsilon_{t+1}$ is clearly iid with mean zero.

Second, what happens if there is Markov switching, but the Markov state s_t is known to the agent at time t when forming the forecast? In that case we have $\mathbb{E}_t x_{t+1} = \varepsilon_t + \omega_{1,t}\mu_1 + \omega_{2,t}\mu_2$ where $\omega_{i,t}$ is the probability of going to state i in period $t + 1$ conditional on the Markov state at time t , hence the time index. While this means that there is time dependence in the forecast, one can show (Hamilton, 2016) that the forecast error due to forecasting the Markov state is a martingale difference sequence, independent of the forecast error due to ε . Hence we do not have any biases in the forecast error in this non-linear, but full information case.

Finally, consider the case where the Markov state at time t is not observed, but instead inferred via Bayes' law. In that case the forecast is $\mathbb{E}_t x_{t+1} = \bar{\omega}_{1,t}\mu_1 + \bar{\omega}_{2,t}\mu_2$ where the weights $\bar{\omega}_{i,t}$ are now functions both of the transition matrix between states (as before), but also of the estimated state probabilities at time t . The forecast error in this case can be decomposed as

$$x_{t+1} - \mathbb{E}_t x_{t+1} = \underbrace{\varepsilon_{t+1} + \mu_{s_{t+1}} - \omega_{1,t} - \mu_1 \omega_{2,t} \mu_2}_{\text{full information forecast error}} - \underbrace{(\bar{\omega}_{1,t} - \omega_{1,t})\mu_1 + (\bar{\omega}_{2,t} - \omega_{2,t})\mu_2}_{\text{part due to learning}} \quad (\text{D.2})$$

The forecast error now contains an extra part due to learning that will generally not be mean 0 because of the dependence on Bayes' law. Note that because we model recurrent regime shifts this effect does not disappear asymptotically, in contrast to scenarios where agents learn about fixed parameters.