Macroeconomic Dynamics, **25**, 2021, 462–488. Printed in the United States of America.
doi:10.1017/S1365100519000257

REGIME-SWITCHING PRODUCTIVITY GROWTH AND BAYESIAN LEARNING IN REAL BUSINESS CYCLES

SAMI ALPANDA
University of Central Florida

Growth in total factor productivity (TFP) in the USA has slowed down significantly since the mid-2000s, reminiscent of the productivity slowdown of the 1970s. This paper investigates the implications of a productivity slowdown on macroeconomic variables using a standard real business cycle (RBC) model, extended with regime-switching in trend productivity growth and Bayesian learning regarding the growth regime. I estimate the Markov-switching parameters using US data and maximum-likelihood methods, and compute the model solution using global projection methods. Simulations reveal that, while adding a regime-switching component to the standard RBC setup increases the volatility in the system, further incorporating incomplete information and learning significantly dampens this effect. The dampening is mainly due to the responses of investment and labor in response to a switch in the trend component of TFP growth, which are weaker in the incomplete information case as agents mistakenly place some probability that the observed decline in TFP growth is due to the transient component and not due to a regime switch. The model offers an objective way to infer slowdowns in trend productivity, and suggests that macroeconomic aggregates in the USA are currently close to their potential levels given observed productivity, while counterfactual simulations indicate that the cost of the productivity slowdown to US welfare has been significant.

Keywords: Productivity Slowdown, Regime-Switching TFP Growth, Real Business Cycles, Bayesian Learning

1. INTRODUCTION

Growth in total factor productivity (TFP) in the USA has slowed down significantly since the mid-2000s, reminiscent of the productivity slowdown of the 1970s. Figure 1 shows the level of TFP in the post-war period, as well as the average growth rate of TFP during the subperiods of 1947–1972, 1973–1994, 1995–2004, and 2005–2017.[1] TFP growth averaged around 2% per year before 1973, while it was reduced to around 0.5% per year following the oil crisis

I thank the editor William A. Barnett, two anonymous referees, Ben Bridgeman, Gino Cateau, Andrew Foerster, Ellen McGrattan, Adrian Peralta-Alva, Edward Prescott, Tatevik Sekhposyan, Alexander Ueberfeldt, and Geoffrey Woglom for suggestions and comments. All remaining errors are my own. Address correspondence to: University of Central Florida, Department of Economics, College of Business Administration, 4336 Scorpius Street, Orlando, FL 32816, USA. Phone: (407) 823-1575. e-mail: sami.alpanda@ucf.edu.

 1469-8056/19

of 1973. A similar pattern was repeated after 1995, with a near 2% TFP growth for about a decade after the mid-1990s followed by the current productivity slowdown since the mid-2000s, where TFP growth has declined to an even slightly lower level than the 0.5% per year observed in the 1970s. The causes behind these productivity slowdowns are likely varied. Gordon (2012, 2016) argues that productivity growth inevitably slows following a period of rapid growth as new innovations falter and spin-off inventions, which in most cases can happen only once, run their course. Productivity slowdowns are also likely related to the recessions or crises events encountered in the beginning or near these episodes, such as the oil crisis of 1973–74 [Alpanda and Peralta-Alva (2010)]. Regardless of their causes, slowdowns in productivity have important implications not only for long-term growth and welfare, but also for the medium-term outlook for the output gap and the appropriate stance of policy since there is significant uncertainty regarding the trend growth rate TFP and, therefore, the potential level of output in real time.

In this paper, I investigate the implications of a productivity slowdown on macroeconomic variables using a standard real business cycle (RBC) model, extended with regime-switching in trend productivity growth and Bayesian learning regarding the growth regime. TFP growth is modeled as the sum of two unobserved components: a trend component which switches between a high and a low growth regime, and a transient component captured by an autoregressive process.[2] Agents in the model observe only the total of these two components, and use the *Hamilton filter* [Hamilton (1989)] to update their beliefs regarding the likelihood of being in either regime before making decisions. This type of long-run uncertainty regarding trend productivity growth is consequential in macroeconomic models since agents are forward-looking. In particular, changes in TFP growth that are expected to be more persistent and affect the trend level of TFP lead to larger responses in consumption and investment, relative to TFP growth shocks that are deemed more transient. If uncertainty regarding the long-run trend accompanies the observed TFP shocks, the responses of agents would depend on how they extract information about the trend given their observations [Kuang and Mitra (2015) and Tortorice (2016)].

I estimate the TFP-related parameters of the model using maximum-likelihood methods as in Hamilton (1989) and TFP data for the USA from Fernald (2012), and show that the smoothed estimates for the posterior probability of being in the low-growth regime are largely consistent with the turning points for productivity growth in the data as indicated in Figure 1. The solution of the theoretical model is computed using the envelope condition method (ECM) of Maliar and Maliar (2013), which is a global projection method intended to better capture the nonlinearities in the model. Simulations reveal that the baseline model with regime-switching TFP growth and Bayesian learning is significantly less volatile than the full information version of the model. The dampening is mainly due to the responses of investment and labor in response to a switch in the trend component of TFP growth. The resulting responses are weaker in the incomplete information

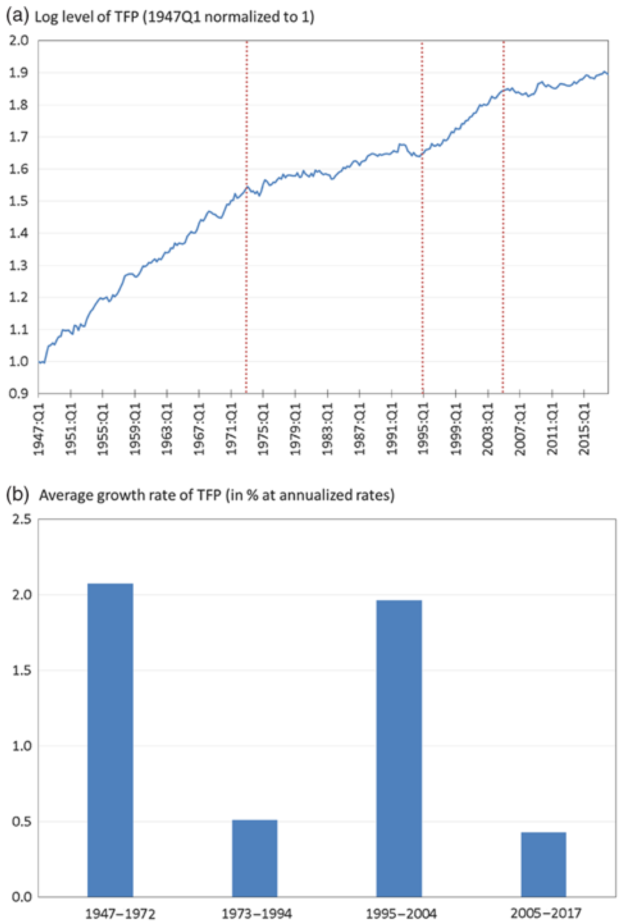

FIGURE 1. Level and growth of TFP in the USA [Source: Fernald (2012)].

case, since agents mistakenly place some probability that the observed low TFP growth is due to the transient component, and not due to a switch in the growth regime. This then necessitates a smaller decline in the steady-state capital stock, thus prompting agents not to reduce their investment and labor supply as much as in the full information case.

The business cycle moments implied by the model are similar to those from the data in many dimensions, but the model does generate slightly lower volatility and significantly lower autocorrelation overall, indicating the possible role of other shocks and frictions in generating and propagating business cycles. Results from different variants of the model imply that, while adding a regime-switching component to the standard RBC setup increases the volatility in the system, further incorporating incomplete information and learning significantly

https://doi.org/10.1017/S1365100519000257 Published online by Cambridge University Press

dampens this effect. Model simulations also indicate that macroeconomic aggregates in the USA are currently close to their potential levels, given the observed productivity series. Furthermore, counterfactual simulations indicate that the cost of the recent productivity slowdown to US output and welfare has been potentially quite large, with the model implying a near 20 pp impact on the level of consumption due to the productivity slowdown between 2005 and 2018.

1.1. Related Literature

In the time-series literature, one of the most common ways of identifying trend information is through the use of the Hamilton filter. Hamilton (1989) represents business cycles as sudden movements in output growth between an expansionary and a contractionary regime, where the growth regime switches between expansions and contractions following a Markov chain. Although the regime is not directly observable due to noise in the observed growth rate, Hamilton (1989) devises a Bayesian filter to assess the conditional probability of the regime given the observed growth. The model has been very successful in replicating the NBER (National Bureau of Economic Research) business cycle dates, often identifying turning points earlier than the NBER's *Business Cycle Committee*, and has supplied economists with a transparent and reproducible methodology to date business cycles [Burns and Mitchell (1946) and Chauvet and Piger (2003)]. The Hamilton filter has also been used for identifying trends in productivity. For instance, French (2001) argues that the Markov-switching model is better suited to capture the infrequent changes in average TFP growth, and identifying trend TFP growth using a variant of the Hamilton filter has several advantages over alternative approaches using the Kalman filter, the Hodrick–Prescott filter, or the bandpass filter.[3]

Incorporating learning about the level or the growth rate of TFP in business cycle models is not new, and dates back at least to the original RBC model of Kydland and Prescott (1982). There, the *level* of observed TFP is modeled as the sum of three unobserved components, namely, a highly persistent autoregressive (AR(1)) component, an *i.i.d*. component, and an *i.i.d*. measurement error, while agents solve a signal extraction problem using the Kalman filter when making decisions.[4] Similarly, Pakko (2002) features an RBC model with shocks to both the *growth* and the *level* of TFP (as well as to the growth and the level of investment-specific technological change). Agents in the model observe only the sum of the growth and level shocks, and infer the trend level of TFP using the Kalman filter. More recently, Lorenzoni (2009) uses a similar framework to show that noise with respect to aggregate TFP shocks can act as demand shocks in a simple New Keynesian setting. Similarly, Boz et al. (2011) use a similar framework with permanent and transitory components in TFP, along with learning through the Kalman filter, in the context of a small open economy model. They show that this model can better account for the business cycle features of emerging market economies, including the higher variability of consumption relative to output, and the countercyclicality of the trade balance.

Note that all the aforementioned models capture the regime component of TFP level or growth as a smoothly transitioning AR(1) process, while I capture changes in trend TFP growth with discrete changes based on a Markov-switching process. This is largely motivated by the discussion in Gordon (2012) which documents bunching in innovations over US history, as well as the persistent nature of the changes in trend growth in the post-war period as depicted in Figure 1. Regime switching in TFP growth has also been recently adopted within the framework with which the Federal Reserve Bank of St. Louis characterizes its macroeconomic and monetary policy outlook for the USA, with a low TFP growth assumption featuring in the current baseline forecasts, while a switch to a higher productivity growth regime seen as an upside risk to this forecast [Bullard (2016)].[5]

There is also a related literature featuring regime switching in dynamic stochastic general equilibrium (DSGE) models, where one or more parameters of the model (e.g., parameters in the monetary or fiscal policy rules) are subject to discrete changes over time following a Markov chain. Farmer et al. (2011), Foerster et al. (2014), and Maih (2015) discuss perturbation methods to compute the solution of these types of models, whereby decision rules are obtained using a first- or higher-order approximation, typically around different steady states implied by each of the different parameter values.[6] Tortorice (2016) uses similar techniques to study the role of uncertainty regarding the long-run trend of TFP and finds that long-run uncertainty amplifies the effects of shocks. Similarly, Bianchi and Melosi (2016) use perturbation techniques to study the effects of learning in Markov-switching DSGE environments, providing an example with unobserved components in TFP growth. In contrast to the above papers, the computation procedure utilized here is nonlinear and global, taking into account the nonlinearities inherent in the model.

The next section introduces the model. Section 3 describes the parameterization and computation of the model, Section 4 presents the results, and Section 5 concludes.

2. THE MODEL

The model extends the standard RBC model with regime-switching TFP growth and Bayesian learning.

2.1. Households

The economy is populated by a unit measure of infinitely lived households, whose p. over consumption, c_t, and labor, n_t, are described by the following expected utility function:

$$E_t \sum_{\tau=t}^{\infty} \beta^{\tau-t} \left\{ \frac{1}{1-\sigma} \left[c_\tau \exp\left(-\frac{n_\tau^{1+\varphi}}{1+\varphi} \right) \right]^{1-\sigma} \right\}, \tag{1}$$

https://doi.org/10.1017/S1365100519000257 Published online by Cambridge University Press

where t indexes time, $\beta < 1$ is the time-discount factor, φ is the inverse of the Frisch elasticity of labor supply, and σ is the coefficient of relative risk aversion (and the inverse of the intertemporal elasticity of substitution). When $\sigma = 1$, the utility function reduces to the more common, separable form as

$$E_t \sum_{\tau=t}^{\infty} \beta^{\tau-t} \left(\log c_\tau - \frac{n_\tau^{1+\varphi}}{1+\varphi} \right). \tag{2}$$

The households' period budget constraint is given by

$$c_t + i_t = w_t n_t + r_t k_{t-1}, \tag{3}$$

where i_t and k_t denote investment and capital, respectively, w_t is the wage rate, and r_t is the rental rate of capital. The law of motion of capital accumulation is described by

$$k_t = (1-\delta)\, k_{t-1} + i_t, \tag{4}$$

where δ is the depreciation rate of capital.

The households' objective is to maximize utility subject to their budget constraint. Their optimality conditions with respect to consumption, labor supply and capital accumulation are standard and are, respectively, given by

$$c_t^{-\sigma} \exp\left(\frac{\sigma-1}{1+\varphi} n_t^{1+\varphi} \right) = \lambda_t, \tag{5}$$

$$c_t n_t^{\varphi} = w_t, \tag{6}$$

$$\lambda_t = \beta E_t \left[\lambda_{t+1} \left(r_{t+1} + 1 - \delta \right) \right], \tag{7}$$

where λ_t is the Lagrange multiplier on the budget constraint.

2.2. Production

There is a representative firm whose technology is described by the following production function:

$$y_t = e^{(1-\alpha)a_t} k_{t-1}^{\alpha} n_t^{1-\alpha}, \tag{8}$$

where y_t denotes output, a_t captures the level of *log* TFP, and α is the share of capital in overall production. The output good can be used for consumption or investment:

$$c_t + i_t = y_t. \tag{9}$$

In equilibrium, each input is paid its marginal product, and the firm earns zero profits; hence

$$w_t = (1-\alpha) \frac{y_t}{n_t} \text{ and } r_t = \alpha \frac{y_t}{k_{t-1}}. \tag{10}$$

https://doi.org/10.1017/S1365100519000257 Published online by Cambridge University Press

2.2.1. Markov-switching TFP growth with unobserved components. Let $\Delta a_t = a_t - a_{t-1}$ denote the first difference of the TFP factor. Δa_t is the sum of two *unobserved* components; a regime-switching component, $\mu_t \in \{\mu_H, \mu_L\}$, and a transient component, z_t:

$$\Delta a_t = \mu_t + z_t. \tag{11}$$

The regime component, μ_t, switches between μ_H (i.e., the high growth regime) and μ_L (i.e., the low growth regime) following a 2×2 transition probability matrix:

$$\Pi_\mu = \begin{bmatrix} \pi_{HH} & \pi_{HL} \\ \pi_{LH} & \pi_{LL} \end{bmatrix}, \tag{12}$$

where $\pi_{ij} = \Pr\left(\mu_t = \mu_j \mid \mu_{t-1} = \mu_i\right)$ for $i, j \in \{H, L\}$, and each row of Π_μ adds up to 1. The transient component, z_t, follows an AR(1) process:

$$z_t = \rho z_{t-1} + \varepsilon_t, \tag{13}$$

where $0 \leq \rho < 1$ is the persistence parameter, and ε_t denotes normally distributed shock innovations with mean 0 and variance σ_ε^2.

2.2.2. Agents' information set and timing. Let p_{t-1} denote the conditional probability that agents assign to being in the high regime at the end of period $t-1$, given all the information up to and including that period, I_{t-1}; hence,

$$p_{t-1} = \Pr\left(\mu_{t-1} = \mu_H \mid I_{t-1}\right) \text{ and } 1 - p_{t-1} = \Pr\left(\mu_{t-1} = \mu_L \mid I_{t-1}\right). \tag{14}$$

The unobserved components of TFP growth, μ_t and z_t, are realized at the beginning of period t. Agents observe only the sum of these two components, Δa_t, and update their prior on the likelihood of the period t regime using Bayes' rule, before they make their period t decisions.

The updated probability of the high regime, $p_t = \Pr\left(\mu_t = \mu_H \mid I_t\right)$, is given by

$$p_t = \frac{p_{t-1}\pi_{HH} f\left(\varepsilon_t^{HH}\right) + \left(1 - p_{t-1}\right)\pi_{LH} f\left(\varepsilon_t^{LH}\right)}{p_{t-1}\left[\pi_{HH} f\left(\varepsilon_t^{HH}\right) + \pi_{HL} f\left(\varepsilon_t^{HL}\right)\right] + \left(1 - p_{t-1}\right)\left[\pi_{LH} f\left(\varepsilon_t^{LH}\right) + \pi_{LL} f\left(\varepsilon_t^{LL}\right)\right]}, \tag{15}$$

where f denotes the probability distribution function (pdf) of the normal distribution with mean 0 and variance σ_ε^2. ε_t^{HH}, ε_t^{LH}, ε_t^{HL}, and ε_t^{LL} denote the four possible realizations for the innovations of the AR(1) process, conditional on the observed TFP growth factors in periods $t-1$ and t; that is,

$$\varepsilon_t^{ij} = \left(\Delta a_t - \mu_j\right) - \rho\left(\Delta a_{t-1} - \mu_i\right) \text{ for } i, j \in \{H, L\}. \tag{16}$$

Along with the capital stock brought from the previous period, agents' state vector before period t decisions are taken can thus be summarized as $(k_{t-1}, \Delta a_t, p_t)$. Agents then choose consumption, c_t, labor, n_t, investment, i_t, etc.,

https://doi.org/10.1017/S1365100519000257 Published online by Cambridge University Press

TABLE 1. Structural parameters

	Symbol	Value
Discount factor	β	0.99
Inverse labor supply elasticity	ϑ	1
Risk aversion (or inverse of intertemp. elast. of subst.)	σ	1
Capital share in production	α	0.33
Depreciation rate	δ	0.025

and the capital stock carried over to the next period, k_t, is determined by the law of motion of capital. In the *full information* version of the model, agents observe μ_t and z_t separately in the beginning of period t, and thus their state vector is given by (k_{t-1}, z_t, μ_t), when they make their period t decisions.

3. PARAMETERIZATION AND COMPUTATION

In this section, I briefly discuss the parameterization of model parameters, including the estimation of TFP-related parameters using maximum-likelihood methods, as well as the computation strategy to obtain the policy functions that describe agents' optimal behavior. Further details regarding the data used in the estimation and the computation algorithm used to obtain the policy rules are provided in Appendices A and B, respectively.

3.1. Parameterization

The structural parameters are set to fairly standard values commonly used in the related literature (see Table 1). The time-discount factor of households, β, is set to 0.99, implying an average annualized real net return on capital of 6%.[7] The inverse of the Frisch elasticity of labor supply, ϑ, is set to 1, implying unit labor supply elasticity as in King et al. (1988). The inverse of the intertemporal rate of substitution, σ, is set to 1 in the baseline case following Cooley and Prescott (1995). The share parameter in the production function, α, is set to 0.33, implying the share of labor in total income is close to two thirds. The depreciation rate of capital, δ, is set to 0.025, implying an annual depreciation rate of 10%. The remaining parameters are related to the TFP processes, and are estimated using maximum-likelihood methods following Hamilton (1989), as explained in the next subsection.

3.1.1. Maximum-likelihood estimation of TFP-related parameters. Let Θ denote the set of parameters related to the TFP shocks

$$\Theta = \{\mu_L, \mu_H, \pi_{LL}, \pi_{HH}, \rho, \sigma_\varepsilon\}, \tag{17}$$

https://doi.org/10.1017/S1365100519000257 Published online by Cambridge University Press

TABLE 2. Maximum-likelihood estimates of TFP-related parameters

	Symbol	Value	Asymp. st. error
Low-regime trend growth rate	μ_L	0.0030282	0.0025606
High-regime trend growth rate	μ_H	0.0075836	0.0020136
Prob. of staying in same growth regime	$\pi_{LL} = \pi_{HH}$	0.9813080	0.0275817
Persistence of transient comp. of TFP (z)	ρ	0.0000000	0.0607141
Standard deviation of innovations to z	σ_ε	0.0120658	0.0005327

and let $L(\Theta|\Delta a)$ denote the likelihood function for Θ given a time-series for Δa of length T such that

$$\log L(\Theta|\Delta a) = \sum_{t=1}^{T} \log l(\Delta a_t, \Theta), \tag{18}$$

where

$$\begin{aligned} l(\Delta a_t, \Theta) = {} & p_{t-1}\left[\pi_{HH} f\left(\varepsilon_t^{HH}\right) + \pi_{HL} f\left(\varepsilon_t^{HL}\right)\right] \\ & + (1 - p_{t-1})\left[\pi_{LH} f\left(\varepsilon_t^{LH}\right) + \pi_{LL} f\left(\varepsilon_t^{LL}\right)\right]. \end{aligned} \tag{19}$$

The estimation procedure maximizes the likelihood function (18) over the space of all possible parameter values, given (19) and the Bayes' rule (15).

For the data series on TFP growth used in the estimation, Δa_t, I use the (utilization-adjusted) quarterly TFP series for the US Business Sector produced by John Fernald of the Federal Reserve Bank of San Francisco [Fernald (2012)], for the period between 1947Q2 and 2018Q1.[8] For better identification of parameters in the estimation, I imposed the restriction that the duration of the low and high TFP growth regimes are equal; that is, $\pi_{LL} = \pi_{HH}$. This restriction is by and large consistent with the data presented in Figure 1, which indicates roughly equal durations for the two regimes.[9] Another restriction that was imposed in the estimation was to ensure that the persistence parameter for the transient component of the TFP shock was between 0 and 1; that is, $0 \leq \rho < 1$.

The maximum-likelihood estimates of the TFP-related parameters, along with their standard errors, are given in Table 2.[10] The estimate for $\pi_{LL} = \pi_{HH}$ is near 0.98, implying an average duration of 13.4 years for the TFP growth regimes.[11] The quarterly trend growth rates under the low and high regimes, μ_L and μ_H, are estimated to be 0.3% and 0.76%, respectively, with significant uncertainty for the former figure. Given these values and equal durations for the two regimes, the implied ergodic mean of Δa_t is 0.0053, which largely conforms with the data. The estimation results also indicate that the autocorrelation in the TFP growth process is exclusively due to the regime component, and the transient component's persistent parameter, ρ, is estimated to be equal to 0.

Note that a switch from the high- to low-growth regimes generates a decline in the TFP growth factor, Δa_t, of size 0.46 percentage points (pp) per quarter,

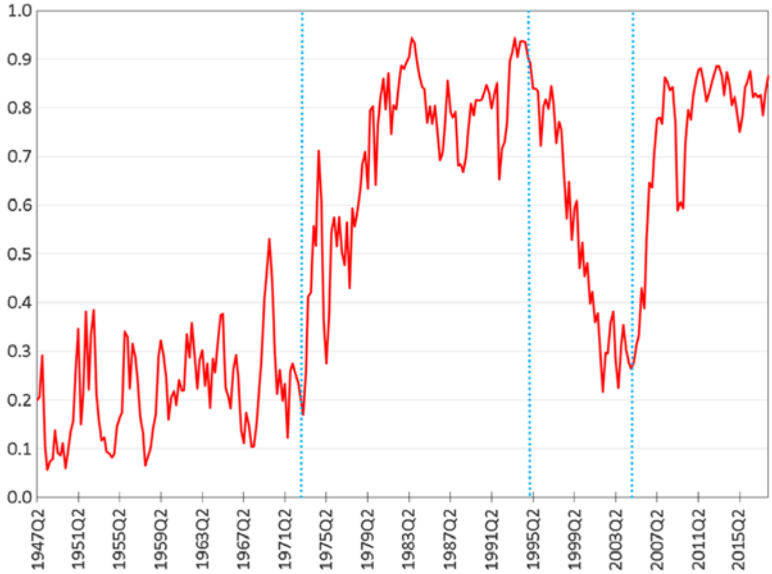

FIGURE 2. Probability that economy is in the low TFP growth regime (solid line), $1 - p_t$, along with the cutoff dates for TFP growth depicted in Figure 1.

which, at the impact period, would be comparable to a −0.38 standard deviations innovation in the transient component, ε_t. Since ε_t is normally distributed, such a magnitude lies within a highly probable band of the distribution of the innovation; hence, the agents in the model would not be able to readily tell if the shock is to the regime component or the transient component. The key distinction would occur in periods following the impact however, given the sizable difference in the durations of the trend versus the transient components of TFP growth. Agents would start increasing their beliefs regarding a regime switch only if they observe very low or negative Δa_t's for a number of consecutive periods.

3.1.2. Probability of low-growth regime in post-war USA. The Hamilton filter used above also allows us to assess the probability of being in the low TFP growth regime for the US economy during the post-war period. Figure 2 plots these estimated probabilities (i.e., $1 - p_t$), along with the cutoff dates for TFP growth regime dates depicted in Figure 1. The probability estimates line up with these cutoff dates fairly well. In particular, the model points to a significant increase in the probability of a productivity slowdown during the 1970s, followed by a decrease in this probability starting from the mid-1990s indicating a possible reversion to the high-growth regime, and then again a significant increase in the probability of a productivity slowdown since the mid-2000s. Note that the probability of a low-growth regime has remained significantly high since the 2007–2008 financial crisis, and is estimated to be around 87% at the end of the sample in

https://doi.org/10.1017/S1365100519000257 Published online by Cambridge University Press

2018Q1. This potentially points to the role of financial factors in reducing trend TFP growth through inefficient capital allocation following periods of financial disintermediation [Gourio (2013)].[12]

3.2. Computation

I first transform the model variables by defining

$$\widehat{c}_t = \frac{c_t}{e^{a_t}},\ \widehat{i}_t = \frac{i_t}{e^{a_t}},\ \widehat{y}_t = \frac{y_t}{e^{a_t}},\ \widehat{k}_t = \frac{k_t}{e^{a_t}},\ \text{and}\ \widehat{w}_t = \frac{w_t}{e^{a_t}}. \tag{20}$$

This transformation renders the model stationary in Δa_t and allows the households' problem to be written as a dynamic programming problem. I then derive the corresponding first-order and envelope conditions of this problem [see Stokey et al. (1989)], and write the other equilibrium conditions of the model (i.e., firm's labor and capital demand expressions, and the feasibility condition) consistent with this recursive notation.[13] Appendix B.1 provides a list of these expressions.

The solution of the model is computed using the ECM of Maliar and Maliar (2013), which is a global projection method intended to better capture the nonlinearities in the model.[14] ECM iterates on the value function derivatives to find the policy functions. In the context of the model presented in Section 2, this updating expression is given by

$$V_k\left(\widehat{k}, \Delta a, p\right) = \beta\,(r + 1 - \delta)\, e^{-\Delta a} E\left[V_k\left(\widehat{k}', \Delta a', p'\right)\right], \tag{21}$$

where V_k denotes the derivative of the households' value function with respect to capital. The above expression can be obtained by combining the household's first-order condition with respect to end-of-period capital, k', and the corresponding envelope condition. The ECM methodology relies on the fact that once V_k is known, all the other variables of the model can be solved easily as a function of the current state variables, $\widehat{k}$, Δa, and p. To find V_k, I approximate it by a Chebyshev polynomial, and iterate on the parameters of this polynomial using equation (21) until convergence.[15] For further details regarding the computational procedure, see Appendix B.2.

4. RESULTS

In this section, I first analyze the impulse response functions generated by the model due to a switch from the high- to low-growth regime or due to a negative shock to the transient component of TFP growth. I then compute the business cycle statistics generated by the model, and compare them to corresponding moments from the full information and the single-regime variants of the model and the data. Finally, I run counterfactuals to assess the impact of the productivity slowdown on the performance of the US economy since the financial crisis.

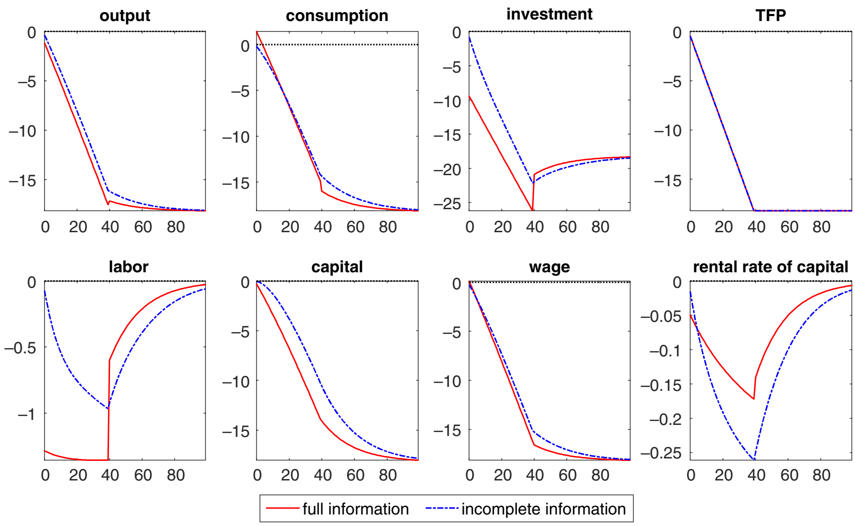

FIGURE 3. Impulse responses (in percent, relative to staying in high regime with no temporary shock) to an unexpected switch from the high to the low TFP growth regime lasting 40 periods in the baseline case with unobserved components and Bayesian learning (incomplete information) versus the full information case.

4.1. Impulse Responses

Figure 3 plots the impulse responses of model variables following an unexpected switch from the high- to the low-growth regime lasting 40 periods (i.e., 10 years), both for the baseline case with incomplete information and the alternative case with full information.[16] The regime switch creates standard recessionary dynamics, whereby output, consumption, investment, as well as the quantity of factor inputs and their returns all decline. Note that, since the shock is to the *growth* rate of TFP, it leads to a *permanent* decline in the *levels* of the nonstationary variables (such as output and consumption). The decline in TFP growth reduces the marginal product of labor and capital, reducing firms' demand for these inputs, which leads to a decline in the wage rate and the rental rate of capital. The decline in the rental rate reduces the investment activity of households, while the overall decline in their lifetime income leads to a reduction in their consumption. The labor supply of households is subject to both substitution and income effects given the utility representation in (1), whereby the decline in wages reduces labor supply, while the decline in lifetime income increases it. The net impact of these is a reduction of labor supply in the short run, while in the long run, the two effects cancel each other out. Note that, in the impact period and beyond, the effect of the regime switch on capital accumulation and labor supply, and therefore aggregate output, are significantly stronger in the full information case. With unobserved components, agents are not fully sure that the decline in TFP growth is due to the regime component, and therefore their reaction is more muted in the short run.

https://doi.org/10.1017/S1365100519000257 Published online by Cambridge University Press

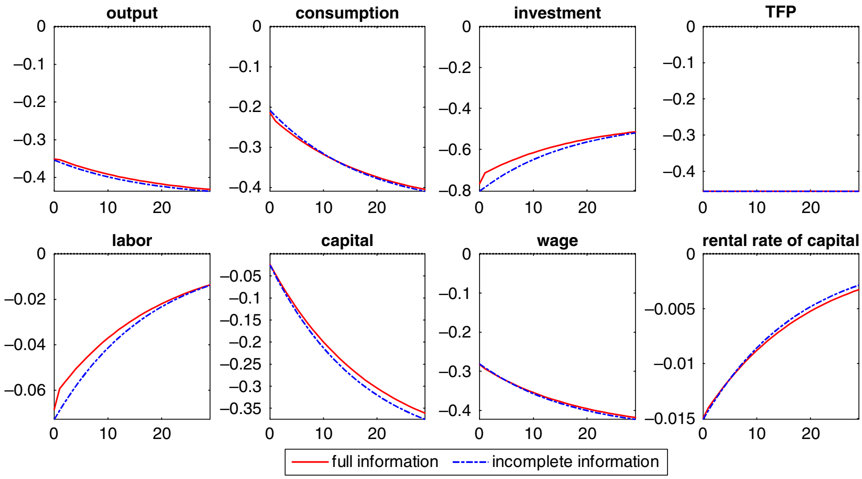

FIGURE 4. Impulse responses (in percent, relative to staying in high regime with no temporary shock) to a negative shock in the transient component of TFP growth of size 0.46 pp in the baseline case with unobserved components and Bayesian learning (incomplete information) versus the full information case.

As I show in the next subsection, these differences are quantitatively important for the implied business cycle moments generated from different versions of the model. In particular, the full information version of the regime-switching model is significantly more volatile than the baseline version with incomplete information and Bayesian learning.

Figure 4 plots the impulse responses of model variables to an innovation in the transient component of TFP growth under full information versus incomplete information and learning. I scale the size of the shock to -0.34 standard deviations of σ_ε (i.e., 0.46 pp), so that at the impact period, the total TFP decline observed by agents is equivalent to what they would observe if instead there was a switch from the high- to low-growth regime. Since agents cannot at first differentiate whether the shock is to the regime component or the transient component, their reaction at the impact period is the same as in Figure 3 under the incomplete information case. After the impact period, they realize fairly quickly that the ultimate decline in the level of TFP is going to be much smaller, and the shock likely is to the transient component of TFP growth rather than a switch in the growth regime.[17] Therefore, while the transitional dynamics of the economy to the new (and lower) balanced growth path in Figure 4 is qualitatively similar to those in Figure 3, the responses are quantitatively much smaller. Furthermore, the relative strengths of the impulse responses under the incomplete information versus the full information cases are now reversed with respect to Figure 3. In particular, the responses of investment and labor in the incomplete information case are now slightly stronger relative to the full information case, since agents place some probability that the economy has switched to the low TFP growth

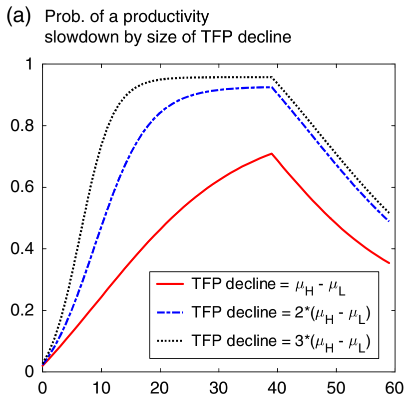

FIGURE 5. Probability of a productivity slowdown varying (a) the size and (b) the duration of the TFP growth decline.

regime. The latter, if true, would necessitate a stronger decline in the steady-state capital stock, and therefore agents are prompted to reduce investment and labor more than in the full information case. The transient growth shock thus leads to an overreaction by agents at the impact period, but even as the shock dies down and agents realize that the regime switch has likely not occurred, the discrepancies between the full information and incomplete information cases continue to linger on, given the endogenous persistence in the model due to capital accumulation and the smoothing incentives of agents.

In Figure 5, I explore the pace of learning in the model based on the size and the duration of a productivity slowdown. In particular, panel (a) plots the probability of a low-growth regime based on the size of the TFP growth shock, where the shock duration is fixed at 10 years (i.e., 40 quarters) as in Figure 3 while the shock sizes are scaled as multiples of 0.46 pp decline in observed TFP growth, Δa_t, equal to the difference between the trend growth rate in the high- and low-growth regimes (i.e., multiples of minus $\mu_H - \mu_L$). As the figure illustrates, the learning process is fairly slow, especially for a lower size shock, since the agent has reason to believe that the shock can be of the transient kind. As the shock size increases, the pace of learning increases and the probability placed by agents to a productivity slowdown is also higher. Panel (b) repeats the same exercise for a shock size of 0.46 pp decline in observed TFP growth, but for different durations for the shock. As expected, the probability placed on the low-growth regime by agents increases as the decline in the observed TFP growth lasts longer.

4.2. Business Cycle Statistics

In this subsection, I compute the business cycle statistics generated by the baseline model with unobserved components and learning versus the full information case, and compare these to the corresponding moments from the data as well as moments generated from a single-regime variant of the model. To generate business cycle moments from each model, I average over 200 simulations, feeding in artificial shocks to the regime and the transient components of TFP growth [Cooley and Prescott (1995)]. I start all simulations from the deterministic steady-state conditional on the high-growth regime (with $p_0 = 1$), simulate for 300 periods, and burn the initial 100 periods. The simulations result in an average of 3.99 switches in the TFP regime, μ_t, implying an average of 2.37 productivity slowdowns per 200 periods (i.e., 50 years), largely consistent with the data.[18]

Note that the model generates time-series for detrended variables as explained in the previous section; I thus transform these detrended variables back to their nonstationary levels by multiplying them with e^{a_t} when appropriate. I then log-difference these level series to render them stationary, and compute their corresponding business cycle moments.[19] Table 3 summarizes these results, comparing (ii) the business cycle moments implied by the US data between 1948Q2 and 2018Q1 to the corresponding model moments under (ii) the baseline case with incomplete information and learning, (iii) under full information, and (iv) under a single-regime version of the model, where both μ_H and μ_L are set to 0.0053, equal to the ergodic mean of TFP growth in the regime-switching version of the model.[20] Since the persistence parameter of the transient component of TFP is estimated to be 0, the single-regime variant of the model essentially features a unit-root TFP process with a drift, similar to King et al. (1988).

The data moments in panel (i) indicate that consumption growth is less than half as volatile as output growth, while investment growth is more than twice as volatile. In terms of the production inputs and factor payments, the growth rates of the capital stock and the rental rates show little cyclical variation over the business cycle, while growth in labor and real wages tend to be relatively more volatile, along with TFP growth. Furthermore, most variables exhibit moderate levels of autocorrelation and cross-correlation with output at various leads and lags. The baseline model featuring regime switching and learning generates business cycle moments that are comparable to the data in many, but not all, dimensions (see panel ii). The model generates slightly lower volatility in output growth relative to the data, indicating the possible role of other shocks and frictions in generating and propagating business cycles that are excluded from the model here, such as labor indivisibility and capital utilization [King and Rebelo (1999)].[21] Consumption growth is found to be about two thirds as volatile as output growth, while investment growth is about twice as volatile, consistent with the findings in the RBC literature featuring stochastic growth [King et al. (1988)]. Given the relatively low values for the labor supply elasticity and capital depreciation rate parameters, the factor inputs are also significantly less volatile than output, while labor growth is relatively more volatile, consistent with the data. In general, the model variables are highly correlated with output growth contemporaneously, but show low rates of autocorrelation and cross-correlation with output growth at leads and lags different than 0. The sensitivity analysis exercise conducted in Table 4 indicates that increasing the elasticity of labor supply (i.e., reducing φ) or

TABLE 3. Business cycle moments of log-differenced variables in % (data vs. model)

	mean	stdev	autocorr	Cross-correlation with Δ log(y_t)				
				−2	−1	0	+1	+2
(i) Data (US 1948Q2–2018Q1)								
$\Delta\ log(y_t)$	0.46	1.30	43.89	24.15	44.08	100.00	44.06	24.69
$\Delta\ log(c_t)$	0.49	0.55	26.73	20.19	39.35	51.82	32.27	21.87
$\Delta\ log(i_t)$	0.42	3.60	33.66	20.90	37.03	96.26	39.96	20.74
$\Delta\ log(n_t)$	−0.02	0.90	63.27	1.98	35.78	73.53	61.52	46.56
$\Delta\ log(k_t)$	0.41	0.29	93.82	−16.28	−3.05	30.43	43.09	49.31
$\Delta\ a_t$	0.46	1.21	−3.97	8.42	−10.28	−0.27	−17.24	−3.02
$\Delta\ log(w_t)$	0.44	0.69	8.28	7.57	0.39	−6.41	−10.08	−7.53
$\Delta\ r_t$	0.00	0.12	42.13	28.71	45.90	97.50	35.81	14.91
(ii) Regime-switching model with unobserved components and Bayesian learning								
$\Delta\ log(y_t)$	0.54	0.92	4.33	3.66	4.33	100.00	4.29	3.63
$\Delta\ log(c_t)$	0.54	0.63	9.33	4.38	5.09	99.35	7.60	6.89
$\Delta\ log(i_t)$	0.54	1.85	0.84	2.80	3.39	99.18	0.61	0.02
$\Delta\ log(n_t)$	0.00	0.15	−0.30	1.92	2.46	97.00	−3.00	3.51
$\Delta\ log(k_t)$	0.54	0.21	95.63	6.80	7.41	35.54	34.47	33.30
$\Delta\ a_t$	0.54	1.22	2.11	3.22	3.87	99.82	2.24	1.63
$\Delta\ log(w_t)$	0.54	0.78	6.18	3.96	4.65	99.89	5.66	4.97
$\Delta\ r_t$	0.00	0.04	−0.54	1.84	2.41	97.32	−3.69	−4.16
(iii) Regime-switching model with full information								
$\Delta\ log(y_t)$	0.55	1.20	−2.33	2.49	−2.38	100.00	−2.26	−2.46
$\Delta\ log(c_t)$	0.55	0.95	−7.11	1.52	11.63	3.16	64.95	7.39
$\Delta\ log(i_t)$	0.55	6.18	−35.68	1.30	−6.87	85.72	−34.46	−1.99
$\Delta\ log(n_t)$	0.00	0.73	−27.34	1.27	−8.31	82.18	−17.66	−4.40
$\Delta\ log(k_t)$	0.54	0.28	75.95	5.98	0.21	57.34	30.67	28.56
$\Delta\ a_t$	0.55	1.23	2.79	2.16	0.15	96.09	2.93	2.29
$\Delta\ log(w_t)$	0.55	0.71	15.85	2.82	3.78	82.96	19.25	7.70
$\Delta\ r_t$	0.00	0.05	−11.36	0.94	−3.66	97.23	−14.70	−4.90
(iv) Single-regime version of the model								
$\Delta\ log(y_t)$	0.53	0.94	0.77	0.17	0.81	100.00	0.84	0.20
$\Delta\ log(c_t)$	0.53	0.55	7.55	0.16	0.80	99.00	5.93	4.99
$\Delta\ log(i_t)$	0.53	2.19	−2.14	0.15	0.78	99.34	−3.15	3.55
$\Delta\ log(n_t)$	0.00	0.20	−2.86	0.10	0.72	98.12	−6.15	−6.34
$\Delta\ log(k_t)$	0.53	0.18	91.99	−1.10	−0.98	37.59	35.00	32.42
$\Delta\ a_t$	0.53	1.20	−0.63	0.15	0.80	99.91	−0.71	1.26
$\Delta\ log(w_t)$	0.53	0.74	2.90	0.17	0.81	99.86	2.73	1.98
$\Delta\ r_t$	0.00	0.04	−2.85	0.08	0.71	98.14	−6.17	−6.37

Note: The cross-correlations refer to $corr\ (x_{t+s}, \Delta \log y_t)$ for $s = -2, \ldots, +2$, where x denotes a variable in the first column.

TABLE 4. Sensitivity analysis on business cycle moments (logdifferenced data vs. model)

				Cross-correlation with Δ log($\mathbf{y}_t$)				
	mean	stdev	autocorr	−2	−1	0	+1	+2
Unobserved components and Bayesian learning (sigma = 1, vartheta = 0.5)								
Δ log(y_t)	0.55	0.97	4.76	3.87	4.71	100.00	4.80	3.81
Δ log(c_t)	0.55	0.66	10.27	4.59	5.55	99.25	8.42	7.34
Δ log(i_t)	0.55	1.99	1.15	3.03	3.73	99.12	0.91	0.02
Δ log(n_t)	0.00	0.22	0.01	2.19	2.75	96.94	2.77	−3.56
Δ log(k_t)	0.54	0.22	95.50	6.75	7.46	35.96	34.98	33.72
Δ a_t	0.55	1.23	2.79	3.49	4.32	99.86	2.97	2.03
Δ log(w_t)	0.55	0.76	7.72	4.30	5.21	99.75	6.90	5.86
Δ r_t	0.00	0.04	−0.23	2.05	2.72	97.25	3.40	−4.17
Unobserved components and Bayesian learning (sigma = 0.75, vartheta = 1)								
Δ log(y_t)	0.54	1.00	4.85	3.59	4.69	100.00	4.70	3.56
Δ log(c_t)	0.54	0.50	22.55	5.39	6.53	96.39	14.84	13.33
Δ log(i_t)	0.53	2.61	−0.42	2.28	3.34	98.63	−1.47	2.38
Δ log(n_t)	0.00	0.27	−1.38	1.47	2.44	96.74	−5.30	6.02
Δ log(k_t)	0.54	0.25	94.15	7.11	7.79	40.31	38.39	36.32
Δ a_t	0.54	1.22	2.83	3.24	4.32	99.89	2.87	1.80
Δ log(w_t)	0.54	0.75	9.73	4.26	5.38	99.59	8.18	6.90
Δ r_t	0.00	0.04	−1.35	1.50	2.45	96.76	−5.32	−6.02

increasing the intertemporal elasticity of substitution (i.e., reducing σ) can bring the baseline model's moments somewhat closer to the data, but do not improve much in terms of the autocorrelation of variables and their cross-correlation with output at leads and lags.[22]

Comparison of panels (ii) and (iii) in Table 3 indicates that the regime switching model with incomplete information and learning generates significantly less volatility relative to the full information variant. The main differences are due to the dampened volatility of investment and labor, consistent with the impulse responses presented in the last subsection. In particular, with unobserved components, agents are not fully sure that the changes in TFP growth is due to the regime component, and therefore their reaction in terms of labor supply and investment is far more muted in the short run in response to a switch in the regime. In fact, the regime-switching model with full information comes closer to replicating the volatility in output growth in the data through TFP shocks only, but this comes at the expense of significantly overpredicting the volatility in investment growth. The full information model also generates significantly negative autocorrelation coefficients for investment and labor growth in contrast with the data, while adding Bayesian learning to the setup largely fixes this sign issue (but still generates significantly lower autocorrelation coefficients than in the data).

Not surprisingly, the results from the single-regime variant of the model (panel iv) are similar to those from King et al. (1988), since the TFP process of the regime switching model reduces to a unit root with a drift, as in the aforementioned paper. More interestingly, the moments generated by the single-regime model are by and large very similar to those from the regime-switching model with unobserved components and learning. Comparison of panels (iii) and (iv) indicates that adding regime-switching TFP growth into the standard RBC model significantly increases the volatility in the system, especially with respect to labor and investment, but also for consumption and overall output. Further adding the unobserved components and learning features to the model however largely offsets this increase in volatility, and renders the regime-switching model more similar to the single-regime model. In effect, the learning feature under incomplete information smooths the effects of discrete jumps caused by the regime switch in the full information case. As a result, while adding a regime-switching component to the standard RBC model increases the overall volatility in the system, incorporating incomplete information and learning dampens volatility, and with similar magnitudes.

4.3. Counterfactual regarding the Recent Productivity Slowdown

In this subsection, I investigate the effects of the productivity slowdown on the US economy since 2005. Figure 6 simulates the baseline model with regime-switching and learning using the observed TFP growth series, Δa_t, and compares the implied log levels of model variables to the data and to a counterfactual simulation of the model in the absence of the productivity slowdown.[23] For the latter, I again use the baseline model with regime-switching and learning, but now assume that agents observe an aggregate TFP growth equal to the actual TFP growth in the data plus the difference between the trend growth rates in the high and low regimes (i.e., $\mu_H - \mu_L$).

Given the slowdown in the observed TFP growth and the resulting increase in the probability that agents place on the low-growth regime, the model predicts only modest increases in consumption and output in the post-2005 period, along with a stagnant investment profile. Note that, perhaps not surprisingly, the model is not able to capture the large decline in investment and output during the Great Recession following the financial crisis of 2008, which were likely caused by financial and demand-type shocks that are absent from the model. In fact, the productivity series of Fernald (2012) indicates a slight increase in productivity after 2008, as can be seen in Figure 6. The finding that productivity shocks cannot account for the output decline during the Great Recession is also consistent with the "business cycle accounting" exercise conducted by Brinca et al. (2016) for this period. In particular, they find that the decline in US output and labor during this period can primarily be captured by the "labor wedge," while the efficiency wedge plays a far smaller role.[24] The model here captures the efficiency wedge through TFP shocks, but abstracts from financial or labor market frictions that can

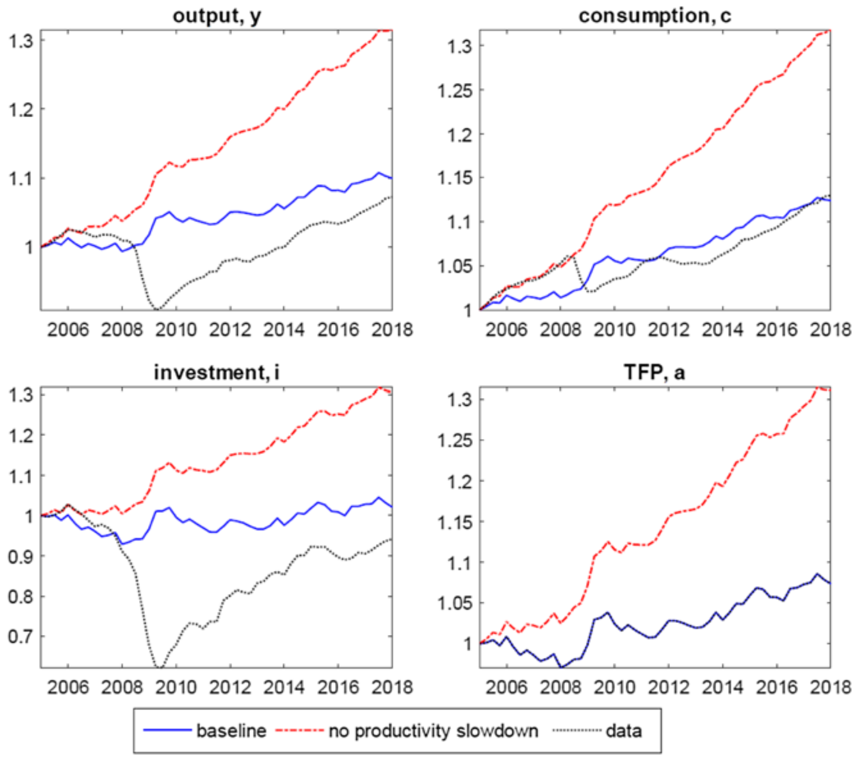

FIGURE 6. Log level of model variables under observed TFP since 2005 and a counterfactual with no productivity slowdown. 2005Q1 figures are normalized to 1.

give rise to changes in the labor wedge, and thus, is unable to fully capture the decline in output that occurred during the Great Recession.

A comparison of the baseline simulation with the data suggests that by 2018Q1 macroeconomic aggregates have essentially returned to, or have come close to, their "potential" levels, based on the observed productivity series. The counterfactual simulation in Figure 6 also indicates that the cost of the productivity slowdown to US welfare has potentially been quite large, given the near 0.46 pp decline in quarterly TFP growth, Δa_t, relative to the high-growth regime. In particular, by 2018Q1, the impact on the level of consumption due to the productivity slowdown has reached near 20 pp.

5. CONCLUSION

The pattern of US productivity growth can be viewed from the perspective of the regime-switching approach of Hamilton (1989), with significant slowdowns in trend TFP growth in the mid-1970s and mid-2000s. In this paper, I extend the standard RBC model with regime-switching in productivity growth and Bayesian

https://doi.org/10.1017/S1365100519000257 Published online by Cambridge University Press

learning and investigate the implications of productivity slowdowns on macroeconomic variables. Model simulations reveal that learning dampens the volatility in the model by reducing the responses of investment and labor to a switch in the trend component of TFP growth, since agents mistakenly place some probability that the observed low TFP growth is due to the transient component and not due to regime switching. The model is largely consistent with US business cycle facts, and since it imbeds the Hamilton filter, it offers an objective way to infer slowdowns in trend productivity, a major issue for policymakers trying to assess the medium-term outlook for output gap in real time. The model also implies that macroeconomic aggregates in the USA are currently close to their potential levels given the observed productivity series, while counterfactual simulations indicate that the cost of the recent productivity slowdown to US welfare has been significantly high.

In future work, the model could be extended to include regime-switching in shock variances, a feature that is now commonly used in the time-series literature [e.g., McConnell and Perez-Quiros (2000)], in order to incorporate issues related to the *Great Moderation* in volatility between 1984 and 2007.[25] The model can also be extended to investigate the asset pricing implications of long-run uncertainty [e.g., Tortorice (2018)], booms and busts in housing markets and household debt [e.g., Alpanda and Ueberfeldt (2016)], or the appropriate monetary framework and stance of policy when there is uncertainty regarding the output gap and potential output growth.

NOTES

1. These figures were constructed using the utilization-adjusted quarterly TFP series for the US Business Sector, produced by John Fernald of the Federal Reserve Bank of San Francisco [Fernald (2012)].

2. As in the standard RBC setup, changes in TFP growth, including those that relate to the regime component, are assumed to occur exogenously here. The literature on endogenous growth indicates that research and development activities, as well as profit incentives, are important to understand the creation and the diffusion of new technologies that lead to this sustained, but uneven, productivity growth [Romer (1990)]. One could extend the model here to make the regime probabilities endogenous, as a way to take this endogeneity into account.

3. Also see Kahn and Rich (2007), who adopt a multivariable version of the Hamilton filter, proposed by Kim and Murray (2002) and Kim and Piger (2002), to study regime-switching in trend growth, which they identify using data on productivity along with consumption and labor compensation.

4. Early RBC models typically featured deterministic long-run TFP growth or abstracted from long-run growth altogether [Prescott (1986) and Cooley and Prescott (1995)]. Based on the time-series evidence that many macroeconomic series contain stochastic trends [Nelson and Plosser (1982)], King et al. (1988) introduced a stochastic trend into the standard RBC framework using a unit root process with a drift in the level of TFP. In the more recent literature on DSGE models, both deterministic and stochastic TFP trends are commonly used [Smets and Wouters (2007)].

5. Note that discrete-valued Markov processes can approximate autoregressive processes fairly well [Tauchen (1986)]. Thus, the adoption of a regime-switching framework for trend TFP growth in

this paper (and in others as cited above) provides an alternative, yet similar, characterization for the data generating process of TFP, as the use of a persistent AR(1) process for the trend component [such as in Pakko (2002)].

6. Also see Davig (2004) and Richter and Throckmorton (2015), who analyze the effects of fiscal policy in DSGE models with regime switching in the government's debt target. Similarly, Dufrénot and Khayat (2017) consider a regime-switching Taylor rule for the European Central Bank.

7. Note that at the steady state of the model, the real net return on capital is given by $r - \delta = e^{\sigma \Delta a}/\beta - 1$, where Δa denotes the ergodic mean of Δa_t.

8. Note that the underlying TFP growth data in Fernald (2012) corresponds to $400\,(1 - \alpha)\,\Delta a_t$ in equation (8) of the model; thus, I render the series consistent with the model's Δa_t by dividing the series by $400\,(1 - \alpha)$, where α is set to 0.33 as in the model parameterization.

9. In the absence of this restriction, the duration of the low TFP regimes is estimated to be significantly shorter, while the high-TFP regime lasts longer than what is implied in Figure 1. One could alternatively try to embed the information in Figure 1 as a prior distribution on the π_{LL} and π_{HH} parameters, and estimate the parameters using Bayesian methods.

10. The estimates for the asymptotic standard errors in Table 2 are constructed using the outer product of gradients approach, which utilizes the fact that the expected second derivatives matrix is equal the covariance matrix of the first derivatives vector; thus, one can use the outer product of the gradient matrix as an estimate for the information matrix [Greene (2000), p. 132]. The initial prior, p_0, is set to 0.8 for 1947Q2, capturing the significant likelihood of being in the high-growth regime in the beginning of the post-war period. Sensitivity analysis on p_0 suggests that the parameter estimates are not affected significantly by the choice of this initial prior. In particular, setting $p0$ to 1 results in virtually the same parameter estimates.

11. The average duration (in quarters) for a growth regime is given by $1/(1 - \pi_{ii})$ for $i \in \{L, H\}$.

12. Also see Brinca et al. (2016) who conduct a "business cycle accounting" exercise for the Great Recession period in the USA, and show that financial frictions such as leverage constraints or input-financing constraints, which arguably were important in the propagation of shocks during the Great Recession, can affect the "efficiency wedge", as well as labor or investment wedges within the standard RBC setup. I discuss the relation between my results and the findings in Brinca et al. (2016) further in Section 4.3.

13. Note that the first welfare theorem holds for this setup, and one can thus solve the model using the social planner's problem instead.

14. Maliar and Maliar (2013) show that the ECM methodology is convenient to implement since envelope conditions are simpler to analyze numerically relative to first-order conditions. The methodology is also faster than the conventional value function iteration method, and is similar in accuracy and speed relative to the endogenous grid method [Carroll (2006)]. See Arellano et al. (2016) for a more recent use of the ECM methodology to solve a challenging default risk model with a kink in the value and policy functions.

15. For the Chebyshev polynomial approximation, I use the Matlab routines in the *CompEcon* toolbox accompanying Miranda and Fackler (2002).

16. Note that the model simulations generate time-series for detrended variables as explained in the previous section; I thus transform these detrended variables back to their nonstationary levels by multiplying them with e^{a_t} when appropriate. The variables are presented in log levels, except for the rental rate of capital, which is not logged. To generate the impulse responses, I start the simulations from the *deterministic* steady-state conditional on the high-growth regime, and assume agents place 80% probability to being in the high-growth regime right before the impact period of the shock in the incomplete information case. Note that agents' assessment of the regime probability changes over the impulse horizon, and their decision rules take into account of the probability of a regime switch during the transition path. To generate each impulse response, I simulate the model twice, with and without the shock, and calculate the (log) difference. Note that, as I am starting the simulations from the deterministic rather than the stochastic steady state, the model has transitional dynamics even in

the absence of any shocks. The "shock minus control" approach used here in generating the impulse responses eliminates the effects of these dynamics and retains the pure effects from the shock itself.

17. Note that the transient component of TFP growth, z_t, is a temporary shock to the *growth* rate of TFP, consistent with the treatment in Hamilton (1989). Thus, it still has a *permanent* impact on the level of TFP, although a much smaller one than implied by the regime component.

18. The timing for switches in the regime is determined using random draws from a uniform distribution and comparing these to a cutoff value, and the ε_t's are drawn from a normal distribution. The resulting series for μ_t and z_t for $t = 1, 2, \ldots$ are then constructed conditional on $\mu_0 = \mu_H$ and $z_0 = 0$. Since I burn a third of the generated sample series in each simulation, the impact of the initial conditions on the moment estimates is largely eliminated.

19. Note that the series for a_t and r_t are not logged before differencing.

20. See Appendix A for a discussion of data sources and the construction of model-consistent series from the data.

21. Increasing the number of shocks in the model would also likely result in a better fit of the baseline model to the data, while potentially decreasing the variation of model variables captured by TFP shocks alone. For example, Smets and Wouters (2007) find that demand shocks or changes to investment-specific technological change can help account for the short-term variation in real variables, while TFP shocks are more important to capture variation in the long run.

22. In the DSGE literature, this deficiency in the propagation mechanisms within the basic RBC models is typically overcome through the use of real rigidities in the model, such as habit formation in consumption and adjustment costs in capital accumulation [Christiano et al. (2005) and Smets and Wouters (2007)]. I do not include these features in order to focus exclusively on the regime-switching and learning aspects of the model. Excluding these features also ease computation by limiting the size of the state space.

23. Thus, the TFP series generated from the baseline model is equivalent to that from the data, which is from Fernald (2012).

24. The business cycle accounting framework was developed by Chari et al. (2006). The "efficiency wedge" shows up as time-varying TFP, while the "labor wedge" shows up as a time-varying labor income tax within the standard RBC setup. The latter drives a wedge between the households' consumption-leisure marginal rate of substitution and the marginal product of labor in production, and thus, captures distortions related to both labor supply and labor demand.

25. There are also papers in the DSGE literature that incorporate time variation in variances [Justiniano and Primiceri (2006), Fernández-Villaverde et al. (2015), and Lhuissier (2018)].

26. The vintage of data used was produced on May 3, 2018, by John Fernald and Neil Gerstein.

REFERENCES

Alpanda, S. and A. Peralta-Alva (2010) Oil crisis, energy-saving technological change and the stock market crash of 1973-74. *Review of Economic Dynamics* 13, 824–842.

Alpanda, S. and A. Ueberfeldt (2016) Should Monetary Policy Lean Against Housing Market Booms? Bank of Canada Working Paper: No. 2016–19.

Arellano, C., L. Maliar, S. Maliar and V. Tsyrennikov (2016) Envelope condition method with an application to default risk models. *Journal of Economic Dynamics & Control* 69, 436–459.

Bianchi, F. and L. Melosi (2016) Modeling the Evolution of Expectations and Uncertainty in General Equilibrium. *International Economic Review* 57, 717–756.

Boz, E., C. Daude and C. B. Durdu (2011) Emerging market business cycles: Learning about the trend. *Journal of Monetary Economics* 58, 616–631.

Brinca, P., V. V. Chari, P. J. Kehoe and E. McGrattan (2016) Accounting for business cycles. In: J. B. Taylor and H. Uhlig (eds.), *Handbook of Macroeconomics* vol. 2, pp. 1013-1063. Amsterdam: Elsevier.

Bullard, J. (2016) The St. Louis Fed's New Characterization of the Outlook for the U.S. Economy. Federal Reserve Bank of St. Louis, June 17, 2016.

Burns, A. F. and W. C. Mitchell (1946) *Measuring Business Cycles*. New York: National Bureau of Economic Research.

Carroll, C. D. (2006) The method of endogenous grid points for solving dynamic stochastic optimal problems. *Economics Letters* 91, 312–320.

Chari, V. V., P. Kehoe and E. R. McGrattan (2006) Business cycle accounting. *Econometrica* 75, 781–836.

Chauvet, M. and J. M. Piger (2003) Identifying business cycle turning points in real time. *The Federal Reserve Bank of St. Louis Review* March/April, 47–61.

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

Cooley, T. F. and E. C. Prescott (1995) Economic growth and business cycles. In: T. F. Cooley (ed.), *Frontiers of Business Cycle Research*, pp. 1-38. Princeton: Princeton University Press.

Davig, T. (2004) Regime-switching debt and taxation. *Journal of Monetary Economics* 51, 837–859.

Dufrénot, G. and G. A. Khayat (2017) Monetary policy switching in the euro area and multiple steady states: An empirical investigation. *Macroeconomic Dynamics* 21, 1175–1188.

Farmer, R. E. A., D. F. Waggoner and T. Zha (2011) Minimal state variable solutions to Markov switching rational expectations models. *Journal of Economics Dynamics and Control* 35, 2150–2166.

Fernald, J. G. (2012) A Quarterly, Utilization-Adjusted Series on Total Factor Productivity. Federal Reserve Bank of San Francisco Working Paper 2012-19 (updated March 2014).

Fernández-Villaverde, J., P. Guerrón-Quintana, K. Kuester and J. Rubio-Ramírez (2015) Fiscal volatility shocks and economic activity. *American Economic Review* 105, 3352–3384.

Foerster, A., J. Rubio-Ramirez, D. F. Waggoner and T. Zha (2014) Perturbation Methods for Markov-Switching Models. National Bureau of Economic Research, Inc., NBER Working Paper: No. 20390.

French, M. W. (2001) Estimating Changes in Trend Growth of Total Factor Productivity: Kalman and H-P Filters Versus a Markov-Switching Approach. Federal Reserve Board Finance and Economics Discussion Series No. 2001-44.

Gordon, R. J. (2012) Is U.S. Economic Growth Over? Faltering Innovation Confronts the Six Headwinds. National Bureau of Economic Research, Inc., NBER Working Paper: No. 18315.

Gordon, R. J. (2016) *The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War*. Princeton: Princeton University Press.

Gourio, F. (2013) Credit risk and disaster risk. *American Economic Journal: Macroeconomics* 5, 1–34.

Greene, W. H. (2000). *Econometric Analysis*. New Jersey: Prentice Hall.

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

Justiniano, A. and G. E. Primiceri (2006) The Time Varying Volatility of Macroeconomic Fluctuations. National Bureau of Economic Research, Inc., NBER Working Paper: No. 12022.

Kahn, J. A. and R. W. Rich (2007) Tracking the new economy: Using growth theory to detect changes in trend productivity. *Journal of Monetary Economics* 54, 1670–1701.

Kim, C.-J. and C. J. Murray (2002) Permanent and transitory components of recessions. *Empirical Economics* 27, 163–184.

Kim, C.-J. and J. M. Piger (2002) Common stochastic trends, common cycles, and asymmetry in economic fluctuations. *Journal of Monetary Economics* 49, 1189–1211.

King, R. G., C. I. Plosser and S. T. Rebelo (1988) Production, growth and business cycles II: New directions. *Journal of Monetary Economics* 21, 309–341.

King, R. G. and S. T. Rebelo (1999) Resuscitating real business cycles. In M. Woodford and J. Taylor (eds.), *Handbook of Macroeconomics* vol. 1, pp. 927–1007. Amsterdam: Elsevier.

Kuang, P. and K. Mitra (2015) Long-Run Growth Uncertainty. University of Birmingham, Department of Economics Discussion Paper: No. 15-07.

Kydland, F. E. and E. C. Prescott (1982) Time to build and Aggregate Fluctuations. *Econometrica* 50, 1345–1370.

Lhuissier, S. (2018) The regime-switching volatility of euro area business cycles. *Macroeconomic Dynamics* 22, 426–469.

Lorenzoni, G. (2009) A theory of demand shocks. *American Economic Review* 99, 2050–2084.

Maih, J. (2015) Efficient Perturbation Methods for Solving Regime-Switching DSGE Models. Norges Bank Working Paper: No. 01-2015.

Maliar, L. and S. Maliar (2013) Envelope condition method versus endogenous grid method for solving dynamic programming problems. *Economics Letters* 120, 262–266.

McConnell, M. M. and G. Perez-Quiros (2000) Output fluctuations in the United States: What has changed since the early 1980's?. *American Economic Review* 90, 1464–1476.

Miranda, M. J. and P. L. Fackler (2002) *Applied Computational Economics and Finance*. Cambridge, MA: MIT Press.

Nelson, C. R. and C. R. Plosser (1982) Trends and random walks in macroeconomic time series: Some evidence and implications. *Journal of Monetary Economics* 10, 139–162.

Pakko, M. R. (2002) What happens when the technology growth trend changes? Transition dynamics, capital growth, and the "new economy". *Review of Economic Dynamics* 5, 376–407.

Prescott, E. C. (1986) Theory ahead of business cycle measurement. *Federal Reserve Bank of Minneapolis Quarterly Review* 10, 9–22.

Richter, A. and N. A. Throckmorton (2015) The consequences of an unknown debt target. *European Economic Review* 78, 76–96.

Romer, P. M. (1990) Endogenous technological change. *Journal of Political Economy* 98, S71–S102.

Smets, F. and R. Wouters (2007) Shocks and frictions in US business cycles: A Bayesian DSGE approach. *American Economic Review* 97, 586–606.

Stokey, N. L. and R. E. Lucas, Jr., with E. C. Prescott (1989). *Recursive Methods in Economic Dynamics*. Cambridge, MA: Harvard University Press.

Tauchen, G. (1986) Finite state Markov-chain approximations to univariate and vector autoregressions. *Economics Letters* 20, 177–181.

Tortorice, D. L. (2016) The business cycles implications of fluctuating long run expectations. Brandeis University Working Paper: No. 100.

Tortorice, D. L. (2018) Equity return predictability, time varying volatility and learning about the permanence of shocks. *Journal of Economic Behavior and Organization* 148, 315–343.

APPENDIX A: DATA APPENDIX

This appendix describes data sources for the data series used in the estimation and in the construction of the data moments presented in Tables 3 and 4. The series for the growth rate of TFP, Δa_t, is the utilization-adjusted quarterly TFP series for the US Business Sector, produced by John Fernald of the Federal Reserve Bank of San Francisco (Fernald (2012)), for the period between 1948Q2 and 2018Q1.[26] Note that the underlying TFP growth data in Fernald (2012) corresponds to $400\,(1-\alpha)\,\Delta a_t$ in equation (8) of the model; thus, I render the series consistent with the model by dividing it by $400\,(1-\alpha)$, where α is set to 0.33, consistent with the parameterization of the model.

Consumption, c_t, refers to per-capita nominal nondurable goods consumption plus services, deflated by the GDP deflator. Similarly, investment, i_t, refers to per-capita nominal investment plus consumer durable-goods consumption, deflated by the GDP deflator. Output series, y_t, is proxied by the sum of c_t and i_t. Consumption and investment series are from the National Income and Product Accounts (NIPA) of the Bureau of Economic Analysis (BEA), and the population series used to obtain per-capita figures is the *Civilian Noninstitutional Population* of the Census Bureau. Labor data refers to the *Hours index for the Business sector*, which is part of the *Major Sector Productivity and Costs* indices published by the Bureau of Labor Statistics (BLS). All series are quarterly, seasonally adjusted, and cover the period 1948Q2–2018Q1 after log differencing.

Wage income refers to compensation of employees, plus the labor share of proprietor's income, similar to Cooley and Prescott (1995). The nominal series are obtained from the income accounts in NIPA are deflated using the GDP deflator, and divided by population to obtain per-capita figures. The wage rate, w, is then obtained by dividing the total wage income with the labor series, n_t, constructed above. The real rental rate of capital is obtained as a residual from the data using the following model expression:

$$c_t + i_t = w_t n_t + r_t k_{t-1}. \tag{A1}$$

APPENDIX B: COMPUTATIONAL APPENDIX

This appendix provides an ECM algorithm (Maliar and Maliar 2013) to compute the solution of the model in Section 2.

APPENDIX B.1: SOLUTION OF THE MODEL IN RECURSIVE FORM

The aggregate state space for the stationary model is denoted by $S = \left(\widehat{k}, \Delta a, p\right)$, where $\widehat{k}$ denotes the (lagged) detrended level of capital (i.e., $\widehat{k}_{t-1} = k_{t-1}/e^{a_{t-1}}$). With a slight abuse of notation, I denote individual households' state variables with the same letters as the aggregate ones in what follows below. The functional equation of the representative household after detrending can be written as

$$V\left(\widehat{k}, S\right) = \max\left\{\frac{1}{1-\sigma}\left[\widehat{c}\exp\left(-\frac{n^{1+\varphi}}{1+\varphi}\right)\right]^{1-\sigma} + \beta E\left[e^{(1-\sigma)\Delta a'} V\left(\widehat{k}', S'\right)\right]\right\}$$
$$\text{s.t.}\quad \widehat{c} + \widehat{k}' - (1-\delta)\, e^{-\Delta a}\,\widehat{k} = \widehat{w}n + re^{-\Delta a}\,\widehat{k}. \tag{B1}$$

The first-order and envelope conditions of this problem, evaluated at a symmetric equilibrium, are given by

$$\widehat{c}^{-\sigma}\exp\left(\frac{\sigma-1}{1+\varphi}n^{1+\varphi}\right) = \widehat{\lambda}, \tag{B2}$$

$$\widehat{c}n^{\varphi} = \widehat{w}, \tag{B3}$$

$$\beta E\left[e^{(1-\sigma)\Delta a'} V_k\left(S'\right)\right] = \widehat{\lambda}, \tag{B4}$$

$$V_k\left(S\right) = \widehat{\lambda}\left(r + 1 - \delta\right) e^{-\Delta a}, \tag{B5}$$

where V_k (.) denotes the derivative of the households' value function with respect to capital. The last two conditions can be combined to obtain the following recursive equation:

$$V_k\left(\widehat{k}, \Delta a, p\right) = \beta\left(r + 1 - \delta\right) e^{-\Delta a} E\left[e^{(1-\sigma)\Delta a'} V_k\left(\widehat{k}', \Delta a', p'\right)\right]. \tag{B6}$$

The ECM method basically iterates on the above equation to obtain the policy functions.

https://doi.org/10.1017/S1365100519000257 Published online by Cambridge University Press

The first-order conditions of the representative firm can be written in recursive form as

$$\widehat{w} = (1-\alpha)\,\frac{\widehat{y}}{n}, \tag{B7}$$

$$r = \alpha\,\frac{\widehat{y}}{e^{-\Delta a}\,\widehat{k}}, \tag{B8}$$

$$\widehat{y} = Ae^{-\alpha(\Delta a)}\widehat{k}^{\alpha} n^{1-\alpha}. \tag{B9}$$

Similarly, the feasibility conditions can be written as

$$\widehat{c} + \widehat{i} = \widehat{y}, \tag{B10}$$

$$\widehat{i} = \widehat{k}' - (1-\delta)\,e^{-\Delta a}\,\widehat{k}. \tag{B11}$$

Finally, the TFP shocks are characterized by $\Delta a' = \mu' + z'$, where $\mu' \in \{\mu_H, \mu_L\}$ with a transition probability matrix $\Pi_\mu = [\pi_{HH}\ \pi_{HL}; \pi_{LH}\ \pi_{LL}]$, and $z' = \rho z + \varepsilon'$ with $\varepsilon' \sim N\left(0, \sigma_\varepsilon^2\right)$.

APPENDIX B.2: INITIALIZATION FOR THE ITERATION

(i) Define a Chebyshev polynomial approximation for V_k with parameter vector θ

$$\widehat{V}_k\left(\widehat{k}, \Delta a, p; \theta\right) \approx V_k\left(\widehat{k}, \Delta a, p\right). \tag{B12}$$

(ii) Create a meshgrid for the state variables, $\left\{\widehat{k}_m, \Delta a_m, p_m\right\}_{m=1,\ldots,M}$, using Chebyshev nodes for each state.

(iii) Make initial guess for the parameters of the polynomial, $\theta^{(1)}$.

(iv) Construct uniformly spaced nodes for integration over $\left\{\Delta a_j'\right\}_{j=1,\ldots,J}$ with each increment equal to $\Delta a_{inc}'$. Note that, for each j, there are four possible values that the innovations in the AR(1) process can take, and these can be pre-solved as

$$\varepsilon_{j,XY}' = \left(\Delta a_j' - \mu_Y\right) - \rho\,(\Delta a - \mu_X)\,, \text{ for all } X, Y \in \{H, L\}. \tag{B13}$$

The weights, ω_j, and the posterior, p_j', at each node can also be pre-solved as

$$\omega_j = \left\{p\left[\pi_{HH} f\left(\varepsilon_{j,HH}'\right) + \pi_{HL} f\left(\varepsilon_{j,HL}'\right)\right] + (1-p)\left[\pi_{LH} f\left(\varepsilon_{j,LH}'\right) + \pi_{LL} f\left(\varepsilon_{j,LL}'\right)\right]\right\}\Delta a_{inc}', \tag{B14}$$

and

$$p_j' = \frac{p\pi_{HH} f\left(\varepsilon_{j,HH}'\right) + (1-p)\,\pi_{LH} f\left(\varepsilon_{j,LH}'\right)}{p\left[\pi_{HH} f\left(\varepsilon_{j,HH}'\right) + \pi_{HL} f\left(\varepsilon_{j,HL}'\right)\right] + (1-p)\left[\pi_{LH} f\left(\varepsilon_{j,LH}'\right) + \pi_{LL} f\left(\varepsilon_{j,LL}'\right)\right]}. \tag{B15}$$

APPENDIX B.3: ITERATION

(i) For each point in the meshgrid, evaluate $\widehat{V}_k\left(\widehat{k}_m, \Delta a_m, p_m; \theta^{(i)}\right)$ and solve for all model variables including r_m and the implied $\widehat{k}_m'$.

(ii) Compute d_m as

$$d_m = \beta\,(r_m + 1 - \delta)\,e^{-\Delta a_m} E\left[V_k\left(\widehat{k}_m', \Delta a', p'\right)\right], \tag{B16}$$

where the expectation on the right-hand side can be evaluated as

$$E\left[V_k\left(\widehat{k}_m', \Delta a', p'\right)\right] = \sum_{j=1}^{J} \omega_j V\left(\widehat{k}_m', \Delta a_j', p_j'\right). \tag{B17}$$

(iii) Run a regression to find the $\widehat{\theta}$ that best-fits the d vector for the given nodes:

$$\widehat{\theta} = \arg\min_{\theta} \sum_{m=1}^{M} \left\| d_m - \widehat{V}_k \left(\widehat{k}_m, \Delta a_m, p_m; \theta \right) \right\|, \tag{B18}$$

and update θ with damping

$$\theta^{(i+1)} = (1-\xi)\,\theta^{(i)} + \xi \widehat{\theta}. \tag{B19}$$

(iii) Stop iterations when convergence in the capital policy function is achieved; that is, when

$$\frac{1}{\xi M} \sum_{m=1}^{M} \left| \frac{\left(k_m'\right)^{(i+1)} - \left(k_m'\right)^{(i)}}{\left(k_m'\right)^{(i)}} \right| < 10^{-10}. \tag{B20}$$

https://doi.org/10.1017/S1365100519000257 Published online by Cambridge University Press