Macroeconomic Dynamics, **18**, 2014, 418–437. Printed in the United States of America.
doi:10.1017/S1365100512000454

IMPERFECT TRANSMISSION OF TECHNOLOGY SHOCKS AND THE BUSINESS CYCLE CONSEQUENCES

HAMILTON B. FOUT
Federal National Mortgage Association
and
Kansas State University

NEVILLE R. FRANCIS
University of North Carolina

We investigate the business cycle effects of imperfect transmission of technology shocks within a basic real business cycle (RBC) model along two dimensions. First, we assume that agents cannot distinguish a temporary increase in productivity growth from a sustained increase in the underlying growth rate of productivity and instead must conduct signal extraction exercises and update beliefs about the source of aggregated shocks. Second, we propose a technology adjustment cost resulting in the slow diffusion of technological innovations into the production process. Both of these impediments to the transmission of technology result in a large initial wealth effect, increasing investment and hours less, relative to the usual RBC model without these frictions. Furthermore, each of these features is capable of producing a decline in hours on impact of the technology shock matching the negative response in hours found in the data by such works as Gali [*American Economic Review* 89(1), 249–271 (1999)].

Keywords: Technology Shocks, Learning, Kalman Filter, Impulse Responses

1. INTRODUCTION

In standard models of the macroeconomy, technology and other shocks are assumed to be perfectly observable and immediately transmitted through the economy. This theoretical approach stands in contrast to the transmission of technology shocks through the actual economy. As several works in the macroeconomics literature have shown, informational rigidities and other limitations on the diffusion

We would like to thank Richard Froyen, Michael Owyang, Julio Rotemberg, Michael Salemi and Tatevik Sekhposyan, as well as other seminar participants at Kansas State University, the University of Nebraska, the University of North Carolina, the Western Economic Society Meetings, the St.Louis Federal Reserve's Learning Week and the North American Econometric Society Meetings for helpful suggestions. This paper was first circulated under the title "Macroeconomic Learning and the Propagation of Technology Shocks" in May 2008. The analysis, opinions, and conclusions set forth within this paper are to be attributed solely to the authors and do not reflect the belief or opinion of their affiliated institutions. Address correspondence to: Hamilton B. Fout, Federal National Mortgage Association (Fannie Mae), 3900 Wisconsin Avenue, Washington, DC 20016, USA; e-mail: hamfout@gmail.com.

© 2012 Cambridge University Press 1365-1005/12

https://doi.org/10.1017/S1365100512000454 Published online by Cambridge University Press

of technology are key empirical features in the transmission of technology shocks; see for example Rogers (1995), Rotemberg (2003), and Edge et al. (ELW, 2007). In this paper, we investigate the business cycle consequences of these rigidities within a basic general equilibrium framework.

Previous papers that have incorporated informational frictions or limitations on the transmission of technology shocks through real business cycle (RBC) models of the macroeconomy have focused on the long-run consequences of these frictions. For instance, Bullard and Duffy (2004) analyze the effects of adaptive learning under parameter uncertainty in an effort to explain the historical path of U.S. output, consumption, and investment. ELW analyze the medium- to long-run effects of Kalman filter learning on impulse responses to level and trend shocks under uncertainty of the source of the productivity shock. Rotemberg (2003) introduces a technology process that is slowly incorporated into productivity to explain the long-run time series properties of U.S. GDP.[1]

We extend these works by introducing alternative mechanisms that limit the immediate transmission of technological changes in order to investigate the short-run consequences of these frictions within an otherwise standard RBC model. We first modify the ELW framework by including alternative informational assumptions, leading to significantly different conclusions about the effects of uncertainty and learning. Second, we introduce a slow diffusion of technology, in the spirit of Rogers (1995) and Rotemberg (2003), by modeling adjustment costs to implementing technological innovations. The effects of limited transmission of technology shocks on the short-run predictions of the model are particularly important in light of the recent hours debate in the macroeconomic literature [see for example Gali (1999), Christiano et al. (2003), Francis and Ramey (2005), Pesavento and Rossi (2005), Basu et al. (2006) and Gil-Alana and Moreno (2009)].[2] In this debate, the empirical response of hours to a level shock is used to test the validity of the RBC paradigm. Some researchers, such as Christiano et al. (2003), find this permanent technology shock to be expansionary for the labor input in support of the RBC hypothesis, whereas others, such as Gali (1999) and Francis and Ramey (2005), find technology shocks to be contractionary for labor, in contrast to the RBC hypothesis.

The standard model of Kydland and Prescott (1982) provides the benchmark for the RBC paradigm and predicts that hours will rise after a positive shock to technology. To generate a fall in the labor input several adjustments have to be made to the standard theory. Gali (1999) and Basu et al. (2006) both suggest adding nominal rigidities (i.e., sticky prices or wages). On the other hand, Francis and Ramey (2005) suggest a model with the real rigidities of habit in consumption and investment adjustment costs or an economy with a Leontief production technology. As there remains disagreement about the extent of nominal [Bils and Klenow (2004) and Nakamura and Steinsson (2008)] and real [Dynan (2000) and Groth and Khan (2010)] rigidities in actual economies, we instead focus on the underlying processes assumed for technology shocks. In particular, we relax the assumption that shocks are perfectly observable and fully implementable upon impact.

https://doi.org/10.1017/S1365100512000454 Published online by Cambridge University Press

In our first setup, we assume that agents cannot immediately distinguish frequent temporary changes from persistent but infrequent changes in productivity growth. In our model, agents cannot directly observe the underlying innovation to technology and instead must use signal extraction learning to update beliefs about its source.[3] In our second setup, we allow agents to directly observe the underlying technological process but assume a cost for implementing technological innovations into production, resulting in a gradual diffusion of technology over time. We view these two specifications not necessarily as competitors but as each having individual merit.

We begin with a learning environment similar to that of ELW in which agents cannot directly observe the underlying components of the productivity process. Using data on real-time forecasts of long-run productivity growth, ELW provide empirical support for this approach by highlighting the inability of economic agents to distinguish frequent temporary changes from infrequent but persistent changes in productivity growth. Our approach departs from that of ELW in several important ways with fundamentally different implications for the responses of the endogenous variables. Like ELW, we assume that agents in these models have to learn the source of the technology shock hitting the economy by using a Kalman filter to update their forecasts of these underlying components. Therefore, households and firms in these models make allocation decisions based on incomplete information regarding stochastic processes. Unlike ELW, however, we assume that agents use all contemporaneously available information when learning and take into account their underlying learning process when forming expectations about the future. Under the ELW framework, learning has no immediate effect on the response of model variables to technology shocks, whereas under our approach, learning matters for these impact responses.

We next introduce an alternative feature limiting the transmission of technology through the economy by assuming that changes in technology are costly and thus not fully included in the production process on impact. Instead, these changes are adopted into production slowly over time. Rogers (1995) and Rotemberg (2003) provide various historical examples of technological innovations that were slow to diffuse despite evidence of the effectiveness of the new technology. Examples include a hybrid variety of corn that after 10 years still had only diffused to 50% of farms even though it generated 20% higher yields than previous varieties, and the DVORAK keyboard, which despite enabling faster typing is not as widely used as the less ergonomic incumbent QWERTY keyboard. Oliner and Sichel (1994, 2000) analyze the delayed effect of computer technologies that were introduced in the 1980s but did not have a significant impact on productivity until the 1990s. We model these initial periods of slow technological dissemination as arising from costs associated with implementing new technologies.

We find that including these frictions fundamentally alters the predictions of the model. In particular, we find that each of these frictions generates a substantial wealth effect on the impact of a level shock, causing investment and hours to rise by less than in the standard RBC model. We also find that including either of these

https://doi.org/10.1017/S1365100512000454 Published online by Cambridge University Press

limits on the transmission of technology, under a broad set of parameterizations, is sufficient to deliver a fall in labor following a level shock, even in the absence of the commonly used rigidities of sticky prices, habit persistence, and investment adjustment costs. Thus these frictions provide a potential theoretical explanation for the findings of Gali (1999) and Francis and Ramey (2005), with only a slight change in the RBC model.

The rest of the paper proceeds as follows. Section 2 presents the frictionless RBC model. Section 3 introduces signal extraction learning and reports the impulse responses to productivity shocks in this environment. Section 4 introduces technology adjustment costs and reports impulse responses. Section 5 discusses and concludes.

2. THE RBC MODEL

2.1. Frictionless Case

Here we present the standard RBC model in the frictionless case, with technology fully observable and immediately implementable into production. The representative agent is assumed to maximize

$$\sum_{t=0}^{\infty} \beta^t \left[\ln C_t + \theta \ln(1 - l_t)\right] \tag{1}$$

by choice of consumption C_t and hours l_t at each time t, subject to the following constraints:

$$C_t + I_t \leq Y_t, \tag{2}$$

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

$$Y_t = X_t^{1-\alpha} K_t^{\alpha} l_t^{1-\alpha}, \tag{4}$$

$$X_t = X_{t-1} e^{\gamma_t} e^{\epsilon_t^x}, \tag{5}$$

and

$$\gamma_t = (1 - \rho_\gamma)\gamma + \rho_\gamma \gamma_{t-1} + \epsilon_t^\gamma. \tag{6}$$

Output Y_t is produced using labor and capital K_t and is divided between consumption and investment I_t. The evolution of capital is governed by equation (3). The household is endowed with one unit of time to divide between labor and leisure. Labor-augmenting technology X_t follows the unit root process described in equation (5). Equation (6) describes the process γ_t, which is the trend process for the X_t series. The white noise disturbance terms ϵ_t^γ and ϵ_t^x are the trend and level shocks, respectively, and are assumed to be orthogonal to one another, with respective variances σ_γ^2 and σ_x^2. The inclusion of two technology shocks here

https://doi.org/10.1017/S1365100512000454 Published online by Cambridge University Press

is intended to capture the distinction between persistent but infrequent changes in the underlying trend of productivity growth (ϵ_t^{γ}) and transitory but frequent fluctuations in productivity growth (ϵ_t^x). The parameter $\beta \in (0, 1)$ represents the discount factor for the household, $\theta > 0$ measures the relative value of leisure to the household, $\delta \in (0, 1)$ measures the rate of capital depreciation, $\alpha \in (0, 1)$ measures the capital share of output, and $\rho_\gamma \in (0, 1)$ measures the persistence of the trend shock. When $\rho_\gamma = 1$, productivity growth $\log(X_t/X_{t-1})$ can be decomposed into a temporary component ϵ_t^x and a time-varying mean process γ_t. The first-order conditions and the detrended, steady state, and log-linearized versions of this model are standard and widely available in the literature and thus not reported here.

In parameterizing the model, we set $\alpha = 0.36$, $\beta = 0.98$, $\delta = 0.10$, $\theta = 3.00$, and $\gamma = \log(1.0226)$, standard values in the macroeconomic literature and identical to the values used by ELW.[4] We set the coefficient ρ_γ governing the duration of the trend shock to 0.95 so that ϵ_t^{γ} approximates a permanent change to the growth rate of productivity, while maintaining the steady state features of the model (i.e., γ_t has a steady state value). Thus our productivity growth specification, like those of ELW, Stock and Watson (1998), and Harvey (1985), consists of a noisy temporary level shock ϵ_t^x and a persistent infrequently changing trend shock ϵ_t^{γ}.

In the frictionless setup, with agents observing both the trend and level components of the shock (and immediately able to implement the entire innovation), the standard RBC results hold. In particular, after a level shock, agents realize that the level of technology has increased and unlike a trend shock will not continue to grow over time. Agents thus respond by immediately increasing investment, hours, and consumption to take full advantage of the shock. Here, we obtain the standard result that a positive innovation in the level of technology causes output, all demand-side components of output, and all factor inputs, including hours, to rise on the impact of the shock.

Following the onset of a trend shock, agents realize that because of the nature of the shock, productivity will continue to grow over time. Agents desire to smooth their consumption profiles, consistent with the permanent income hypothesis, and realizing that their consumption opportunities will be higher in the future, they significantly increase current consumption. This increase in consumption is accomplished by initially reducing investment, resulting in a short-run decline in the capital stock relative to its trend. This increase in consumption also lowers the marginal utility of consumption, which then lowers the shadow value of income, and when coupled with the short-run decline in capital results in an initial decline in hours.

Therefore, as these results suggest, one potential way to generate a fall in hours following a technology shock is to assume an innovation process with a sufficiently persistent autoregressive shock process.[5] In fact, in recent work Linde (2009) assumes an autoregressive shock process in an otherwise standard RBC model and is able to generate negative hours responses following this trend shock.

https://doi.org/10.1017/S1365100512000454 Published online by Cambridge University Press

The recent hours debate in the macro literature, however, centers on the hours responses following a level shock, not a trend shock. In particular, as Gali (1999) and Francis and Ramey (2005) explain, the standard RBC model with just a level shock cannot generate the empirically estimated negative hours responses found in these works. In the current work, we thus focus on the impulse responses to the level shock and in particular on frictions in the transmission of technology shocks that can potentially result in a decline in hours following these Gali-type level shocks.

2.2. Real and Nominal Frictions Case

As highlighted in Francis and Ramey (2005), the standard RBC model can be augmented with habit persistence and investment adjustment costs to produce a fall in the labor input on the impact of a level shock. This occurs because a positive technological innovation increases the wealth of the representative household. The presence of investment adjustment costs limits the amount by which investment can be increased on impact, whereas habit persistence in consumption imposes similar restrictions on the initial increase in consumption. Therefore, following the level shock, and with limits on consumption and investment, agents are forced to use the remainder of their wealth from the technology improvement to consume more leisure, the only remaining normal good in the model. Gali (1999) and Basu et al. (2006) show that including nominal rigidities leads to a fall in hours following a level shock. In this case, incomplete adjustment of real wages following the technology shock makes leisure relatively more attractive to the agent than in the frictionless case.

Several works have called into question the empirical relevance of habit persistence and investment adjustment costs [Dynan (2000) and Groth and Khan (2010)], and there remains disagreement on the extent of nominal rigidities such as the degree of price stickiness [Bils and Klenow (2004) and Nakamura and Steinsson (2008)] in real-world economies. We avoid modeling these nominal and real frictions and instead focus on relaxing the assumptions that technology shocks are immediately observable and implementable on impact. In the next two sections, we are particularly interested in the impact responses of macroeconomic variables to Gali's level technology shock under these different technology specifications.

3. IMPERFECT OBSERVABILITY OF TECHNOLOGY SHOCKS

In this section we analyze the impulse-response functions in the case of learning in an environment of uncertainty about the source of the productivity shock, when the productivity shock consists of a trend and a level component.[6] In each time period the household can observe the composite labor-augmenting technology process X_t, but not its individual components γ_t and ϵ_t^x. Agents must use the following

https://doi.org/10.1017/S1365100512000454 Published online by Cambridge University Press

Kalman-filter-updating equation to disentangle any aggregate shock:

$$\hat{\gamma}_{t+1|t} = \rho_\gamma \hat{\gamma}_{t|t-1} + \kappa(\hat{x}_t - \hat{\gamma}_{t|t-1}). \tag{7}$$

Here, $\hat{\gamma}_{t|t-1}$ represents the forecast of the unobservable time t trend component $\hat{\gamma}_t = \gamma_t - \gamma$ based on observable information as of time $t-1$. The relevant information for agents when they update forecasts is the log-linearized first difference of the observable aggregate productivity series $\hat{x}_t$. The parameter κ represents the optimal steady state Kalman gain, and the key parameter of the Kalman gain is the signal-to-noise ratio $\phi = \left(\sigma_\gamma^2/\sigma_x^2\right)$.[7] This ratio reflects the likelihood of observing a persistent change (trend shock) relative to a temporary change (level shock) in productivity growth. We estimate this parameter to be 0.05 using annual U.S. data from 1960 to 1995 for the growth rate of nonfarm business output per hour and the median unbiased estimation (MUE) approach of Stock and Watson (1998). The small value of this ratio reflects the economic reality that temporary changes in productivity growth are much more likely in the data than changes in the mean of the productivity growth series. This implies a Kalman gain value of $\kappa = 0.15$ when $\rho_\gamma = 0.95$. This is higher than the value of $\kappa = 0.11$ used by ELW. This difference arises from a difference in sample periods, as our sample (1960–1995) is smaller than that of ELW (1948–2004).[8] Notice that the latter period contains two potential structural breaks in the productivity growth series [the slowdown of the mid-1970s and the speedup of the middle to late 1990s; see Bullard and Duffy (2004)]. The MUE approach, which assumes only one structural break in the sample, may not be appropriate for the entire postwar period, and we thus concentrate on the shorter period.

The informational assumptions we use here are also fundamentally different from those of ELW. In particular, ELW assume that agents are unaware of their own learning process and also ignore available information when updating forecasts. These assumptions pin down the initial responses of all model variables to both level and growth shocks, and thus the presence of learning has no effect on the initial response of hours to either of the technology shocks. This is because under these assumptions, agents believe any shock to be a level shock on impact (i.e., $\hat{\gamma}_{1|0} = 0$), and because their expectations for the future do not take into account the fact that they are learning, all responses, whether to growth or level shocks, result in identical responses to the level shock under full information. Thus in the ELW framework, including informational frictions and learning has no effect on the impact responses of technology shocks (see Figures 3 and 4 in ELW).[9]

Because we are primarily interested in the short- and medium-run effects of learning, we move away from the assumptions of ELW. We first assume that agents use all observable information to form forecasts. In particular, because agents are allowed to make decisions by observing the contemporaneous value X_t, we also allow them to use this information when forming forecasts. This means that, in contrast to ELW, agents can begin learning in the period of the shock. Furthermore, unlike ELW, we assume that agents take their underlying learning process into

https://doi.org/10.1017/S1365100512000454 Published online by Cambridge University Press

account when forming expectations about the other endogenous variables. In particular, notice that the log-linearized consumption Euler equation can be written as follows:

$$\hat{c}_t = E_t\hat{c}_{t+1} - [1 - \beta e^{-\gamma}(1-\delta)]E_t\hat{y}_{t+1} + [1 - \beta e^{-\gamma}(1-\delta)]E_t\hat{k}_{t+1} + \beta e^{-\gamma}(1-\delta)\hat{x}_{t+1|t}. \quad (8)$$

Here, $\hat{c}_t$ is the log-linearized detrended consumption series, with the log-linearized forms of output and capital defined analogously. The term $\hat{x}_{t+1|t}$ represents the Kalman filter forecast of $\hat{x}_{t+1}$ based on observation of $\hat{x}_t$ and is given by

$$\hat{x}_{t+1|t} = \hat{\gamma}_{t+1|t}, \quad (9)$$

with $\hat{\gamma}_{t+1|t}$ representing the one-step-ahead forecast of $\hat{\gamma}_{t+1}$ based on observation of $\hat{x}_t$ and calculated using equation (7). When agents in our model choose consumption, they do so according to (9), based on their calculated forecasts of $\hat{x}_{t+1|t}$ as well as their rationally formed expectations of consumption, output, and capital. The rational expectation of consumption is based on (8) and the learning process given by (7) and (9), whereas ELW assume that agents receive only the forecast value of technology without accounting for the learning mechanism when forming rational expectations. Therefore, in our setting, there are two channels for signal extraction learning to affect the impact response of the endogenous variables of the model—(i) immediate observation and learning about the technology process and (ii) inclusion of the learning process in the formation of rational expectations of other future endogenous variables, neither of which exist in the ELW framework.[10]

The impulse responses for the case of signal extraction learning are shown in Figures 1 and 2. The initial magnitudes of the trend and level shocks in this paper have been constrained so that the level of productivity permanently rises in the long run by 1%. Figure 1 presents the impulse responses to an increase in the trend of technology. The trend shock acts to increase the growth rate of productivity, with ρ_γ governing the persistence of the shock. In this case, given the signal-to-noise ratio, agents put considerably higher weight on a productivity disturbance being a level shock as opposed to a trend shock. Believing they are likely facing a level shock, agents smooth consumption less than they would otherwise (i.e., consumption rises by less than in the frictionless case with agents directly observing the change in trend). This results in an increase in investment (the opposite of the frictionless case) and a slightly negative initial hours response (in contrast to the large fall on impact in the frictionless case). Overall, the results here indicate that the trend responses under learning are driven closer to the level responses under the frictionless specification because of the small variance of the trend of productivity relative to the variance of the level component.

The impulse responses to a shock to the level of productivity under learning are reported in Figure 2. Even though agents believe that any shock they face is more

https://doi.org/10.1017/S1365100512000454 Published online by Cambridge University Press

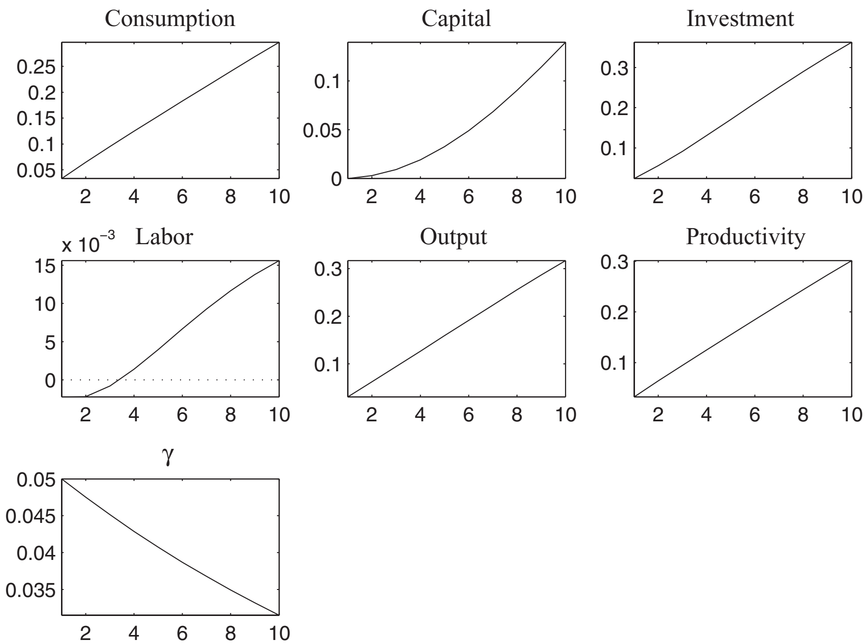

FIGURE 1. Impulse responses to a trend shock in productivity with uncertainty and Kalman filter learning. The solid line represents the impulse responses following a trend shock to productivity in the model in which agents observe a composite productivity shock that they use to create forecasts of the underlying trend and level components and the future composite productivity shock.

likely to be a level shock, the persistence parameter ρ_γ and signal-to-noise ratio are high enough to create a large benefit of consumption smoothing in the unlikely event of a trend shock. Thus, in response to a level shock, agents raise consumption by more under learning than they do in the frictionless case. This causes a rise in investment that is smaller than in the frictionless case and a decline in the response of hours on impact, in contrast to the frictionless case. The muted initial response of investment is supported by several works in the empirical literature, such as Altig et al. (2010) and Francis et al. (2008), and the decline in hours matches the empirical results of Gali (1999) and Francis and Ramey (2005).

Huang et al. (2009) also present results for labor responses to neutral technology shocks, but they assume an adaptive learning framework in which agents learn the rational expectations equilibrium of the model. Huang et al. find that under learning, hours rise by more than in the full information case, in direct contrast to our results. Their results are driven by the backward-looking nature of the adaptive learning framework, in which agents form expectations based on past observations, resulting in a dampening of the wealth effect, with consumption rising less and hours more than under the full information case.

https://doi.org/10.1017/S1365100512000454 Published online by Cambridge University Press

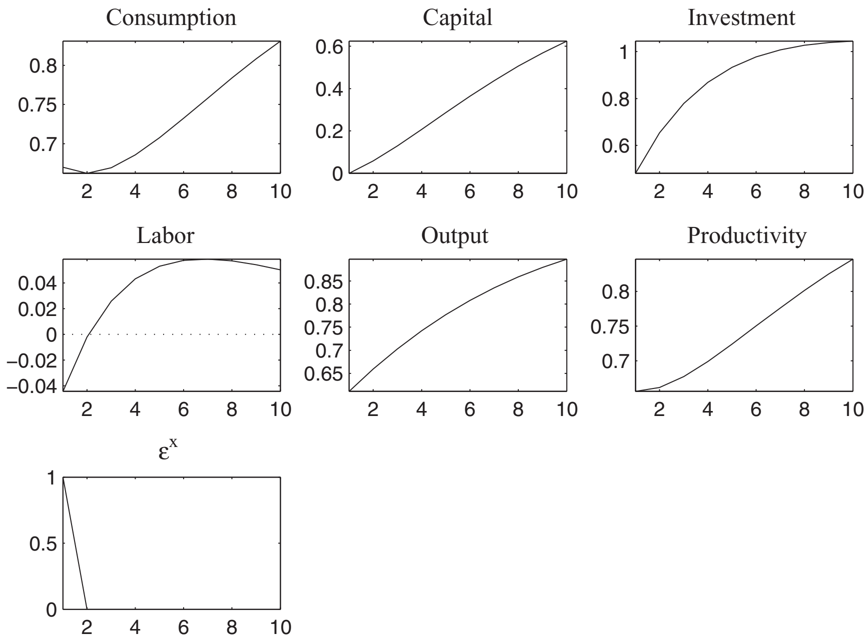

FIGURE 2. Impulse responses to a level shock in productivity with uncertainty and Kalman filter learning. The solid line represents the impulse responses following a level shock to productivity in the model in which agents observe a composite productivity shock that they use to create forecasts of the underlying trend and level components and the future composite productivity shock.

3.1. Learning Parameterization

The response of hours in the learning framework is sensitive to the values of the signal-to-noise ratio (ϕ) and the persistence parameter (ρ_γ). In particular, higher values of either of these parameters will increase the weight agents put on the trend shock in the signal extraction problem and further depress the response of hours to a level shock, whereas lower values will result in less weight put on the trend shock in the learning problem and at low enough values will result in an increase of hours in response to a level shock. When setting the value of ρ_γ (and thus the implied Kalman gain), we face a tradeoff. In particular, our model of productivity shocks and the MUE approach used to estimate ϕ [Stock and Watson (1998)] assume that ρ_γ should be unity, whereas for the steady state of our model to exist ρ_γ should be strictly less than unity.

Notice that our technology process can be written as

$$\hat{x}_t = \hat{\gamma}_t + \epsilon_t^x = \rho_\gamma \hat{\gamma}_{t-1} + \epsilon_t^\gamma + \epsilon_t^x. \tag{10}$$

https://doi.org/10.1017/S1365100512000454 Published online by Cambridge University Press

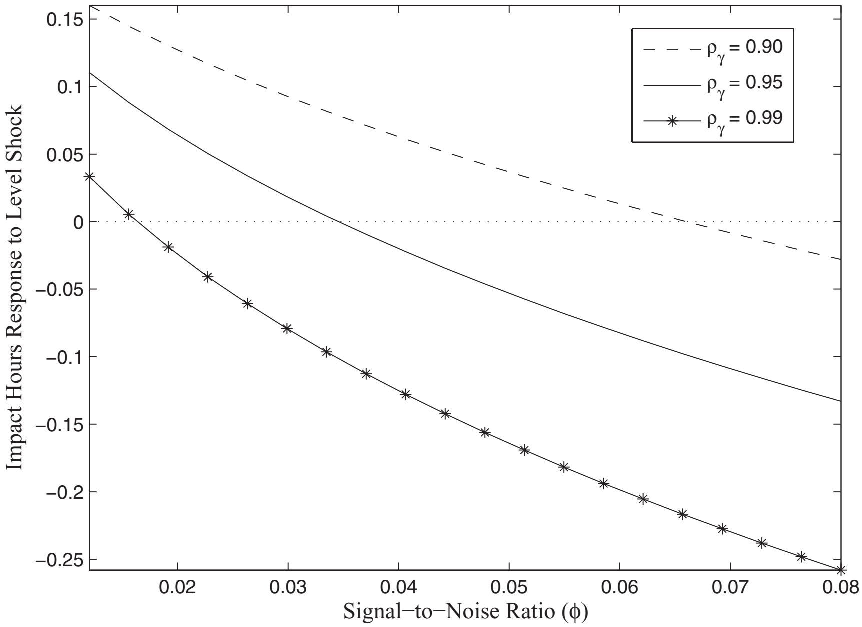

FIGURE 3. Impact hours response to a level shock for alternative learning parameterizations. The dashed line represents the impact hours response for the case in which $\rho_\gamma = 0.90$. The solid line represents the impact hours response for the case in which $\rho_\gamma = 0.95$. The solid starred line represents the impact hours response for the case when $\rho_\gamma = 0.99$.

Estimating this equation is not straightforward, as it involves two error terms and the unobservable term $\hat{\gamma}_t$. One potential approach is to use the Kalman filter, but as Stock and Watson (1998) point out, when σ_γ^2 is small, as it is in this application, the signal-to-noise ratio ϕ can be mistaken for zero; this is known as the pile-up problem. To circumvent this problem, we follow Stock and Watson and ELW, who like us are interested in estimating the signal-to-noise ratio for U.S. productivity growth when productivity growth has a temporary and permanent component, in applying the MUE approach to estimate ϕ. We set ρ_γ at 0.95 in our benchmark calibration to satisfy the steady state assumptions of our model. This value is close enough to unity to maintain consistency with our model of technology shocks and the MUE approach used to estimate ϕ. This value is also consistent with the theoretical and empirical literature on productivity growth [e.g., Stock and Watson, ELW, and Harvey (1985)].

With this in mind, we are interested in exploring how persistent the trend shock needs to be to deliver a negative hours response. Figure 3 presents impact hours responses for alternative values of the persistence parameter ($\rho_\gamma = 0.90, 0.95, 0.99$) and the signal-to-noise ratio ($\phi \in [0.012, 0.08]$). Higher values of ρ_γ result in impact responses of hours that are more negative in magnitude for a given value of ϕ.

https://doi.org/10.1017/S1365100512000454 Published online by Cambridge University Press

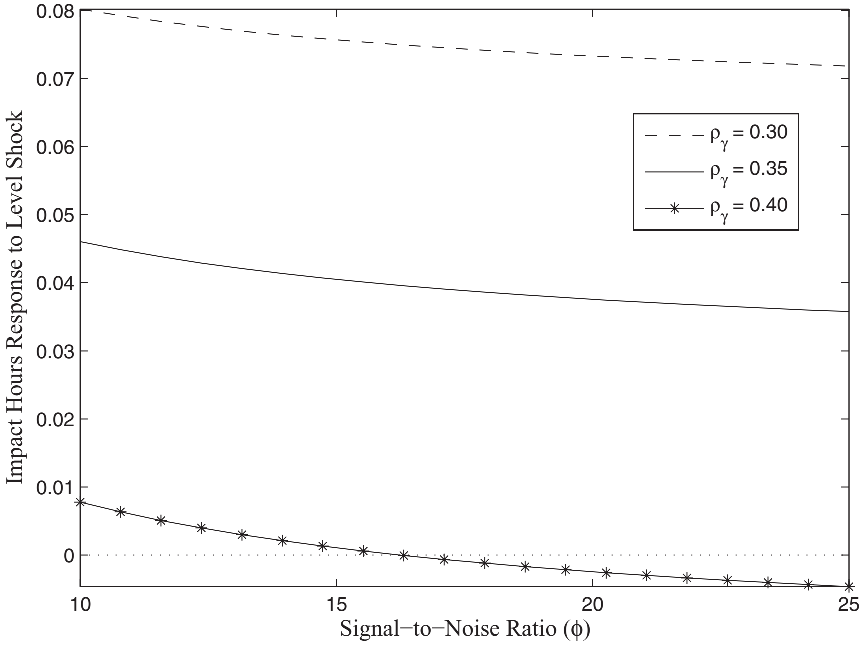

FIGURE 4. Impact hours response to a level shock for alternative learning parameterizations when ρ_γ is small. The dashed line represents the impact hours response for the case in which $\rho_\gamma = 0.30$. The solid line represents the impact hours response for the case in which $\rho_\gamma = 0.35$. The solid starred line represents the impact hours response for the case in which $\rho_\gamma = 0.40$.

Negative hours responses are also possible with values of ρ_γ as low as 0.4, as depicted in Figure 4, but require a value of the signal-to-noise ratio of approximately 16.[11] As ρ_γ approaches zero, the responses of the level and trend shocks converge. At values lower than 0.4, following a trend shock in the full information case, hours rise on impact. Thus, at such low values of ρ_γ, there is no value of the signal-to-noise ratio for which learning can produce an impact decline in hours following a level shock. This is because with both the level and trend shocks resulting in positive impact hours responses under full-information, there is no sign tradeoff for the learning model to exploit.

We also investigate the impulse-response functions arising from alternative parameterizations for the more general utility function $\ln C_t + \frac{\theta(1-l_t)^{1+\chi}}{1+\chi}$. We check values of χ ranging between -1.5 and 2 and find that the impact hours response to a level shock (not shown) remains negative for all values of χ within this range (holding ϕ and ρ_γ at our original values), becoming more negative as χ rises.[12] We find that when $\chi = 0.5$, values of ϕ as low as 0.012 can generate negative hours responses to level shocks with ρ_γ near unity, owing to the additional weight agents now place on leisure.

https://doi.org/10.1017/S1365100512000454 Published online by Cambridge University Press

4. TECHNOLOGY ADJUSTMENT COSTS

In this section, we explore another friction that, like the imperfect observability of technology shocks, limits the immediate effect of technology shocks. Similarly to Rotemberg (2003), we model a technology process that is slow to diffuse through the economy; however, we employ a fundamentally different modeling approach. Whereas Rotemberg assumes a framework of heterogeneous firms, each with a continuum of capital types that absorb innovations at varying rates, we assume a direct cost when implementing innovations that causes agents to slowly incorporate the technology over time.

We assume an underlying *potential* level of technology $\bar{X}_t$ that is slowly adopted into the *installed* level of technology X_t; the latter enters the production function according to (4). We further assume that the agent faces an adjustment cost when implementing changes in potential technology into installed technology. The structure we use is similar to that of Moscoso Boedo (2010), who assumes that firms can only increase stocks of physical capital and skilled workers by incurring a cost. In our case, the equations governing $\bar{X}_t$ and X_t are as follows:

$$\bar{X}_t = \bar{X}_{t-1} e^{\gamma_t} e^{\epsilon_t^x}, \tag{11}$$

$$X_t = \bar{X}_t \left[1 - G\left(\frac{\bar{X}_t}{\bar{X}_{t-1}}, \frac{\bar{X}_{t-1}}{X_{t-1}} \right) \right]. \tag{12}$$

Thus, only a fraction $1 - G(\cdot)$ of the potential technology is implemented, with the following functional form governing adjustment costs:

$$G\left(\frac{\bar{X}_t}{\bar{X}_{t-1}}, \frac{\bar{X}_{t-1}}{X_{t-1}} \right) = e^{\omega_1\left(\frac{\bar{X}_t/\bar{X}_{t-1}}{e^{\gamma}} - 1\right) + \omega_2\left(\frac{\bar{X}_{t-1}}{X_{t-1}} - 1\right)} - 1. \tag{13}$$

We assume that adjustment costs rise when potential technology rises faster than steady state growth ($\bar{X}_t/\bar{X}_{t-1} > e^{\gamma}$) and when installed output falls below potential output ($\bar{X}_{t-1}/X_{t-1} > 1$).[13] Furthermore, in the steady state, $\bar{X}_t = X_t = e^{\gamma}\bar{X}_{t-1}$, and thus adjustment costs in this case are zero. The log-linearized system with adjustment costs is the same as that in the frictionless case, with the exceptions of the log-linearized equations for X_t and $\bar{X}_t$.[14] The log-linearized version of (11) is given by

$$\hat{\bar{x}}_t = \hat{\gamma}_t + \epsilon_t^x. \tag{14}$$

Here, $\hat{\bar{x}}_t = \log(\bar{X}_t/\bar{X}_{t-1}) - \gamma$ represents the log-linear transformation of $\frac{\bar{X}_t}{\bar{X}_{t-1}}$. Substituting (13) into (12) and log-linearizing results in the following equation:

$$\hat{x}_t = (1 - \omega_1)\hat{\bar{x}}_t + \omega_2 \hat{x}_{t-1} + (\omega_1 - \omega_2)\hat{\bar{x}}_{t-1}. \tag{15}$$

Thus the parameters ω_1 and ω_2 represent the fundamental parameters in the adjustment costs model. We investigate two different cases—the first with $\omega_2 = 0$ and the second with $\omega_1 = \omega_2 = \omega$. In the first of these cases, with $\omega_2 = 0$, agents observe an innovation, but because of the adjustment cost, they can only implement

https://doi.org/10.1017/S1365100512000454 Published online by Cambridge University Press

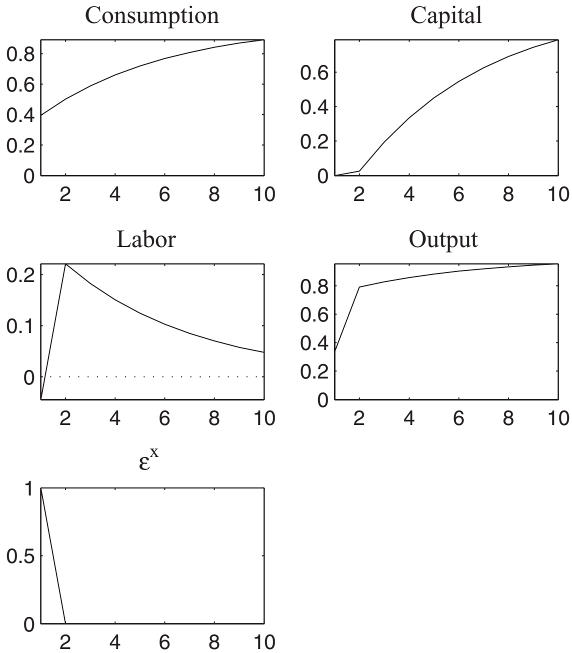

FIGURE 5. Impulse responses to a level shock with technology adjustment costs $\omega_1 = 0.43, \omega_2 = 0$. The solid line represents the impulse responses following a level shock to productivity in which agents directly observe the level shock but face an adjustment cost in implementing the technological innovation.

$1 - \omega_1$ of the innovation on impact and the remainder in the subsequent period. We calibrate the parameter $\omega_1 = 0.43$ in order to match the impact hours response following a level shock under the imperfect observability of shocks in the previous section. This allows us to understand the degree of adjustment costs necessary to match the impact dynamics of a comparable model. When $\omega = 0.43$, only 57% of the innovation is absorbed into production on impact. Figure 5 presents the impulse responses to a level shock in this case. With the costly adjustment of technology, agents respond to the level shock by immediately raising consumption, knowing that productivity will be higher next period when all of the innovation is absorbed. The rise in consumption causes investment to fall below trend and results in a negative labor response (in contrast to the frictionless case) because of the fall in the marginal utility of consumption. The increase in consumption, however, is smaller than in the frictionless case, because the impact effect of the technological change on production is smaller with adjustment costs.[15]

In the second case we set $\omega_1 = \omega_2 = \omega = 0.47$, with the calibration again chosen to match the impact response of hours following a level shock under imperfect observability of shocks.[16] The responses of a level shock in this case are observationally equivalent to those of a growth shock in the frictionless case when

https://doi.org/10.1017/S1365100512000454 Published online by Cambridge University Press

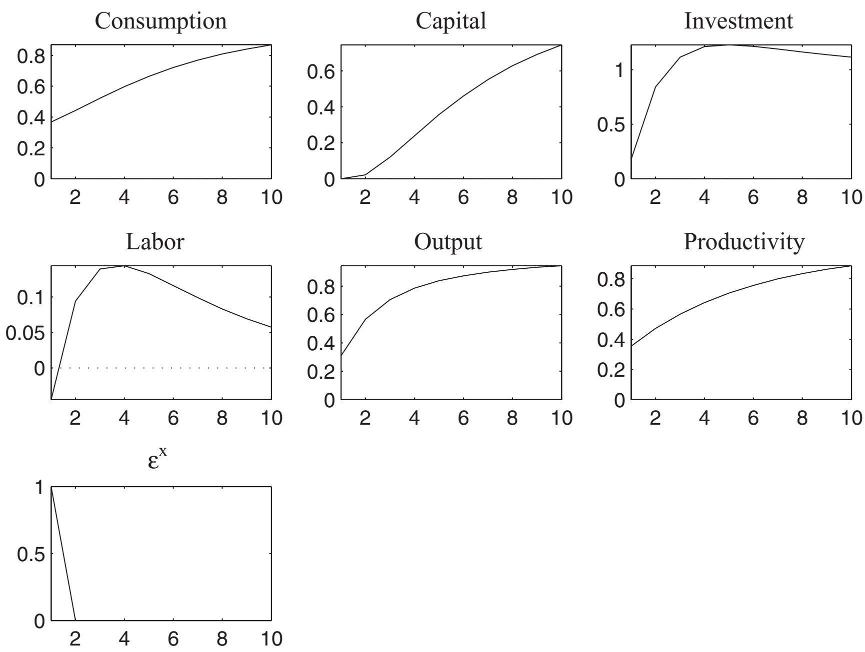

FIGURE 6. Impulse responses to a level shock with technology adjustment costs $\omega_1 = \omega_2 = 0.47$. The solid line represents the impulse responses following a level shock to productivity in which agents directly observe the level shock but face an adjustment cost in implementing the technological innovation.

$\omega = \rho_\gamma$. To see this, notice that abstracting from level shocks, the frictionless technology series can be expressed as

$$\hat{x}_t = \hat{\gamma}_t = \rho_\gamma \hat{\gamma}_{t-1} + \epsilon_t^\gamma = \rho_\gamma \hat{x}_{t-1} + \epsilon_t^\gamma. \tag{16}$$

The adjustment cost technology equation for a level shock in this case can be expressed as

$$\hat{x}_t = \omega \hat{x}_{t-1} + (1 - \omega)\epsilon_t^x. \tag{17}$$

Thus, if $\omega = \rho_\gamma$, a unit shock to ϵ_t^x in the adjustment costs model is identical to a growth shock of $1 - \rho_\gamma$ in the frictionless model. Figure 6 further illustrates this equivalence, showing the impulse responses to the level shock for the adjustment cost case. Agents now respond exactly as they would following a growth shock in the frictionless case by raising consumption on impact (to account for the expected permanent increase in income) and lowering hours worked. A trend shock in an environment of adjustment costs (not shown) causes consumption to rise by much more than output. Here, agents realize that the effect of the shock will increase over time for two reasons: (i) technology takes time to fully diffuse and (ii) a shock to trend also means higher levels of productivity in the future, causing a

https://doi.org/10.1017/S1365100512000454 Published online by Cambridge University Press

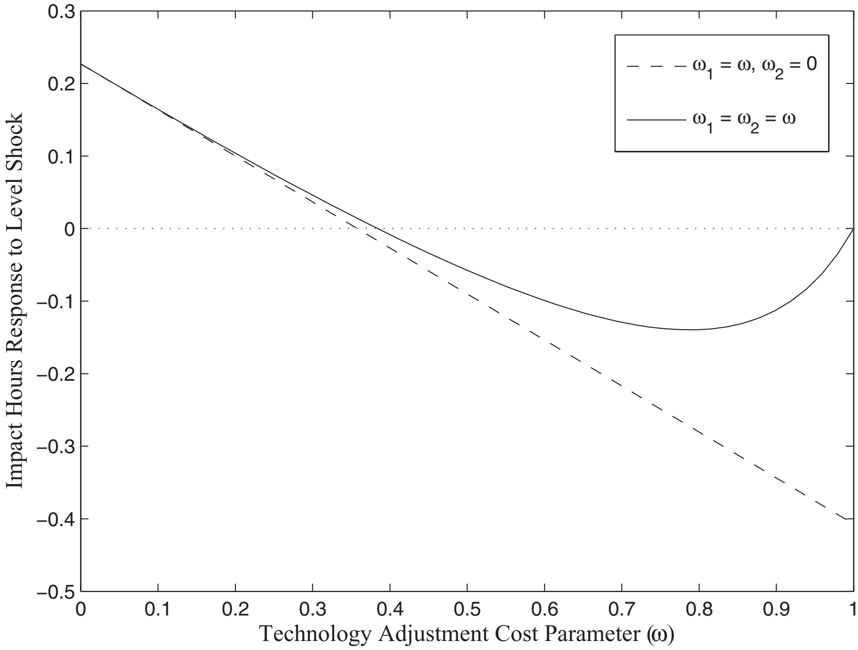

FIGURE 7. Impact hours response to level shock for alternative adjustment cost parameterizations. The dashed line represents the impact hours response for the case in which $\omega_1 = \omega$ and $\omega_2 = 0$. The solid line represents the impact hours response for the case in which $\omega_1 = \omega_2 = \omega$.

large initial rise in consumption. Investment falls relative to its unshocked path to accommodate this increase in consumption. This limits the short-run rise in capital, which, when coupled with the fall in the marginal value of consumption, causes hours to fall, a finding similar to that in the frictionless case.

The degree of rigidity we impose in both cases here is much less restrictive than in many of the cases analyzed by Oliner and Sichel (1994, 2000), Rogers (1995), and Rotemberg (2003), who find examples of technological innovations taking several years before they reach even half of their long-run impact. The parameterizations we investigate all result in over half of the innovation being absorbed on impact. Figure 7 presents the impact response for hours for various values of adjustment costs for both cases we examine here. When adjustment costs are zero, we recover the positive hours response of the frictionless case. As ω rises, the degree of adjustment cost increases, the size of future output relative to current output grows, and the desire to smooth consumption leads to a fall in hours relative to the frictionless case. When both ω_2 and ω_1 are unity, we get no response of hours, because adjustment costs completely prohibit the diffusion of technology.

The results of Section 2 show that with a sufficiently high autoregressive parameter ρ_γ, the RBC model can generate a fall in hours following a growth shock.

https://doi.org/10.1017/S1365100512000454 Published online by Cambridge University Press

The hours debate of Gali (1999), Francis and Ramey (2005), and Christiano et al. (2003), however, revolves around the level shock. The adjustment cost environment presented here provides a structural framework for understanding the link between the level and growth shocks. In particular, by delaying the full long-run value of the level innovation for only one period ($\omega_2 = 0$), we can generate a decline in hours on impact. In this case, the shock has its entire effect spread over two periods and is arguably much closer to the Gali-type level shock than the asymptotically lived autoregressive trend shock. As we increase the duration of the adjustment costs ($\omega_2 > 0$), the level shock now behaves more like the frictionless growth shock and is identical to this growth shock when $\omega_1 = \omega_2 = \rho_\gamma$. We thus move away from the shock of Gali and the hours debate literature as we increase the degree of adjustment costs.[17] Overall, by varying the adjustment cost parameter, we are able to recover the range of investment and hours responses in the literature. For instance, when adjustment costs are higher, we can match Gali (1999), Francis and Ramey (2005), Francis et al. (2008), and Altig et al. (2010).[18] When these costs are lower, we can match the standard RBC results of Christiano et al. (2003).

5. CONCLUSION

In this paper, we have included two empirically supported frameworks of limited transmission of technology shocks in the RBC model that fundamentally alter the short-run predictions of the model. In the first framework, agents are unable to observe the source of the productivity shock and thus unable to distinguish temporary from sustained increases in productivity growth. Instead agents are assumed to observe an aggregate productivity series, which they use to form forecasts of the underlying trend and level components. These agents use signal extraction to update these beliefs about the source of the productivity shock. We find that responses to permanent level shocks in this model differ substantially from those in the standard frictionless case. In particular, if sufficient weight is put on the trend shock in the learning process, which amounts to sufficiently high values of the signal-to-noise ratio and persistence of technology shocks, investment and hours rise by less relative to the frictionless case.

We also examine a second framework that assumes that technological changes diffuse slowly through the economy. Here we introduce an adjustment cost to technology, causing agents to adopt the underlying technology change slowly and implement it into production. We find that this modification is able to produce a variety of impulse responses by simply varying the degree of adjustment costs. When costs are high, again investment and hours rise by less than in the frictionless model. We thus find that including costly adjustment of technology fundamentally alters the predictions of the RBC model and provides a compelling alternative to the use of the real rigidities of habit persistence and investment adjustment costs in the RBC model, for instance in matching the decline in hours found in such works as Gali (1999) and Francis and Ramey (2005).

https://doi.org/10.1017/S1365100512000454 Published online by Cambridge University Press

NOTES

1. As we explain later, ELW are interested in the medium- to long-run effects of uncertainty and learning, and their informational assumptions ensure that including these features has no impact effect on the RBC model's predictions relative to the frictionless case. Rotemberg, on the other hand, mentions the short-run impact of the slow diffusion technology, but his major aim is to explain the long-run trend of U.S. GDP.

2. Gali and Rabanal (2005) provide a comprehensive review of this literature.

3. We adopt the term "learning" used by ELW to describe the process of agents updating beliefs about the underlying source of technology shocks based on observation of the aggregate technology shock. This use of "learning" differs from that of works in the adaptive learning literature [Bullard and Duffy (2004)], which assume that the agent learns the rational expectations equilibrium by updating a perceived law of motion. Both versions of learning, however, involve agents updating beliefs based on their own forecast errors.

4. We parameterize the model to describe U.S. annual data. We extend this parameterization to the model with technology frictions in the next sections.

5. As the value of ρ_γ declines, the responses to the trend shock move closer to those of the level shock, and when $\rho_\gamma = 0$, the trend shock is identical to the level shock.

6. In the impulse-response analysis, we report all responses in terms of the deviations of logged variables from their unshocked paths.

7. The equation of the Kalman gain and the details of the Kalman filtering approach are presented in the appendix to this paper. Van Nieuwerburgh and Veldkamp (2006) present a model in which the signal-to-noise ratio is procyclical, varying with the state of the economy. Agents in their model get more precise signals about productivity during economic booms and more noisy signals about productivity during recessions. We abstract from any such considerations in our simulations, holding the signal-to-noise ratio constant.

8. ELW also assume a modeling environment in which agents cannot fully observe the individual components of aggregate productivity. ELW report a MUE value of $\phi = 0.12$ and a gain value of $\kappa = 0.11$ (assuming $\rho_\gamma = 1$). We find that when $\phi = 0.12$, κ should equal 0.29 in this case, and we verify this with both their formula and our formula. We suspect that the value ELW intend to report is $\phi = 0.012$, resulting in a value of $\kappa = 0.1037$.

9. To see this clearly, note that the initial responses under learning in Figures 3 and 4 of ELW are identical to the responses under full information in Figure 3 for all variables (with slight scaling differences due to a difference in magnitude of the two shocks). Also, these figures reveal that changing the learning parameter λ in ELW's framework does not change the impact full information responses and thus does not change the impact learning responses. Therefore, learning does not affect the impact responses in ELW's model.

10. Kydland and Prescott (1982) also model imperfect observability of technology shocks. In their approach, the aggregate productivity series itself is unobservable to the representative agent when choosing labor and capital. Agents instead must form forecasts of the aggregate series based on a noisy signal of its underlying components. Kydland and Prescott rely on a counterfactually high signal-to-noise ratio ($\phi = 25$) to match empirically observed correlations of macro variables.

11. Higher values of the signal-to-noise ratio imply higher frequency of occurrence of trend shocks. This is a natural assumption at lower values of the persistence parameter, because the long-run effects of trend shocks are smaller and thus presumably more likely to occur in the data.

12. The results shown in Figures 1 and 2 correspond with the case of $\chi = -1$.

13. We abstract away from technological regress and assume only technological progress.

14. Because X_t converges to $\bar{X}_t$, in the steady state $(X_t/\bar{X}_t)^{ss} = 1$.

15. The level shock causes productivity to rise permanently by 1% in both frameworks, but the impact increase to the production function is smaller because of the adjustment costs relative to the frictionless case.

https://doi.org/10.1017/S1365100512000454 Published online by Cambridge University Press

16. The long-term adjustment costs in this case limit the cumulative increase in output, the rise in consumption and thus the fall in marginal utility. Thus, a larger cost on impact is necessary to generate an identical fall in hours when costs only affect the innovation on impact.

17. Gali's structural VAR identification restriction that the technology shock is the only shock able to permanently affect labor productivity does not actually rule out the growth shock as the empirically identified technology shock. The theoretical RBC model discussed in the hours debate, however, assumes that the relevant technology shock is the level shock.

18. Altig et al. (2010) actually find a positive but insignificant impact response of hours. They also find a muted increase in investment immediately following the level shock, matching the predictions of the adjustment cost frameworks presented here.

REFERENCES

Altig, David, Lawrence J. Christiano, Martin Eichenbaum, and Jesper Linde (2011) Firm-specific capital, nominal rigidities and the business cycle. *Review of Economic Dynamics* 14(2), 225–247.

Basu, Susanto, John G. Fernald, and Miles S. Kimball (2006) Are technology improvements contractionary? *American Economic Review* 96(5), 1418–1448.

Bils, Mark and Peter J. Klenow (2004) Some evidence on the importance of sticky prices. *Journal of Political Economy* 112(5), 947–985.

Bullard, James and John Duffy (2004) Learning and Structural Change in Macroeconomic Data. Working paper, Federal Reserve Bank of St. Louis, 1–45.

Christiano, Lawrence J., Martin Eichenbaum, and Robert Vigfusson (2003) What Happens after a Technology Shock? NBER working paper 9819, 1–54.

Dynan, Karen (2000) Habit formation in consumer preferences: Evidence from panel data. *American Economic Review* 90(3), 391–406.

Edge, Rochelle M., Thomas Laubach, and John C. Williams (2007) Learning and shifts in long-run productivity growth. *Journal of Monetary Economics* 54(8), 2421–2438.

Francis, Neville R., Michael T. Owyang, and Jennifer E. Roush (2008) A Flexible Infinite-Horizon Alternative to Long-Run Restrictions with an Application to Technology Shocks. Working paper, Federal Reserve Bank of St. Louis, 1–39.

Francis, Neville R. and Valerie A. Ramey (2005) Is the technology-driven real business cycle hypothesis dead? *Journal of Monetary Economics* 52(8), 1379–1399.

Gali, Jordi (1999) Technology, employment, and the business cycle: Do technology shocks explain aggregate fluctuations? *American Economic Review* 89(1), 249–271.

Gali, Jordi and Pau Rabanal (2005) Technology shocks and aggregate fluctuations: How well does the real business cycle model fit postwar U.S. data? *NBER Macroeconomics Annual* 19, 225–288.

Gil-Alana, Luis Alberiko and Antonio Moreno (2009) Technology shocks and hours worked: A fractional integration perspective. *Macroeconomic Dynamics* 13, 580–604.

Groth, Charlotta and Hashmat Khan (2010) Investment adjustment costs: An empirical assessment. *Journal of Money, Credit and Banking* 42(8), 1469–1494.

Harvey, A.C. (1985) Trends and cycles in macroeconomic time series. *Journal of Business and Economic Statistics* 3(3), 216–227.

Huang, Kevin X.D., Zheng Liu, and Tao Zha (2009) Learning, adaptive expectations and technology shocks. *Economic Journal* 119(536), 377–405.

Kydland, Finn and Edward Prescott (1982) Time to build and aggregate fluctuations. *Econometrica* 50(6), 1345–1370.

Linde, Jesper (2009) The effects of permanent technology shocks on hours: Can the RBC-model fit the VAR evidence? *Journal of Economic Dynamics and Control* 33(3), 597–613.

Moscoso Boedo, Hernan J. (2010) Optimal technology and development. *Journal of Macroeconomics* 32(2), 617–634.

Nakamura, Emi and Jon Steinsson (2008) Five facts about prices: A reevaluation of menu cost models. *Quarterly Journal of Economics* 123(4), 1415–1464.

https://doi.org/10.1017/S1365100512000454 Published online by Cambridge University Press

Oliner, Stephen D. and Daniel E. Sichel (1994) Computers and output growth revisited: How big is the puzzle? *Brookings Papers on Economic Activity* 2, 273–334.
Oliner, Stephen D. and Daniel E. Sichel. (2000) The resurgence of growth in the late 1990s: Is information technology the story? *Journal of Economic Perspectives* 14(4), 3–22.
Pesavento, Elena and Barbara Rossi (2005) Do technology shocks drive hours up or down? A little evidence from an agnostic procedure. *Macroeconomic Dynamics* 9, 478–488.
Quandt, R.E. (1960) Tests of the hypothesis that a linear regression system obeys two separate regimes. *Journal of the American Statistical Association* 55, 324–330.
Rogers, Evertt M. (1995) *Diffusion of Innovations*, 4th ed., New York: Free Press.
Rotemberg, Julio J. (2003) Stochastic technical progress, smooth trends, and nearly distinct business cycles. *American Economic Review* 93(5), 1543–1559.
Stock, James H. and Mark W. Watson (1998) Median unbiased estimation of coefficient variance in a time-varying parameter model. *Journal of the American Statistical Association* 93(2), 349–358.
Van Nieuwerburgh, Stijn and Laura Veldkamp (2006) Learning asymmetries in real business cycles. *Journal of Monetary Economics* 53(4), 753–772.

TECHNICAL APPENDIX

The value of the steady state Kalman gain in equation (7) is given as

$$\kappa = \frac{\rho_\gamma \left[\sqrt{(1-\rho_\gamma^2-\phi)^2+4\phi} - (1-\rho_\gamma^2-\phi)\right]}{2+\sqrt{(1-\rho_\gamma^2-\phi)^2+4\phi} - (1-\rho_\gamma^2-\phi)}. \quad \textbf{(A.1)}$$

We follow the MUE approach of Stock and Watson (1998) and assume that the signal-to-noise ratio ϕ can be expressed as

$$\phi = \frac{\sigma_\gamma^2}{\sigma_x^2} = \left(\frac{\lambda}{T}\right)^2. \quad \textbf{(A.2)}$$

Here, T is the sample size of $\hat{x}_t$ and λ is calculated using the tables in Stock and Watson (1998, p. 353) and the Quandt (1960) likelihood ratio statistic (QLR_T):

$$\mathrm{QLR}_T = \sup_{s\in(s_0,s_1)} F_T(s), \quad \textbf{(A.3)}$$

where

$$F_T(s) = \frac{\mathrm{SSR}_{1,T} - \mathrm{SSR}_{1,[Ts]} - \mathrm{SSR}_{[Ts]+1,T}}{(\mathrm{SSR}_{1,[Ts]} + \mathrm{SSR}_{[Ts]+1,T})/(T-1)}. \quad \textbf{(A.4)}$$

SSR_{t_0,t_1} is the sum of squared residuals from the regression of productivity growth on a constant for the time period from t_0 to t_1, $s \in (0, 1)$, and the potential break dates are given by $[Ts]$, where the operator $[\cdot]$ rounds its argument to the next lower integer.

We use annual values of the growth of nonfarm business output per hour to measure $\hat{x}_t$ (available from the Bureau of Labor Statistics) for the period 1960–1995 in order to estimate λ. We estimate QLR_T to be 8.8834, implying an interpolated value of $\lambda = 7.8197$, $\phi = 0.0472$, and a Kalman gain value [equation (A.1)] of $\kappa = 0.1949$ when we assume that $\hat{\gamma}_t$ follows a random walk ($\rho_\gamma = 1$).

https://doi.org/10.1017/S1365100512000454 Published online by Cambridge University Press