The Effects of Permanent and Transitory Shocks under Imperfect Information

Andrew Foerster and Pierre-Daniel Sarte

The conventional framework in modern macroeconomics assumes that households know all information pertinent to the trade-offs they face and have rational expectations. Within this framework, they optimally adjust their behavior in response to disturbances in the economy in a way that is forward-looking. Thus, the persistence of shocks matters. Moreover, in an economy driven by shocks that differ in their degree of persistence, households can perfectly distinguish between these shocks. Permanent shocks move the economy to a new steady state while transitory shocks have no effect in the long run.

The idea that persistent and transitory shocks have different effects in a setting where agents are forward-looking is well-documented. Blanchard and Quah (1989), for example, use this fact to identify demand shocks as those that only have temporary effects on unemployment and output but supply shocks as those that have permanent effects on output. In a similar vein, King et al. (1991) identify permanent productivity shocks to the common trend in output, consumption, and investment. They find that permanent shocks account for over two-thirds of output fluctuations over the two- to five-year horizon. However, they also show that including nominal variables decreases the explanatory power of balanced-growth shocks on output fluctuations.

The underlying class of models used, for example, by either Blanchard and Quah (1989) or King et al. (1991), is one in which agents

have perfect information and know all relevant shocks. With imperfect information, for instance in a setting where agents may not be able to distinguish between transitory and persistent shocks, the long-run identifying restrictions used by these papers remain, but the effects of shocks differ in the short and medium run.[1] Households have to solve a signal extraction problem in which they try to distinguish between the different types of shocks they face. This signal extraction problem, in turn, can have significant implications for the way in which the economy reacts to shocks at short or medium horizons (see, for example, Edge et al. (2007) and Blanchard et al. (2013)).

In this paper, we use the one-sector neoclassical growth model with permanent and temporary shocks to productivity growth to explore the implications of imperfect information on household behavior. We describe how to solve the model when households cannot distinguish between permanent and transitory shocks in real time. We show that the solution method involves recasting the driving process in terms of the estimates of exogenous states and households' forecast errors implied by the signal extraction problem. We further show that the calculation determining the stability of the model's dynamics is independent from this signal extraction problem. Thus, standard linear rational expectations toolkits may be readily used to solve the imperfect information model. Given this information friction, we then highlight the fact that real variables not only respond to fundamental shocks to productivity growth but also to errors made in forecasting these shocks.

Section 1 presents evidence that shocks to total factor productivity (TFP) growth in the postwar United States are characterized by both a permanent and a transitory component. Section 2 then describes a model environment in which households cannot distinguish between these components and presents the solution method. Section 3 shows how the model's dynamic responses to productivity shocks differ from the complete information case, as well as how inferences about which shock (i.e., temporary or permanent) occurred affect how quickly households adjust their behavior in response to the shock. Section 4 concludes.

---

[1] Imperfect information in macroeconomic models has increasingly been used to explain economic fluctuations. Collard et al. (2009), for example, use a New Keynesian model to show how several different forms of informational frictions featured in recent literature improve their model's explanatory power over the same model under perfect information.

1. AN UNOBSERVED COMPONENT MODEL OF TREND PRODUCTIVITY GROWTH

Let $\Delta \ln z_t$ denote the annualized growth rate ($400 \times$ the difference of the logarithm) of quarterly measurements of aggregate TFP at date t. We allow TFP growth to be driven by both permanent and temporary shocks. Thus, the univariate unobserved component (UC) model expresses $\Delta \ln z_t$ as the sum of a permanent component, τ_t, and a transitory component, ε_t. Specifically,

$$\Delta \ln z_t = \tau_t + \varepsilon_t, \tag{1}$$

where

$$\tau_t = \tau_{t-1} + \eta_t, \tag{2}$$

and

$$\eta_t \sim N\left(0, \sigma_\eta^2\right),$$

$$\varepsilon_t \sim N\left(0, \sigma_\varepsilon^2\right).$$

The UC model is essentially a state-space model with an unobserved state, τ_t, an observation equation, (1), a state transition equation, (2), and unknown variances, σ_η^2 and σ_ε^2. The relative size of these variances, to be estimated below, tell us about the importance of permanent and temporary shocks in driving TFP growth.

Figure 1 shows the data along with the estimated trend. The data are taken from Fernald (2014), who constructs a quarterly, utilization-adjusted measure of TFP for the United States since 1947.

We estimate the empirical model using Bayesian methods and rely on a Gibbs sampler to draw repeatedly from the posterior distribution, a combination of the likelihood and a prior. The use of priors allows us to nudge the posterior estimates to be economically meaningful. In this case, our estimates of σ_η^2 and σ_ε^2 imply that the state, τ_t, evolves at relatively low frequency while still capturing some of the variation in the original series. Table 1 shows our estimates of the variances, σ_η^2 and σ_ε^2. The median estimates, as well as their 68 percent credible intervals, suggest that most of the variation in the data is explained by the transitory component, ε_t. At the same time, the permanent component also has some variation and plays a role since the 68 percent credible interval associated with σ_η^2 does not contain zero.

The fact that a substantial part of the variation in the data is attributed to the transitory component rather than the trend is also clear in Figure 1. Interestingly, the trend in TFP growth declines over the full sample period, 1947–2018, but shows an uptick during the information technology revolution of the mid-1990s to mid-2000s.

Figure 1 U.S. Total Factor Productivity Growth
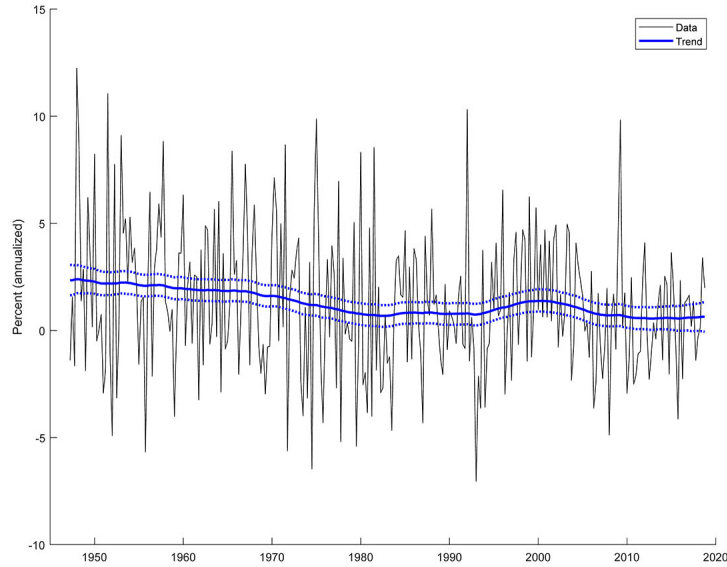
Notes: The figure shows annualized growth rates for quarterly utilization-adjusted TFP from Fernald (2014) and the estimate of the trend from the unobserved component model. Dotted lines indicate the 68 percent credible interval.

Table 1 Variance Estimates

	Median	16th	84th		Median	16th	84th
σ_η^2	0.160	0.125	0.209	σ_ε^2	3.319	3.180	3.461

2.   THE ONE-SECTOR GROWTH MODEL WITH IMPERFECT INFORMATION

To the extent that shocks affecting productivity growth may be permanent or temporary, these have different effects on agents' decisions. For convenience, it is generally assumed in macroeconomic models that the representative household has full information and knows which type of shock is affecting the economy. In practice, however, households would

not be able to distinguish between permanent and temporary shocks from simply observing realized TFP growth.[2] The following model, therefore, lays out how to solve the household's problem in the presence of this informational friction. Relative to the standard setup, the household in this case faces an additional filtering problem in which it must infer estimates of permanent and transitory shocks in deciding how much to consume and save.

We write the household's problem as:

$$\max_{c_t, k_{t+1}} \widetilde{\mathbb{E}}_0 \sum_{t=0}^{\infty} \beta^t \phi \ln c_t \tag{P1}$$

subject to

$$c_t + k_{t+1} - (1-\delta)k_t = \underbrace{A_t k_t^\alpha \ell_t^{1-\alpha}}_{y_t}, \ \alpha \in [0,1], \tag{3}$$

$$k_0 > 0 \text{ given,}$$

where c_t, k_t, and ℓ_t denote consumption, capital, and labor input, respectively. Capital depreciates at rate δ and, in a competitive environment, $1-\alpha$ captures the share of labor payments in value added.

The "tilde" over the expectations operator accounts for the fact that the information set available to households does not distinguish between persistent and transitory shocks. This implies that households face a signal extraction problem. We show below that the implied filtering exercise that needs to be solved does not affect the standard calculation of eigenvalues determining whether the dynamic equilibrium is saddle-path stable or, alternatively, unique. This filtering problem, however, does lead to Markov decision rules that differ from those in the perfect information case insofar as expectations of future shocks now involve estimates of the unobserved exogenous states.

Let

$$z_t = A_t \ell_t^{1-\alpha},$$

where, for ease of presentation, we set the level of exogenous labor, ℓ_t, equal to 1. As in Section 1, we assume that the growth rate of z_t is driven by both persistent shocks, η_t, and transitory shocks, ε_t,

$$\Delta \ln z_t = \tau_t + \varepsilon_t, \tag{4}$$

---

[2] Moreover, there may be other types of shocks, unrelated to fundamentals and that effectively constitute noise, that households are unable to distinguish from fundamental shocks affecting the economy. See, for example, Blanchard et al. (2013).

$$\tau_t = (1 - \rho)g_z + \rho\tau_{t-1} + \eta_t. \tag{5}$$

There is an apparent tension here between the UC model in Section 1, which includes a random walk term, and the model presented in this section, which requires stationarity of the growth rates, $\rho < 1$. We resolve this tension by assuming that the τ term follows a highly persistent but stationary AR(1) model with root local to unity. This process is approximated as a random walk in the empirical description of z_t.

The first-order conditions associated with problem (P1) are

$$\frac{\phi}{c_t} = \lambda_t,$$

$$\lambda_t = \beta\widetilde{\mathbb{E}}_t\lambda_{t+1}\left(\alpha z_{t+1}k_{t+1}^{\alpha-1} + 1 - \delta\right),$$

where the λ_t are Lagrange multipliers. The first-order conditions can be summarized as

$$1 = \beta\widetilde{\mathbb{E}}_t\frac{c_t}{c_{t+1}}\left(\alpha z_{t+1}k_{t+1}^{\alpha-1} + 1 - \delta\right). \tag{6}$$

Balanced Growth

In the model we have just presented, there exists a long-run steady state in which all variables grow at the same constant rate. Specifically, consider the case where all variables are growing at a constant rate along a nonstochastic, steady-state path, $\varepsilon_t = \eta_t = 0$ and $\tau_t = \tau_{t-1}$ $\forall t$. Then, equations (4) and (5) imply that

$$\Delta \ln z_t = g_z$$

so that

$$\frac{z_t}{z_{t-1}} = e^{g_z} \approx 1 + g_z \tag{7}$$

for reasonable growth rates.

Let g_c, g_k, and g_y represent the growth rates of consumption, the capital stock, and GDP, respectively. The resource constraint (3) then implies that $(1 + g_c) = (1 + g_k) = (1 + g_y)$. In addition, the technology implies that $(1 + g_y) = (1 + g_z)(1 + g_k)^\alpha$. Thus, along a balanced growth path,

$$(1 + g_c) = (1 + g_k) = (1 + g_y) = (1 + g_z)^{\frac{1}{1-\alpha}}.$$

Making the Model Stationary

Since variables in the model grow at a constant rate in the long run, solving for the model's transitional dynamics requires that these vari-

ables be detrended in such a way that the detrended variables are constant in the steady state. Let $\mu_t = z_t^{\frac{1}{1-\alpha}}$. Hence, we normalize the model's variables as follows, $\widetilde{c}_t = \frac{c_t}{\mu_t}$, $\widetilde{k}_{t+1} = \frac{k_{t+1}}{\mu_t}$, $\widetilde{y}_t = \frac{y_t}{\mu_t}$, and $\widetilde{z}_t = \frac{z_t}{z_{t-1}}$. Under this normalization, note that detrended GDP is given by $\widetilde{y}_t = \widetilde{z}_t^{\frac{-\alpha}{1-\alpha}} \widetilde{k}_t^{\alpha}$.

In detrended form, the resource constraint (3) becomes

$$\widetilde{c}_t \mu_t + \widetilde{k}_{t+1}\mu_t - (1-\delta)\widetilde{k}_t \mu_{t-1} = \widetilde{z}_t z_{t-1} \widetilde{k}_t^{\alpha} \mu_{t-1}^{\alpha},$$

or

$$\widetilde{c}_t + \widetilde{k}_{t+1} - (1-\delta)\widetilde{z}_t^{\frac{-1}{1-\alpha}} \widetilde{k}_t = \widetilde{z}_t^{\frac{-\alpha}{1-\alpha}} \widetilde{k}_t^{\alpha}. \qquad (8)$$

The Euler equation (6) becomes

$$1 = \beta \widetilde{\mathbb{E}}_t \frac{\widetilde{c}_t \mu_t}{\widetilde{c}_{t+1}\mu_{t+1}} \left(\alpha \widetilde{z}_{t+1} z_t \widetilde{k}_{t+1}^{\alpha-1} \mu_t^{\alpha-1} + 1 - \delta \right),$$

or

$$1 = \beta \widetilde{\mathbb{E}}_t \left(\frac{\widetilde{c}_t}{\widetilde{c}_{t+1}} \right) (\widetilde{z}_{t+1})^{\frac{-1}{1-\alpha}} \left(\alpha \widetilde{z}_{t+1} \widetilde{k}_{t+1}^{\alpha-1} + 1 - \delta \right). \qquad (9)$$

Finally, the process driving productivity in (4) and (5) becomes

$$\ln \widetilde{z}_t = \tau_t + \varepsilon_t, \qquad (10)$$

$$\tau_t = (1-\rho)g_z + \rho \tau_{t-1} + \eta_t. \qquad (11)$$

The Steady State of the Stationary Model

Consider a steady state where $\varepsilon_t = \eta_t = 0 \; \forall t$. Then, $\tau_t = \tau_{t\,1} = g_z$ and the steady state of the stationary model comprises a system of three equations,

$$\ln \widetilde{z} = g_z \implies \widetilde{z} = 1 + g_z,$$

$$1 = \beta \left(1 + g_z\right)^{\frac{-1}{1-\alpha}} \left(\alpha \left(1 + g_z\right) \widetilde{k}^{\alpha-1} + 1 - \delta \right),$$

$$\widetilde{c} + \widetilde{k} \left(1 - (1-\delta)(1 + g_z)^{\frac{-1}{1-\alpha}} \right) = (1 + g_z)^{\frac{-\alpha}{1-\alpha}} \widetilde{k}^{\alpha},$$

in three unknowns, $\widetilde{c}$, $\widetilde{k}$, and $\widetilde{z}$. Denote the solution to this system of equations by $\widetilde{c}^*$, $\widetilde{k}^*$, and $\widetilde{z}^*$, with implied values for $\widetilde{y}^*$.

Dynamics of the System

Given the model's normalized equations such that detrended variables are constant in the steady state, we are now ready to solve for the model's (linearized) transitional dynamics. In what follows, the "hat" notation stands for percent deviation from steady state for the normalized variables (i.e., for some variable x, and with some abuse of notation, $\widehat{x}_t = \log(\widetilde{x}_t/x^*) = \widehat{\widetilde{x}}_t$). In linearized form, our system of equations becomes

$$0 = \widehat{c}_t - \widetilde{\mathbb{E}}_t\widehat{c}_{t+1} - \frac{1}{1-\alpha}\widetilde{\mathbb{E}}_t\widehat{z}_{t+1} + \beta(1+g_z)^{\frac{-\alpha}{1-\alpha}}\alpha\left(\widetilde{k}^*\right)^{\alpha-1}\widetilde{\mathbb{E}}_t\widehat{z}_{t+1}$$

$$+ \beta(1+g_z)^{\frac{-\alpha}{1-\alpha}}\alpha(\alpha-1)\left(\widetilde{k}^*\right)^{\alpha-1}\widetilde{\mathbb{E}}_t\widehat{k}_{t+1},$$

$$s_c\widehat{c}_t + s_k\widehat{k}_{t+1} - (1-\delta)(1+g_z)^{\frac{-1}{1-\alpha}}s_k\widehat{k}_t + \frac{(1-\delta)(1+g_z)^{\frac{-1}{1-\alpha}}s_k}{1-\alpha}\widehat{z}_t = \alpha\widehat{k}_t - \frac{\alpha}{1-\alpha}\widehat{z}_t,$$

where $s_c = \widetilde{c}^*/\widetilde{y}^*$ and $s_k = \widetilde{k}^*/\widetilde{y}^*$. We summarize these equations into the following system,

$$\underbrace{\begin{bmatrix} 1 & \beta\alpha(1-\alpha)(1+g_z)^{\frac{-\alpha}{1-\alpha}}\left(\widetilde{k}^*\right)^{\alpha-1} \\ 0 & s_k \end{bmatrix}}_{A}\widetilde{\mathbb{E}}_t\begin{bmatrix} \widehat{c}_{t+1} \\ \widehat{k}_{t+1} \end{bmatrix}$$

$$= \underbrace{\begin{bmatrix} 1 & 0 \\ -s_c & \alpha+(1-\delta)(1+g_z)^{\frac{-1}{1-\alpha}}s_k \end{bmatrix}}_{B}\begin{bmatrix} \widehat{c}_t \\ \widehat{k}_t \end{bmatrix} \qquad (12)$$

$$+ \underbrace{\begin{bmatrix} 0 \\ \frac{-\left(\alpha+(1-\delta)(1+g_z)^{\frac{-1}{1-\alpha}}s_k\right)}{1-\alpha} \end{bmatrix}}_{C}\widehat{z}_t + \widetilde{\mathbb{E}}_t\underbrace{\begin{bmatrix} \frac{-1}{1-\alpha}+\beta(1+g_z)^{\frac{-\alpha}{1-\alpha}}\alpha\left(\widetilde{k}^*\right)^{\alpha-1} \\ 0 \end{bmatrix}}_{D}\widehat{z}_{t+1}.$$

Aside from the expectations operator, which conditions on an information set that does not allow households to distinguish between permanent and transitory shocks, the system of equations (12) is otherwise expressed in a format identical to that studied in King and Watson (2002). In the next section, we show that despite the information friction, the dynamics of the above system, along with the driving process,

$$\widehat{z}_t = \ln\left(\frac{\widetilde{z}_t}{1+g_z}\right) = \ln(\widetilde{z}_t) - g_z = \underbrace{\tau_t - g_z}_{\widehat{\tau}_t} + \varepsilon_t, \qquad (13)$$

$$\widehat{\tau}_t = \rho\widehat{\tau}_{t-1} + \eta_t, \qquad (14)$$

may nevertheless be expressed in a way that allows it to be solved using standard linear rational expectations solution toolkits.

Solving the Model with Imperfect Information

Since the matrix A in (12) is invertible, the system of equations may alternatively be expressed as

$$\widetilde{\mathbb{E}}_t \begin{bmatrix} \widehat{c}_{t+1} \\ \widehat{k}_{t+1} \end{bmatrix} = W \begin{bmatrix} \widehat{c}_t \\ \widehat{k}_t \end{bmatrix} + \Psi(\mathbf{F})\widetilde{\mathbb{E}}_t \widehat{z}_t, \tag{15}$$

where $W = A^{-1}B$ and $\Psi(\mathbf{F})$ is defined similarly, where $\mathbf{F}$ denotes the forward operator. Moreover, consider a decomposition of W such that $W = V^{-1}\mu V$, where μ is a diagonal matrix containing the eigenvalues of W. Equation (15) may then be expressed as

$$\widetilde{\mathbb{E}}_t d_{t+1} = \begin{bmatrix} \mu_s & 0 \\ 0 & \mu_u \end{bmatrix} d_t + \Psi^d(\mathbf{F})\widetilde{\mathbb{E}}_t \widehat{z}_t,$$

where $d_t = \begin{bmatrix} d_{1,t} \\ d_{2,t} \end{bmatrix} = V \begin{bmatrix} \widehat{c}_t \\ \widehat{k}_t \end{bmatrix}$ and $\Psi^d(\mathbf{F}) = V\Psi(\mathbf{F})$. In other words, the information problem households face does not affect the calculation of eigenvalues that determine whether the dynamic equilibrium is saddle-path stable or, alternatively, unique. Put another way, in this context, the determination of the initial conditions for the costate variable, $\widehat{c}_0$ (or alternatively $\widehat{\lambda}_0$), still follows Blanchard and Kahn (1980).[3] Solving the model with imperfect information, therefore, reduces as a first step to solving a signal extraction problem involving the term $\Psi(\mathbf{F})\widetilde{\mathbb{E}}_t \widehat{z}_t$ in equation (15). Markov decision rules will differ from those of the perfect information case insofar as expectations of future shocks then depend on estimates of unobserved exogenous states, $\widehat{\tau}_t$. We obtain these estimates using a Kalman filter, which offers an algorithmic approach to predicting current-state variables and updating these estimates given new information. Consistent with the Kalman filter, we denote by $\widehat{\tau}_{t|t-1}$ and $\widehat{\tau}_{t|t}$ estimates of $\widehat{\tau}_t$ conditional on information available at date $t-1$ and t, respectively.

Given the driving process, (13) and (14), standard Kalman filtering gives the following prediction equations,

$$\widehat{\tau}_{t|t-1} = \rho\widehat{\tau}_{t-1|t-1},$$

---

[3] For the problem at hand, the equilibrium is indeed unique as the system contains one stable eigenvalue, $\mu_s < 1$, and one unstable eigenvalue, $\mu_u > 1$.

and

$$\widehat{z}_{t|t-1} = \widehat{\tau}_{t|t-1} = \rho\widehat{\tau}_{t-1|t-1},$$

where $\widehat{z}_{t|t-1}$ is the household's forecast of productivity, $\widehat{z}_t$, conditional on all information available up to date $t-1$. Furthermore, given a steady-state Kalman gain matrix, K, the Kalman updating equation is

$$\widehat{\tau}_{t|t} = \widehat{\tau}_{t|t-1} + Ka_t, \tag{16}$$

where a_t is the forecast error,

$$a_t = \widehat{z}_t - \widehat{z}_{t|t-1},$$

observed in period t. The Kalman gain, $0 \leq K \leq 1$, controls the extent to which the household updates its estimate of the state upon realization of its forecast error.

The imperfect information problem introduces estimates of the unobserved exogenous states, $\widehat{\tau}_{t|t}$ and $\widehat{\tau}_{t-1|t-1}$, as state variables (rather than the states themselves) and modifies the driving process faced by the representative household. However, given this modified process, the model may then be solved using standard linear rational expectations solution toolkits. For example, in the formulation of King and Watson (2002), the system of equations to be solved continues to be given by (12), but in place of equations (13) and (14), as the relevant driving process, we instead have that

$$\widehat{z}_t = \begin{bmatrix} 0 & \rho & 1 \end{bmatrix} \begin{bmatrix} \widehat{\tau}_{t|t} \\ \widehat{\tau}_{t-1|t-1} \\ a_t \end{bmatrix},$$

and

$$\begin{bmatrix} \widehat{\tau}_{t|t} \\ \widehat{\tau}_{t-1|t-1} \\ a_t \end{bmatrix} = \begin{bmatrix} \rho & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \widehat{\tau}_{t-1|t-1} \\ \widehat{\tau}_{t-2|t-2} \\ a_{t-1} \end{bmatrix} + \begin{bmatrix} K \\ 0 \\ 1 \end{bmatrix} a_t.$$

It follows, therefore, that the key difference in solving the model with and without perfect information lies in the treatment of the driving process, not in any other aspect of the solution method. Under the conventional perfect information setup, the representative household is assumed to know all shocks affecting the economy so that the driving process, (13) and (14), include the actual exogenous state, $\widehat{\tau}_t$, and disturbances, η_t and ε_t. In contrast, in the imperfect information case, the household is unable to distinguish between shocks so that the relevant driving process includes contemporaneous estimates of the exogenous state, $\widehat{\tau}_{t|t}$ and $\widehat{\tau}_{t-1|t-1}$, as well as the date t forecast error, $a_t = \widehat{z}_t - \widehat{z}_{t|t-1}$, via the Kalman filter.

Intuitively, since the representative household cannot distinguish between shocks under the imperfect information setup, those shocks cannot be driving its decisions. Instead, in each period t, the household is able to observe the current realization of productivity, $\widehat{z}_t$, relative to the forecast it made based on all available information up to $t-1$, $\widehat{z}_{t|t-1}$, and hence knows its forecast error, $a_t = \widehat{z}_t - \widehat{z}_{t|t\ 1}$. The forecast error then, in turn, becomes the object that drives decisions as the household adjusts its behavior, depending on the magnitude and direction of the error, in accordance with the Kalman gain K. Details of the computation of the forecast errors and steady-state Kalman gain, given the properties of the underlying shocks, are provided in the Appendix.

3.   COMPARING THE SOLUTIONS WITH PERFECT AND IMPERFECT INFORMATION

Figure 2 shows the response of GDP growth to a 1 percentage point increase in TFP growth under perfect information. The left panels depict the responses to a permanent shock in TFP growth, while in the right panels, the shock is temporary (in this case, the shock lasts exactly one period). When the shock to TFP growth is permanent, as illustrated in the top left-hand panel, GDP growth initially increases to 3 percent from its prior steady-state value of 2 percent. After the initial burst, growth slows slightly but then increases toward its new steady state. In the case of temporary shocks, GDP growth notably increases at the time the shock materializes but then slowly reverts back to its unchanged steady state (here, 2 percent).

To get a sense of how the GDP growth response (to variations in TFP growth) changes under the imperfect information setup, we consider special cases of the Kalman gain, K. In particular, recall that the Kalman gain governs the degree to which the representative household adjusts to its forecast errors. At one extreme, when $K = 0$, the household never adjusts its estimate of the state, $\widehat{\tau}_{t|t} = \widehat{\tau}_{t\ 1|t\ 1}$ in equation (16), which is to say that it treats every shock as temporary. At the other extreme, $K = 1$, the household reacts fully to every forecast error, $\widehat{\tau}_{t|t} = \widehat{\tau}_{t\ 1|t\ 1} + a_t$ in equation (16), which is to say that it treats every disturbance as permanent. More generally, the optimal value of K is determined by the estimates of the variances of the shocks, meaning the representative household updates forecast errors based on the likelihood of receiving permanent and transitory shocks.

Figure 3 depicts the economy's response to both permanent and temporary increases in TFP growth when the representative household treats every disturbance as permanent. This is for illustrative purposes only since, in practice, estimates of the driving process and its prop-

Figure 2 Impulse Responses under Perfect Information
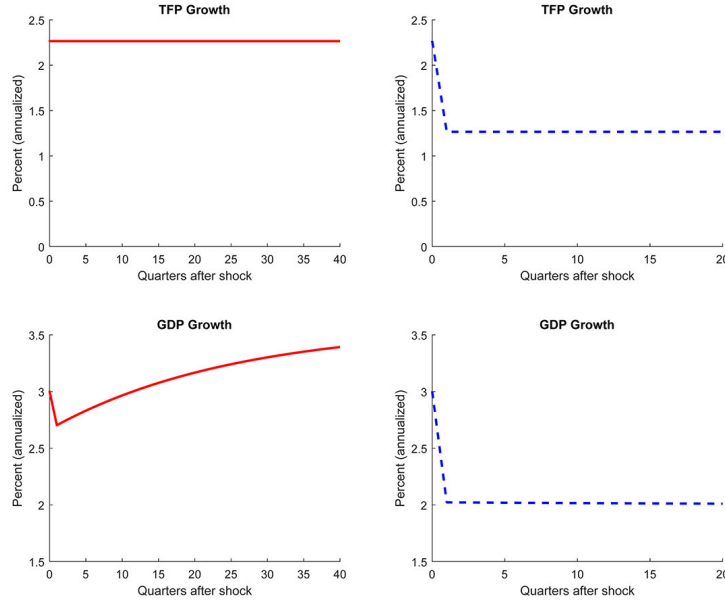
Notes: The figure shows the responses of GDP growth to a 1 percentage point increase in TFP growth The shock to TFP growth in the left panels is permanent, while the shock in the right panels lasts exactly one quarter.

erties will determine the size of the steady-state Kalman gain. This gain, in turn, will lie strictly between 0 and 1 as long as the history of TFP growth reflects both permanent and temporary shocks, made evident by nonzero estimates of their respective variances, as shown in the Appendix.
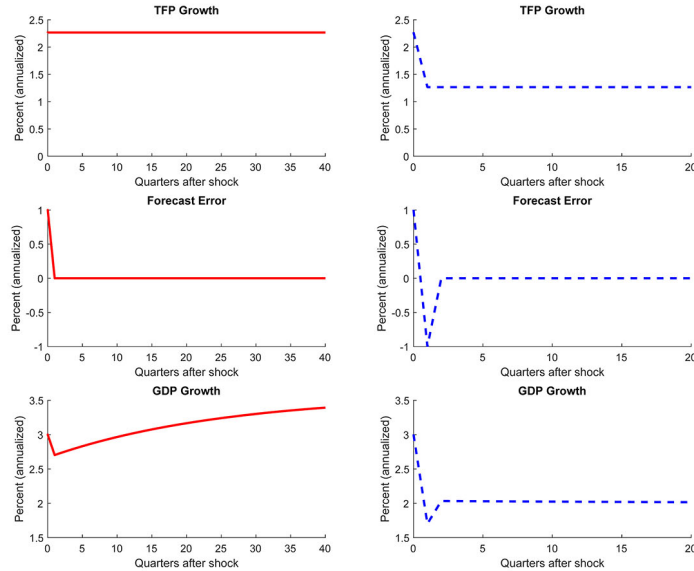
In the extreme case, where K is pinned down to 1, the relevant driving process for the representative household follows

$$\widehat{\tau}_{t|t} = \widehat{\tau}_{t|t-1} + a_t,$$

$$a_t = z_t - \rho\widehat{\tau}_{t-1|t-1},$$

where the true underlying shocks are still given by equations (13) and (14).

**Figure 3 Impulse Responses under Imperfect Information, Where $K = 1$**

Notes: See notes to Figure 2. The middle row shows the date t forecast error, the difference between the current realization of productivity and the household estimate of productivity conditional on information available up to date $t-1$.

As the left panels of Figure 3 illustrate, when the increase in TFP growth is permanent, the household makes a one-period forecast error at the time of the shock, equivalent to the surprise it would have observed under perfect information, but no further forecast errors after that date. This is because the household treats every disturbance as permanent and thus, in this case, accurately assesses the nature of the shock. It follows that the response of GDP growth to a permanent increase in TFP growth, depicted in the bottom left-hand panel of Figure 3, is identical to that in the perfect information case.

When the increase in TFP growth is temporary, shown in the right-hand panels of Figure 3, the household makes a forecast error at the time the shock materializes as before but also in the following period when the shock dies out while it is expected to last indefinitely. In this

**Figure 4 Impulse Responses under Imperfect Information,
        Where $K = 1/2$**
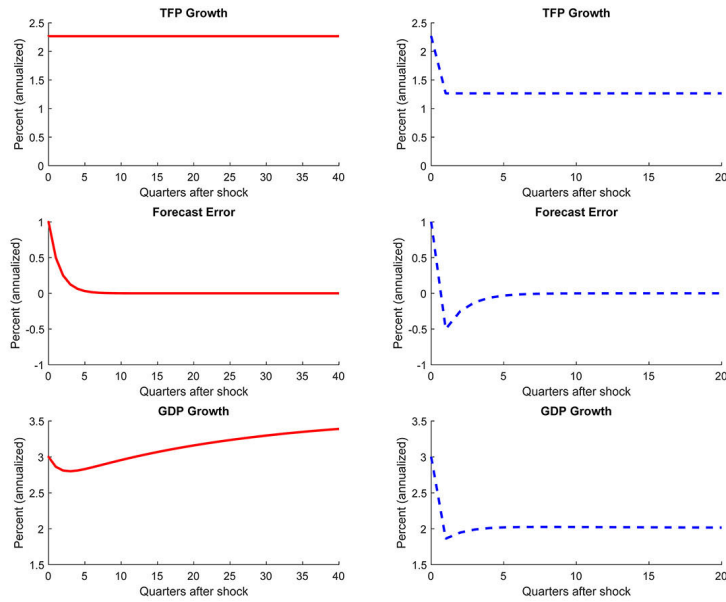
Notes: See notes to Figure 2 and Figure 3.

case, GDP growth increases at the time of the shock as in the perfect information case, shown in the bottom right-hand panel of Figure 3, but unlike the perfect information case, GDP growth falls below its steady state in the following period as the household realizes that the shock was only temporary. Because the household turned out to be overly optimistic, the mistake in this case is associated with a negative forecast error.

The stark nature of the driving process, where disturbances are either i.i.d. or fully permanent, combined with extreme values of the Kalman gain, implies that forecast errors are short-lived. In an intermediate case, where the representative household thinks it equally likely that shocks are permanent as they are temporary, $K = 1/2$, forecast errors persist beyond two periods and the effects of having imperfect information, therefore, are more drawn out relative to the perfect in-

formation environment. This is shown in Figure 4, where the impulse responses for both shocks are now smoother or more persistent despite the one-time nature of the disturbances. This effect would be even more pronounced in an economic environment where, as in Blanchard et al. (2013), both temporary and permanent shocks to the growth rate of TFP are themselves persistent.

Ultimately, the contrast between Figure 2 and Figure 4 highlights the possibility of real economic responses, in this case highlighted by GDP growth, to errors made by households and firms in reading the economic environment rather than actual changes in fundamentals. In such a world, the correlation between economic variables and the processes thought to drive them would be lower than expected, and assessments of the economy would be more challenging to make from a policymaker's perspective.

4.   CONCLUDING REMARKS

While the conventional framework in modern macroeconomics assumes that households have perfect information and rational expectations, the economic environment is complicated enough in practice that they are unlikely to be able to distinguish between different types of shocks in real time. One important such distinction is that between temporary and persistent shocks to fundamentals. In this paper, therefore, we use the one-sector neoclassical growth model, driven by permanent and temporary shocks to productivity growth, to explore key implications of imperfect information on household behavior.

We describe a method for solving the model under this information friction and show that it involves a recasting of the driving process. In particular, whereas the driving process is expressed in terms of the fundamental states and shocks under perfect information, it is instead expressed in terms of estimates of the exogenous states and households' forecast errors under imperfect information. Importantly, we also show that the stability of the model's dynamics is independent of the signal extraction problem from which these estimates and forecast errors are derived. Therefore, standard linear rational expectations toolkits readily apply. We highlight the fact that, under imperfect information, real variables not only respond to fundamental shocks, but also to errors made in forecasting these shocks. In this sense, variations in economic activity do not have to primarily reflect fundamental changes to the economic environment.

REFERENCES

Blanchard, Olivier J., and Charles M. Kahn. 1980. "The Solution of
     Linear Difference Models under Rational Expectations."
     Econometrica 48 (July): 1305–11.

Blanchard, Olivier J., Jean-Paul L'Huillier, and Guido Lorenzoni.
     2013. "News, Noise, and Fluctuations: An Empirical Exploration."
     American Economic Review 103 (December): 3045–70.

Blanchard, Olivier J., and Danny Quah. 1989. "The Dynamic Effects
     of Aggregate Demand and Supply Disturbances." *American
     Economic Review* 79 (September): 655–73.

Collard, Fabrice, Harris Dellas, and Frank Smets. 2009. "Imperfect
     Information and the Business Cycle." *Journal of Monetary
     Economics* 56 (October): S38–S56.

Edge, Rochelle M., Thomas Laubach, and John C. Williams. 2007.
     "Learning and Shifts in Long-Run Productivity Growth." *Journal
     of Monetary Economics* 54 (8): 2421–38.

Fernald, John. 2014. "A Quarterly, Utilization-Adjusted Series on
     Total Factor Productivity." Federal Reserve Bank of San
     Francisco Working Paper Series No. 2012-19.

King, Robert G., Charles I. Plosser, James H. Stock, and Mark W.
     Watson. 1991. "Stochastic Trends and Economic Fluctuations."
     American Economic Review 81 (September): 819–40.

King, Robert G., and Mark W. Watson. 2002. "System Reduction
     and Solution Algorithms for Singular Linear Difference Systems
     under Rational Expectations." *Computational Economics* 20
     (October): 57–86.

APPENDIX

The observation and state equations in the standard Kalman filter problem are respectively,

$$z_t = H\tau_t + u_t, \ u_t \sim N(0, R),$$

and

$$\tau_t = F\tau_{t-1} + v_t, \ v_t \sim N(0, Q),$$

where, in our case, $H = 1$ and $F = \rho$, $u_t = \varepsilon_t$, and $v_t = \eta_t$ (for simplicity, we leave out the "ˆ" notation over the variables). The Kalman filter proceeds in two stages: a prediction stage and an updating stage. During the prediction step, the filter produces estimates of the state, $\tau_{t|t-1}$, and its error covariance, $P_{t|t-1}$, at time t using information available up to time $t-1$. From these estimates, we further obtain a forecast of $z_{t|t\ 1}$ and its associated error covariance, $S_{t|t\ 1}$. In the updating step, we use information from period t to improve the state and error covariance estimates.

The prediction equations are:

- $\tau_{t|t-1} = F\tau_{t-1|t-1}$,

- $P_{t|t-1} = E[(\tau_t - \tau_{t|t-1})(\tau_t - \tau_{t|t-1})'|z^{t-1}] = FP_{t-1|t-1}F' + Q$,

- $z_{t|t-1} = H\tau_{t|t-1}$,

- $S_{t|t-1} = E[(z_t - z_{t|t-1})(z_t - z_{t|t-1})'|z^{t-1}] = HP_{t|t-1}H' + R.$

The updating equations are:

- $\tau_{t|t} = \tau_{t|t-1} + \underbrace{P_{t|t-1}H'S_{t|t-1}^{-1}}_{K_t}(z_t - z_{t|t-1}),$

- $P_{t|t} = P_{t|t-1} - \underbrace{P_{t|t-1}H'S_{t|t-1}^{-1}}_{K_t}HP_{t|t-1}.$

The Kalman gain, K_t, reveals the degree to which the estimates change once z_t is realized. Intuitively, a low gain means the new information will have little impact on the estimates, while a high gain will give the new information more weight. The steady-state Kalman gain is then given by $K = PH'S^{-1}$, where P and S, respectively, solve

$$P = F(P - PH'S^{-1}HP)F' + Q$$

and

$$S = HPH' + R.$$

We obtain Q and R by computing covariances of the shocks from the UC model,

$$Q = \sigma_\eta^2,$$

$$R = \sigma_\varepsilon^2.$$

We then compute the steady-state Kalman gain by iterating on the following equations, where, given iteration $j - 1$,

$$P_1^{(j)} = F P_0^{(j-1)} F' + Q,$$

$$S = H P_1^{(j)} H' + R,$$

$$K^{(j)} = P_1^{(j)} H' S^{-1},$$

and

$$P_0^{(j)} = P_1^{(j)} - K^{(j)} H P_1^{(j)}.$$

We stop the iterations when $\left\| P_0^{(j)} - P_0^{(j-1)} \right\| < 10^{-8}$. The steady-state Kalman gain is then $K^{(j)}$.

With $H = 1$ and $F = \rho \approx 1$, $v_t = \eta_t$ and $u_t = \varepsilon_t$, convergence as $j \to \infty$ implies that

$$P_1^{(\infty)} = P_0^{(\infty)} + \sigma_\eta^2,$$

$$S = P_1^{(\infty)} + \sigma_\varepsilon^2,$$

$$K^{(\infty)} = P_1^{(\infty)} S^{-1},$$

$$P_0^{(\infty)} = P_1^{(\infty)} - K^{(\infty)} P_1^{(\infty)}.$$

Note that if $K^{(\infty)} = 0$, then it must be that $P_0^{(\infty)} = P_1^{(\infty)}$, and hence the equations imply $\sigma_\eta^2 = 0$. On the other hand, if $K^{(\infty)} = 1$, then $P_0^{(\infty)} = 0$, and, therefore, $\sigma_\varepsilon^2 = 0$.

Thus, we have that the households' estimates of the states and their forecast errors evolve according to:

$$a_t = (\rho \widehat{\tau}_{t-1} + v_t + \varepsilon_t) - \rho \widehat{\tau}_{t-1|t-1},$$

$$\widehat{\tau}_t = \rho \widehat{\tau}_{t-1} + \eta_t.$$

$$\begin{bmatrix} \widehat{\tau}_{t|t} \\ a_t \\ \widehat{\tau}_t \end{bmatrix} = \begin{bmatrix} \rho\left(1-K\right) & 0 & K\rho \\ -\rho & 0 & \rho \\ 0 & 0 & \rho \end{bmatrix} \begin{bmatrix} \widehat{\tau}_{t-1|t-1} \\ a_{t-1} \\ \widehat{\tau}_{t-1} \end{bmatrix} + \begin{bmatrix} K & K \\ 1 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} \eta_t \\ \varepsilon_t \end{bmatrix}$$