

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

■ The views expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Richmond, the Federal Reserve Bank of San Francisco, or the Federal Reserve System. We thank Mark Watson for highlighting the imperfect information problem discussed herein and suggesting the solution we describe. We also thank Reiko Laski for outstanding research assistance.

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,

$$\Delta \ln z_t = \tau_t + \varepsilon_t, \quad (1)$$

where

$$\tau_t = \tau_{t-1} + \eta_t, \quad (2)$$

and

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

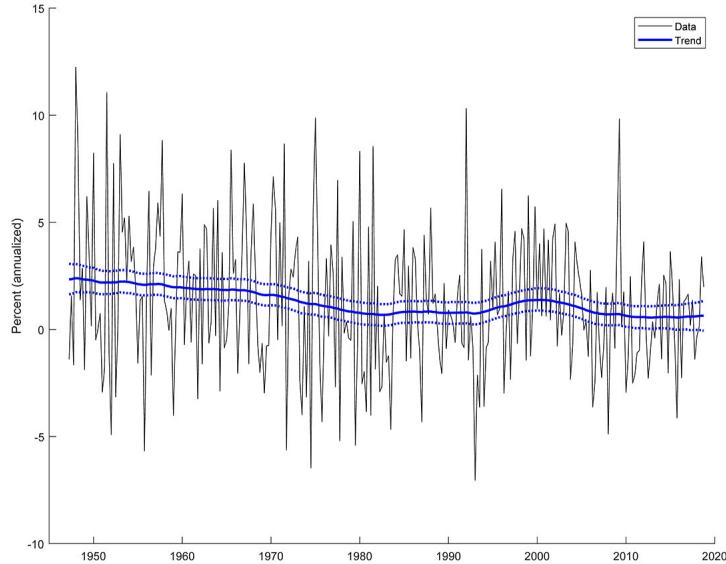
$$\varepsilon_t \sim N(0, \sigma_\varepsilon^2).$$

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:

$$\max_{c_t, k_{t+1}} \tilde{\mathbb{E}}_0 \sum_{t=0}^{\infty} \beta^t \phi \ln c_t \quad (\text{P1})$$

subject to

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

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

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 ,

$$\Delta \ln z_t = \tau_t + \varepsilon_t, \quad (4)$$

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).

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

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

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

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

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

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

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

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

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, $\tilde{c}_t = \frac{c_t}{\mu_t}$, $\tilde{k}_{t+1} = \frac{k_{t+1}}{\mu_t}$, $\tilde{y}_t = \frac{y_t}{\mu_t}$, and $\tilde{z}_t = \frac{z_t}{z_{t-1}}$. Under this normalization, note that detrended GDP is given by $\tilde{y}_t = \tilde{z}_t^{\frac{-\alpha}{1-\alpha}} \tilde{k}_t^\alpha$.

In detrended form, the resource constraint (3) becomes

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

or

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

The Euler equation (6) becomes

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

or

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

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

$$\ln \tilde{z}_t = \tau_t + \varepsilon_t, \quad (10)$$

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

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,

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

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

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

in three unknowns, \tilde{c} , \tilde{k} , and \tilde{z} . Denote the solution to this system of equations by \tilde{c}^* , \tilde{k}^* , and \tilde{z}^* , with implied values for \tilde{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, $\hat{x}_t = \log(\tilde{x}_t/x^*) = \tilde{\hat{x}}_t$). In linearized form, our system of equations becomes

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

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

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

$$\underbrace{\begin{bmatrix} 1 & \beta\alpha(1-\alpha)(1+g_z)^{\frac{-\alpha}{1-\alpha}} \left(\tilde{k}^* \right)^{\alpha-1} \\ 0 & s_k \end{bmatrix}}_A \tilde{\mathbb{E}}_t \begin{bmatrix} \hat{c}_{t+1} \\ \hat{k}_{t+1} \end{bmatrix} \\ = \underbrace{\begin{bmatrix} 1 & 0 \\ -s_c & \alpha + (1-\delta)(1+g_z)^{\frac{-1}{1-\alpha}} s_k \end{bmatrix}}_B \begin{bmatrix} \hat{c}_t \\ \hat{k}_t \end{bmatrix} \quad (12) \\ + \underbrace{\begin{bmatrix} 0 \\ -\left(\alpha + (1-\delta)(1+g_z)^{\frac{-1}{1-\alpha}} s_k \right) \end{bmatrix}}_{C} \tilde{z}_t + \underbrace{\tilde{\mathbb{E}}_t \begin{bmatrix} \frac{-1}{1-\alpha} + \beta(1+g_z)^{\frac{-\alpha}{1-\alpha}} \alpha \left(\tilde{k}^* \right)^{\alpha-1} \\ 0 \end{bmatrix}}_D \tilde{z}_{t+1}.$$

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,

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

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

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

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

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

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

where $d_t = \begin{bmatrix} d_{1,t} \\ d_{2,t} \end{bmatrix} = V \begin{bmatrix} \hat{c}_t \\ \hat{k}_t \end{bmatrix}$ 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, \hat{c}_0 (or alternatively $\hat{\lambda}_0$), 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 $\Psi(\mathbf{F})\tilde{\mathbb{E}}_t \hat{z}_t$ 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, $\hat{\tau}_t$. 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 $\hat{\tau}_{t|t-1}$ and $\hat{\tau}_{t|t}$ estimates of $\hat{\tau}_t$ 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,

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

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

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

where $\hat{z}_{t|t-1}$ is the household's forecast of productivity, \hat{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

$$\hat{\tau}_{t|t} = \hat{\tau}_{t|t-1} + K a_t, \quad (16)$$

where a_t is the forecast error,

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

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, $\hat{\tau}_{t|t}$ and $\hat{\tau}_{t-1|t-1}$, 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

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

and

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

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, $\hat{\tau}_t$, 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, $\hat{\tau}_{t|t}$ and $\hat{\tau}_{t-1|t-1}$, as well as the date t forecast error, $a_t = \hat{z}_t - \hat{z}_{t|t-1}$, via the Kalman filter.

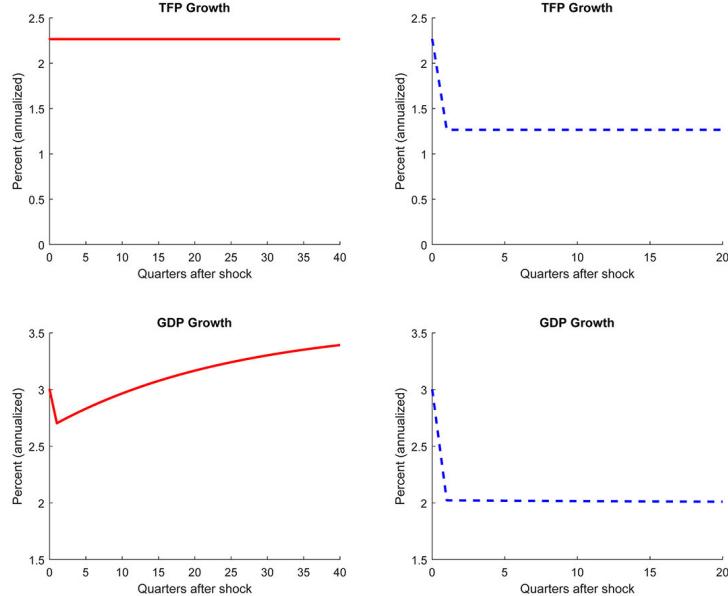
Intuitively, since the representative household cannot distinguish between shocks under the imperfect information setup, those shocks cannot be driving its decisions. Instead, in each period t , the household is able to observe the current realization of productivity, \hat{z}_t , relative to the forecast it made based on all available information up to $t-1$, $\hat{z}_{t|t-1}$, and hence knows its forecast error, $a_t = \hat{z}_t - \hat{z}_{t|t-1}$. 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, $\hat{\tau}_{t|t} = \hat{\tau}_{t-1|t-1}$ 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, $\hat{\tau}_{t|t} = \hat{\tau}_{t-1|t-1} + a_t$ 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.

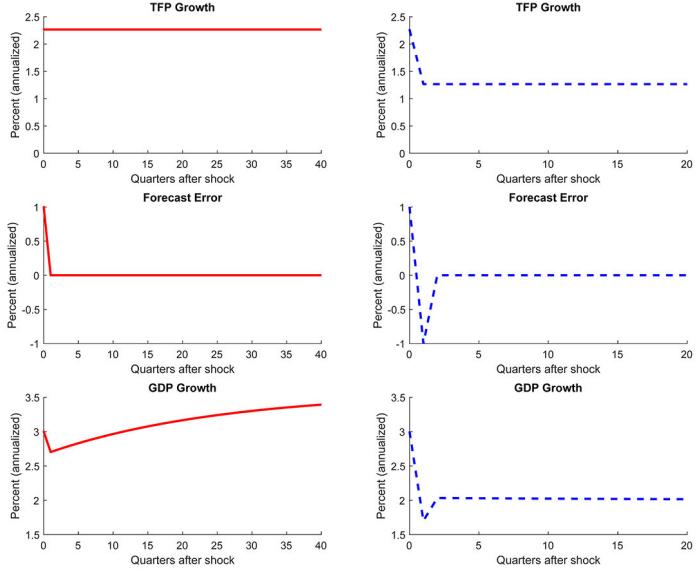
In the extreme case, where K is pinned down to 1, the relevant driving process for the representative household follows

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

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

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

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

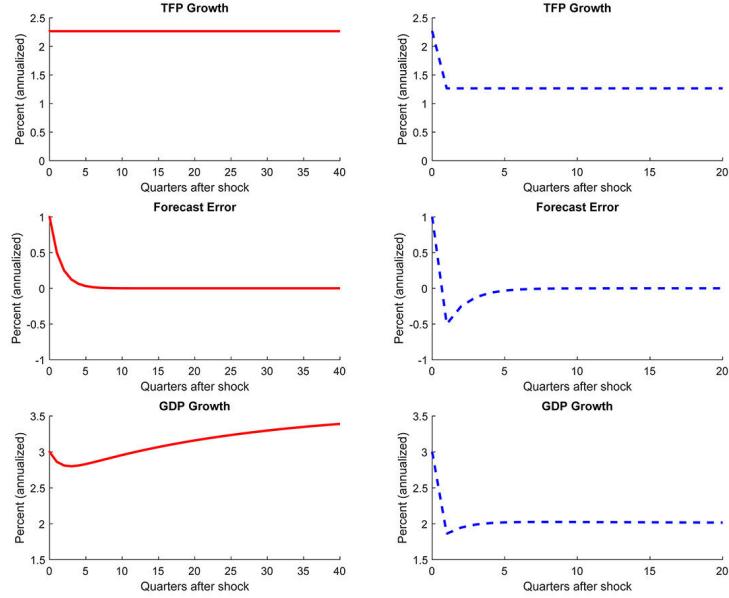


Notes: See notes to Figure 2. The middle row shows the date t forecast error, the difference between the current realization of productivity and the household estimate of productivity conditional on information available up to date $t - 1$.

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,

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

and

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

where, in our case, $H = 1$ and $F = \rho$, $u_t = \varepsilon_t$, and $v_t = \eta_t$ (for simplicity, we leave out the “ $\hat{}$ ” 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:

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

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 $j - 1$,

$$P_1^{(j)} = FP_0^{(j-1)}F' + Q,$$

$$S = HP_1^{(j)}H' + R,$$

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

and

$$P_0^{(j)} = P_1^{(j)} - K^{(j)}HP_1^{(j)}.$$

We stop the iterations when $\|P_0^{(j)} - P_0^{(j-1)}\| < 10^{-8}$. 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 $j \rightarrow \infty$ 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:

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

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

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

Computing Dynamic Heterogeneous-Agent Economies: Tracking the Distribution

Grey Gordon

The evolution of prices in dynamic heterogeneous-agent economies typically depends on the state of every agent, thereby requiring that a distribution be a state variable. The contribution of this paper is to introduce a method for computing equilibrium in these models by including an entire distribution if finite-dimensional—or a fine approximation of it if infinite-dimensional—as a state variable. The insight of Krusell and Smith (1997, 1998) is that this approach is not necessary if a model features *quasi-aggregation*, the condition where prices can be accurately forecasted using just a few state variables. However, not all economies feature quasi-aggregation, and I show that the method presented in this paper is capable of accurately computing equilibrium in at least one of these: Huffman’s (1987) overlapping-generations (OLG) economy paired with an extreme calibration used in Krueger and Kubler (2004). Even when quasi-aggregation obtains, including a distribution as a state variable may be desirable from a conceptual or purely pragmatic perspective. I show that the method accurately computes equilibrium in an economy of this type also: a version of Krusell and Smith’s (1998) (KS) economy where households face

I would like to thank Dirk Krueger for many helpful conversations on this project. Also, thanks to Jesús Fernández-Villaverde and Aaron Hedlund for helpful comments. Comments and questions should be sent to greygordon@gmail.com. Computer code for this paper can be found at sites.google.com/site/greygordon. This article is a slightly revised version of a chapter of my 2012 dissertation, and the first draft was from spring of 2011. Any errors are my own. The views expressed do not necessarily represent those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

occasionally binding constraints. The method is feasible for these two economies with equilibrium for both computed in just a few minutes in Matlab.¹ As discussed momentarily, Smolyak's (1963) sparse-grid interpolation algorithm introduced to economics by Krueger and Kubler (2004) makes this possible.

Smolyak's algorithm is a projection method that uses collocation on a very sparse grid.² The algorithm approximates a function by interpolating its value at a set of predefined gridpoints (collocation points) using weighted sums of polynomials. The fineness of the approximation is controlled by using different "levels of approximation." For the lowest level of approximation, which is the only one used in this paper, the number of gridpoints grows only *linearly* in dimension. More specifically, given a function of dimension d , Smolyak's algorithm gives $2d+1$ points that the function must be evaluated at in order to approximate it. In contrast, linear interpolation or any tensor-product interpolation method would require at least 2^d points. To see the difference this makes, consider that the distributions (and hence state spaces) used in this paper have up to 200 elements: to approximate a function of this dimension using linear interpolation would require more than 10^{48} trillion function evaluations compared with only 401 for Smolyak interpolation. Not only is the Smolyak algorithm computationally efficient, but Barthelmann, Novack, and Ritter (2000) prove the approximation has nearly optimal error bounds for smooth functions.³ The disadvantage of Smolyak's algorithm is that approximations to nonsmooth functions may be quite poor.

The application of Smolyak's algorithm presented in this paper leverages the strengths of the Smolyak algorithm, computational efficiency and accuracy for smooth functions, while avoiding its main weakness, poor approximation of nonsmooth functions. While many heterogeneous-agent models feature policy functions that are kinked in individual wealth or income and hence are not smooth, as long as they are smooth in the aggregate state, they can be approximated well by the Smolyak algorithm. This is accomplished through *indexing* policy functions by individual states and constructing a Smolyak approxima-

¹ Carroll's (2006) endogenous gridpoints method is used to solve the household problem. Value function iteration is also feasible, just slower and less accurate than Carroll's Euler-equation-based method.

² For an excellent introduction to projection methods, including projection methods that use collocation, the reader is referred to Judd (1998). For an accessible exposition of the Smolyak algorithm, the reader is referred to Krueger, Kubler, and Malin (2011). Section 1 of this paper also discusses the algorithm.

³ The error bounds depend not only on the number of times a function is continuously differentiable, but also on how little curvature a function has. The term "smooth" is used atypically here to cover both of these properties.

tion to *each* indexed policy function. For example, given a capital policy function $k'(k, \mu)$, where k is a household's current capital holdings and μ is a distribution of holdings across households, the Smolyak approximation to $k'(k, \mu)$ would likely be poor if k' were kinked in k . However, if k' is fairly smooth in μ for fixed k , then the indexed policy function $k'_k(\mu)$ could be accurately approximated using Smolyak interpolation.⁴ By indexing policies in this way, the resulting Smolyak approximations may be accurate even if the policies are “not smooth.” I refer to this approach as the Smolyak method.

Recognizing the computational challenge posed by solving a model where the distribution was part of the state space, Krusell and Smith (1997, 1998) found a clever way of circumventing it. By replacing the distribution with a few aggregate statistics and assuming that households perceive prices to be functions of only these statistics (and the aggregate shocks), a law of motion for them enables households to predict current and future prices and hence to optimize. Given optimal household policies, it is then possible to check the accuracy of the perceived prices and law of motion through simulation. If a small set of statistics can be found that results in an accurate law of motion and accurate price forecasts, then quasi-aggregation is said to obtain, in which case it is hoped that the computed bounded-rationality equilibrium is close to the equilibrium of the full-rationality model. Equilibrium has typically been computed by guessing on a law of motion, solving the household problem, simulating the economy, and updating the law of motion using data from the simulation. I refer to this approach as the KS method.

The Smolyak method has three advantages over the KS method. First, the Smolyak method does not rely on quasi-aggregation, an equilibrium property that is not known *a priori*. Second, there is no need to simulate the economy to compute the solution. Not only can this result in substantial time savings, but it also means the computed solution is not a random variable. Third, for certain classes of models, namely those where the distribution is finite-dimensional, the solution can be regarded as a full rational-expectations equilibrium.⁵

While the Smolyak method has several advantages over the KS method, this paper is in no way a critique of it. When the KS method works, that is when quasi-aggregation obtains, it is extremely powerful.

⁴ Of course, approximating each function requires that the number of values k takes on is finite. Section 2 discusses how to apply the Smolyak algorithm even if k takes on an infinite number of values.

⁵ Even if the distribution has infinite dimension but is represented by a finite number of elements, as is the case when using the method of Ríos-Rull (1997), the solution may be regarded as an approximate full rational-expectations equilibrium.

Indeed, whereas the Smolyak method has gridpoints growing linearly in the dimension of the underlying state space, the KS method's gridpoints need not grow at all! Moreover, quasi-aggregation has obtained in a wide variety of models. The KS method is robust, conceptually simple, and easy to program, and so is a powerful tool.

Yet there are cases where the KS method does not work well. As already mentioned, I present one such OLG economy that has a known solution due to Huffman (1987) and calibration due to Krueger and Kubler (2004) (KK).⁶ In the most extreme case, where there are only three generations, a linear forecasting rule for the aggregate capital stock results in an R^2 statistic of 0.676 and a maximum percent error of 3.17. In contrast, the Smolyak method's forecast performs very well resulting in an R^2 statistic of 0.999 and a maximum percent error of 0.07. While the performance of the KS method could be improved by adding more moments, here that would mean the distribution could be completely summarized because only the oldest two generations have positive capital holdings. Moreover, the Smolyak method is already faster in this case than the KS method. The Euler-equation errors, a measure of household optimization error, are similar across the two models with the KS method better in terms of maximum errors and the Smolyak method better in terms of average errors.

Even when the KS method does work well, the Smolyak method may achieve a similar level of accuracy and possibly be even faster to run. As for accuracy, in the modified Krusell and Smith (1998) economy studied, I find the computed equilibria are virtually identical across methods, both in terms of optimization errors and simulated aggregate moments: the Euler-equation errors for the Smolyak method (-2.22 maximum and -4.97 average) are slightly smaller than those of the KS method (-1.89 maximum and -4.79 average), and the simulated capital series are at most 0.28 percent apart. As for speed, the KS method will typically be faster. However, this depends on the cost of household optimization relative to the cost of simulation. In the KS economy, the former cost dominates and the Smolyak method takes 6.4 minutes compared with 3.3 minutes for the KS method. However, in the OLG economy, the latter cost dominates when the number of generations is less than fifty, making the Smolyak method faster.⁷

⁶ A more realistic example where quasi-aggregation does not obtain comes from the equity-premium literature. Chien, Cole, and Lustig (2009) find that including even five moments of the distribution results in an R^2 of only 0.50 to 0.75 when forecasting the pricing kernel (cf. Table 2 of their paper).

⁷ Parallelization of the household problem, which was not used, could shift this balance substantially in favor of the Smolyak method. Value function iteration in particular is known to parallelize well (see Aldrich et al. 2011).

The Smolyak method may also be more intuitive than the KS method for certain classes of models. For instance, dynamic models of voting do not typically have a natural “sufficient statistic” representation.⁸ Instead, researchers have typically used coarse histograms to summarize the distribution, as is done in Krusell and Ríos-Rull (1999). While there is usually some way of approximating the aggregate state space with just a few statistics, the Smolyak method provides an alternative that may be both feasible and accurate, as it is in two non-trivial economies.

Since the seminal papers of Krusell and Smith (1997, 1998), many methods have been developed to solve dynamic heterogeneous-agent models. For a thorough review of current methods, the reader is referred to the January 2010 special issue of the *Journal of Economic Dynamics and Control* “Computational Suite of Models with Heterogeneous Agents: Incomplete Markets and Aggregate Uncertainty” (volume 34, issue 1), the comparison paper Terry (2017), and Childers (2018). I highlight a few of these that are most closely related to the Smolyak method. The first approach is the “backward induction” method of Reiter (2010). As in the KS method, the aggregate state space is a small set of statistics. A distinguishing aspect of his approach is that these statistics map into a specific “proxy” distribution that agents use to make forecasts. A qualitatively similar approach is due to Algan, Allais, and den Haan (2010): like Reiter (2010), they link a few moments to a specific distribution but do so in a different way. A third approach is due to den Haan and Rendahl (2010). Roughly speaking, they construct an approximation to the true policy function that results in exact aggregation. While these methods have many merits, they place special structure on either the distribution (in the case of Reiter 2010 and Algan et al. 2010) or on the policy functions (in the case of den Haan and Rendahl 2010) to construct a law of motion. The Smolyak method does neither of these.

In the same OLG economy studied here, Krueger and Kubler (2004) use Smolyak interpolation to solve for a full rational-expectations equilibrium with a large state space. Their approach differs from the one I present because they make no distinction between individual and aggregate states.⁹ This means that the same approximation must be used

⁸ An exception to this is Azzimonti, de Francisco, and Krusell (2006), where the mean and median of wealth are proven to be sufficient statistics.

⁹ For KK, the state space is a distribution of capital holdings, \mathbf{k} , and a particular generation’s capital holdings is just “read off” this distribution. However, by expanding the state space to (k, \mathbf{k}) , where k is a particular generation’s capital holdings, it is possible to use a fine approximation for the individual state k and a coarse approximation for the aggregate state \mathbf{k} . Essentially this separates the role of prices, which are determined from \mathbf{k} , from the role of individual wealth, which is proportional to k .

for both, and so a higher-order Smolyak approximation (one that grows quadratically or cubically in dimension) is required to achieve sufficient accuracy. Consequently, their method can only handle economies with relatively small state-space dimensions.¹⁰

Smolyak interpolation is not the only method that could be used to include the distribution as a state variable. In particular, the recently developed cluster-grid projection method of Judd, Maliar, and Maliar (2010) is capable of handling problems of very large dimension. Relative to Smolyak interpolation, their method provides greater flexibility in terms of where gridpoints are placed and which basis functions are used. This comes at the cost of using weakly more gridpoints else equal.¹¹ Whether the cluster-grid projection method provides a feasible and accurate alternative to Smolyak interpolation in this context is left as a question for future research.

This paper is organized as follows. Section 1 describes the Smolyak algorithm. Section 2 discusses the Smolyak method, i.e., the application of the Smolyak algorithm used to approximate equilibrium. Section 3 presents the OLG and KS models and their calibrations. Section 4 discusses implementation details specific to the models. Section 5 analyzes the performance of the Smolyak and KS methods. Section 6 concludes. The appendix examines alternative implementations of the Smolyak method.

1. THE SMOLYAK ALGORITHM

This section describes how functions can be approximated using the Smolyak algorithm. To distinguish the algorithm from its application to approximating equilibrium, the latter is referred to as the Smolyak method.

Let f be an arbitrary function mapping \mathbb{R}^d to \mathbb{R} with typical element x . The Smolyak algorithm is best described in three steps, which I present as a “black box.” See Krueger, Kubler, and Malin (2011)

While in equilibrium, k must be consistent with \mathbf{k} . This only matters when simulating the economy and is trivial to enforce.

¹⁰ KK report that the Euler-equation errors and computation times increase rapidly in the number of generations. They conclude their algorithm can only be applied if the number of generations is less than thirty (p. 19). The Smolyak method presented in this paper can easily handle one hundred generations, and the maximum errors appear to asymptotically approach -2.46 (roughly a one-dollar mistake for every \$300 spent). See Table 4 of this paper.

¹¹ Judd et al. (2010) find the method works best (both in terms of accuracy and numerical stability) when the number of gridpoints is larger than the number of basis functions. The authors argue that a 20 percent increase in the number of points (relative to collocation that has the fewest possible number) has a “sizable effect on accuracy” (p. 30).

(KKM) for a careful exposition of all the necessary steps. The code provided is organized similarly to the description given here.¹² Attention is restricted to the lowest level of approximation.¹³

Step 1 – Set Up

Fix bounds \underline{x} and \bar{x} in \mathbb{R}^d on the state space such that $\underline{x} < \bar{x}$. These bounds define a hypercube. The Smolyak algorithm then provides $n := 2d+1$ collocation points $\{x^i\}_{i=1}^n$ within this hypercube. The advantage of the Smolyak algorithm lies in the construction of these points, whose number grows only linearly in dimension.

Step 2 – Polynomial Construction

Evaluate f at each of the n collocation points. The Smolyak algorithm then provides polynomial coefficients θ . The coefficients θ implicitly define an approximating polynomial \hat{f} .

Step 3 – Polynomial Evaluation

Given coefficients θ , the Smolyak algorithm provides a way to evaluate \hat{f} at arbitrary x (inside or outside of the hypercube).

The collocation points and interpolating polynomial are constructed in such a way that the following conditions are guaranteed:

1. \hat{f} agrees with f at each collocation point, i.e. $\hat{f}(x^i) = f(x^i)$ for all $i \in \{1, \dots, n\}$.
2. If f is a linear combination of the polynomials x_j^2 , x_j , and 1 for $j \in \{1, \dots, d\}$, then \hat{f} agrees with f *everywhere* in the hypercube.¹⁴
3. If f is not perfectly reproduced but is at least continuous, then the polynomial \hat{f} is an almost optimal approximation in a certain

¹² The code is available at sites.google.com/site/greygordon. There are several alternatives to my code. In particular “spinterp” is a free Matlab sparse-grid interpolation toolbox available at www.ians.uni-stuttgart.de/spinterp/. This toolbox has more features than what I provide. Additionally, KKM provide Fortran routines.

¹³ This is the only one that’s feasible for very large distributions. However, for small- to medium-sized distributions, a higher level of approximation may be feasible. It is easy to try a higher level of approximation when using the provided code.

¹⁴ Unfortunately, there are no cross terms for this level of approximation. However, this does not prevent obtaining an accurate solution for the two nontrivial economies considered in this paper.

sense.¹⁵ In general, the less curvature f has, the better \hat{f} will be as an approximation.

For additional details on the algorithm and its properties, the interested reader is referred to KKM and Barthelmann et al. (2000).

2. THE SMOLYAK METHOD

This section describes the Smolyak method, that is, the application of the Smolyak algorithm to approximating equilibrium. First, a typical definition of equilibrium is redefined as a set of functions of *only* the aggregate state. Second, the algorithm is used to approximate these functions.

Redefining Equilibrium

Let f represent a typical policy function, value function, price function, or law of motion. Without loss of generality, assume that f is a function of some “individual” state $x \in X$ and an “aggregate” state $\omega \in \Omega$ that is common across functions.¹⁶ For notational convenience, also assume X is shared by all functions and is non-empty. A typical definition of equilibrium is then a possibly uncountable collection of functions

$$\{f(x; \omega)\} \tag{1}$$

that satisfy conditions that are not explicitly stated, such as optimality, budget balance, market clearing, and consistency of a law of motion. Consider a new definition of equilibrium comprised of indexed functions

$$\mathcal{F} := \{f^x(\omega) | f^x(\omega) = f(x; \omega) \forall x \in X, \omega \in \Omega\} \tag{2}$$

that satisfy implicitly the same conditions as before. Now the original equilibrium has been represented as a (large) collection of functions of only the aggregate state.

A worked example

To fix ideas, consider the following worked example. Suppose there are two agents $i = 1, 2$ in the economy, each with some bonds b , and

¹⁵ Barthelmann et al. (2000) show it is not the best (in the sense of minimizing the sup norm) interpolating polynomial, but it is close to it in that its error bounds are the same up to a logarithmic factor in the number of collocation points. See Theorem 2 and Remark 4 in their paper.

¹⁶ Note that correspondences can be treated as a possibly uncountable collection of functions. Also note that if a function does not depend on the aggregate state, it can just be regarded as a trivial function of the aggregate state.

the budget constraint $c + q(\omega)b' = b$, where the equilibrium bond price q depends on the distribution of wealth ω . In this simple case, we can literally track the exact asset holdings of each agent, $\omega = (b_1, b_2)$, but this will not be the case in the more complicated KS model. The bond is in zero net supply, so the equilibrium price $q(\omega)$ is such that total demand for the bond is zero. In choosing optimal bonds and consumption, the agents must know not only what the price is today, given by $q(\omega)$, but also the price at every point in the future. This is done by assuming a law of motion Γ that maps the current aggregate state to the next period aggregate state: $\omega' = \Gamma(\omega)$. With that in hand, prices in one period will be $q(\Gamma(\omega))$, prices in two periods will be $q(\Gamma(\Gamma(\omega)))$, and so on. Supposing flow utility is $\log(c)$ and the time discount factor is β , optimality requires the consumption policy function $c(b; \omega)$ and savings policy $g(b; \omega)$ satisfy the budget constraint and a Euler equation.¹⁷

Equilibrium can be stated as follows. A recursive competitive equilibrium is a price $q(\omega)$ and consumption and savings policies, $c(b; \omega)$ and $g(b; \omega)$, such that all the following hold:

$$\begin{aligned} c(b; \omega) + q(\omega)g(b; \omega) &= b && (BC) \\ \frac{1}{c(b; \omega)}q(\omega) &= \beta \frac{1}{c(g(b; \omega); \Gamma(\omega))} && \forall b, \omega && (EE) \\ 0 &= \sum_i g(\omega_i; \omega) && \forall \omega && (MC) \\ (g(\omega_1; \omega), g(\omega_2; \omega)) &= \Gamma(\omega) && \forall \omega && (LOM). \end{aligned} \tag{3}$$

Here, ω_i means the i th component of ω , which is the bond holdings of agent i . When BC and EE hold, the policy functions solve the household problem. When MC holds, market clearing obtains. And when LOM holds, the law of motion is consistent with individual policies.

To rewrite equilibrium in terms of indexed functions, we define—for every b —functions $c^b(\omega) = c(b; \omega)$ and $g^b(\omega) = g(b; \omega)$. Then we write

$$\begin{aligned} c^b(\omega) + q(\omega)g^b(\omega) &= b && \forall \omega && (BC^b) \\ \frac{1}{c^b(\omega)}q(\omega) &= \beta \frac{1}{c^{g^b(\omega)}(\Gamma(\omega))} && \forall \omega && (EE^b) \\ 0 &= \sum_i g^{\omega_i}(\omega) && \forall \omega && (MC) \\ (g^{\omega_1}(\omega), g^{\omega_2}(\omega)) &= \Gamma(\omega) && \forall \omega && (LOM), \end{aligned} \tag{4}$$

where now instead of one BC and one EE equilibrium condition, we have as many BC and EE conditions as there are b values.

¹⁷ A transversality condition or no-Ponzi condition is also required.

Applying the Smolyak Algorithm

With equilibrium redefined, it is now straightforward to approximate it using the Smolyak algorithm. First consider the easiest case where Ω is a subset of R^d for some $d < \infty$ and X is a finite set.

1. Fix bounds $\underline{\omega}$ and $\bar{\omega}$ on the aggregate state space such that $\underline{\omega} < \bar{\omega}$.
2. Use the Smolyak algorithm to generate collocation points $\Omega^c := \{\omega^i\}_{i=1}^n$, where $n = 2d + 1$.
3. Make a guess on $f^x(\omega)$ for all $x \in X$ and for each $\omega \in \Omega^c$ for each $f^x \in \mathcal{F}$. Alternatively, make a guess on only a subset of \mathcal{F} , but a subset that is sufficient to construct all the other functions through equilibrium conditions.¹⁸
4. Use the Smolyak algorithm to construct approximations $\hat{f}^x(\omega)$. If the guesses were made for a subset of \mathcal{F} , then approximations will only be explicitly constructed for this subset with the other functions approximated implicitly.
5. Determine whether the approximated functions nearly satisfy all the equilibrium conditions. If they do, stop. If they do not, proceed from step (3) with new guesses. Alternatively, change the bounds and proceed from (1), explicitly approximate other functions in (3) and (4), or pursue different definitions of equilibrium functions or the aggregate state space (e.g., using logs instead of levels).¹⁹

While in abstract this is complicated, the process is simple. Basically, guess on function values at the collocation points, construct Smolyak approximations, and check whether the approximated functions constitute an approximate equilibrium.

The preceding algorithm assumed that X was a finite set and that Ω had finite dimension. If X is not a finite set, then $f(x; \omega)$ for fixed ω must be approximated by its values in a finite set \tilde{X} . This set will typically just be the nodes used for an interpolation, projection, or quadrature method. If $\omega \in \Omega$ has infinite dimension, then it must be approximated using a vector $\tilde{\omega}$ in a subset $\tilde{\Omega}$ of R^d for some $d < \infty$. If ω is a distribution, then a natural way to accomplish this is with the method of either Ríos-Rull (1997) or Young (2010).²⁰ If ω is not a

¹⁸ For instance, one could explicitly approximate consumption and price functions with the savings function being given implicitly through the budget constraint.

¹⁹ One could also check whether a higher level of approximation is feasible.

²⁰ These methods handle distributions of the type $\mu(a, s)$, where $a \in [\underline{a}, \bar{a}]$ and $s \in S$, where S is a finite set. If S is infinite, it must be discretized.

distribution, then some other method must be used, which will depend on the application. Using \tilde{X} and $\tilde{\Omega}$ in place of X and Ω , the algorithm above can then be applied.

The worked example continued

What does this look like in practice? Reconsider the worked example. In this case, the set X is all possible bond values—which is \mathbb{R} —and hence infinite. Because of this, some interpolation or projection scheme has to be used for the individual state b , not just the aggregate state ω . For simplicity, suppose one wanted to represent c and g by a linear interpolant in b conditional on ω . To be extra simple, suppose we only had two interpolation nodes, $b = \mathfrak{b}_1$ and $b = \mathfrak{b}_2$ with $\mathfrak{b}_1 < \mathfrak{b}_2$. Now take \tilde{X} as the interpolation nodes, $\tilde{X} = \{\mathfrak{b}_1, \mathfrak{b}_2\}$.

To apply the Smolyak method, step 1 requires the choice of bounds $\underline{\omega}$ and $\bar{\omega}$. Here, those might naturally be $\underline{\omega} = (\mathfrak{b}_1, \mathfrak{b}_1)$ and $\bar{\omega} = (\mathfrak{b}_2, \mathfrak{b}_2)$. With those in place, step 2 of the method produces collocation points Ω^c . Then, in step 3, for each $\omega \in \Omega^c$, one has to supply a guess on $c^{\mathfrak{b}_1}(\omega)$, $c^{\mathfrak{b}_2}(\omega)$, $g^{\mathfrak{b}_1}(\omega)$, $g^{\mathfrak{b}_2}(\omega)$, and $q(\omega)$. This is five numbers for each ω , and since the dimensionality of ω is two in this example, there are $n = 5$ (since $n = 2d + 1$) values of ω in Ω^c . Consequently, the guess on the approximate equilibrium functions can be summarized by a length 25 vector \mathbf{x} . Specifically, the 25 values in \mathbf{x} are $\{c^{\mathfrak{b}_1}(\omega), c^{\mathfrak{b}_2}(\omega), g^{\mathfrak{b}_1}(\omega), g^{\mathfrak{b}_2}(\omega), q(\omega) | \omega \in \Omega^c\}$. Then with \mathbf{x} , the Smolyak method produces approximations $\{\hat{c}^{\mathfrak{b}_i}(\cdot), \hat{g}^{\mathfrak{b}_i}(\cdot), \hat{q}(\cdot) | i \in \{1, 2\}\}$ in step 4.

At this point, if b had only been allowed to reside in a finite set, one would immediately check the equilibrium conditions in (4). Here, with the added complication of b allowed to be any \mathbb{R} value, in *general* one must take the functions of aggregate states ω for each individual state $\mathfrak{b}_1, \mathfrak{b}_2$, and construct the policy and price functions for any combination of (b, ω) by applying the linearity in b :

$$\begin{aligned} c(b; \omega) &= (1 - \frac{b - \mathfrak{b}_1}{\mathfrak{b}_2 - \mathfrak{b}_1})\hat{c}^{\mathfrak{b}_1}(\omega) + \frac{b - \mathfrak{b}_1}{\mathfrak{b}_2 - \mathfrak{b}_1}\hat{c}^{\mathfrak{b}_2}(\omega) \\ g(b; \omega) &= (1 - \frac{b - \mathfrak{b}_1}{\mathfrak{b}_2 - \mathfrak{b}_1})\hat{g}^{\mathfrak{b}_1}(\omega) + \frac{b - \mathfrak{b}_1}{\mathfrak{b}_2 - \mathfrak{b}_1}\hat{g}^{\mathfrak{b}_2}(\omega) \\ q(\omega) &= \hat{q}(\omega) \end{aligned}$$

Given these policy and price functions, the equilibrium conditions in (3) can be checked and the guess \mathbf{x} modified until the error is sufficiently small or cannot be made smaller. However, a natural choice for minimizing the error would be to make it zero at all twenty-five values of $(b, \omega) \in \tilde{X} \times \Omega^c$, i.e., to use collocation for both the individual and aggregate states. This is equivalent to forcing the equilibrium condi-

tions in (4)— $BC^{\mathfrak{b}_1}, BC^{\mathfrak{b}_2}, EE^{\mathfrak{b}_1}, EE^{\mathfrak{b}_2}, MC, LOM$ —to hold exactly at each of the five $\omega \in \Omega^c$.

3. MODELS AND CALIBRATIONS

This section describes the OLG and KS models and calibrations. In the case of the OLG economy, an analytic solution is also given. The OLG model is set up to be qualitatively similar to the KS model so that both feature capital, inelastic labor supply, production, total factor productivity shocks, and log utility. The model calibrations are similar in several respects but differ drastically with respect to time discounting and depreciation.

OLG Economy

The OLG economy is very similar to Krueger and Kubler (2004) and based on Huffman (1987). The model is set up in sequential rather than recursive form to simplify notation.

A neoclassical production firm operates a production technology $z_t F(K_t, N_t) = z_t K_t^\alpha N_t^{1-\alpha}$ with $\alpha \in (0, 1)$ that uses as inputs capital K_t rented at rate r_t and labor N_t hired at wage w_t and is subject to a productivity shock z_t that evolves according to a Markov chain. Capital depreciates at a stochastic rate δ_t that also evolves according to a Markov chain. The firm takes prices as given and so the equilibrium rental and wage rates are $r_t = z_t \alpha (K_t/N_t)^{\alpha-1}$ and $w_t = z_t (1 - \alpha) (K_t/N_t)^\alpha$, respectively.

Households consist of generations 1 through $T < \infty$ with no intragenerational heterogeneity. The measure of households is constant across generations with the total measure of households normalized to T . It is assumed, and this is key for tractability, that households have log utility, a strictly positive labor endowment in their first period of life, and no labor endowment for the rest of their life. The time t labor endowment of the youngest generation denoted l_t^1 and normalized to T is supplied inelastically resulting in total labor supply $N_t = 1 (= l_t^1/T)$. The time t labor endowment for generation i in $2, \dots, T$ is denoted l_t^i and is equal to zero. At time $t = 0$, households are endowed with capital holdings denoted by a vector $\mathbf{k}_0 = (k_0^1, k_0^2, \dots, k_0^T)$, where k_t^i denotes the capital holdings of generation i at time t . The resulting time 0 aggregate capital endowment is $K_0 = \sum \mathbf{k}_0 / T$. Assume that newborn households have zero capital holdings.

Households maximize expected discounted lifetime utility subject to a budget constraint, nonnegative consumption, and a natural borrowing limit (equal to zero). The budget constraint at time t is given

by

$$c_t^i + k_{t+1}^{i+1} = (1 + r_t - \delta_t)k_t^i + w_t l_t^i \quad (5)$$

for generations $i \in \{1, \dots, T-1\}$ and

$$c_T^i = (1 + r_t - \delta_t)k_T^i + w_t l_t^i \quad (6)$$

for generation $i = T$. Utility of a household beginning life in period t is given by

$$_t \sum_{j=1}^T \beta^{j-1} \log(c_{t+j-1}^j), \quad (7)$$

where $\beta \in (0, 1)$ is the time discount factor.

The necessary and sufficient first-order condition of an age $i < T$ household at time t is given by

$$1/c_t^i = \beta \mathbb{E}_t(1 + r_{t+1} - \delta_{t+1})/c_{t+1}^{i+1}. \quad (8)$$

Using backward induction, the solution to the household problem is shown to be

$$\begin{aligned} k_{t+1}^{i+1} &= \gamma^i (1 + r_t - \delta_t) k_t^i \\ k_{t+1}^2 &= \gamma^1 w_t l_t^1 \\ \gamma^i &= \frac{\beta \sum_{j=0}^{T-1-i} \beta^j}{\sum_{j=0}^{T-i} \beta^j} \end{aligned} \quad (9)$$

for all i in $\{1, \dots, T-1\}$ and for all t . Note that γ^i is the marginal propensity of generation i to save and is constant.

With this solution to the household problem, it is straightforward to calculate the law of motion. Let the time t distribution of capital holdings across generations be given by the vector $\mathbf{k}_t = (0, k_t^2, \dots, k_t^T)$. Then the time t capital stock is $K_t = \sum \mathbf{k}_t / T$, and since total labor supply equals one, the marginal product of capital is $r_t = z_t \alpha K_t^{\alpha-1}$ and the marginal product of labor is $w_t = z_t (1 - \alpha) K_t^\alpha$. Using (9), the time $t+1$ distribution of capital holdings is shown to be

$$\mathbf{k}_{t+1} = (0, \gamma^1 w_t l_t^1, \gamma^2 k_t^2 (1 + r_t - \delta_t), \dots, \gamma^{T-1} k_t^{T-1} (1 + r_t - \delta_t)), \quad (10)$$

which is a correspondence of only the time t aggregate shocks (δ_t, z_t) and distribution \mathbf{k}_t . This law of motion will be used to check the forecast accuracy of both the KS and Smolyak methods. Equilibrium is given by the capital policies in (9), the law of motion in (10), and competitive factor prices. Goods market clearing is ensured by Walras's law.

KS Economy

The KS economy used is a slightly modified version of the original KS (1998) model and is laid out in den Haan, Judd, and Juillard (2010) (dHJJ). The only difference between the two is dHJJ add unemployment insurance so that the zero-borrowing constraint is sometimes binding. The model is set up in recursive form to save on notation.

A neoclassical production firm operates a production technology $zF(K, N) = zK^\alpha N^{1-\alpha}$ with $\alpha \in (0, 1)$ that uses as input capital K rented at rate r and labor N hired at wage w and is subject to a productivity shock z . The productivity shock z takes on one of two values $z \in \{g, b\}$ and evolves according to a Markov chain $\Pi_{zz'}$. Capital depreciates at a constant rate δ . Perfect competition ensures $r = z\alpha(K/N)^{\alpha-1}$ and $w = z(1 - \alpha)(K/N)^\alpha$.

Households have stochastic employment status s taking on one of two values $s \in \{1, 0\}$ with $s = 1$ representing employment and $s = 0$ representing unemployment. Employed workers receive a labor endowment \bar{e} that they supply inelastically to the firm for labor income $w\bar{e}$. Unemployed workers receive unemployment insurance equal to $w\bar{u}$ from the government. Unemployment insurance is funded by the government which levies labor income tax τ on employed workers and runs a balanced budget.

Employment status evolves with the productivity shock according to a Markov chain $\Pi_{ss',zz'}$. The (exogenous) stock of unemployed workers U is assumed to be a function of only the current shock and so is denoted U_z .²¹ The employment process implies total labor supply is known as a function of z with $N_z = (1 - U_z)\bar{e}$. For the government budget to balance, τ must be a function of z with $\tau_z = \frac{\bar{u}}{\bar{e}} \frac{U_z}{(1-U_z)}$. Households seek to maximize the expected discounted lifetime log-utility of consumption discounted at rate β .

The problem of the household is

$$V(k, s; z, \mu) = \max_{c, k'} \log(c) + \beta \sum_{s' z'} \Pi_{ss', zz'} V(k', s'; z', \mu') \quad (11)$$

subject to

$$\begin{aligned} c + k' &= (1 + r - \delta)k + sw\bar{e}(1 - \tau_z) + (1 - s)w\bar{u} \\ c &\geq 0 \\ k' &\in [0, \bar{k}] \\ \mu' &= \Gamma_{zz'}(\mu), \end{aligned} \quad (12)$$

²¹ For ease of exposition, I set up the model as being “initialized” from a long-run distribution. In general, U as well as N and τ must be determined from the distribution.

Table 1 OLG Calibration

Parameter	Value	Parameter	Value
β	0.70	z	[1.05, 1.05, 0.95, 0.95]
α	0.36	$\Pi_{\delta z}$	1/4
δ	[0.9, 0.5, 0.9, 0.5]		

where $r = r(z, \mu)$ and $w = w(z, \mu)$, μ is a joint distribution of capital holdings and employment status across households (giving K and N), and \bar{k} is an exogenous upper bound on possible capital choices (chosen large enough so as to not be binding in equilibrium). Equilibrium is a collection of policy, value, and price functions c, k', V, r, w , together with a law of motion $\Gamma_{zz'}$ (for each z, z') such that V , c , and k' solve the household problem taking r, w , and $\Gamma_{zz'}$ as given, factor prices r and w are competitive, and the law of motion $\Gamma_{zz'}$ is consistent with individual policies and exogenous transition probabilities. Goods market clearing is ensured by Walras's law. Unfortunately, there is no known solution for this model.²²

Calibration

For the OLG economy, I focus on the extreme calibration presented by KK in which quasi-aggregation fails for small T . Depreciation takes on one of two values $\delta \in \{0.9, 0.5\}$, and the productivity shock takes on one of two values $z \in \{1.05, 0.95\}$. Both of these are iid and the four combinations of δ and z occur with equal probability $\Pi_{\delta z} = 1/4$. The discount factor β is taken to be 0.7. The parameters are summarized in Table 1.

For the KS economy, the calibration is the same as in dHJJ, which is only a slight modification of the original KS calibration. The calibration matches select business-cycle statistics at a quarterly frequency. Relative to the OLG calibration, households are much more patient with a discount factor of 0.99, capital depreciates much more slowly at 0.025, the productivity shocks are somewhat smaller at 1.01 and 0.99, and the productivity shock is not iid but has persistence $\Pi_{gg} = \Pi_{bb} =$

²² Interestingly though, there is a solution for a “nearby” economy. If all households are unemployed, there is no unemployment insurance, and the production technology is zK^α , then the equilibrium capital policy function is $k' = \beta(1 + r - \delta)k$. This is the limiting case of the Huffman (1987) example.

Table 2 KS Calibration

Parameter	Value	Parameter	Value
β	.99	$\Pi_{00,gg}$	$\frac{1}{3}\Pi_{gg}$
α	.36	$\Pi_{00,bb}$	$\frac{3}{5}\Pi_{bb}$
δ	.025	$\Pi_{00,gb}$	$\frac{5}{4}\frac{\Pi_{00,bb}}{\Pi_{bb}}\Pi_{gb}$
\bar{e}, \bar{u}	10/9, .15	$\Pi_{00,bg}$	$\frac{3}{4}\frac{\Pi_{00,gg}}{\Pi_{gg}}\Pi_{bg}$
U_g, U_b	.04, .1	$\Pi_{10,zz'}$	$(U_{z'} - U_z \frac{\Pi_{00,zz'}}{\Pi_{zz'}})/(1 - U_z)\Pi_{zz'}$
g, b	1.01, .99	$\Pi_{01,zz'}$	$\Pi_{zz'} - \Pi_{00,zz'}$
Π_{gg}, Π_{bb}	7/8, 7/8	$\Pi_{11,zz'}$	$\Pi_{zz'} - \Pi_{10,zz'}$

7/8. All the parameters, including the employment process parameters, are listed in Table 2.

4. IMPLEMENTATION

This section discusses implementation issues specific to solving the OLG and KS economies using both the Smolyak and KS methods. In abstract, the procedure for computing equilibrium is the same across both economies and both methods. Fixing a law of motion, backward induction along with Carroll's (2006) endogenous gridpoints method is used to solve the household problem. The household capital policies are then used to update the law of motion. This procedure is repeated until the change in the consumption policy and law of motion is less than 10^{-7} in levels. The rest of this section discusses the solution procedures in more detail.

Specific Implementation for OLG

There are several implementation choices to be made when using the Smolyak method to compute the OLG economy. One pertains to representing the distribution of capital holdings. In particular, the distribution can be represented, as in the theoretical model, by using the levels of capital holdings across agents, $\mathbf{k} = (0, k^2, \dots, k^T)$. However, one can instead use (K, \mathbf{s}) , where \mathbf{s} represents the capital-holding *shares* $\mathbf{s} = \mathbf{k}/\sum \mathbf{k}$. Of course, it is possible to switch between the state spaces using $\mathbf{k} = \mathbf{s}KT$, but the two state spaces will result in different numerical solutions. It was found that using (K, \mathbf{s}) produced less error in both the forecasted capital stock and in the Euler equations, and so this is adopted as the benchmark method. The appendix presents accuracy numbers for the other state-space representation. The state-space

bounds were taken to be ± 20 percent of the (nonstochastic) steady-state capital stock and ± 40 percent of the steady-state share distribution. In general, it is a good idea to place state-space bounds as $\pm X$ percent of the steady-state values as this will cluster the collocation points around the steady-state values.

Another implementation choice applies only if using shares in the state space and regards handling of the restriction $\sum \mathbf{s} = 1$. The Smolyak algorithm is not designed to handle this case because it gives collocation points $\{(K, \mathbf{s})\} \subset \mathbb{R}^{T+1}$ that in general will not satisfy this restriction. The method I adopt is to use a mapping from the hypercube $[0, 1]^T$ into the unit-simplex $\Delta(T-1) \subset [0, 1]^T$. In particular, given a collocation point $(K, \tilde{\mathbf{s}})$ with $\sum \tilde{\mathbf{s}} \neq 1$, the mapping $\mathbf{s} = \tilde{\mathbf{s}} / \sum \tilde{\mathbf{s}}$ is used to recover (K, \mathbf{s}) with $\sum \mathbf{s} = 1$. For the reverse mapping, $\tilde{\mathbf{s}} = \mathbf{s}$ is used. The appendix explores a different mapping that is more uniform in a probabilistic sense but produces a worse approximation.

A final implementation choice regards simulating the economy. One can construct an approximation to the law of motion and use it to find the distribution of capital in the next period. Alternatively, one can construct approximations to the capital policy functions and use these to find the distribution. In solving the model, this is a nonissue because the two agree at the collocation points. It was found that approximating the capital policies produced less error, and so this is adopted in the benchmark. The appendix presents accuracy numbers for the other method.

To solve for equilibrium using the KS method, the law was updated by nonstochastically simulating the economy for 5,000 periods, discarding the first 1,000 periods, and using least squares regression to obtain a new law of motion (no relaxation was used).²³ The grid for aggregate capital was set to cover ± 60 percent of the steady-state capital stock with eleven evenly spaced points. A linear, rather than log-linear, functional form for the law of motion was assumed, but the two result in nearly identical approximations.²⁴

Specific Implementation for KS

To solve for equilibrium in the KS economy using the Smolyak method, the following implementation was used. After choosing a set \mathcal{K} of capital gridpoints, the infinite-dimensional distribution $\mu(k, s)$ over $k \in$

²³ No relaxation was used for the Smolyak method either.

²⁴ The Smolyak method's law of motion was also represented using levels, so this makes for a straightforward comparison. In Section 5, I argue no functional form will result in quasi-aggregation for small T .

$[0, \bar{k}], s \in \{0, 1\}$ was approximated by a discrete distribution $\tilde{\mu}(k, s)$ over $k \in \mathcal{K}, s \in \{0, 1\}$ using the method of Young (2010). The capital grid was constructed using one hundred gridpoints resulting in a distribution of dimension 200. Typically the population would be normalized to unity implying the restriction $\sum \tilde{\mu}(k, s) = 1$ in which case the collocation points would not all satisfy this restriction. However, all that matters for prices is the capital-labor ratio, and so I did not impose this.²⁵ The bounds on the state space were taken to be ± 100 percent of the steady-state distribution.²⁶ Instead of iterating to convergence on the household problem every time before updating the law, it was found that iterating only ten times converged to arbitrary precision, did not require relaxation, and was fast (this process was repeated until both the law of motion and household policies fully converged).

To solve for equilibrium using the KS method, the law was updated by nonstochastically simulating the economy for 5,000 periods, discarding the first 1,000 periods, and obtaining new coefficients through least squares regression. The law was only updated after iterating to convergence on the household problem. When updating the law, a relaxation parameter of 0.5 was used because a looser value of 0.25 did not converge. The grid for aggregate capital was set to cover ± 30 percent of the steady-state capital stock using eleven evenly spaced points. A linear functional form for the law of motion was assumed.

5. PERFORMANCE

This section analyzes the performance of the Smolyak and KS methods in computing the OLG and KS economies.

OLG Economy

To evaluate the accuracy of the Smolyak and KS solution methods for the OLG economy, I focus on capital-stock forecast errors and Euler-equation errors along a long simulated path. The path is simulated using the *true* law of motion. The simulation length is set to 15,000 periods, and the first 1,000 periods are discarded.

²⁵ Originally, the collocation points $\{\hat{\mu}\}$ were mapped into the simplex using the transformation $\hat{\mu} = \hat{\mu}/\sum \hat{\mu}$. However, it was found that not using this mapping resulted in smaller Euler-equation errors and a more stable solution.

²⁶ This isn't entirely true as a small constant 10^{-6} was added to the upper bound to ensure the hypercube had positive volume. Also, in general, a lower bound of zero could cause problems, but this won't be the case when the level of approximation is one as it is here.

Table 3 Error in the Law of Motion

T	Max % Error		Minimal R^2		Best Minimal R^2	
	KS	Smolyak	KS	Smolyak	Linear	Log-Lin
3	3.17	0.07	.67568	.99940	.67587	.67635
6	1.03	0.27	.99505	.99987	.99520	.99766
10	1.05	0.65	.99784	.99982	.99792	.99982
25	1.04	0.85	.99807	.99981	.99816	.99984
50	1.04	0.85	.99807	.99980	.99816	.99984
100	1.04	0.85	.99807	.99980	.99816	.99984

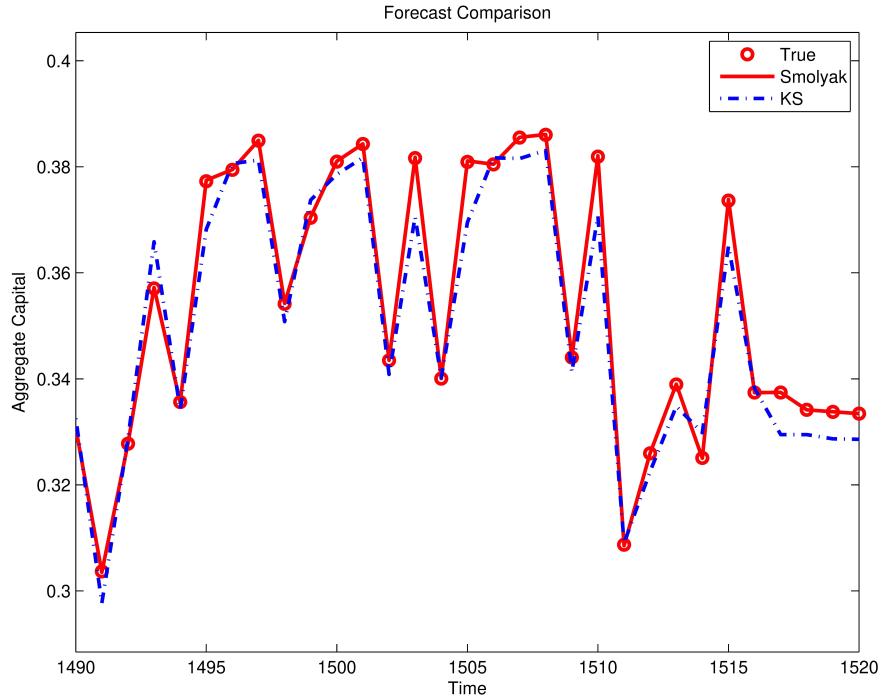
To assess the accuracy of the approximate law of motion, the one-step-ahead, capital-stock forecasts are compared with the realized values in several ways. One measure of the accuracy is given by the largest forecast error $|\hat{K}' - K'|/K'$ observed during the simulation, where \hat{K}' is the forecasted value and K' the actual. Another measure is the R^2 statistic, which indicates how much of the variation in K' is explained by \hat{K}' .²⁷ As there is a separate approximate law of motion for each (δ, z) pair, there are four R^2 statistics, and the worst of these is referred to as “minimal R^2 .” For the KS method, an upper bound on the minimal R^2 value is found by running a linear regression *ex post* on the simulated data. As a robustness check, a log-linear regression is also run to calculate the best minimal R^2 if a log-linear law of motion were assumed for the KS method. The maximum error, minimal R^2 , and best-possible minimal R^2 values are reported in Table 3.

When the number of generations is small, the Smolyak method performs much better than the KS method. The KS method produces maximum errors as large as 3.17 percent and an R^2 value as low as 0.676. In contrast, the Smolyak method’s maximum error is only 0.07 percent and its minimal R^2 statistic is 0.999. That the Smolyak method performs better in this case is confirmed visually in Figure 1, which plots the capital-stock forecasts made by both the Smolyak and KS methods with the true values.

While the Smolyak forecasts and true values are virtually indistinguishable, the KS forecasts deviate noticeably.

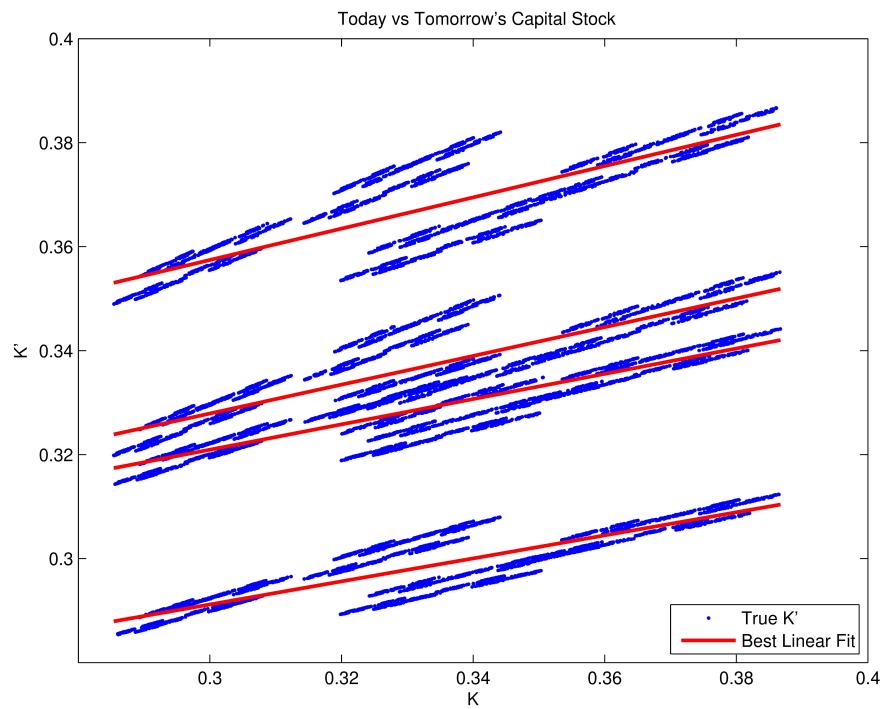
The reason quasi-aggregation fails to obtain for small T is clear. When there are only three generations, marginal propensities to save (which are roughly 0.54, 0.41, and 0 for generations 1, 2, and 3, respectively) differ substantially. Moreover, a generation’s share of total cap-

²⁷ Formally, this is computed as $R^2 = 1 - \sum(\hat{K}' - K'^2) / \sum(K' - \bar{K}'^2)$, where the summation is over a particular sample and \bar{K}' is the sample mean of the next period’s capital stock.

Figure 1 Capital-Stock Forecast Comparison for $T = 3$ 

ital and labor income fluctuates greatly because of large depreciation shocks that only affect capital-rich generations. When the youngest generation holds most of the income, the aggregate propensity to save is roughly 0.54. If instead the middle-aged or oldest generation holds most of the income, the aggregate propensity to save is closer to 0.41 or 0, respectively. Hence, what matters here is not just aggregate income (given by the capital stock), but also the share of income held by each generation, which varies substantially with the history of aggregate shocks. As argued in KK, if either the marginal propensities to save were similar *or* the distribution did not vary much, quasi-aggregation would obtain.

Furthermore, this failure of quasi-aggregation for small T is not due to the chosen functional form of the law of motion. This is made clear in Figure 2, where a scatter plot of today versus tomorrow's capital stock is contrasted against the best linear rules (one for each pair of shocks) one could have.

Figure 2 Today Versus Tomorrow's Capital Stock for $T = 3$ 

While a linear rule does not work well, this figure also demonstrates that any forecast rule that is a *function* of today's shocks and capital stock will fail to produce a good fit because the capital-stock “clouds” are stacked one on another.

For a larger number of generations, the KS and Smolyak methods result in similar performance. For example, when there are one hundred generations, the maximum observed error is 1.04 percent for the KS method and 0.85 percent for the Smolyak method with minimal R^2 values of 0.9981 for the KS method and 0.9998 for the Smolyak method. The KS method's performance noticeably improves as T is increased, while the Smolyak method's performance improves by one measure and worsens by another.

The reason for the KS method's improved performance is clear. In the limiting economy, as T goes large, γ^i , the marginal propensity to save of generation i , converges to β for any fixed i . Hence, nearly all

households have nearly the same marginal propensity to save resulting in quasi-aggregation. Since quasi-aggregation obtains, the KS method performs well.

While quasi-aggregation obtains for large T , the Smolyak method's performance does not noticeably improve, in fact, it gets worse by one measure. However, the advantage of the Smolyak method is that, by keeping track of the entire distribution, its performance is not tied to quasi-aggregation. Rather, the method's performance hinges on the polynomial structure of the law of motion, which does not fundamentally change as T increases.

To test the accuracy of the household policy functions, both maximum and average Euler-equation errors are computed along the simulated path. The errors are computed following Judd (1992) as

$$EEE_t^i(\omega_t) = \log_{10} \left| 1 - \frac{u'^{-1}(\beta \mathbb{E} u_{t+1}^{i+1}(k_{t+1}^i; \omega_t); \hat{\omega}_{t+1})) R(\hat{\omega}_{t+1})}{c_t^i(k_t^i; \omega_t)} \right|, \quad (13)$$

where $R(\hat{\omega}_{t+1}) = 1 + z_{t+1}\alpha\hat{K}_{t+1}^{\alpha-1} - \delta_{t+1}$ and $u(\cdot) = \log(\cdot)$. For the Smolyak method, $\omega_t = (z_t, \delta_t, K_t, \mathbf{s}_t)$, $\hat{\omega}_{t+1} = (z_{t+1}, \delta_{t+1}, \hat{K}_{t+1}, \hat{\mathbf{s}}_{t+1})$, and $(\hat{K}_{t+1}, \hat{\mathbf{s}}_{t+1})$ is the aggregate-state in the next period, according to the perceived law of motion. For the KS method, \mathbf{s}_t and $\hat{\mathbf{s}}_{t+1}$ are simply dropped from the definition of ω_t and $\hat{\omega}_{t+1}$. The interpretation of these errors, derived from Judd and Guu (1997), is that a one-dollar mistake in optimization is made for every $10^{-EEE_t^i}$ dollars spent. For example, if EEE_t^i is -3 , then a one-dollar mistake is made for every \$1,000 spent. Note that, as has typically been done in the literature, the Euler errors are measured with respect to the *perceived* state in the next period. In this sense, they isolate household optimization error conditional on a law of motion from error in the law of motion.

Table 4 reports the maximum and average errors (across both generations and time). For the most part, the optimization errors of the two methods are comparable. Whereas the KS method results in smaller maximum errors, the Smolyak method results in smaller average errors. For large T , the maximum percent errors for the Smolyak method are noticeably larger than those for the KS method and result in a one-dollar mistake for every \$290 spent compared with \$830 for the KS method. The performance of both methods tends to decrease as T increases but appears to level off for $T \geq 25$.

Despite using a fairly coarse approximation for the aggregate state, the Smolyak method produces both optimization errors and errors in the law of motion that are quite small. Having an analytic solution makes it possible to see why this is the case. The chosen implementation of the Smolyak method effectively constructs an approximation of

Table 4 Euler-Equation Errors in the OLG Economy

T	Max Euler Errors		Avg Euler Errors	
	KS	Smolyak	KS	Smolyak
3	-3.15	-3.28	-4.17	-4.42
6	-3.08	-2.99	-3.98	-4.15
10	-2.91	-2.65	-3.81	-4.18
25	-2.92	-2.46	-3.82	-4.14
50	-2.92	-2.46	-3.82	-4.13
100	-2.92	-2.46	-3.82	-4.13

the function

$$k^{i+1,k^i,z,\delta}(K, \mathbf{s}) = \gamma^i(1 + z\alpha K^{\alpha-1} - \delta)k^i \quad (14)$$

for each $i > 1$, k^i in a grid, and (z, δ) combination ($i = 1$ is similar). Because α is in $(0, 1)$, this function is not a polynomial. Hence, away from the collocation points, the approximation is not perfect. However, note that this is a function of only *one* variable, K , and that the polynomial basis used has terms K and K^2 . Because of this, the approximation is quite good for any generation and any level of capital holdings.

It is also possible to see what indexing the policy functions and separating the individual from the aggregate state accomplishes. If the policy function were not indexed, then the Smolyak approximation would be applied to

$$k^{i+1,z,\delta}(k^i; K, \mathbf{s}) = \gamma^i(1 + z\alpha K^{\alpha-1} - \delta)k^i, \quad (15)$$

which has a term k^i and a cross term $k^i K^{\alpha-1}$. To capture the impact of this cross term, one would need a finer Smolyak approximation. If in addition the aggregate and individual states were combined, the Smolyak approximation would be applied to

$$k^{i+1,z,\delta}(K, \mathbf{s}) = \gamma^i(1 + z\alpha K^{\alpha-1} - \delta)\mathbf{s}_i K T, \quad (16)$$

which has cross terms $\mathbf{s}_i K$ and $\mathbf{s}_i K^\alpha$. This also would require a higher level of approximation. Indexing policy functions and separating the individual and aggregate states makes a fairly coarse Smolyak approximation accurate.

Now the running times of the two methods are briefly considered. It is important to remember that while the Smolyak method's aggregate state space grows linearly in T , the KS method's aggregate state space does not grow at all. Hence, for large T , it is guaranteed that the KS method will be faster. However, for small T , the Smolyak method may be faster because it doesn't need to simulate the economy in order to

Table 5 Running Times in Minutes for the OLG Economy

<i>T</i>	KS	Smolyak
3	0.12	0.004
6	0.20	0.02
10	0.25	0.05
25	0.59	0.29
50	1.21	1.18
100	2.29	4.97

update the law of motion. Table 5 reports the running times. For $T < 50$, the Smolyak method is faster than the KS method; for $T = 50$ the times are roughly even, and for $T > 50$ the KS method is faster. Note that while computation time for the KS method grows linearly (roughly) in the number of generations, it grows quadratically for the Smolyak method.

KS Economy

Because the KS economy does not have an analytic solution, it is difficult to assess the accuracy of the KS and Smolyak methods. This is especially true for the Smolyak method. To test the law of motion, researchers have typically compared simulated series generated using only household policies with series generated using an approximate law of motion. This is not really applicable to the Smolyak method: if the law of motion is not explicitly approximated but rather given implicitly by the policy functions, then there is no disagreement between the series. In other words, the Smolyak method has an R^2 of 1 in this case. However, that does not mean there is no error in the law of motion because interpolating the policies is not typically an error-free process.

In light of this, the accuracy of the Smolyak method is assessed in three ways. The first is to compute Euler-equation errors with respect to the *realized* aggregate state in the next period rather than the *perceived* aggregate state. This measure gives an idea of how much error in the law of motion translates into error in household optimization. The second is to compare the capital-stock series from the Smolyak and KS simulations. In addition to the convincing argument made by KS that their computed equilibrium must be close to the true equilibrium, many different solution methods have computed nearly the same equilibrium as the KS method (den Haan 2010). Hence, the KS method's solution can be used as an accuracy check. The third and final test is the comparison of a simulated series generated by explicitly

Table 6 Euler-Equation Errors in the KS Economy

Max Euler Errors		Avg Euler Errors	
KS	Smolyak	KS	Smolyak
-1.89	-2.22	-4.79	-4.97

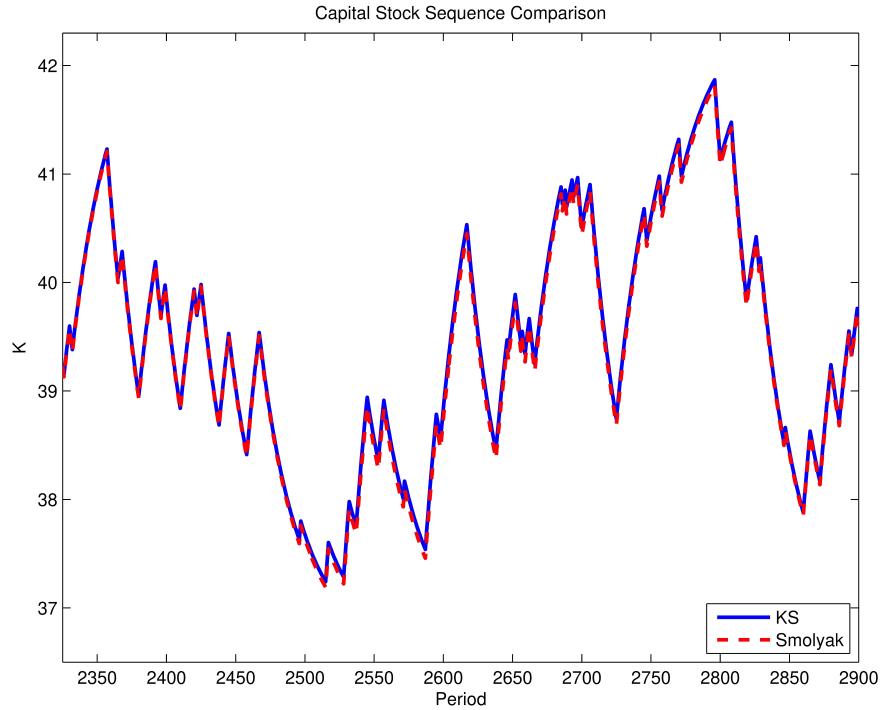
approximating the capital policies with a simulated series generated by explicitly approximating the law of motion. Because it is possible to simulate the economy with either approximation, this may be helpful in assessing how well the Smolyak method is working.²⁸ These three tests are conducted using a simulation length of 5,000 periods with the first 1,000 periods discarded. The accuracy of the KS method is assessed with the first test and also the typical comparison of the forecasted and realized capital-stock sequences.

First, Euler-equation errors are computed along the simulated path and both the maximum and average errors are reported in Table 6. The Euler errors are calculated analogously to (11) except that the rental rate and consumption in the next period are found using the realized next-period moment or distribution and the errors are only counted if the capital choice is strictly positive.²⁹ The errors for the Smolyak method are slightly smaller in both maximum and average terms, but in this measure, the KS and Smolyak methods are roughly equivalent.

Second, the capital sequence generated by the KS method is compared with that of the Smolyak method. Table 7 reports the maximum and average differences between the Smolyak and KS aggregate capital series, and Figure 3 plots them. Visually, the series are almost indistinguishable, although at times, the Smolyak series lies slightly below the KS one. Over the entire simulation, the series exhibit a maximum difference of only 0.28 percent and an average difference of 0.14 percent. However, while this difference is small, it is systematic with the average *nonabsolute* difference also being 0.14 percent (measured using the Smolyak series subtracted from the KS series) confirming what was noticed visually. As the KS method is likely very close to the truth,

²⁸ This is not sure to be helpful, however, because the capital policies and law of motion have different properties. For instance, typically the law of motion will vary with the distribution even if the capital policy does not. In this case, it would be better to approximate the capital policies and compute the law of motion indirectly.

²⁹ When the capital choice equals zero, the no-borrowing constraint is almost certainly binding, in which case the Euler-equation error is not useful.

Figure 3 Simulated Capital Sequence Comparison**Table 7 Simulated Capital Sequence Comparison**

Max Abs (%)	Mean Abs (%)	Mean (%) (KS-Smolyak)
0.278	0.141	0.141

the proximity of these two series confirms the accuracy of the Smolyak method for this economy.

Third and finally, the maximum and average capital-stock “forecast errors” at one, twenty-five, and one hundred steps ahead are examined. As already mentioned, the errors for the Smolyak method are not forecast errors in the typical sense, but rather the discrepancy in capital-stock series constructed by two simulation methods. For the KS method, I proceed as usual, comparing the forecasts made by the perceived and actual law of motion. Table 8 reports the errors, which

Table 8 Law of Motion Forecast Errors in the KS Economy

	1-Step Forecast		25-Step Forecast		100-Step Forecast	
	KS	Smolyak	KS	Smolyak	KS	Smolyak
R^2	.999999	.999998	.999780	.999918	.999474	.999888
Max Abs (%)	.0073	.0204	.1007	.1533	.1530	.1882
Avg Abs (%)	.0021	.0026	.0346	.0454	.0554	.0769
Avg (%)	-.0013	.0026	-.0211	.0454	-.0349	.0769

are clearly small for both methods: even one hundred steps ahead, the capital-stock forecast is off by at most 0.18 percent for the Smolyak and 0.15 percent for the KS method.

The KS method outperforms the Smolyak method in all measures except the R^2 , but its performance tends to deteriorate faster as the forecast length increases. Of note is the systematic bias in the Smolyak's forecasts as seen in the mean absolute errors being the same as the mean errors. While the KS method is guaranteed to be correct on average (because it makes unbiased forecasts by construction), the Smolyak method is guaranteed to be correct only as the level of approximation goes large. It appears, however, that the lowest level of approximation produces small errors in the KS economy just as it does in the OLG economy.

Since the Smolyak method places no restrictions on the functional form of the law of motion and appears to be accurate, it is interesting to ask whether the Smolyak-approximated economy exhibits quasi-aggregation. I test this by running linear and log-linear regressions *ex post* on the simulated capital stock series from the Smolyak method. As expected, the fit is extremely good with all the R^2 values exceeding 0.999999. This verifies directly, inasmuch as the Smolyak solution approximates the truth, that KS's argument for quasi-aggregation was indeed correct: only the mean matters for this model.

The running times of the two methods are now briefly considered. While the KS method (paired with Carroll's 2006 endogenous gridpoints method) is very fast at 3.25 minutes, the Smolyak method is also quite fast at 6.37 minutes. The Smolyak method performs quite well in this regard as it has 401 collocation points compared with the KS method's eleven moment gridpoints. Its comparative advantage lies in avoiding the simulation step, where much of the computation time is spent for the KS method.

6. CONCLUSION

The Smolyak method is a promising technique for computing equilibrium in dynamic heterogeneous-agent economies.

While including a distribution as a state variable massively increases the dimensionality of the state space, the Smolyak sparse-grid interpolation algorithm makes this increase manageable. This technique developed in Smolyak (1963) and Barthelmann et al. (2000) shows great promise for economic applications as Krueger and Kubler (2004) first illustrated. The application of the Smolyak algorithm here results not only in tractability, but also in very good accuracy. In the KS economy, the Smolyak method produces errors similar to those of the KS method. In the OLG economy, the Smolyak method performs much better than the KS method when the number of generations is small because it does not rely on quasi-aggregation which fails in this case. Moreover, for models where the distribution is finite-dimensional, as in the OLG model, the method can be regarded as solving for a full rational-expectations equilibrium. In models where the distribution is infinite-dimensional, as in the KS model, the method comes very close to a full rational-expectations equilibrium.

Keeping track of the distribution does come at a cost: the Smolyak method is not as easy to implement nor as fast as the KS method in most cases. With regard to implementation, the code provided with this paper is meant to reduce the programming costs as much as possible. With this code in hand, the Smolyak algorithm is not much more complicated than any other interpolation scheme: values at a few pre-defined points define an interpolating function. With regards to speed, this paper has shown the Smolyak method need not be much slower than the KS method. For models with larger aggregate state spaces or where Carroll's (2006) efficient solution method cannot be used, parallelization (which has not been used in this paper) could prove very helpful as the work by Aldrich et al. (2011) has shown.

While the application of the Smolyak algorithm here has been to approximate full-rationality equilibrium as closely as possible, the Smolyak algorithm could also be useful in solving for bounded-rationality equilibrium quickly and accurately. While many methods could benefit from the Smolyak algorithm, of particular promise is the explicit aggregation technique of den Haan and Rendahl (2010). This method avoids the simulation step of the KS method by explicitly aggregating “auxiliary” policy functions. In its most basic implementation, the auxiliary policy functions are constructed to be linear in asset holdings but as close as possible to the original policies. Because the auxiliary policies are linear, they aggregate perfectly and the minimal aggregate state space is average capital holdings for each type of agent. Having more

types of agents or more curvature in the auxiliary policy functions requires having more moments, but Smolyak interpolation is well-adapted to handling this increase in aggregate moments. This pairing of den Haan and Rendahl's (2010) method with the Smolyak algorithm could achieve some of the benefits of the Smolyak method (no simulation, high accuracy) while being extremely efficient. The exploration of this idea is left for future research.

REFERENCES

- Aldrich, Eric M., Jesús Fernández-Villaverde, A. Ronald Gallant, and Juan F. Rubio-Ramírez. 2011. “Tapping the Supercomputer under Your Desk: Solving Dynamic Equilibrium Models with Graphics Processors.” *Journal of Economic Dynamics and Control* 35 (March): 386–93.
- Aruoba, S. Borağan, Jesús Fernández-Villaverde, and Juan F. Rubio-Ramírez. 2006. “Comparing Solution Methods for Dynamic Equilibrium Economies.” *Journal of Economic Dynamics and Control* 30 (December): 2477–508.
- Azzimonti, Marina, Eva de Francisco, and Per Krusell. 2006. “Median-voter Equilibria in the Neoclassical Growth Model under Aggregation.” *Scandinavian Journal of Economics* 108 (December): 587–606.
- Barthelmann, Volker, Erich Novak, and Klaus Ritter. 2000. “High Dimensional Polynomial Interpolation on Sparse Grids.” *Advances in Computational Mathematics* 12 (March): 273–88.
- Carroll, Christopher D. 2006. “The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems.” *Economics Letters* 91 (June): 312–20.
- Chien, Yili, Harold Cole, and Hanno Lustig. 2009. “A Multiplier Approach to Understanding the Macro Implications of Household Finance.” Manuscript.
- Childers, David. 2018. “Solution of Rational Expectations Models with Function Valued States.” Manuscript.
- Den Haan, Wouter J. 2010. “Comparison of Solutions to the Incomplete Markets Model with Aggregate Uncertainty.” *Journal of Economic Dynamics and Control* 34 (January): 4–27.
- Den Haan, Wouter J., Kenneth L. Judd, and Michel Juillard. 2010. “Computational Suite of Models with Heterogeneous Agents: Incomplete Markets and Aggregate Uncertainty.” *Journal of Economic Dynamics and Control* 34 (January): 1–3.
- Den Haan, Wouter J., and Pontus Rendahl. 2010. “Solving the Incomplete Markets Model with Aggregate Uncertainty Using Explicit Aggregation.” *Journal of Economic Dynamics and Control* 34 (January): 69–78.

- Huffman, Gregory W. 1987. "A Dynamic Equilibrium Model of Asset Prices and Transaction Volume." *Journal of Political Economy* 95 (February): 138–59.
- Judd, Kenneth L. 1992. "Projection Methods for Solving Aggregate Growth Models." *Journal of Economic Theory* 58 (December): 410–52.
- Judd, Kenneth L. 1998. *Numerical Methods in Economics*. Cambridge, Mass.: MIT Press.
- Judd, Kenneth L., and Sy-Ming Guu. 1997. "Asymptotic Methods for Aggregate Growth Models." *Journal of Economic Dynamics and Control* 21 (June): 1025–42.
- Judd, Kenneth L., Lilia Maliar, and Serguei Maliar. 2010. "A Cluster-Grid Projection Method: Solving Problems with High-Dimensionality." NBER Working Paper No. 15965, May.
- Krueger, Dirk, Felix Kubler. 2004. "Computing Equilibrium in OLG models with Stochastic Production." *Journal of Economic Dynamics and Control* 28 (April): 1411–36.
- Krueger, Dirk, Felix Kubler, and Benjamin Malin. 2011. "Computing Stochastic Dynamic Economic Models with a Large Number of State Variables: A Description and Application of a Smolyak-Collocation Method." *Journal of Economic Dynamics and Control* 35 (February): 229–39.
- Krusell, Per, and José-Victor Ríos-Rull. 1999. "On the Size of U.S. Government: Political Economy in the Neoclassical Growth Model." *American Economic Review* 89 (December): 1156–81.
- Krusell, Per, and Anthony A. Smith Jr. 1997. "Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns." *Macroeconomic Dynamics* 1 (June): 387–422.
- Krusell, Per, and Anthony A. Smith Jr. 1998. "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy* 106 (October): 867–96.
- Reiter, Michael. 2010. "Solving the Incomplete Markets Model with Aggregate Uncertainty by Backward Induction." *Journal of Economic Dynamics and Control* 34 (January): 28–35.
- Ríos-Rull, José-Victor. 1997. "Computation of Equilibria in Heterogeneous Agent Models." Federal Reserve Bank of Minneapolis Staff Report 231, pp. 238–64.
- Smolyak, S. A. 1963. "Quadrature and Interpolation Formulas for

- Tensor Products of Certain Classes of Functions.” *Soviet Mathematics, Doklady* 4: 240–43.
- Terry, Stephen J. 2017. “Alternative Methods for Solving Heterogeneous Firm Models.” *Journal of Money, Credit and Banking* 49 (September): 1081–111.
- Young, Eric R. 2010. “Solving the Incomplete Markets Model with Aggregate Uncertainty Using the Krusell-Smith Algorithm and Non-stochastic Simulations.” *Journal of Economic Dynamics and Control* 34 (January): 36–41.

APPENDIX

This appendix explores alternative implementations of the Smolyak method in the context of the OLG economy (which has a known solution). Four different implementations are considered and their descriptions are given below.

Sm1 – Benchmark Implementation

The state space uses shares \mathbf{s} ; the mapping from $\tilde{\mathbf{s}}$ in the cube to \mathbf{s} in the simplex is $\mathbf{s} = \tilde{\mathbf{s}} / \sum \tilde{\mathbf{s}}$, and the reverse mapping is $\tilde{\mathbf{s}} = \mathbf{s}$; the distribution is forecasted using an approximated capital policy function (not an approximate law of motion). Because of the loss in dimension in mapping to the simplex, there are many reverse mappings into the cube.³⁰

Sm2 – Alternative Mapping

The state space uses shares \mathbf{s} ; the mapping from $\tilde{\mathbf{s}}$ in the cube to \mathbf{s} in the simplex is $\mathbf{s} = -\log(\tilde{\mathbf{s}}) / \sum -\log(\tilde{\mathbf{s}})$, and the reverse mapping is $\tilde{\mathbf{s}} = e^{(10 \log(.5)\mathbf{s})}$; the distribution is forecasted using approximated policy functions. The mapping is motivated by a method of drawing uniformly from a unit simplex (which is accomplished by drawing from the Dirichlet distribution with concentration parameter equal to 1). Several reverse mappings were tried, but the one used worked best.³¹

Sm3 – Simpler State Space

The state space uses levels of capital stock for each generation \mathbf{k} ; there is no mapping; and the distribution is forecasted using an approximated capital policy function.

³⁰ It's unclear which of these reverse mappings is optimal, but the identity mapping is a natural choice.

³¹ In some sense, the reverse mapping $\tilde{\mathbf{s}} = e^{(T \log(.5)\mathbf{s})}$ should be optimal because most of the time, $\sum -\log(\tilde{\mathbf{s}})$ is equal to $-T \log(.5)$. Numerically, however, raising values to such a large power creates instability. I found $\tilde{\mathbf{s}} = e^{(10 \log(.5)\mathbf{s})}$ works fairly well.

Table 9 Accuracy of the Law of Motion

T	Max % Error					Minimal R^2				
	KS	Sm1	Sm2	Sm3	Sm4	KS	Sm1	Sm2	Sm3	Sm4
3	3.17	0.07	0.07	0.22	0.37	.676	.999	.999	.997	.996
6	1.03	0.27	0.27	0.84	1.25	.995	1.000	1.000	.996	.992
10	1.05	0.65	0.66	1.42	1.85	.998	1.000	1.000	.995	.989
25	1.04	0.85	0.86	1.66	2.14	.998	1.000	1.000	.994	.987
50	1.04	0.85	0.86	1.66	2.14	.998	1.000	1.000	.994	.987
100	1.04	0.85	0.86	1.66	2.14	.998	1.000	1.000	.994	.987

Sm4 – Simpler State Space, Alternative Simulation Method

The state space uses levels of capital stock for each generation \mathbf{k} ; there is no mapping; and the distribution is forecasted using an approximated law of motion.

The accuracy numbers for the laws of motion are displayed in Table 9. For space, the R^2 values are rounded to three decimal places. Alternative implementations of the Smolyak method display similar characteristics to the benchmark implementation: each outperforms the KS method for small T (which is when quasi-aggregation breaks down), and the performance of each implementation decreases as the number of periods increases. Sm1 and Sm2 display similar errors, which are less than the errors of Sm3 and Sm4. This suggests that the precise mapping of shares may not matter as much as the choice of whether or not to use shares. That Sm1 and Sm2 perform better than Sm3 and Sm4 is likely due to the ability of Smolyak interpolation to achieve high accuracy in one dimension relative to accuracy in several dimension as discussed in the main text. Because Sm1 and Sm2 include K in the state space directly rather than indirectly through $K = \sum \mathbf{k}/T$, Sm1 and Sm2 exploit this feature of the Smolyak algorithm and so capture more of the general equilibrium effects.

The Euler-equation errors are reported in Table 10. Again, the alternative implementations display similar patterns to the benchmark one: the errors are smaller than the KS method errors for $T = 3$ but increase in the number of generations. Again, Sm1 and Sm2 fare better than Sm3 and Sm4, with Sm1 outperforming Sm2 most of the time. While the KS method tends to produce less error, the errors of the Smolyak implementations would not typically be considered large, at least in terms of average errors. Note that Sm3 and Sm4 have the exact same solution but different simulations (and so result in different maximum and average errors along the simulation path).

Table 10 Accuracy of the Policy Functions

T	KS	Max Euler Errors				KS	Avg Euler Errors			
		Sm1	Sm2	Sm3	Sm4		Sm1	Sm2	Sm3	Sm4
3	-3.15	-3.28	-3.28	-3.27	-3.27	-4.17	-4.42	-4.41	-4.19	-4.16
6	-3.08	-2.99	-3.05	-2.24	-2.14	-3.98	-4.15	-4.11	-3.25	-3.27
10	-2.91	-2.65	-2.24	-1.94	-1.82	-3.81	-4.18	-3.71	-3.06	-3.08
25	-2.92	-2.46	-2.12	-1.86	-1.72	-3.82	-4.14	-3.73	-3.01	-3.03
50	-2.92	-2.46	-2.20	-1.86	-1.72	-3.82	-4.13	-3.82	-3.01	-3.03
100	-2.92	-2.46	-2.28	-1.86	-1.72	-3.82	-4.13	-3.92	-3.01	-3.03

Table 11 Running Times in Minutes

T	KS	Sm1/Sm2		Sm3/Sm4	
3	0.12		0.00		0.00
6	0.20		0.02		0.02
10	0.25		0.05		0.05
25	0.59		0.29		0.34
50	1.21		1.18		1.40
100	2.29		4.97		6.00

The running times for the various implementations are reported in Table 11. As running times are virtually the same for Sm1 as Sm2 and Sm3 as Sm4, I only report joint Sm1/Sm2 and Sm3/Sm4 times. The Sm3 and Sm4 implementations take slightly longer to run than the Sm1 and Sm2 implementations. This is not because the actual interpolation takes any longer, but because extra iterations are required for the law of motion to converge, which is itself due to larger errors in the law of motion.