

ESTIMATING DYNAMIC EQUILIBRIUM ECONOMIES: LINEAR VERSUS NONLINEAR LIKELIHOOD

JESÚS FERNÁNDEZ-VILLAVERDE^{a*} AND JUAN F. RUBIO-RAMÍREZ^b

^a *Department of Economics, University of Pennsylvania, USA*

^b *Research Department, Federal Reserve Bank of Atlanta, USA*

SUMMARY

This paper compares two methods for undertaking likelihood-based inference in dynamic equilibrium economies: a sequential Monte Carlo filter and the Kalman filter. The sequential Monte Carlo filter exploits the nonlinear structure of the economy and evaluates the likelihood function of the model by simulation methods. The Kalman filter estimates a linearization of the economy around the steady state. We report two main results. First, both for simulated and for real data, the sequential Monte Carlo filter delivers a substantially better fit of the model to the data as measured by the marginal likelihood. This is true even for a nearly linear case. Second, the differences in terms of point estimates, although relatively small in absolute values, have important effects on the moments of the model. We conclude that the nonlinear filter is a superior procedure for taking models to the data. Copyright © 2005 John Wiley & Sons, Ltd.

1. INTRODUCTION

Recently, a growing literature has focused on the formulation and estimation of dynamic equilibrium models using a likelihood-based approach. Examples include the seminal paper of Sargent (1989) and, more recently, Bouakez *et al.* (2002), DeJong *et al.* (2000), Dib (2001), Fernández-Villaverde and Rubio-Ramírez (2003), Hall (1996), Ireland (2002), Kim (2000), Landon-Lane (1999), Lubik and Schorfheide (2003), McGrattan *et al.* (1997), Moran and Dolar (2002), Otrok (2001), Rabanal and Rubio-Ramírez (2003), Schorfheide (2000), Smets and Wouters (2003), to name just a few. Most of these papers have used the Kalman filter to estimate a linear approximation to the original model.

This paper studies the effects of estimating the nonlinear representation of a dynamic equilibrium model instead of working with its linearized version. We document how the estimation of the nonlinear solution of the economy substantially improves the empirical fitting of the model. The marginal likelihood of the economy, i.e., the probability that the model assigns to the data, increases by two orders of magnitude. This is true even for our application, the neoclassical growth model, which is nearly linear. We also report that, although the effect of linearization on point estimates is small, the impact on the moments of the model is of first-order importance. This finding is key for applied economists because quantitative models are widely judged by their ability to match data moments.

Dynamic equilibrium models have become a standard tool in quantitative economics (see Cooley, 1995 for a summary of applications). These models can be described as a likelihood function for observables, given the model's structural parameters—those characterizing preferences and

* Correspondence to: Jesús Fernández-Villaverde, Department of Economics, 160 McNeil Building, 3718 Locust Walk, University of Pennsylvania, Philadelphia, PA 30309, USA. E-mail: jesusfv@econ.upenn.edu

technology. The advantage of thinking about models as a likelihood function is that, once we can evaluate the likelihood, inference is a direct exercise. In a classical environment we maximize this likelihood function to get point estimates and standard errors. A Bayesian researcher can use the likelihood and priors about the parameters to find the posterior. The advent of Markov chain Monte Carlo algorithms has facilitated this task. In addition, we can compare models by likelihood ratios (Vuong, 1989) or Bayes factors (Geweke, 1998) even if the models are misspecified and nonnested.

The previous discussion points out the need to evaluate the likelihood function. The task is conceptually simple, but its implementation is more cumbersome. Dynamic equilibrium economies do not have a 'paper and pencil' solution. This means that we can study only an approximation to them, usually generated by a computer. The lack of a closed form for the solution of the model complicates the process of finding the likelihood.

The literature shows how to write this likelihood analytically only in a few cases (see Rust, 1994 for a survey). Outside those, Sargent (1989) proposed an approach that has become popular. Sargent noticed that a standard procedure for solving dynamic models is to linearize them. This can be done either directly in the equilibrium conditions or by generating a quadratic approximation to the utility function of the agents. Both approaches imply that the optimal decision rules are linear in the states of the economy. The resulting linear system of difference equations can be solved with standard methods (see Anderson *et al.*, 1996; Uhlig, 1999 for a detailed explanation).

For estimation purposes, Sargent emphasized that the resulting system has a linear representation in a state-space form. If, in addition, we assume that the shocks exogenously hitting the economy are normal, we can use the Kalman filter to evaluate the likelihood. It has been argued (for example, Kim *et al.*, 2003) that this linear solution is likely to be accurate enough for fitting the model to the data.

However, exploiting the linear approximation to the economy can be misleading. For instance, linearization may be an inaccurate approximation if the nonlinearities of the model are important or if we are travelling far away from the steady state of the model. Also, accuracy in terms of the policy function of the model does not necessarily imply accuracy in terms of the likelihood function. Finally, the assumption of normal innovations may be a poor representation of the dynamics of the shocks in the data.

An alternative to linearization is to work instead with the nonlinear representation of the model and to apply a nonlinear filter to evaluate the likelihood. This is possible thanks to the recent development of sequential Monte Carlo methods (see the seminal paper of Gordon *et al.*, 1993 and the review of the literature by Doucet *et al.*, 2000 for extensive references). Fernández-Villaverde and Rubio-Ramírez (2004) build on this literature to show how a sequential Monte Carlo filter delivers a consistent and efficient evaluation of the likelihood function of a nonlinear and/or nonnormal dynamic equilibrium model.

The presence of the two alternative filters begets the following question: how different are the answers provided by each of them? We study this question with the canonical stochastic neoclassical growth model with leisure choice. We estimate the model using both simulated and real data and compare the results obtained with the sequential Monte Carlo filter and the Kalman filter.

Why do we choose the neoclassical growth model for our comparison? First, this model is the workhorse of modern macroeconomics. Since any lesson learned in this paper is conditional on our particular model, we want to select an economy that is the foundation of numerous applications. Second, even if the model is nearly linear for the standard calibration, the answers provided by

each of the filters are nevertheless quite different. In this way, we make our point that linearization has a nontrivial impact on estimation in the simplest possible environment.

Our main finding is that, while linearization may have a second-order effect on the accuracy of the policy function given some parameter values, it has a first-order impact on the model's likelihood function. Both for simulated and for real data, the sequential Monte Carlo filter generates an overwhelmingly better fit of the model as measured by the marginal likelihood. This is true even if most differences in the point estimates of the parameters generated by each of the two filters are small.

Why is the marginal likelihood so much higher for the sequential Monte Carlo? Because this filter delivers point estimates for the parameters that imply model's moments closer to the moments of the data. This result is crucial in applied work because models are often judged by their ability to match empirical moments.

Our finding is not the first in the literature that suggests accounting for nonlinearities substantially improves the measures of fit of a model. For example, Sims and Zha (2002) report that the ability of a structural VAR to account for the dynamics of the output and monetary policy increases by several orders of magnitude when they allow the structural equation variances to change over time. A similar finding is emphasized by the literature on regime switching (Kim and Nelson, 1999) and by the literature on the asymmetries of the business cycle (Kim and Piger, 2002).

The rest of the paper is organized as follows. In Section 2 we discuss the two alternatives to evaluate the likelihood of a dynamic equilibrium economy. Section 3 presents our application. Section 4 describes the estimation algorithm, and Section 5 reports our main findings. Section 6 concludes.

2. TWO FRAMEWORKS TO EVALUATE THE LIKELIHOOD

In this section we describe the nonlinear and linear filters used to evaluate the likelihood function of a dynamic equilibrium economy. First, we present the state-space representation of a dynamic equilibrium model solved by nonlinear and linear methods. Second, we show how to use a sequential Monte Carlo filter to evaluate the likelihood of the nonlinear state-space representation of the economy. Finally, we do the same with the Kalman filter.

2.1. The State-Space Representation

Assume that we observe $y^T = \{y_t\}_{t=1}^T$, a realization of the n -dimensional random variable $Y^T = \{Y_t\}_{t=1}^T \in R^{nT}$. The researcher is interested in evaluating the likelihood function of the observable y^T implied by a dynamic equilibrium economy at any given γ , $L(y^T; \gamma) = p(y^T; \gamma)$, where $\gamma \in \Upsilon$ is the vector collecting the parameters that characterize preferences, information and technology in the model.

Unfortunately, in general, it is not possible to compute this function. Part of the reason is that most dynamic models do not have a closed-form solution. Consequently, just to solve the model before any estimation, we need to approximate the equilibrium path using numerical techniques. This approximation is going to affect the characterization of the likelihood.

There are two main routes to evaluate the likelihood function. If we opt to solve the model nonlinearly, we can use a sequential Monte Carlo. If we linearize the model, we can apply the Kalman filter. We now describe both methodologies.

The Nonlinear Solution of the Model

Dynamic equilibrium economies solved using nonlinear methods have the following state-space representation. The vector of state variables, S_t , evolves over time according to the transition equation

$$S_t = f(S_{t-1}, W_t; \gamma) \quad (1)$$

where $\{W_t\}$ is a sequence of exogenous random variables. The observable y_t is governed by the measurement equation

$$Y_t = g(S_t, V_t; \gamma) \quad (2)$$

where $\{V_t\}$ is a sequence of exogenous independent random variables. The sequences $\{W_t\}$ and $\{V_t\}$ are independent of each other. Assuming independence of $\{W_t\}$ and $\{V_t\}$ is only for notational convenience. Generalization to more involved structures is achieved by increasing the dimension of the state-space. Along some dimension, the function g can be the identity mapping if a state is observed without noise.

The functions f and g depend on the equations that describe the equilibrium of the model—policy functions, laws of motion for variables, resource constraints—and on the nonlinear solution method used to approximate the policy functions.

To ensure that the model is not stochastically singular, we assume that $\dim(W_t) + \dim(V_t) \geq \dim(Y_t)$. We do not impose any restrictions on how those degrees of stochasticity are achieved.

The Linear Solution of the Model

If we pick a linear method to solve the same model, the state-space representation has the following linear form:

$$S_t = E(\gamma) + A(\gamma)S_{t-1} + B(\gamma)W_t \quad (3)$$

$$Y_t = F(\gamma) + C(\gamma)S_t + D(\gamma)V_t \quad (4)$$

where $A(\gamma)$, $B(\gamma)$, $C(\gamma)$, $D(\gamma)$, $E(\gamma)$ and $F(\gamma)$ are matrices with the required dimension that depend on the parameters of the model. Notice how these equations are nothing more than a particular case of (1) and (2). Also, we make the same assumptions regarding stochastic singularity as above.

We have presented two representations of the same economy. Section 2.2 introduces a sequential Monte Carlo filter to evaluate the likelihood function implied by (1) and (2). Section 2.3 exploits the Kalman filter to calculate the likelihood entailed by (3) and (4).

2.2. The Nonlinear Approach: A Sequential Monte Carlo Filter

Fernández-Villaverde and Rubio-Ramírez (2004) propose the following sequential Monte Carlo to evaluate the likelihood function of y^T induced by (1) and (2).

We assume that we can partition $\{W_t\}$ into two separate sequences $\{W_{1,t}\}$ and $\{W_{2,t}\}$, such that $W_t = (W_{1,t}, W_{2,t})$ and $\dim(W_{2,t}) + \dim(V_t) = \dim(Y_t)$. Let $W_i^t = \{W_{i,m}\}_{m=1}^t$, for $i = 1, 2$, $V^t = \{V_m\}_{m=1}^t$ and $S^t = \{S_m\}_{m=0}^t \forall t$. We also define $W_i^0 = \{\emptyset\}$ and $y^0 = \{\emptyset\}$. We could work under weaker assumptions, paying the cost of heavier notation.

Given our assumptions, we factor the likelihood function as follows:

$$\begin{aligned} p(y^T; \gamma) &= \prod_{t=1}^T p(y_t | y^{t-1}; \gamma) \\ &= \prod_{t=1}^T \int \int p(y_t | W_1^t, S_0, y^{t-1}; \gamma) p(W_1^t, S_0 | y^{t-1}; \gamma) dW_1^t dS_0 \end{aligned} \quad (5)$$

Conditional on having N draws of $\{s_0^{t|t-1,i}, w_1^{t|t-1,i}\}_{i=1}^N$ from the sequence of densities $\{p(W_1^t, S_0 | y^{t-1}; \gamma)\}_{t=1}^T$, this likelihood can be approximated by

$$p(y^T; \gamma) \simeq p_{SMC}(y^T; \gamma) = \prod_{t=1}^T \frac{1}{N} \sum_{i=1}^N p(y_t | w_1^{t|t-1,i}, s_0^{t|t-1,i}, y^{t-1}; \gamma) \quad (6)$$

using a law of large numbers. Thus, the problem of evaluating $p_{SMC}(y^T; \gamma)$ is equivalent to the problem of drawing from $\{p(W_1^t, S_0 | y^{t-1}; \gamma)\}_{t=1}^T$. The sequential Monte Carlo filter accomplishes this objective.

Let us fix some additional notation. Let $\{s_0^{t,i}, w_1^{t,i}\}_{i=1}^N$ be a sequence of N i.i.d. draws from $p(W_1^t, S_0 | y^t; \gamma)$ and $\{s_0^{t|t-1,i}, w_1^{t|t-1,i}\}_{i=1}^N$ be a sequence of N i.i.d. draws from $p(W_1^t, S_0 | y^{t-1}; \gamma)$. We call each draw $(s_0^{t,i}, w_1^{t,i})$ a particle and the whole sequence $\{s_0^{t,i}, w_1^{t,i}\}_{i=1}^N$ a swarm of particles.

Fernández-Villaverde and Rubio-Ramírez (2004) prove the following result that shows how to use $p(W_1^t, S_0 | y^{t-1}; \gamma)$ as an important sampling density to draw from $p(W_1^t, S_0 | y^t; \gamma)$.

Proposition 1. *Let $\{s_0^{t|t-1,i}, w_1^{t|t-1,i}\}_{i=1}^N$ be a draw from $p(W_1^t, S_0 | y^{t-1}; \gamma)$. Let the sequence $\{\tilde{s}_0^i, \tilde{w}_1^i\}_{i=1}^N$ be a draw with replacement from $\{s_0^{t|t-1,i}, w_1^{t|t-1,i}\}_{i=1}^N$ where q_t^i , defined as*

$$q_t^i = \frac{p(y_t | w_1^{t|t-1,i}, s_0^{t|t-1,i}, y^{t-1}; \gamma)}{\sum_{i=1}^N p(y_t | w_1^{t|t-1,i}, s_0^{t|t-1,i}, y^{t-1}; \gamma)}$$

is the probability of $(s_0^{t|t-1,i}, w_1^{t|t-1,i})$ being drawn $\forall i$. Then $\{\tilde{s}_0^i, \tilde{w}_1^i\}_{i=1}^N$ is a draw from $p(W_1^t, S_0 | y^t; \gamma)$.

Proposition 1 shows how a draw $\{s_0^{t|t-1,i}, w_1^{t|t-1,i}\}_{i=1}^N$ from $p(W_1^t, S_0 | y^{t-1}; \gamma)$ can be used to get a draw $\{s_0^{t,i}, w_1^{t,i}\}_{i=1}^N$ from $p(W_1^t, S_0 | y^t; \gamma)$. This result is crucial. Given particles at $t-1$, $\{s_0^{t-1,i}, w_1^{t-1,i}\}_{i=1}^N$ distributed according to $p(W_1^{t-1}, S_0 | y^{t-1}; \gamma)$, we can use $p(W_1^t, S_0 | y^t; \gamma)$ to generate proposal draws $\{s_0^{t|t-1,i}, w_1^{t|t-1,i}\}_{i=1}^N$ from $p(W_1^t, S_0 | y^{t-1}; \gamma)$. Then, we can exploit Proposition 1 and resample from $\{s_0^{t|t-1,i}, w_1^{t|t-1,i}\}_{i=1}^N$ a new swarm of particles, $\{s_0^{t,i}, w_1^{t,i}\}_{i=1}^N$, distributed according to $p(W_1^t, S_0 | y^t; \gamma)$. The output of the algorithm $\{s_0^{t|t-1,i}, w_1^{t|t-1,i}\}_{i=1}^N$ is used to compute the likelihood (6).

This is the key step in the filter. A naive extension of Monte Carlo techniques diverges as T grows because only one particle will eventually accumulate all the information. To avoid this problem, we do not carry over all the simulations to the next period. We keep those with higher probability of explaining the data.

The following pseudocode summarizes the description of the algorithm:

-
- Step 0, Initialization:** Set $t \rightsquigarrow 1$. Initialize $p(W_1^{t-1}, S_0|y^{t-1}; \gamma) = p(S_0; \gamma)$.
- Step 1, Prediction:** Sample N values $\{s_0^{t|t-1,i}, w_1^{t|t-1,i}\}_{i=1}^N$ from the conditional density $p(W_1^t, S_0|y^{t-1}; \gamma) = p(W_{1,t}; \gamma)p(W_1^{t-1}, S_0|y^{t-1}; \gamma)$.
- Step 2, Filtering:** Assign to each draw $(s_0^{t|t-1,i}, w_1^{t|t-1,i})$ the weight q_i^t as defined in Proposition 1.
- Step 3, Sampling:** Sample N times with replacement from $\{s_0^{t|t-1,i}, w_1^{t|t-1,i}\}_{i=1}^N$ with probabilities $\{q_i^t\}_{i=1}^N$. Call each draw $(s_0^{t,i}, w_1^{t,i})$. If $t < T$ set $t \rightsquigarrow t + 1$ and go to step 1. Otherwise stop.
-

An important point is that the algorithm can be implemented on a good desktop computer. All programs needed for the computation of the model and the sequential Monte Carlo were coded in Fortran 95 and compiled in Intel Fortran 8.0 to run on Windows-based machines. On a Pentium 4 at 3.00 GHz, each draw from the posterior using the sequential Monte Carlo with 60,000 particles takes around 6.1 s. That implies a total of about 88 h for each simulation of 50,000 draws. This relatively low computational cost opens the nonlinear estimation of dynamic equilibrium models to practitioners, and should be seen as an important selling point of our procedure.¹

It is important to note that we are presenting here only a basic sequential Monte Carlo filter and that the literature has presented several refinements to improve efficiency (see, for example, Kitagawa, 1996; Pitt and Shephard, 1999). The interested reader can find further details, comparison with alternative schemes and a discussion of convergence in Fernández-Villaverde and Rubio-Ramírez (2004).

2.3. The Linear Approach: The Kalman Filter

Now we describe how to evaluate the likelihood function implied by (3) and (4) with the Kalman filter.

To apply this filter, we need to assume that $\{W_t\}$ and $\{V_t\}$ are both normally distributed. Therefore, we can define $\tilde{W}_t = B(\gamma)W_t$ and $\tilde{V}_t = D(\gamma)V_t$ to be normal with distributions $\tilde{W}_t \sim N(0, Q(\gamma))$ and $\tilde{V}_t \sim N(0, R(\gamma))$.

The formulae of the Kalman filter allow us to recursively build a linear forecast of y_t , called $y_{t|t-1}$, based on previous observations. Also, it implies that the difference between y_t and its prediction $y_{t|t-1}$ is normally distributed with zero mean and variance $\Sigma_{t|t-1}$. In the interest of space we refer the reader to Harvey (1989) for details regarding how to compute $y_{t|t-1}$ and $\Sigma_{t|t-1}$. Suffice it to say that we need to keep track of the linear prediction of s_t (called $s_{t|t-1}$) and its variance $P_{t|t-1}$. The output of the Kalman filter can be used to calculate the likelihood function as follows:

$$p(y^T; \gamma) \simeq p_{KF}(y^T; \gamma) = \prod_{t=1}^T \frac{1}{\sqrt{(2\pi)^n |\Sigma_{t|t-1}|}} \exp\left(-\frac{(y_t - y_{t|t-1})' \Sigma_{t|t-1}^{-1} (y_t - y_{t|t-1})}{2}\right) \quad (7)$$

¹ All of the code is available upon request from the corresponding author.

We can also check that, if the transition and measurement equations of the model are linear and the innovations normally distributed, the sequential Monte Carlo and the Kalman filter deliver *precisely* the same value for the likelihood function (up to a numerical error).

3. AN APPLICATION

This section presents an example of how to implement the two alternatives described above to evaluate the likelihood function of a dynamic equilibrium model. We select the neoclassical growth model as our application. The reasons are twofold. First, this environment is the workhorse of quantitative macroeconomics. In this way, we perform our comparison in an application that is ‘representative’ of a large number of papers. Since any lesson learned is conditional on our particular model, we want to deal with a case that can be partially extrapolated to other setups. Second, the application of the two procedures delivers answers that are substantially different even if the model is nearly linear. The neoclassical growth model is a simple environment where we can make our main point. For a more nonlinear model, the disparities are more striking.

The rest of the section is organized as follows. First, we introduce the neoclassical growth model. Second, we discuss our linear and nonlinear approaches to solution methods. Third, we compute $p_{SMC}(y^T; \gamma)$ and $p_{KF}(y^T; \gamma)$.

3.1. The Neoclassical Growth Model

We work with the stochastic neoclassical growth model with leisure. Since this model is well known (see Cooley and Prescott, 1995), we present it only to fix notation.

There is a representative agent in the economy, whose preferences about consumption c_t and leisure l_t are represented by the utility function

$$U = E_0 \sum_{t=1}^{\infty} \beta^{t-1} \frac{(c_t^\theta (1-l_t)^{1-\theta})^{1-\tau}}{1-\tau}$$

where $\beta \in (0, 1)$ is the discount factor, τ controls the elasticity of intertemporal substitution, θ pins down labour supply, and E_0 is the conditional expectation operator.

The only good of this economy is produced according to the production function $e^{z_t} k_t^\alpha l_t^{1-\alpha}$ where k_t is the aggregate capital stock, l_t is the aggregate labour input and z_t is the technology level. z_t follows an AR(1), $z_t = \rho z_{t-1} + \varepsilon_t$ with $\varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon)$. We consider the stationary case (i.e., $|\rho| < 1$). The law of motion for capital is $k_{t+1} = i_t + (1-\delta)k_t$, where i_t is investment. Finally, the economy satisfies the resource constraint $c_t + i_t = e^{z_t} k_t^\alpha l_t^{1-\alpha}$.

A competitive equilibrium can be defined in a standard way. Since both welfare theorems hold, we can solve the equivalent and simpler social planner’s problem. We can think about this problem as finding policy functions for consumption $c(\cdot, \cdot)$, labour $l(\cdot, \cdot)$ and next period’s capital $k'(\cdot, \cdot)$ that deliver the optimal choices as functions of the two state variables, capital and the technology level.

3.2. The Solution Methods

The sequential Monte Carlo filter is independent of the particular nonlinear solution method employed. Aruoba *et al.* (2003) document that the finite elements method delivers an accurate,

fast and stable solution for a wide range of parameter values in a model exactly like the one considered here. Therefore, we choose this method for our nonlinear approach. Details of how to implement the finite elements method are also provided by Aruoba *et al.* (2003). For the linearized approach, the situation is easier, since all the methods existing in the literature (conditional on applicability) deliver *exactly* the same solution. Out of pure convenience, we use the undetermined coefficients procedure described by Uhlig (1999).

3.3. Evaluating $p_{SMC}(y^T; \gamma)$

Let $\gamma^1 \equiv (\theta, \rho, \tau, \alpha, \delta, \beta, \sigma_\varepsilon) \in \Upsilon^1 \subset R^7$ be the parameters of the neoclassical growth model. Since the finite element method requires the shocks be bounded between -1 and 1 , we transform the productivity shock as $\lambda_t = \tanh(z_t)$. Let $S_t = (k_t, \lambda_t)$ be the states of the model and set $W_t = \varepsilon_t$. Define $V_t \sim \mathcal{N}(0, \Sigma)$ as the vector of measurement errors. To economize on parameters we assume that Σ is diagonal with entries σ_1^2, σ_2^2 and σ_3^2 . Define $\gamma^2 = (\sigma_1^2, \sigma_2^2, \sigma_3^2) \in \Upsilon^2 \subset R_+^3$ and $\gamma = (\gamma^1, \gamma^2) \in \Upsilon$. Finally, call the approximated labour policy function $l_{fem}(\cdot, \cdot; \gamma)$ and capital accumulation $k_{fem}(\cdot, \cdot; \gamma)$, where we make the dependence on the structural parameter values explicit.

The transition equation for this model is

$$\begin{aligned} k_{t+1} &= f_1(S_{t-1}, W_t; \gamma) = k_{fem}(k_t, \tanh^{-1}(\lambda_t); \gamma) \\ \lambda_t &= f_2(S_{t-1}, W_t; \gamma) = \tanh(\rho \tanh^{-1}(\lambda_{t-1}) + \varepsilon_t) \end{aligned}$$

We assume that the observed time series, y_t , has three components: output, $gd p_t$; hours worked, $hours_t$; and gross investment, inv_t . The measurement equation is then

$$\begin{aligned} gd p_t &= g_1(S_t, V_t; \gamma) = x_1(S_t; \gamma) + V_{1,t} = e^{\tanh^{-1}(\lambda_t)} k_t^\alpha l_{fem}(k_t, \tanh^{-1}(\lambda_t); \gamma)^{1-\alpha} + V_{1,t} \\ hours_t &= g_2(S_t, V_t; \gamma) = x_2(S_t; \gamma) + V_{2,t} = l_{fem}(k_t, \tanh^{-1}(\lambda_t); \gamma) + V_{2,t} \\ inv_t &= g_3(S_t, V_t; \gamma) = x_3(S_t; \gamma) + V_{3,t} = k_{fem}(k_t, \tanh^{-1}(\lambda_t); \gamma) - (1 - \delta)k_t + V_{3,t} \end{aligned}$$

We comment on two assumptions made for convenience: the observables and the presence of measurement error. First, the selection of observables keeps the dimensionality of the problem low while capturing some of the most important dynamics of the data. Three dimensions will be enough to document the differences between the two filters. Second, we add measurement errors to avoid stochastic singularity. Nothing in our procedure critically depends on the presence of measurement errors. For example, we could instead work with a version of the model with shocks to technology, preferences and depreciation. This alternative environment might be more empirically interesting but it would make the solution of the model much more complicated. Since our goal here is to evaluate the impact of linearization on estimation, we follow the simple route.

Given that we have four sources of uncertainty, we set $\dim(W_{2,t}) = 0$ and $W_{1,t} = W_t = \varepsilon_t$. Drawing from $p(W_{1,t}; \gamma)$ is then equivalent to drawing from a normal with mean zero and standard deviation σ_ε . Since $S_t = g(S_{t-1}, W_t; \gamma)$ the reader can note, first, that $p(y_t | W_{1,t}^t, S_0, y^{t-1}; \gamma) = p(y_t | S_t; \gamma)$ and second, that drawing from $p(W_{1,t}^t, S_0 | y^{t-1}; \gamma)$ is equivalent to sampling from $p(S_t | y^{t-1}; \gamma)$. This allows us to write the likelihood function as

$$p(y^T; \gamma) = \prod_{t=1}^T \int p(y_t | S_t; \gamma) p(S_t | y^{t-1}; \gamma) dS_t \quad (8)$$

But since our measurement equation implies that $p(y_i|S_i; \gamma) = (2\pi)^{-\frac{3}{2}} |\Sigma|^{-\frac{1}{2}} e^{-\frac{\omega(S_i; \gamma)}{2}}$ where $\omega(S_i; \gamma) = (y_i - x(S_i; \gamma))' \Sigma^{-1} (y_i - x(S_i; \gamma)) \forall t$, we can rewrite (8) as

$$p(y^T; \gamma) = (2\pi)^{-\frac{3T}{2}} |\Sigma|^{-\frac{T}{2}} \prod_{t=1}^T \int e^{-\frac{\omega(S_t; \gamma)}{2}} p(S_t|y^{t-1}; \gamma) dS_t$$

The last expression is simple to handle. With the particles $\{s_0^{t|t-1, i}, w_1^{t|t-1, i}\}_{i=1}^N$ coming from our filter, we build the states $\{s_i^t\}_{i=1}^N$ and the prediction error $\{\omega(s_i^t; \gamma)\}_{i=1}^N$. We set $s_{0|0}^i = S_{ss} \forall i$. Therefore, the likelihood is approximated by

$$p(y^T; \gamma) \simeq p_{SMC}(y^T; \gamma) = (2\pi)^{-\frac{3T}{2}} |\Sigma|^{-\frac{T}{2}} \prod_{t=1}^T \frac{1}{N} \sum_{i=1}^N e^{-\frac{\omega(s_i^t; \gamma)}{2}} \tag{9}$$

3.4. Evaluating $p_{KF}(y^T; \gamma)$

Let γ, W_t and V_t be defined as in Section 3.3. The undetermined coefficients method delivers linear functions of output, investment and labour, depending on current capital and the technology level. Then we build the transition equation for the model and, as in the previous case, the measurement equations $gdp_t = y_t + V_{1,t}, hours_t = l_t + V_{2,t}$ and $inv_t = y_t - c_t + V_{3,t}$. Mapping these equations into the notation of equations (3) and (4), we apply the Kalman filter and evaluate $p_{KF}(y^T; \gamma)$ as described in Section 2.3. We initialize the filter to $s_{0|0} = S_{ss}$ and $P_{0|0} = 0$.

4. THE ESTIMATION ALGORITHM

Now we explain how to incorporate the likelihood functions in a Bayesian estimation algorithm. In the Bayesian approach, the main inference tool is the parameters' posterior distribution given the data, $\pi(\gamma|y^T)$. The posterior density is proportional to the likelihood times the prior. Therefore, we need to specify priors on the parameters, $\pi(\gamma)$, and to evaluate the likelihood function.

We specify our priors in Section 5.1, and the likelihood function is evaluated either by (6) or by (7), depending on how we solve the model. Since none of these posteriors have a closed-form, we use a Metropolis–Hasting algorithm to draw from them. We call $\pi_{SMC}(\gamma|y^T)$ the posterior implied by the sequential Monte Carlo filter and $\pi_{KF}(\gamma|y^T)$ the posterior derived from the Kalman filter. To simplify the notation, we let $f_{SMC}(\cdot, \cdot; \gamma_i)$ and $g_{SMC}(\cdot, \cdot; \gamma_i)$ be defined by (1) and (2), and $f_{KF}(\cdot, \cdot; \gamma_i)$ and $g_{KF}(\cdot, \cdot; \gamma_i)$ by (3) and (4).

The algorithm to draw a chain $\{\gamma_i\}_{i=1}^M$ from $\pi_j(\gamma|y^T) \forall j \in \{SMC, KF\}$ is as follows:

Step 0, Initialization: Set $i \sim 0$ and initial γ_i . Compute functions $f_j(\cdot, \cdot; \gamma_i)$ and $g_j(\cdot, \cdot; \gamma_i)$. Evaluate $\pi(\gamma_i)$ and $p_j(y^T; \gamma_i)$ using (6) or (7). Set $i \sim i + 1$.

Step 1, Proposal draw: Get a proposal draw $\gamma_i^p = \gamma_{i-1} + v_i$, where $v_i \sim N(0, \Psi)$.

Step 2, Solving the model: Solve the model for γ_i^p and compute $f_j(\cdot, \cdot; \gamma_i^p)$ and $g_j(\cdot, \cdot; \gamma_i^p)$.

Step 3, Evaluating the proposal: Evaluate $\pi(\gamma_i^p)$ and $p_j(y^T | \gamma_i^p)$ using either (6) or (7).

Step 4, Accept/reject: Draw $\chi_i \sim U(0, 1)$. If $\chi_i \leq \frac{\pi(\gamma_i^p) p_j(y^T | \gamma_i^p)}{\pi(\gamma_{i-1}) p_j(y^T | \gamma_{i-1})}$ set $\gamma_i = \gamma_i^p$, otherwise $\gamma_i = \gamma_{i-1}$. If $i < M$, set $i \rightsquigarrow i + 1$ and go to step 1. Otherwise stop.

We used standard methods to check the convergence of the chain generated by the Metropolis–Hasting algorithm (see Mengersen *et al.*, 1999). Also, we selected the variance of the innovations in the proposals for the parameters to achieve an acceptance rate of proposals of around 40%.

We concentrate in this paper on Bayesian inference. However, we could also perform classical inference, maximizing the likelihood function obtained in the previous section, and building an asymptotic variance–covariance matrix using standard numerical methods. Also, the value of the likelihood function at its maximum would be useful to compute likelihood ratios for model comparison purposes.

5. FINDINGS

We undertake two main exercises. In the first one, we simulate ‘artificial’ data using the nonlinear solution of the model for a particular choice of values of γ^* . Then, we define some priors over γ , and draw from its posterior distribution implied by both $p_{SMC}(y^T; \gamma)$ and $p_{KF}(y^T; \gamma)$. Finally, we compute the marginal likelihood of the ‘artificial’ data implied by each likelihood approximation. This exercise answers the following two questions: (1) How accurate is the estimation of the ‘true’ parameter values, γ^* , implied by each filter? (2) How big is the improvement delivered by the sequential Monte Carlo filter over the Kalman filter? From the posterior means, we address the first question. From the marginal likelihoods, we respond to the second.

Since the difference between the policy functions implied by the finite elements and the linear method depends greatly on γ^* , we perform the described exercise for two different values of γ^* , one with low risk aversion and low variance, γ_l^* , when both policies are close, and another with high risk aversion and high variance, γ_h^* , when the policies are farther away.

Our second exercise uses US data to estimate the model with the sequential Monte Carlo and the Kalman filters. This exercise answers the following question: Is the Sequential Monte Carlo providing a better explanation of the real data?

We divide our exposition into three parts. First, we specify the priors for the parameters. Second, we present results from the ‘artificial’ data experiment. Finally, we present the results with US data.

5.1. The Priors

We postulate flat priors for all 10 parameters, subject to some boundary constraints to make the priors proper. This choice is motivated by two considerations. First, since we are going to estimate our model using ‘artificial’ data generated at some value γ^* , we do not want to bias the results in favour of any alternative by our choice of priors. Second, with a flat prior, the posterior is proportional to the likelihood function (except for the very small issue of the bounded support of the priors). As a consequence, our experiment can be interpreted as a classical exercise in which the mode of the likelihood function is the maximum likelihood estimate. A researcher who prefers more informative priors can always reweight the likelihood to accommodate her priors (see Geweke, 1998).

The parameter governing labour supply, θ , follows a uniform distribution between 0 and 1. That constraint imposes only a positive marginal utility of leisure. The persistence of the technology shock, ρ , also follows a uniform distribution between 0 and 1. The parameter τ follows a uniform distribution between 0 and 100. That choice rules out only risk-loving behaviour and risk aversions that will predict differences in interest rates several orders of magnitude higher than the observed ones. The prior for the technology parameter, α , is uniform between 0 and 1. The prior on the depreciation rate ranges between 0 and 0.05, covering all national accounts estimates of quarterly depreciation. The discount factor, β , is allowed to vary between 0.75 and 1, implying steady-state annual interest rates between 0 and 316%. The standard deviation of the innovation of productivity, σ_ε , follows a uniform distribution between 0 and 0.1, a bound 15 times higher than the usual estimates. We also pick this prior for the three standard deviations of the measurement errors. Table I summarizes our discussion.

5.2. Results with ‘Artificial’ Data

We simulate observations from the model and use them as data for the estimation. We generate data from two different calibrations.

First, to make our experiment as realistic as possible, we present a benchmark calibration of the model. The discount factor $\beta = 0.9896$ matches an annual interest rate of 4.27%. The risk aversion $\tau = 2$ is a common choice in the literature. $\theta = 0.357$ matches the micro evidence of labour supply. We reproduce the labour share of national income with $\alpha = 0.4$. The depreciation rate $\delta = 0.02$ fixes the investment/output ratio and $\rho = 0.95$ and $\sigma = 0.007$ match the historical properties of the Solow residual of the US economy. With respect to the standard deviations of the measurement errors, we set them equal to 0.01% of the steady-state value of output, 0.35% of the steady-state value of hours and 0.2% of the steady-state value of investment based on our priors regarding the relative importance of measurement errors in the National Income and Product Accounts. We summarize the chosen values in Table II.

Table I. Priors for the parameters of the model

Parameter	Distribution	Hyperparameters
θ	Uniform	0, 1
ρ	Uniform	0, 1
τ	Uniform	0, 100
α	Uniform	0, 1
δ	Uniform	0, 0.05
β	Uniform	0.75, 1
σ_ε	Uniform	0, 0.1
σ_1	Uniform	0, 0.1
σ_2	Uniform	0, 0.1
σ_3	Uniform	0, 0.1

Table II. Calibrated parameters

Parameter	θ	ρ	τ	α	δ	β	σ_ε	σ_1	σ_2	σ_3
Value	0.357	0.95	2.0	0.4	0.02	0.99	0.007	1.58×10^{-4}	0.0011	8.66×10^{-4}

The second calibration, which we will call extreme, maintains the same parameters except that it increases τ to 50 (implying a relative risk aversion of 24.5) and σ_ε to 0.035. This high risk aversion and variance introduce a strong nonlinearity to the economy. This particular choice of parameters allows us to check the differences between the sequential Monte Carlo filter and the Kalman filter in a highly nonlinear world while maintaining a familiar framework. We justify our choice, thus, not basing it on empirical considerations but on its usefulness as a 'test' case.

After generating a sample of size 100 for each of the two calibrations,² we apply our priors and our likelihood evaluation algorithms. For the sequential Monte Carlo filter, we use 60,000 particles to get 50,000 draws from the posterior distribution. Since we do not suffer from an attrition problem, we do not replenish the swarm. See Fernández-Villaverde and Rubio-Ramírez (2004) for further details of this issue and of convergence. For the Kalman filter, we also get 50,000 draws. In both cases, we have a long burn-in period.

In Figure 1 we plot the likelihood function in logs of the model, given our simulated data for the sequential Monte Carlo filter (continuous line) and the Kalman filter (discontinuous line). We draw in each panel the likelihood function for an interval around the calibrated value of the structural parameter, keeping all the other parameters fixed at their calibrated values. We can think of each panel then as a transversal cut of the likelihood function. To facilitate the comparison, we show the 'true' value for the parameter corresponding to the direction being plotted with a vertical line, and we do not draw values lower than $-20,000$.

Figure 1 reveals two points. First, both likelihoods have the same shape and are centred on the 'true' value of the parameter, although the Kalman filter delivers a slight bias for four parameters (α , δ , β and θ). Note that, since we are assuming flat priors, none of the curvature of the likelihoods is coming from the prior. Second, there is a difference in level between the likelihood generated by the sequential Monte Carlo filter and the one delivered by the Kalman filter. This is a first proof that the nonlinear model fits the data better even for this nearly linear economy.

Table III conveys similar information: the point estimates are approximately equal regardless of the filter. However, the sequential Monte Carlo delivers estimates that are better in the sense of being closer to their pseudotrue value.³

Table IV reports the logmarginal likelihood differences between the nonlinear and the linear case. We compute the marginal likelihood with Geweke's (1998) harmonic mean proposal. Consequently, we need to specify a bound on the support of the weight density. To ensure robustness, we report the distances for a range of values of the truncation value p from 0.1 to 0.9. All the values convey the same message: the nonlinear solution method fits the data two orders of magnitude better than the linear approximation. A good way to read this number is to use Jeffreys' (1961) rule: if one hypothesis is more than 100 times more likely than the other, the evidence is decisive in its favour. This translates into differences in logmarginal likelihoods of 4.6 or higher. Our value of 73.6 is, then, well beyond decisiveness in favour of nonlinear filtering.

Fernández-Villaverde *et al.* (2004) provide a theoretical explanation for this finding. They show how the bound on the error in the likelihood induced by the linear approximation of the policy

² The results were robust when we used different simulated data from the same model. We omit details because of space considerations.

³ The whole posteriors are available upon request from the authors. We also checked that the numerical errors of the estimates were negligible.

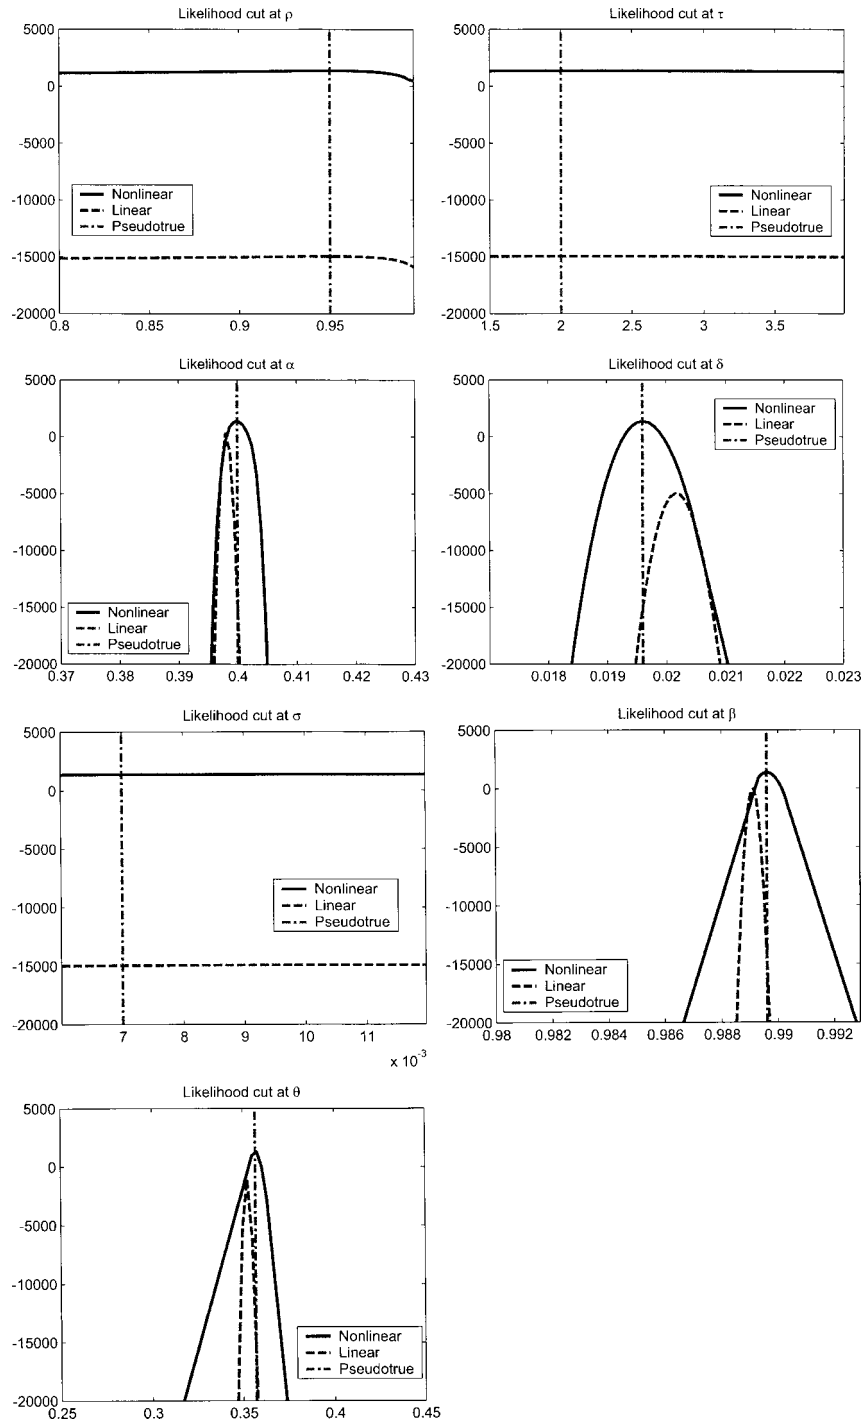


Figure 1. Likelihood function, benchmark calibration

Table III. Nonlinear versus linear posterior distributions benchmark case

Nonlinear (SMC filter)			Linear (Kalman filter)		
Parameter	Mean	s.d.	Parameter	Mean	s.d.
θ	0.357	0.10×10^{-3}	θ	0.374	0.06×10^{-3}
ρ	0.950	0.29×10^{-3}	ρ	0.914	0.24×10^{-3}
τ	2.000	0.92×10^{-3}	τ	3.536	1.17×10^{-3}
α	0.400	0.10×10^{-3}	α	0.443	0.08×10^{-3}
δ	0.020	0.02×10^{-3}	δ	0.030	0.02×10^{-3}
β	0.990	0.02×10^{-3}	β	0.978	0.03×10^{-3}
σ_ε	0.007	0.04×10^{-4}	σ_ε	0.011	0.36×10^{-4}
σ_1	1.58×10^{-4}	0.15×10^{-6}	σ_1	1.86×10^{-4}	4.64×10^{-7}
σ_2	1.12×10^{-3}	0.68×10^{-6}	σ_2	5.55×10^{-4}	2.01×10^{-6}
σ_3	5.64×10^{-4}	0.81×10^{-6}	σ_3	2.42×10^{-3}	1.75×10^{-6}

Table IV. Logmarginal likelihood difference benchmark case

p	Nonlinear vs linear
0.1	73.631
0.5	73.627
0.9	73.603

function gets compounded with the size of the sample. The intuition is as follows. Small errors in the policy function accumulate at the same rate as the sample size grows. This means that, as the sample size goes to infinity, a linear approximation will deliver an approximation of the likelihood that will fail to converge.

We now move to study the results for the extreme calibration. Figure 2 is equivalent to Figure 1 for the extreme case. First, note how the likelihood generated by the sequential Monte Carlo filter is again centred on the 'true' value of the parameter. In comparison, the likelihood generated by the Kalman filter is not. These differences will have an important impact on the marginal likelihood. Table V recasts the same information in terms of means and standard deviations of the posteriors. As in the benchmark case, the sequential Monte Carlo delivers better estimates of the parameters of the model.

Table VI reports the logmarginal likelihood differences between the nonlinear and the linear case for the extreme calibration for different p 's. Again, we can see how the evidence in favour of the nonlinear filter is overwhelming.

As a conclusion, our exercise shows how even for a nearly linear case such as the neoclassical growth model, an estimation that respects the nonlinear structure of the economy improves substantially the ability of the model to fit the data. This may indicate that we greatly handicap dynamic equilibrium economies when we linearize them before taking them to the data and that some empirical rejections of these models may be due to the biases introduced by linearization.

Our results do not imply, however, that we should completely abandon linear methods. We have also shown that their accuracy for point estimates is acceptable. For some exercises where only point estimates are required, the extra computational cost of the sequential Monte Carlo filter

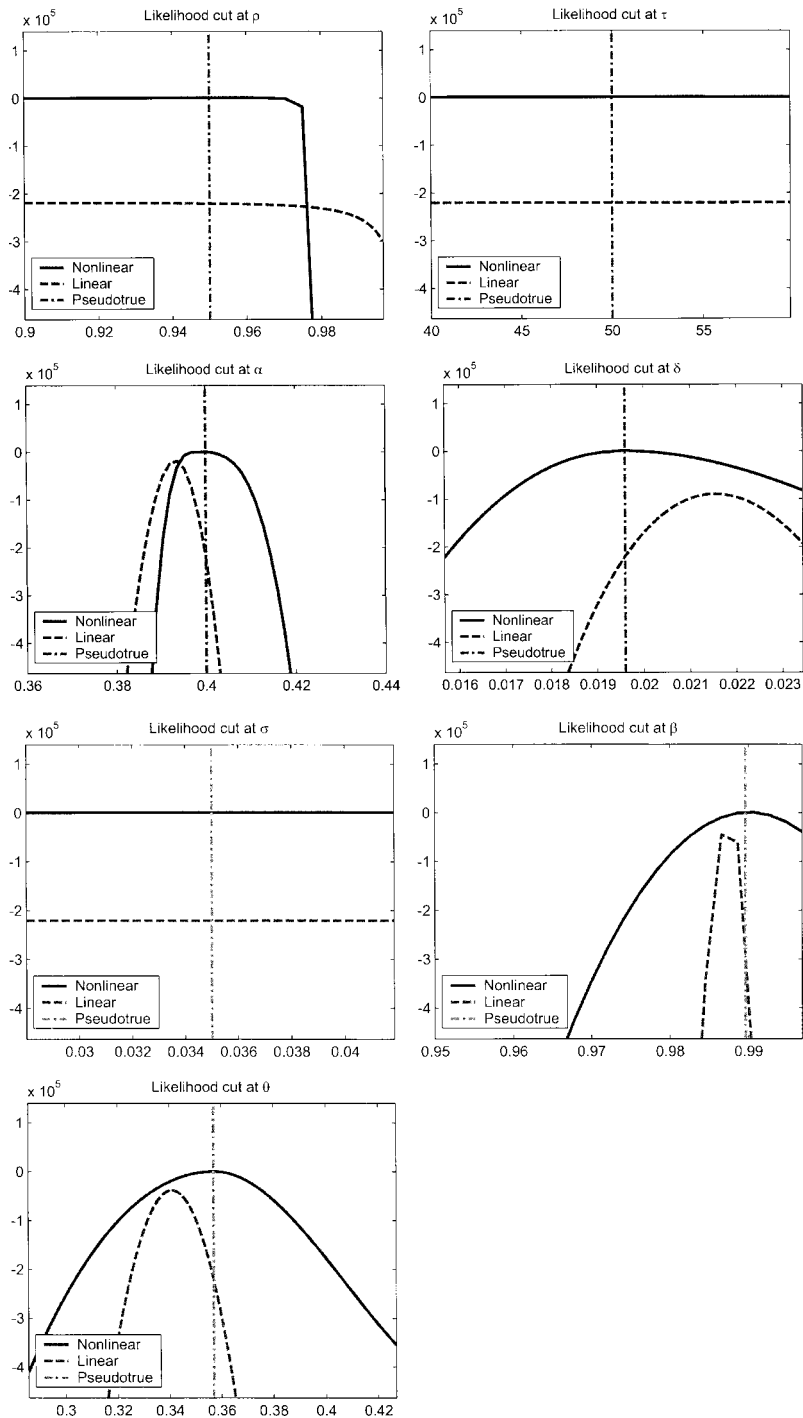


Figure 2. Likelihood function, extreme calibration

Table V. Nonlinear versus linear posterior distributions extreme case

Nonlinear (SMC filter)			Linear (Kalman filter)		
Parameter	Mean	s.d.	Parameter	Mean	s.d.
θ	0.357	0.08×10^{-3}	θ	0.337	0.06×10^{-3}
ρ	0.950	0.17×10^{-3}	ρ	0.894	0.19×10^{-3}
τ	50.000	0.24×10^{-1}	τ	67.70	0.10×10^{-1}
α	0.400	0.05×10^{-3}	α	0.346	0.04×10^{-3}
δ	0.020	0.05×10^{-4}	δ	0.010	0.02×10^{-4}
β	0.990	0.08×10^{-4}	β	0.996	0.08×10^{-4}
σ_e	3.50×10^{-2}	0.03×10^{-4}	σ_e	3.61×10^{-2}	0.09×10^{-4}
σ_1	1.58×10^{-4}	0.06×10^{-6}	σ_1	1.72×10^{-4}	0.05×10^{-6}
σ_2	1.12×10^{-3}	0.05×10^{-5}	σ_2	9.08×10^{-4}	0.03×10^{-5}
σ_3	8.66×10^{-4}	0.02×10^{-5}	σ_3	2.64×10^{-3}	0.04×10^{-5}

Table VI. Logmarginal likelihood difference extreme case

p	Nonlinear vs linear
0.1	117.608
0.5	117.592
0.9	117.564

may not compensate for the reduction in bias. Practitioners should weight the advantages and disadvantages of each procedure in their particular application.

5.3. Results with Real Data

Now we estimate the neoclassical growth model with US quarterly data. We use real output per capita, average hours worked and real gross fixed investment per capita from 1964:Q1 to 2003:Q1. We first remove a trend from the data using an H-P filter. In this way, we do not need to model explicitly the presence of a trend and its possible changes.

Table VII presents the results from the posterior distributions from 100,000 draws for each filter, again after a long burn-in period. The discount factor, β , is estimated to be 0.997 with the nonlinear filter and 0.973 with the Kalman filter. This is an important difference when using quarterly data. The linear model compensates for the lack of curvature induced by its certainty equivalence with more impatience. The parameter controlling the elasticity of substitution, τ , is estimated by the nonlinear filter to be 1.717 and by the Kalman filter to be 1.965. The parameter α is close to the canonical value of one-third in the case of the sequential Monte Carlo, and higher (0.412) in the case of the Kalman filter. Finally, we note how the standard deviation of the parameters is estimated to be much higher when we use the nonlinear filter than when we employ the Kalman filter, indicating that the nonlinear likelihood is more dispersed.

It is difficult to assess whether the differences in point estimates documented in Table VII are big or small. A possible answer is based on the impact of the different estimates on the moments generated by the model. Macroeconomists often use these moments to evaluate the model's ability to account for the data. Table VIII presents the moments of the real data and reports the moments

Table VII. Nonlinear versus linear posterior distributions real data

Nonlinear (SMC filter)			Linear (Kalman filter)		
Parameter	Mean	s.d.	Parameter	Mean	s.d.
θ	0.390	0.11×10^{-2}	θ	0.423	0.19×10^{-3}
ρ	0.978	0.52×10^{-2}	ρ	0.941	0.27×10^{-3}
τ	1.717	0.12×10^{-1}	τ	1.965	0.12×10^{-2}
α	0.324	0.71×10^{-3}	α	0.412	0.37×10^{-3}
δ	0.006	0.36×10^{-4}	δ	0.019	0.79×10^{-5}
β	0.997	0.92×10^{-4}	β	0.973	0.43×10^{-5}
σ_e	0.020	0.11×10^{-3}	σ_e	0.009	0.43×10^{-5}
σ_1	0.45×10^{-1}	0.42×10^{-3}	σ_1	0.11×10^{-3}	0.18×10^{-6}
σ_2	0.15×10^{-1}	0.25×10^{-3}	σ_2	0.83×10^{-2}	0.34×10^{-5}
σ_3	0.38×10^{-1}	0.39×10^{-3}	σ_3	0.26×10^{-1}	0.18×10^{-4}

Table VIII. Nonlinear versus linear moments real data

	Real data		Nonlinear (SMC filter)		Linear (Kalman filter)	
	Mean	s.d.	Mean	s.d.	Mean	s.d.
<i>output</i>	1.95	0.073	1.91	0.129	1.61	0.068
<i>hours</i>	0.36	0.014	0.36	0.023	0.34	0.004
<i>inv</i>	0.42	0.066	0.44	0.073	0.28	0.044

that the stochastic neoclassical growth model generates by simulation when we calibrated it at the mean of the posterior distribution of the parameters given by each of the two filters.

We highlight two observations from Table VIII. First, the nonlinear model matches the data much better than the linearized one. This difference is significant because the moments are nearly identical if we simulate the model using the linear or the nonlinear solution method with the same set of parameter values. The differences come thus from the point estimates delivered by each procedure. The nonlinear estimation nails down the mean of each of the three observables and does a fairly good job with the standard deviations. Second, the estimation by the nonlinear filter implies a higher output, investment and hours worked than the estimation by the linear filter.

The main reason for these two differences is the higher β estimated by the sequential Monte Carlo. The lower discount factor induces a higher accumulation of capital and, consequently, a higher output, investment and hours worked. The differences for the standard deviation of the economy are also important. The nonlinear economy is also more volatile than the linearized model in terms of the standard deviation of output and hours.

Table IX reports the logmarginal likelihood differences between the nonlinear and the linear case. As in the previous cases, the real data strongly support the nonlinear version of the economy with differences in log terms of around 93. The differences in moments discussed above are one of the main driving forces behind the finding. A second force is that the likelihood function generated by the sequential Monte Carlo is less concentrated than that coming from the Kalman filter.⁴

⁴ Another way to think about the marginal likelihood is as a measure of the ability of the model to forecast within sample. The much higher marginal likelihood of the nonlinear model indeed translates into a better forecasting record within the

Table IX. Logmarginal likelihood difference real data

p	Nonlinear vs linear
0.1	93.65
0.5	93.55
0.9	93.55

6. CONCLUSIONS

We have compared the effects of estimating dynamic equilibrium models using a sequential Monte Carlo filter proposed by Fernández-Villaverde and Rubio-Ramírez (2004) and a Kalman filter. The sequential Monte Carlo filter exploits the nonlinear structure of the economy and evaluates the likelihood function of the model by simulation methods. The Kalman filter estimates a linearization of the economy around the deterministic steady state. The advantage of the Kalman filter is its simplicity and speed. We compare both methodologies using the neoclassical growth model. We report two main results. First, both for simulated and for real data, the sequential Monte Carlo filter delivers a substantially better fit of the model to the data. This difference exists even for a nearly linear case. Second, the differences in terms of point estimates, even if small in absolute terms, have quite important effects on the moments of the model. From these two results we conclude that the nonlinear filter is superior as a procedure for taking models to the data.

An additional advantage of the sequential Monte Carlo filter is that it allows the estimation of nonnormal economies. Nonnormalities or the presence of stochastic volatility may be important to account for the dynamics of macro data. Future research will address how much accuracy is gained with the use of a sequential Monte Carlo filter when estimating this class of models.

ACKNOWLEDGEMENTS

We thank Sally Burke, Will Roberds, Chris Sims, Tao Zha, participants at several seminars, the editor and two anonymous referees for useful comments. Valuable assistance was provided by the University of Minnesota Supercomputer Institute. Beyond the usual disclaimer, we must note that any views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Atlanta or of the Federal Reserve System.

REFERENCES

- Anderson EW, Hansen LP, McGrattan ER, Sargent TJ. 1996. On the mechanics of forming and estimating dynamic linear economies. In *Handbook of Computational Economics*, Amman HM *et al.* (eds). Elsevier: Amsterdam.
- Aruoba SB, Fernández-Villaverde J, Rubio-Ramírez J. 2003. Comparing solution methods for dynamic equilibrium economies. Federal Reserve Bank of Atlanta Working Paper 2003–27.

sample (this is also true for simulated data). We omit details regarding this superior forecasting power of the nonlinear procedure because of space constraints.

- Bouakez H, Cardia E, Ruge-Murcia FJ. 2002. Habit formation and the persistence of monetary shocks. Bank of Canada Working Paper 2002–27.
- Cooley TF (ed.). 1995. *Frontiers of Business Cycle Research*, Princeton University Press: Princeton, NJ.
- Cooley TF, Prescott EC. 1995. Economic growth and business cycles. In *Frontiers of Business Cycle Research*, Cooley TF (ed.). Princeton University Press: Princeton, NJ.
- DeJong DN, Ingram BF, Whiteman CH. 2000. A Bayesian approach to dynamic macroeconomics. *Journal of Econometrics* **98**: 203–223.
- Dib A. 2001. An estimated Canadian DSGE model with nominal and real rigidities. Bank of Canada Working Paper 2001–26.
- Doucet A, de Freitas N, Gordon N. 2001. *Sequential Monte Carlo Methods in Practice*. Springer-Verlag: New York.
- Fernández-Villaverde J, Rubio-Ramírez J. 2004. Comparing dynamic equilibrium models to data: a Bayesian approach. *Journal of Econometrics* **123**: 153–187.
- Fernández-Villaverde J, Rubio-Ramírez J. 2004. Estimating nonlinear dynamic equilibrium economies: a likelihood approach. Mimeo, University of Pennsylvania.
- Fernández-Villaverde J, Rubio-Ramírez J, Santos M. 2004. Convergence properties of the likelihood of computed dynamic equilibrium models. Mimeo, University of Pennsylvania.
- Geweke J. 1998. Using simulation methods for Bayesian econometric models: inference, development and communication. Federal Reserve Bank of Minneapolis Staff Report 249.
- Gordon NJ, Salmund DJ, Smith AFM. 1993. Novel approaches to NonLinear/Non-Gaussian Bayesian State Estimation. *IEE Proceedings-F* **140**: 107–113.
- Hall GJ. 1996. Overtime, effort and the propagation of business cycle shocks. *Journal of Monetary Economics* **38**: 139–160.
- Harvey AC. 1989. *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press: Cambridge.
- Ireland P. 2002. Technology shocks in the new Keynesian model. Mimeo, Boston College.
- Jeffreys H. 1961. *Theory of Probability*, 3rd edn. Oxford University Press: Oxford.
- Kim C, Nelson CR. 1999. *State-Space Models with Regime Switching*. MIT Press: Boston, MA.
- Kim C, Piger J. 2002. Nonlinearity and the permanent effects of recessions. Federal Reserve Bank of St. Louis Working Paper 2002-014E.
- Kim J. 2000. Constructing and estimating a realistic optimizing model of monetary policy. *Journal of Monetary Economics* **45**: 329–359.
- Kim J, Kim S, Schaumburg E, Sims C. 2003. Calculating and using second order accurate solutions of discrete time dynamic equilibrium models. Mimeo, Princeton University.
- Kitagawa G. 1996. Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *Journal of Computational and Graphical Statistics* **5**: 1–25.
- Landon-Lane J. 1999. Bayesian comparison of dynamic macroeconomic models. PhD thesis, University of Minnesota.
- Lubik T, Schorfheide F. 2003. Do central banks respond to exchange rates? A structural investigation? Mimeo, University of Pennsylvania.
- McGrattan E, Rogerson R, Wright R. 1997. An equilibrium model of the business cycle with household production and fiscal policy. *International Economic Review* **33**: 573–601.
- Mengersen KL, Robert CP, Guihenneuc-Jouyaux C. 1999. MCMC convergence diagnostics: a ‘reviewww’. In *Bayesian Statistics 6*, Berger J, Bernardo J, Dawid AP, Smith AFM (eds). Oxford Sciences Publications: Oxford.
- Moran K, Dolar V. 2002. Estimated DGE models and forecasting accuracy: a preliminary investigation with Canadian data. Bank of Canada Working Paper 2002–18.
- Otrok C. 2001. On measuring the welfare cost of business cycles. *Journal of Monetary Economics* **47**: 61–92.
- Pitt MK, Shephard N. 1999. Filtering via simulation: auxiliary particle filters. *Journal of the American Statistical Association* **94**: 590–599.
- Rabanal P, Rubio-Ramírez J. 2003. Comparing new Keynesian models of the business cycle: a Bayesian approach. Federal Reserve Bank of Atlanta Working Paper 2003–30.
- Rust J. 1994. Structural estimation of Markov decision processes. In *Handbook of Econometrics*, Vol. 4, Engle RF, McFadden DL (eds). North Holland: Amsterdam.

- Sargent TJ. 1989. Two models of measurements and the investment accelerator. *Journal of Political Economy* **97**: 251–287.
- Schorfheide F. 2000. Loss function-based evaluation of DGSE models. *Journal of Applied Econometrics* **15**: 645–670.
- Sims CA, Zha T. 2002. Macroeconomic switching. Mimeo, Princeton University.
- Smets F, Wouters R. 2003. Shocks and frictions in US business cycle fluctuations: a Bayesian DSGE approach. Mimeo, European Central Bank.
- Uhlig H. 1999. A toolkit for analyzing nonlinear dynamic stochastic models easily. In *Computational Methods for the Study of Dynamic Economies*, Marimon R, Scott A. (eds). Oxford University Press: Oxford.
- Vuong QH. 1989. Likelihood ratio test for model selection and non-nested hypothesis. *Econometrica* **57**: 307–333.