Journal of Econometrics 185 (2015) 216–229

Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

Estimating dynamic equilibrium models with stochastic volatility

Jesús Fernández-Villaverde [a], Pablo Guerrón-Quintana [b,*], Juan F. Rubio-Ramírez [c]

[a] University of Pennsylvania, NBER, and CEPR, United States

[b] Federal Reserve Bank of Philadelphia, United States

[c] Duke University, Federal Reserve Bank of Atlanta, BBVA Research and FEDEA, United States

ARTICLE INFO

Article history:
Received 21 May 2013
Received in revised form
31 January 2014
Accepted 18 August 2014
Available online 18 October 2014

JEL classification:
E10
E30
C11

Keywords:
Dynamic equilibrium models
Stochastic volatility
Parameter drifting
Bayesian methods

ABSTRACT

This paper develops a particle filtering algorithm to estimate dynamic equilibrium models with stochastic volatility using a likelihood-based approach. The algorithm, which exploits the structure and profusion of shocks in stochastic volatility models, is versatile and computationally tractable even in large-scale models. As an application, we use our algorithm and Bayesian methods to estimate a business cycle model of the US economy with both stochastic volatility and parameter drifting in monetary policy. Our application shows the importance of stochastic volatility in accounting for the dynamics of the data.

1. Introduction

This paper develops a particle filtering algorithm to estimate dynamic equilibrium models with stochastic volatility using a likelihood-based approach. The novelty of our algorithm is that it does not require the presence of linear measurement errors to evaluate the likelihood function of the model. In order to do that, we characterize the properties of the solution of these models when approximated with the second-order expansion. As an application of our procedure, we estimate a medium-size business cycle economy.

Our results are useful because, motivated by the findings of Stock and Watson (2002) and Sims and Zha (2006), many recent papers have built dynamic equilibrium models with volatility shocks (also known as uncertainty shocks). Among them, we can highlight Fernández-Villaverde and Rubio-Ramirez (2007), Justiniano and Primiceri (2008), Bloom (2009), and Fernández-Villaverde et al. (2010b). In these models, and in the tradition of stochastic volatility (Shephard, 2008), there are two types of shocks: structural shocks (shock to productivity, to preferences, etc.) and volatility shocks (shocks to the standard deviation of the innovations to the structural shocks).

To fulfill the promise in this literature, we need tools to estimate this class of models. However, the task is complicated by the inherent non-linearity that stochastic volatility generates. Linearization is ill-equipped to handle time-varying volatility because it yields certainty-equivalent policy functions. That is, volatility influences neither the agents' decision rules nor the laws of motion of the aggregate variables. Hence, to consider how stochastic volatility affects those factors, it is imperative to employ at least the second-order approximation to the equilibrium dynamics of the economy and to use simulation-based estimators of the likelihood.

To accomplish that last task, one could, in principle, rely on the baseline particle filter presented in Fernández-Villaverde and Rubio-Ramirez (2007). Unfortunately, that version of the particle filter requires, when estimating models with stochastic volatility, the presence of linear measurement errors in observables. Otherwise, we would be forced to solve a large quadratic system of equations with multiple solutions, an endeavor for which there are no suitable algorithms. Although measurement errors are plausible, they complicate identification in small samples and entangle the interpretation of the results.

* Corresponding author.
E-mail addresses: jesusfv@econ.upenn.edu (J. Fernández-Villaverde), pablo.guerron@phil.frb.org (P. Guerrón-Quintana), juan.rubio-ramirez@duke.edu (J.F. Rubio-Ramírez).

http://dx.doi.org/10.1016/j.jeconom.2014.08.010
0304-4076/© 2014 Elsevier B.V. All rights reserved.

To get around this problem, we show how to write an alternative particle filter that exploits the structure of the second-order approximation to the equilibrium dynamics of an economy with stochastic volatility without the need of linear measurement errors. Second-order approximations accurately capture important implications of stochastic volatility and are convenient because they are not computationally expensive.

We proceed in two steps. First, we characterize the second-order approximation to the decision rules of a dynamic equilibrium model with stochastic volatility. Second, we demonstrate how to use this characterization to write the alternative particle filter. The key is to show how the quadratic problem associated with the evaluation of the approximated measurement density is reduced to a much simpler linear problem that only involves a matrix inversion. After we have evaluated the likelihood, we can combine it with a prior and a Markov chain Monte Carlo (McMc) algorithm to draw from the posterior distribution.

Our characterization of the second-order approximation to the decision rules is also of interest in itself. Among other things, it is useful to analyze the equilibrium of the model, to explore the shape of its impulse response functions, or to calibrate it. More concretely, we prove that:

1. The first-order approximation to the decision rules of the agents (or any other equilibrium object of interest) does not depend on volatility shocks and they are certainty equivalent.
2. The second-order approximation to the decision rules of the agents only depends on volatility shocks on terms where volatility is multiplied by the innovation to its own structural shock. For instance, if we have a productivity shock and a volatility shock to it, the only non-zero term where the volatility shock to productivity appears is the one where the volatility shock multiplies the innovation to the productivity shock. Thus, only a few of the terms in the second-order approximation are non-zero.
3. The perturbation parameter will only appear in a non-zero term where it is raised to a square. This term is a constant that corrects for precautionary behavior induced by risk.

As an application, we estimate a business cycle model of the US economy. The model incorporates stochastic volatility in the shocks that drive its dynamics and parameter drifting in the parameters that control monetary policy. In that way, we include two of the main mechanisms that researchers have highlighted to account for the time-varying volatility of US time series – heteroscedastic shocks and parameter drifting – and let the likelihood decide which of them better accounts for the data. Last, we have a model that is as rich as many of the models employed in modern quantitative macroeconomics. While estimating such a large model is a computational challenge, we wanted to demonstrate that our procedure is of practical use and to make our application a blueprint for the estimation of other dynamic equilibrium models.

Our main empirical findings are as follows. First, the posterior distribution of the parameters puts most of its mass in areas that denote a fair amount of stochastic volatility. Second, a model comparison exercise indicates that, even after controlling for stochastic volatility, the data prefer a specification where monetary policy changes over time. This finding should not be interpreted, though, as implying that volatility shocks did not play a role. It means, instead, that a successful model of the US economy requires the presence of both stochastic volatility and parameter drifting, a result that challenges the results of Sims and Zha (2006). Finally, we document the evolution of the structural shocks, of stochastic volatility, and the parameters of monetary policy. We emphasize the confluence, during the 1970s, of times of high volatility and weak responses to inflation, and, during the 1990s, of positive structural shocks and low volatility even if monetary policy was weaker than often argued. In the appendix, we construct counterfactual histories of the US data by varying some aspect of the model such as shutting down time-varying volatility or imposing alternative monetary policies.

An alternative to our stochastic volatility framework would be to work with Markov regime-switching models such as those of Bianchi (2009) or Farmer et al. (2009). These models provide a promising extra degree of flexibility in modeling aggregate dynamics. In fact, some of the fast changes in policy parameters that we document in our empirical section suggest that discrete jumps could be a good representation of the data. We hope to undertake in the future a more careful assessment of the advantages and disadvantages of stochastic volatility versus Markov regime-switching models.

Finally, even if the motivation for our approach and the application belong to macroeconomics, the tools we present are not specific to that field. One can think about the importance of estimating dynamic equilibrium models with stochastic volatility in many other fields such as finance (Bansal and Yaron, 2004) or international economics (Fernández-Villaverde et al., 2010b).

The rest of the paper is organized as follows. Section 2 introduces a generic dynamic equilibrium model with stochastic volatility to fix notation and discuss how to solve it. Section 3 explains the evaluation of the likelihood of the model. Section 4 compares our approach with continuous-time methods. Section 5 presents our application. Section 6 concludes. An extensive technical appendix includes additional material (see Appendix A).

2. Dynamic equilibrium models with stochastic volatility

2.1. The model

The set of equilibrium conditions of a wide class of dynamic equilibrium models can be written as

$$\mathbb{E}_t f\left(\mathcal{Y}_{t+1}, \mathcal{Y}_t, \mathcal{S}_{t+1}, \mathcal{S}_t, \mathcal{Z}_{t+1}, \mathcal{Z}_t; \gamma\right) = 0, \tag{1}$$

where $\mathbb{E}_t$ is the conditional expectation operator at time t, $\mathcal{Y}_t = (\mathcal{Y}_{1t}, \mathcal{Y}_{2t}, \ldots, \mathcal{Y}_{kt})'$ is the $k \times 1$ vector of observables at time t, $\mathcal{S}_t = (\mathcal{S}_{1t}, \mathcal{S}_{2t}, \ldots, \mathcal{S}_{nt})'$ is the $n \times 1$ vector of endogenous states at time t, $\mathcal{Z}_t = (\mathcal{Z}_{1t}, \mathcal{Z}_{2t}, \ldots, \mathcal{Z}_{mt})'$ is the $m \times 1$ vector of structural shocks at time t, f maps $\mathbb{R}^{2(k+n+m)}$ into $\mathbb{R}^{k+n+m}$, and γ is the $n_\gamma \times 1$ vector of parameters that describe preferences and technology. In this paper, γ is also the vector of parameters to be estimated.

We will consider models where $\mathcal{Z}_{it+1}$ follow a stochastic volatility process of the form

$$\mathcal{Z}_{it+1} = \rho_i \mathcal{Z}_{it} + \Lambda \sigma_i \sigma_{it+1} \varepsilon_{it+1} \tag{2}$$

for all $i \in \{1, \ldots, m\}$, where Λ is a perturbation parameter, σ_i is the mean volatility, and $\log \sigma_{it+1}$, the percentage deviation of the standard deviation of the innovations to the structural shocks with respect to its mean, evolves as

$$\log \sigma_{it+1} = \vartheta_i \log \sigma_{it} + \Lambda \left(1 - \vartheta_i^2\right)^{\frac{1}{2}} \eta_i u_{it+1} \tag{3}$$

for all $i \in \{1, \ldots, m\}$. The combination of levels in (2) and logs in (3) ensures a positive σ_{it+1}. We multiply the innovation in (3) by $\left(1 - \vartheta_i^2\right)^{\frac{1}{2}}$ to normalize its size by the persistence of σ_{it}. It will be clear momentarily why we specify (2) and (3) in terms of the perturbation parameter Λ. It is also convenient to write, for all $i \in \{1, \ldots, m\}$, the laws of motions for $\mathcal{Z}_{it}$ and $\log \sigma_{it}$

$$\mathcal{Z}_{it} = \rho_i \mathcal{Z}_{it-1} + \sigma_i \sigma_{it} \varepsilon_{it} \tag{4}$$

and

$$\log \sigma_{it} = \vartheta_i \log \sigma_{it-1} + \left(1 - \vartheta_i^2\right)^{\frac{1}{2}} \eta_i u_{it}. \tag{5}$$