

Our Research Agenda: Estimating DSGE Models

Jesús Fernández-Villaverde

Duke University, NBER, and CEPR

Juan F. Rubio-Ramírez

Duke University and Federal Reserve Bank of Atlanta

October 5, 2006

1. Introduction

Our research agenda has focused on the estimation of dynamic stochastic general equilibrium (DSGE) models. In particular, we have worked on the likelihood-based approach to inference.

DSGE models are the standard tool of quantitative macroeconomics. We use them to organize our thinking, to measure the importance of different phenomena, and to provide policy prescriptions. However, since Kydland and Prescott's immensely influential 1982 paper, the profession has fought about how to take these models to the data. Three issues are at stake: first, how to determine the values of the parameters that describe preferences and technology (the unfortunately named "structural" parameters); second, how to measure the fit of the model; and third, how to decide which of the existing theories better accounts for the observed data.

Kydland and Prescott proposed to "calibrate" their model, i.e., to select parameter values by matching some moments of the data and by borrowing from microeconomic evidence. Calibration was a reasonable choice at the time. Macroeconomists were unsure about how to compute their models efficiently, a necessary condition to perform likelihood-based inference. Moreover, even if economists had known how to do so, most of the techniques required for estimating DSGE models using a likelihood approach did not exist. Finally, as recalled by Sargent (2005), the early results on estimation brought much disappointment. The models were being blown out of the water by likelihood ratio tests despite the feeling that those models could teach practitioners important lessons. Calibration offered a way out. By focusing only on a very limited set of moments of the model, researchers could claim success and keep developing the theory.

The landscape changed dramatically in the 1990s. There were developments along three fronts. First, macroeconomists learned how to efficiently compute equilibrium models with rich dynamics. There is not much point in estimating very stylized models that do not even have a remote chance of fitting the data well. Second, statisticians developed simulation techniques like Markov chain Monte Carlo (MCMC), which we require to estimate DSGE models. Third, and perhaps most important, computer power has become so cheap and readily available that we can now do things that were unthinkable 20 years ago.

One of the things we can now do is to estimate non-linear and/or non-normal DSGE models using a likelihood approach. This statement begets two questions: 1) Why do we want to estimate those DSGE models? and 2) How do we do it?

2. Why Do We Want to Estimate Non-linear and/or Non-normal DSGE Models?

Let us begin with some background. There are many reasons why the likelihood estimation of DSGE models is an important topic. First of all, a rational expectations equilibrium is a likelihood function. Therefore, if you trust your model, you have to trust its likelihood. Second, the likelihood approach provides a coherent and systematic procedure to estimate all the parameters of interest. The calibration approach may have made sense back in the 1980s when we had only a small bundle of parameters to select values for. However, current models are richly parameterized. Neither a loose application of the method of moments (which is what moment matching in calibration amounts to) nor some disparate collection of microeconomic estimates will provide us with the discipline to quantify the behavior of the model. Parameters do not have a life of their own: their estimated values are always conditional on one particular model. Hence, we cannot import these estimated values from one model to another. Finally, the likelihood yields excellent asymptotic properties and sound small sample behavior.

However, likelihood-based estimation suffers from a fundamental problem: the need to evaluate the likelihood function of the DSGE model. Except in a few cases, there is no analytical or numerical procedure to write down the likelihood.

The standard solution in the literature has been to find the linear approximation to the policy functions of the model. If, in addition, we assume that the shocks to the economy are normally distributed, we can apply the Kalman filter and evaluate the likelihood implied by the approximated policy functions. This strategy depends on the accuracy of the approximation of the exact policy functions by a linear relation and on the presence of normal shocks. Each of those two assumptions is problematic.

2.1. Linear Policy Functions

When we talk about linearization, the first temptation is to sweep it under the rug as a small numerical error. However, the impact of linearization is grimmer than it looks. We explore this assertion in our paper “Convergence Properties of the Likelihood of Computed Dynamic Models”, published in *Econometrica* and coauthored with Manuel Santos. In that paper, we prove that second order approximation errors in the policy function, like those generated by linearization, have first order effects on the likelihood function. Moreover, we demonstrate that the error in the approximated likelihood is compounded with the size of the sample.

Period by period, small errors in the policy function accumulate at the same rate at which the sample size grows. Thus, the approximated likelihood diverges from the exact one as we get more and more observations.

We have documented how those theoretical insights are quantitatively relevant for real-life applications. The main piece of evidence is in our paper “Estimating Dynamic Equilibrium Economies: Linear versus Nonlinear Likelihood”, published in the *Journal of Applied Econometrics*. The paper compares the results of estimating the linearized version of a DSGE model with the results from estimating the non-linear version. In the first case, we evaluate the likelihood of the model with the Kalman filter. In the second case, we evaluate the likelihood with the particle filter (which we will discuss below). Our findings highlight how linearization has a non-trivial impact on inference. First, both for simulated and for U.S. data, the non-linear version of the model fits the data substantially better. This is true even for a nearly linear case. Second, the differences in terms of point estimates, although relatively small in absolute values, have substantive effects on the behavior of the model.

Other researchers have found similar results when they take DSGE models to the data. We particularly like the work of Amisano and Tristani (2005) and An (2005). Both papers investigate New Keynesian models. They find that the non-linear estimation allows them to identify more structural parameters, to fit the data better, and to obtain more accurate estimates of the welfare effects of monetary policies.

2.2. Normal Shocks

The second requirement for applying the Kalman filter to estimate DSGE models is the assumption that the shocks driving the economy are normally distributed. Since nearly all DSGE models make this assumption, this requirement may not look dangerous. This impression is wrong: normality is extremely restrictive.

Researchers put normal shocks in their models out of convenience, not for any substantive reason. In fact, fat tails are such a pervasive feature of the data that normality is implausible. More thoughtful treatments of the shocks deliver huge benefits. For example, the fit of an ARMA process to U.S. output data improves dramatically when the innovations are distributed as student-t’s (a density with fat tails) instead of normal ones (Geweke, 1993 and 1994).

A simple way to generate fat tails, and one that captures the evidence of volatility clustering in the data, is to have time-varying volatility in the shocks. Why macroeconomists have not focused more effort on the topic is a puzzle. After all, Engle (1982), in the first

work on time-varying volatility, picked as his application of the ARCH model the process for United Kingdom inflation. However, that route was not followed. Even today, and beyond our own work on the issue, only Justiniano and Primiceri (2005) take seriously the idea that shocks in a DSGE model may have a richer structure than normal innovations.

Time-varying volatility of the shocks is not only a device to achieve a better fit, it is key to understanding economic facts. Think about the “Great Moderation.” Kim and Nelson (1999), McConnell and Pérez-Quirós (2000), and Stock and Watson (2002) have documented a decline in the variance of output growth since the mid 1980s. Moreover, there is a narrowing gap between growth rates during booms and recessions. What has caused the change in observed aggregate volatility? Was it due to better conducting of monetary policy by the Fed? Or was it because we did not suffer large shocks like the oil crises of the 1970s? We can answer that question only if we estimate structural models where we let both the monetary policy rule and the volatility of the shocks evolve over time. We will elaborate below on how to explore policy change as a particular case of parameter drifting.

There are two possibilities to introduce time-varying variance in shocks. One is stochastic volatility. The other one is Markov regime-switching models. We have worked more on the first approach since it is easier to handle. However, as we will explain below, we are currently exploring the second one.

A common feature of both stochastic volatility and regime-switching models is that they induce fundamental non-linearities and fat tails. Linearization, by construction, precludes any possibility of assessing time-varying volatility. If we linearize the laws of motion for the shocks, as someone who wanted to rely on the Kalman filter would be forced to do, the volatility terms would drop. Justiniano and Primiceri (2005) have got around that problem by pioneering the use of partially linear models in a specially clever way. Unfortunately, there is only so much we can do even with partially linear models. We need a general procedure to tackle non-linear and/or non-normal problems.

3. How Do We Do It?

Our previous arguments point out the need to evaluate the likelihood function of the non-linear and/or non-normal solution of DSGE models. But, how can we do that? This is where our paper, “Estimating Macroeconomic Models: A Likelihood Approach,” comes in. This paper shows how a simulation technique known as the particle filter allows us to evaluate that likelihood function. Once we have the likelihood, we can estimate the parameters of the

model by maximizing the likelihood (if you are a classical econometrician) or by combining the likelihood with a prior density for the model parameter to form a posterior distribution (if you are a Bayesian one). Also, we can compare how well different economies explain the data with likelihood ratio tests or Bayes factors.

The particle filter is a sequential Monte Carlo method that tracks the unobservable distribution of states of a dynamic model conditional on observables. The reason we are keenly interested in tracking such distribution is that, with it, we can obtain a consistent evaluation of the likelihood of the model using a straightforward application of the law of the large numbers.

The particle filter substitutes the population conditional distribution of states, which is difficult if not impossible to characterize, by an empirical distribution generated by simulation. The twist of ingenuity of the particle filter is that the simulation is generated through a device known as sequential importance resampling (SIR). SIR ensures that the Monte Carlo method achieves sufficient accuracy in a reasonable amount of time. Hence, the particle filter delivers the key object that we need to estimate non-linear and/or non-normal DSGE models: an efficient evaluation of the likelihood function of the model.

To illustrate our method, we follow Greenwood, Hercowitz, and Krusell (1997 and 2000). These authors have vigorously defended the importance of technological change specific to new investment goods for understanding postwar U.S. growth and aggregate fluctuations. We estimate a version of their business cycle model. The model has three shocks: to preferences, to neutral technology, and to investment-specific technology. All three shocks display stochastic volatility. Also, there are two unit roots and cointegration relations derived from the balanced growth path properties of the economy. We solve the model using second order approximations and apply the particle filter to evaluate the likelihood function.

The data reveal three facts. First, there is strong evidence for the presence of stochastic volatility in U.S. data. Capturing this phenomenon notably improves the fit of the model. Second, the decline in aggregate volatility has been a gradual trend and not, as suggested by the literature, the result of an abrupt drop in the mid 1980s. The fall in volatility started in the late 1950s, was interrupted in the late 1960s and early 1970s, and resumed around 1979. Third, changes in the volatility of preference shocks account for most of the variation in the volatility of output growth over the last 50 years.

Summarizing, our paper shows how to conduct an estimation of non-linear and/or non-normal DSGE models, that such estimation is feasible in real life, and that it helps us to obtain many answers we could not otherwise generate.

4. Complementary Papers

Parallel to our main line of estimation of non-linear and/or non-normal DSGE models, we have written other papers that complement our work.

The first paper in this line of research is “Comparing Dynamic Equilibrium Economies to Data: a Bayesian Approach,” published in the *Journal of Econometrics*. This paper studies the properties of the Bayesian approach to estimation and comparison of dynamic economies. First, we show that Bayesian methods have a classical interpretation: asymptotically, the parameter point estimates converge to their pseudotrue values, and the best model under the Kullback-Leibler distance will have the highest posterior probability. Both results hold even if the models are non-nested, misspecified, and non-linear. Second, we illustrate the strong small sample behavior of the approach using a well-known example: the U.S. cattle cycle. Bayesian estimates outperform maximum likelihood, and the proposed model is easily compared with a set of Bayesian vector autoregressions.

A second paper we would like to mention is “A,B,C’s (and D)’s for Understanding VARs”, written with Thomas Sargent and Mark Watson. This paper analyzes the connections between DSGE models and vector autoregressions (VARs), a popular empirical strategy. An approximation to the equilibrium of a DSGE model can be expressed in terms of a linear state space system. An associated linear state space system determines a vector autoregression for observables available to an econometrician. We provide a simple algebraic condition to check whether the impulse response of the VAR resembles the impulse response associated with the economic model. If the condition does not hold, the interpretation exercises done with VARs are misleading. Also, the paper describes many interesting links between DSGE models and empirical representations. Finally, we give four examples that illustrate how the condition works in practice.

In “Comparing Solution Methods for Dynamic Equilibrium Economies”, published in the *Journal of Economic Dynamics and Control* and joint with Boragan Aruoba, of the University of Maryland, we assess different solution methods for DSGE models. This comparison is relevant because when we estimate DSGE models, we want to solve them quickly and accurately. In the paper, we compute and simulate the stochastic neoclassical growth model with leisure choice by implementing first, second, and fifth order perturbations in levels and in logs, the finite elements method, Chebyshev polynomials, and value function iteration for several calibrations. We document the performance of the methods in terms of computing time, implementation complexity, and accuracy, and we present some conclusions and pointers for future research.

This paper motivated us to think about the possibility of developing new and efficient solution techniques for dynamic models. A first outcome of this work has been “Solving DSGE Models with Perturbation Methods and a Change of Variables,” also published in the *Journal of Economic Dynamics and Control*. This paper explores the changes of variables technique to solve the stochastic neoclassical growth model with leisure choice. We build upon Kenn Judd’s idea of changing variables in the computed policy functions of the economy. The optimal change of variables for an exponential family reduces the average absolute Euler equation errors of the solution of the model by a factor of three. We demonstrate how changes of variables can correct for variations in the risk level of the economy even if we work with first-order approximations to the policy functions. Moreover, we can keep a linear representation of the laws of motion of the model if we employ a nearly optimal transformation. We finish by discussing how to employ our results to estimate DSGE models

5. What is Next?

The previous paragraphs were just a summary of the work we have done on the estimation of DSGE models. But there is plenty of work ahead of us.

Currently, we are working on a commissioned article for the NBER Macroeconomics Annual. This paper will study the following question: How stable over time are the so-called “structural parameters” of DSGE models? At the core of these models, we have the parameters that define the preferences and technology that describe the environment. Usually, we assume that these parameters are structural in the sense of Hurwicz (1962): they are invariant to interventions, including shocks by nature. Their invariance permits us to exploit the model fruitfully as a laboratory for quantitative analysis. At the same time, the profession is accumulating more and more evidence of parameter instability in dynamic models. We are undertaking the first systematic analysis of parameter instability in the context of a “state of the art” DSGE model. One important application of this research is that we can explore changes in monetary policy over time. If you model monetary policy as a feedback function, you can think about the policy change as a change in the parameters of that feedback function, i.e., as one particular example of parameter drifting.

A related project is our work on semi-nonparametric estimation of DSGE models. The recent DSGE models used by the profession are complicated structures. They rely on many parametric assumptions: utility function, production function, adjustment costs, structure of stochastic shocks, etc. Some of those parametric choices are based on restrictions imposed

by the data on theory. For example, the fact that labor income share has been relative constant since 1950s suggests a Cobb-Douglas production function. Unfortunately, many other parametric assumptions are not. Researchers choose parametric forms for those functions based only on convenience. How dependent are our findings on the previous parametric assumptions? Can we make more robust assumptions? Our conversations with Xiaohong Chen have convinced us that this is a worthwhile avenue of improvement. We are pursuing the estimation of DSGE models when we relax parametric assumptions along certain aspects of the model with the method of Sieves, which Xiaohong has passionately championed.

We would also like to better understand how to compute and estimate models with Markov regime-switching. Those models are a nice alternative to stochastic volatility models. They allow for less variation in volatility, hence gaining much efficiency. Also, they may better capture phenomena such as the abrupt break in U.S. interest rates in 1979. Regime-switching models present interesting challenges in terms of computation and estimation.

Finally, we are interested in the integration of microeconomic heterogeneity within estimated DSGE models. James Heckman has emphasized again and again that individual heterogeneity is the defining feature of micro data (see Browning, Hansen, and Heckman, 1999, for the empirical importance of individual heterogeneity and its relevance for macro-economists). Our macro models need to move away from the basic representative agent paradigm and include richer configurations. The work of Victor Ríos-Rull in this area has been pathbreaking. Of course, this raises the difficult challenge of how to effectively estimate these economies. We expect to tackle some of those difficulties in the near future.

References

- [1] An, S. (2005). “Bayesian Estimation of DSGE Models: Lessons from Second Order Approximations.” *Mimeo*, University of Pennsylvania.
- [2] Amisano, G. and O. Tristani (2005). “Euro Area Inflation Persistence in an Estimated Nonlinear DSGE Model.” *Mimeo*, ECB.
- [3] Aruoba, S.B., J. Fernández-Villaverde and J. Rubio-Ramírez (2006). “Comparing Solution Methods for Dynamic Equilibrium Economies.” *Journal of Economic Dynamics and Control* 30, 2447-2508.
- [4] Browning, M., L.P. Hansen, and J.J. Heckman (1999). “Micro Data and General Equilibrium Models.” in: J.B. Taylor & M. Woodford (eds.), *Handbook of Macroeconomics*, volume 1, chapter 8, pages 543-633 Elsevier.
- [5] Fernández-Villaverde, J. and J. Rubio-Ramírez (2004). “Comparing Dynamic Equilibrium Models to Data: A Bayesian Approach.” *Journal of Econometrics* 123, 153-187.
- [6] Fernández-Villaverde, J. and J. Rubio-Ramírez (2005a). “Estimating Dynamic Equilibrium Economies: Linear versus Nonlinear Likelihood.” *Journal of Applied Econometrics*, 20, 891-910.
- [7] Fernández-Villaverde, J. and J. Rubio-Ramírez (2005b). “Estimating Macroeconomic Models: A Likelihood Approach.” *NBER Technical Working Paper* T0308.
- [8] Fernández-Villaverde, J. and J. Rubio-Ramírez (2006). “Solving DSGE Models with Perturbation Methods and a Change of Variables.” *Journal of Economic Dynamics and Control* 30, 2509-2531.
- [9] Fernández-Villaverde, J., J. Rubio-Ramírez, T.J. Sargent, and M. Watson (2006). “A,B,C’s (and D)’s for Understanding VARs.” *Mimeo*, Duke University.
- [10] Fernández-Villaverde, J., J. Rubio-Ramírez, and M.S. Santos (2006). “Convergence Properties of the Likelihood of Computed Dynamic Models.” *Econometrica* 74, 93-119.
- [11] Geweke, J.F. (1993). “Bayesian Treatment of the Independent Student-t Linear Model.” *Journal of Applied Econometrics* 1993, 8, S19-S40.
- [12] Geweke, J.F. (1994). “Priors for Macroeconomic Time Series and Their Application.” *Econometric Theory* 10, 609-632.
- [13] Greenwood, J, Z. Hercowitz, and P. Krusell (1997). “Long-Run Implications of Investment-Specific Technological Change.” *American Economic Review* 87, 342-362.
- [14] Greenwood, J, Z. Hercowitz, and P. Krusell (2000). “The Role of Investment-Specific Technological Change in the Business Cycle.” *European Economic Review* 44, 91-115.

- [15] Hurwicz, L. (1962). “On the Structural Form of Interdependent Systems”. In E. Nagel, P. Suppes, and A. Tarski (eds.), *Logic, Methodology and Philosophy of Science*. Stanford University Press.
- [16] Justiniano A. and G.E. Primiceri (2005). “The Time Varying Volatility of Macroeconomic Fluctuations.” *Mimeo*, Northwestern.
- [17] Kim, C. and C.R. Nelson (1999) “Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle.” *Review of Economics and Statistics* 81, 608-616.
- [18] McConnell, M.M. and G. Pérez-Quirós (2000). “Output Fluctuations in the United States: What Has Changed Since the Early 1980’s?” *American Economic Review* 90, 1464-1476.
- [19] Sargent, T.J. (2005). “An Interview with Thomas J. Sargent by George W. Evans and Seppo Honkapohja.” *Macroeconomic Dynamics* 9, 561-583.
- [20] Stock, J.H. and M.W. Watson (2002). “Has the Business Cycle Changed, and Why?” *NBER Macroeconomics Annual* 17, 159-218.