Abstract

This essay discusses how Bayesian methods can be used to cope with challenges that arise in the econometric analysis of dynamic stochastic general equilibrium models and vector autoregressions.

Bayesian Methods in Macroeconometrics

Macroeconometrics encompasses a large variety of probability models for macroeconomic time series as well as estimation and inference procedures to study the determinants of economic growth, to examine the sources of business cycle fluctuations, to understand the propagation of shocks, generate forecasts, and to predict the effects of economic policy changes. Bayesian methods are a collection of inference procedures that permit researchers to combine initial information about models and their parameters with sample information in a logically coherent manner by use of Bayes' Theorem. Both prior and post-data information is represented by probability distributions.

Unfortunately, the term 'macroeconometrics' is often narrowly associated with large-scale system-of-equations models in the Cowles Commission tradition that were developed from the 1950s to the 1970s. These models came under attack on academic grounds in the mid 1970s. Lucas (1976) argued that the models are unreliable tools for policy analysis because they are unable to predict the effects of policy regime changes on the expectation formation of economic agents in a coherent manner. Sims (1980) criticized that many of the restrictions that are used to identify behavioral equations in these models are inconsistent with dynamic macroeconomic theories and proposed the use of vector autoregressions (VAR) as an alternative. Academic research on econometric models in the Cowles tradition reached a trough in the early 1980s and never recovered. The state-of-the-art is summarized in a monograph by Fair (1994).

I am adopting a modern view of macroeconometrics in this essay and will portray an active research area that is tied to modern dynamic macroeconomic theory. Reviewing Bayesian methods in macroeconometrics is a short essay is a difficult task. My review is selective and not representative of Bayesian time series analysis in general. I have chosen some topics that I believe are important, but the list is by no means exhaustive. I will focus on the question how Bayesian methods are used to address some of the challenges that arise in the econometric analysis of dynamic stochastic general equilibrium (DSGE) models and vector autoregressions (VAR). A more extensive treatment can be found in the survey article by An and Schorfheide (2006).

DSGE Models

The term DSGE model is often used to refer to a broad class of dynamic macroeconomic models that spans the standard neoclassical growth model discussed in King, Plosser, and Rebelo (1988) as well as the monetary model with numerous real and nominal frictions developed by Christiano, Eichenbaum, and Evans (2005).

A common feature of these models is that decision rules of economic agents are derived from assumptions about preferences and technologies by solving intertemporal optimization problems. Moreover, agents potentially face uncertainty with respect to, for instance, total factor productivity or the nominal interest rate set by a central bank. This uncertainty is generated by exogenous stochastic processes or shocks that shift technology or generate unanticipated deviations from a central bank's interest-rate feedback rule. Conditional on distributional assumptions for the exogenous shocks, the DSGE model generates a joint probability distribution for the endogenous model variables such as output, consumption, investment, and inflation.

What are the goals?

While macroeconometric methods are used to address many different questions, several issues stand out. Business cycle analysts are interested in identifying the sources of fluctuations, for instance, how important are monetary policy shocks for movements in aggregate output? We would like to understand the propagation of shocks, e.g., what happens to aggregate hours worked in response to a technology shock? Moreover, researchers ask questions about structural changes in the economy: has monetary policy changed in the early 1980s? Why did the volatility of many macroeconomic time series drop in the mid 1980s? Macroeconometricians are also interested in forecasting the future: how will inflation and output growth rates evolve over the next eight quarters? Finally, an important aspect of macroeconometrics is to predict the effect of policy changes: how will output and inflation respond to an unanticipated change in the nominal interest rate? Is it desirable to adopt an inflation targeting regime?

What are the challenges?

In principle one could proceed as follows: specify a DSGE model that is sufficiently rich to address the substantive economic question of interest; derive its likelihood function and fit the model to historical data; answer the questions based on the estimated DSGE model. Unfortunately, this is easier said than done. A trade-off between theoretical coherence and empirical fit poses the first challenge to macroeconometric analysis.

Under certain regularity conditions DSGE models can be well approximated by VARs that satisfy particular cross-coefficient restrictions. The DSGE model is misspecified if these restrictions are at odds with the data and the model has difficulties to track and forecast historical time series. Misspecification was quite apparent for the first generation of DSGE models and has led Kydland, Prescott, and their followers since the early 1980's to abandon formal econometric procedures and advocate a calibration approach, outlined for instance, in Kydland and Prescott (1996). Recent Bayesian and non-Bayesian research, however, has resulted in formal econometric tools that are general enough to explicitly account for misspecification problems that arise in the context of DSGE models. Examples of Bayesian approaches are Canova (1994), Dejong, Ingram, and Whiteman (1996), Geweke (1999), Schorfheide

(2000), Del Negro and Schorfheide (2004), and Del Negro, Schorfheide, Smets, and Wouters (2006).

The presence of misspecification might suggest that we should simply ignore the cross-coefficient restrictions implied by dynamic economic theories in the empirical work and try to answer the questions posed above directly by VARs. Unfortunately, there is no free lunch. VARs have many free parameters and without restrictions on their coefficients can lead to poor forecasts. VARs do not provide a tight economic interpretation of economic dynamics in terms of the behavior of rational, optimizing agents. Moreover, it is difficult to predict the effects of rare policy regime changes on the expectation formation and the behavior of economic agents since these are not explicitly modelled. While the most recent generation of DSGE models comes much closer to matching the empirical fit of VARs, as documented in Smets and Wouters (2003), a trade-off between theoretical coherence and empirical fit remains.

A second challenge is identification. The parameters of a model are identifiable if no two parameterizations of that model generate the same probability distribution for the observables. In VARs the mapping between the one-step-ahead forecast errors of the endogenous variables and the underlying structural shocks is not unique, and additional restrictions are necessary to identify, say, a monetary policy or a technology shock. Many of the popular identification schemes and the controversies surrounding them are surveyed in Cochrane (1994), Christiano and Eichenbaum (1999) and Stock and Watson (2001).

DSGE models can be locally approximated by linear rational expectations (LRE) models. While tightly parameterized compared to VARs, LRE models can generate delicate identification problems. Suppose a model implies that $y_t = \theta I\!E_t[y_{t+1}] + u_t$, where u_t is an independently distributed random variable with mean zero. If $0 \leq \theta < 1$ then the only stable law of motion for y_t that satisfies the rational expectations restrictions is $y_t = u_t$, which means that θ is not identifiable. More elaborate examples are discussed in Lubik and Schorfheide (2004, 2006), Beyer and Farmer (2004), and Canova and Sala (2005). Unfortunately, it is in many cases difficult to detect identification problems in DSGE models, since the mapping from the structural parameters into the autoregressive law of motion for y_t is highly nonlinear and typically can only be evaluated numerically.

Many regularities of macroeconomic time series are indicative of nonlinearities, for instance, the rise and fall of inflation in the 1970s and early 1980s and time-varying volatility of many macroeconomic time series, e.g., Cogley and Sargent (2005), Sargent, Williams, and Zha (2005), and Sims and Zha (2005). In VARs nonlinear dynamics are typically generated with time-varying coefficients, whereas most DSGE models are nonlinear and only for convenience approximated by linear rational expectations models. Conceptually the analysis of nonlinear models is very similar to the analysis of linear models, but the implementation of the computations is often more cumbersome and poses a third challenge.

How can Bayesian analysis help?

Bayesian analysis is conceptually straightforward. Pre-sample information about parameters is summarized by a prior distribution $p(\theta)$. We can also assign discrete probabilities to distinct models although the distinction between models and parameters is somewhat artificial. The prior is combined with the conditional distribution of the data given the parameters (likelihood function) $p(Y|\theta)$. The application of Bayes theorem yields the posterior model probabilities and parameter distributions $p(\theta|Y)$. Markov-Chain-Monte-Carlo methods can be used to generate θ draws from the posterior. Based on these draws one can numerically approximate the relevant moments of the posterior and make inference about taste and technology parameters as well as the relative importance and the propagation of the various shocks.

The literature on Bayesian estimation of DSGE models began with work by Landon-Lane (1998), DeJong, Ingram, and Whiteman (2000), Schorfheide (2000), and Otrok (2001). DeJong, Ingram, and Whiteman (2000) estimate a stochastic growth model and examine its forecasting performance, Otrok (2001) fits a real business cycle with habit formation and time-to-build to the data to assess the welfare costs of business cycles, and Schorfheide (2000) considers cash-in-advance monetary DSGE models. The Bayesian analysis of VAR dates at least back to Doan, Litterman, and Sims (1984).

Since DSGE models are to some extent micro-founded, macroeconomists require their parameterization to be consistent with microeconometric evidence on, for instance, labor supply elasticities and the frequency with which firms adjust their prices. If information in the estimation sample were abundant and model misspecification were not a concern, then there would be little need for a prior distribution that summarizes information contained in other data sets. However, in the estimation of DSGE model this additional information plays an important role.

The prior is used to down-weigh the likelihood function in regions of the parameter space that are inconsistent with out-of-sample information and in which the structural model becomes un-interpretable. The shift from prior to posterior can be an indicator of tensions between different sources of information. If the likelihood function peaks at a value that is at odds with, say, the micro-level information that has been used to construct the prior distribution then marginal data density $\int p(Y|\theta)p(\theta)d\theta$ will be low. If two models have equal prior probabilities then the ratio of their marginal data densities determine the posterior model odds. Hence, in a posterior odds comparison a DSGE model will automatically be penalized for not being able to reconcile two sources of information with a single set of parameters.

Identification problems manifest themselves through ridges and multiple peaks of equal height in the likelihood function. While Bayesian inference is based on the same likelihood function as classical maximum likelihood estimation, it can bring to bear additional information that may help to discriminate between different parameterizations of a model. If, for instance, the likelihood function is invariant to a subvector θ_1 of θ then the posterior distribution of θ_1 will simply equal to the prior distribution. Hence, a comparison of priors and posteriors can provide important insights about the extent to which the data provide information about the parameters of interest. Regardless, the posterior provides a coherent summary of pre-sample and sample information and can be used for inference and decision making. This insight has been used, for instance, by Lubik and Schorfheide (2004) to assess whether monetary policy in the 1970's was conducted in a way that would allow expectations to be self-fulfilling and cause business cycle fluctuations unrelated to fundamental shocks.

Bayesian inference is well suited for model comparisons. Under a loss function that is zero if the correct model is chosen and one otherwise, is optimal to select the model that has the highest posterior probability. However, in many applications, in particular related to the comparison of two possibly misspecified DSGE models, this zero-one loss function is not very attractive because it does provide little insights about the dimensions along which the structural models should be improved. Schorfheide (2000) provides a framework for the comparison of two or more potentially misspecified DSGE models. A VAR plays the role of a reference model. If the DSGE models are indeed misspecified the VAR will attain the highest posterior probability and the model comparison is based on the question: given a particular loss function, which DSGE model best mimics the dynamics captured by the VAR?

VARs typically have many more parameters than DSGE models and the role of prior distributions is mainly to reduce the effective dimensionality of this parameter space to avoid over-fitting. More interestingly, if one interprets the DSGE model as a set of restrictions on the VAR then the DSGE model induces a degenerate prior for the VAR coefficients. If the researcher is concerned about potential misspecification of the DSGE model, a natural approach is to relax the DSGE model restrictions and construct a non-degenerate prior distribution that concentrates most of its mass near the restrictions. This approach was originally proposed by Ingram and Whiteman (1994) and has been further developed by Del Negro and Schorfheide (2004) who provide a framework for the joint estimation of VAR and DSGE model parameters. The framework generates a continuum of intermediate specification that differ according to the degree by which the restrictions are relaxed. This degree is measured by a hyperparameter and the posterior distribution of the hyperparameter can be interpreted as a measure of fit.

Incorporating model and parameter uncertainty into a decision is straightforward in a Bayesian setup. Levin, Onatski, Williams, and Williams (2006), for instance, study the effect of optimal monetary policy under parameter uncertainty in the context of an estimated DSGE model. Let δ denote a decision, such as the choice of a monetary policy rule or a tax rate and $L(\delta, \theta)$ be a loss function that is used to evaluate the decision. The optimal choice minimizes the posterior risk $\int L(\delta, \theta) p(\theta|Y) d\theta$. The calculation of the risk is facilitated by Markov-Chain-Monte-Carlo methods that enable a numerical evaluation of expected losses. If the parameter θ in the loss function is replaced by a future observation y' and $p(\theta|Y)$ is replaced by the predictive distribution $p(y'|Y)$ the decision-theoretic framework can also be used to generate forecasts from the Bayes model.

Finally, with respect to the analysis of nonlinear models Bayesian methods are in some instances very helpful. Data-augmentation techniques let researchers efficiently deal with numerical complications that arise in models with latent state variables, such as regime-switching models or VARs with time-varying coefficients as in Cogley and Sargent (2005) and Sims and Zha (2005). On the other hand, the need to compute a likelihood function can create serious obstacles. For instance, the computation of the likelihood function for a DSGE model solved with a nonlinear solution method requires a computational-intensive particle filter as in Fernández-Villaverde and Rubio-Ramírez (2004).

Conclusion

The Bayesian paradigm provides a rich framework for inference and decision making with modern macroeconometric models such as DSGE models and VARs. The econometric methods can be tailored to cope with the challenges in this literature: potential model misspecification and a trade-off between theoretical coherence and empirical fit, identification problems, and estimation of models with many parameters based on relatively few observations. Advances in Bayesian computations let the researcher efficiently deal with numerical complications that arise in models with latent state variables, such as regime-switching models, or nonlinear state-space models.

Frank Schorfheide

See also Bayesian econometrics, likelihood function, Markov-Chain-Monte Carlo methods, vector autoregressions.

Bibliography

An, S. and F. Schorfheide 2006. Bayesian analysis of DSGE models. Econometric Reviews, forthcoming.

Beyer, A. and R. Farmer 2004. On the indeterminacy of New Keynesian economics. ECB Working Paper 323.

Canova, F. 1994. Statistical inference in calibrated models. Journal of Applied Econometrics 9, S123-S144.

Canova, F. and L. Sala 2005. Back to square one: identification issues in DSGE models. Manuscript, IGIER and Bocconi University.

Christiano, L. and M. Eichenbaum 1999. Monetary policy shocks: what have we learned and to what end? Handbook of Macroeconomics.

Christiano, L.J., M. Eichenbaum, and C. Evans 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. Journal of Political Economy 113, 1-45.

Cochrane, J. 1994. Shocks. Carnegie Rochester Conference Series 41, 295-364.

Cogley, T. and T. Sargent 2005. Drifts and volatilities: monetary policies and outcomes in the post WWII U.S. Review of Economic Dynamics 8, 262-302.

DeJong, D., B. Ingram and C. Whiteman 1996. A Bayesian approach to calibration. Journal of Business Economics and Statistics 14, 1-9.

DeJong, D., B. Ingram and C. Whiteman 2000. A Bayesian approach to dynamic macroeconomics. Journal of Econometrics 98, 201-223.

Del Negro, M. and F. Schorfheide 2004. Priors from equilibrium models for VARs. International Economic Review 45, 643-673.

Del Negro, M., F. Schorfheide, F. Smets, and R. Wouters 2006. On the Fit of New Keynesian Models. Journal of Business and Economic Statistics, forthcoming.

Doan, T., R. Litterman, and C. Sims 1984. Forecasting and conditional projections using realistic prior distributions. Econometric Reviews 3, 1-100.

Fair, R. 1994. Testing Macroeconomic Models. Cambridge: Harvard University Press.

Fernández-Villaverde, J. J. Rubio-Ramírez 2004. Estimating macroeconomic models: a likelihood approach. Manuscript, University of Pennsylvania.

Geweke, J. 1999. Computational experiments and reality. Manuscript, University of Iowa.

Ingram, B. and C. Whiteman 1994. Supplanting the Minnesota prior – Forecasting macroeconomic time series using real business cycle model priors. Journal of Monetary Economics 34, 497-510.

Levin, A., A. Onatski, J. Williams, and N. Williams 2006. Monetary policy under uncertainty in micro-founded macroeconometric models. In NBER Macroeconomics Annual 2005, ed. M. Gertler and K. Rogoff, Cambridge: MIT Press, 229-312.

King, R., C. Plosser, and S. Rebelo 1988. Production, growth, and business cycles: I. neoclassical model. Journal of Monetary Economics 81, 819-40.

Kydland F. and E. Prescott 1996. The computational experiment: an econometric tool. Journal of Economic Perspectives 10, 69-85.

Landon-Lane, J. 1998. Bayesian comparison of dynamic macroeconomic models. Ph.D. Dissertation, University of Minnesota.

Lubik, T. and F. Schorfheide 2004. Testing for indeterminacy: an application to U.S. monetary policy. American Economic Review 94, 190-217.

Lubik, T. and F. Schorfheide 2006. A Bayesian look at New Open Economy Macroeconomics. In NBER Macroeconomics Annual 2005, ed. M. Gertler and K. Rogoff, Cambridge: MIT Press, 313-382.

Lucas, R.E., Jr. 1976. Econometric policy evaluation: a critique. In: The Phillips Curve and Labor Markets, ed. K. Brunner and A.H. Meltzer, Amsterdam: North-Holland.

Otrok, C. 2001. On measuring the welfare cost of business cycles. Journal of Monetary Economics 47, 61-92.

Sargent, T., N. Williams, and T. Zha 2005. Shocks and government beliefs: the rise and fall of American inflation. American Economic Review, forthcoming.

Schorfheide, F. 2000. Loss function-based evaluation of DSGE models. Journal of Applied Econometrics 15, 645-670.

Sims, C.A. 1980. Macroeconomics and reality. Econometrica 48, January, 1-48.

Sims, C.A. and T. Zha 2005. Were there regime switches in U.S. monetary policy? American Economic Review, forthcoming.

Smets, F. and R. Wouters 2003: An estimated stochastic dynamic general equilibrium model of the Euro Area. Journal of the European Economic Association 1, 1123-75.

Stock, J.H. and M.W. Watson 2001. Vector Autoregression. Journal of Economic Perspectives 15, 101-115.