Journal of Economic Dynamics & Control 36 (2012) 294–313

Contents lists available at SciVerse ScienceDirect

Journal of Economic Dynamics & Control

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

Bayesian prior elicitation in DSGE models: Macro- vs micropriors

Marco J. Lombardi [a], Giulio Nicoletti [a,b,*]

[a] European Central Bank, Kaiserstrasse 29, 63660 Frankfurt am Main, Germany
[b] Banca d'Italia, Via Nazionale 91, 00184, Rome, Italy

ARTICLE INFO

Article history:
Received 11 January 2011
Received in revised form
21 September 2011
Accepted 21 September 2011
Available online 18 October 2011

JEL classification:
C11
C51
E30

Keywords:
DSGE models
Bayesian estimation
Prior distribution
Impulse response function

ABSTRACT

Bayesian approaches to the estimation of DSGE models are becoming increasingly popular. Prior knowledge is normally formalized either directly on deep parameters' values ('microprior') or indirectly, on macroeconomic indicators, e.g. moments of observable variables ('macroprior'). We introduce a non-parametric macroprior which is elicited from impulse response functions and assess its performance in shaping posterior estimates. We find that using a macroprior can lead to substantially different posterior estimates. We probe into the details of our result, showing that model misspecification is likely to be responsible of that. In addition, we assess to what extent the use of macropriors is impaired by the need of calibrating some hyperparameters.

1. Introduction

Since the seminal contributions of DeJong et al. (2000) and Schorfheide (2000), Bayesian estimation methods have gained ground as a very attractive alternative over frequentist approaches in the field of dynamic stochastic general equilibrium (DSGE) models, among both academicians and practitioners; a much notable empirical example being Smets and Wouters (2003). Unlike in the frequentist approach, a Bayesian researcher bases his inference on his prior knowledge in addition to the available data, in order to provide the so-called posterior estimates. Therefore, an important aspect of the Bayesian methodology relates to how the researcher elicits prior information and the way this influences posterior estimates.

Prior knowledge takes the form of a probability distribution for the model's parameters. In the case of DSGE models, prior knowledge is normally expressed as independent probability distributions placed on each of the structural parameters. Hence, the resulting joint prior distribution has independent components. Often such independent distributions are informed by using previous microeconometric studies. In the following, we will refer to this approach as 'microprior'.

The microprior selection scheme is very simple and ready to implement, but suffers from several shortcomings. First, a priori independence of the deep parameters entails implicit assumptions for the a priori beliefs on data moments and the responses of macroeconomic variables after some shock (e.g. impulse response functions, IRFs heretofore) which may indeed be at odds with prior knowledge drawing from macroeconometric studies. Second, for some parameters,

* Corresponding author at: European Central Bank, Kaiserstrasse 29, 63660 Frankfurt am Main, Germany.
E-mail addresses: marco.lombardi@ecb.europa.eu (M.J. Lombardi), giulio.nicoletti@bancaditalia.it, giulio.nicoletti@ecb.int (G. Nicoletti).

0165-1889/$ - see front matter © 2011 Published by Elsevier B.V.
doi:10.1016/j.jedc.2011.09.010

microeconometric information can be abundant (e.g. the frequency at which firms adjust their prices, a direct measure of nominal rigidities in the economy), while for others it can be scant (e.g. the average size of some shock or its persistence). At the same time, the researcher might possess substantial a priori knowledge about the behaviour of macroeconomic variables, and may possibly want to use it to inform the prior at least for a subset of the deep parameters. We will label this approach as 'macroprior' in general, or 'hybrid', when only a subset of parameters is informed by a macroprior.

Among the possible 'macroeconomic objects' that can be used to elicit prior information about the structural parameters, a special role is occupied by impulse response functions—i.e. the response of macroeconomic variables after some shock occurs. A key advantage of eliciting prior knowledge from IRFs is that, differently from the deep parameters and, although to a lesser extent, from sample moments, researchers and policy makers often have clear prior views on how the response of the economic system to certain shocks should look like, even without explicitly referring to a data-based pre-sample.

The contribution of our paper is two-fold. First, we introduce a new type of macroprior, elicited in a non-parametric way from impulse response functions (IRF-prior heretofore) and provide posterior estimates for a simple DSGE model. Second, we highlight the implications for deep parameter estimates of combining micro- and macroinformation.

Our main result is that when an IRF-prior is applied to exogenous state parameters, e.g. the persistence of exogenous shocks and their standard deviations, posterior estimates can differ with respect to the case of the independent prior case by a little amount, especially when the overall variance of the prior is large. If the IRF-prior is applied to a different set of deep parameters, posterior estimates change more widely. By recurring to simulated data, we highlight the fact that this feature is likely to be due to model misspecification.

Mostly related to our work, Del Negro and Schorfheide (2008) – DS heretofore – pioneered the approach to form prior distributions on the basis of a priori beliefs on moments of the macrovariables. More specifically, DS propose a hybrid approach in which a macroprior is used for the subset of parameters concerning the exogenous states and a microprior for the other parameters such as those related to the amount of nominal rigidities in the economy. The difference between our approach and the one of DS is that they inform their macroprior with information about (unconditional) moments of the data rather than IRFs. Our opinion is that IRFs have an advantage in that they allow the researcher to draw prior information only from the shocks whose effects are most known, even if they contribute to the overall data moments by little; a good example of that is a monetary policy shock. In general, IRFs are the most studied objects in macroeconomics and should hence entail more detailed information in terms of prior knowledge. DS report large differences in their posterior estimates, compared to the case of independent priors: we probe into this discrepancy of results and find that the difference is likely to be produced by the fact that, when combining macro with micropriors, some parameters need to be calibrated parameters. The posterior estimates produced by the DS prior appear to be very sensitive to this calibration, whereas our IRF-prior seems to yield more robust results.

The paper is structured as follows: in Section 2 we describe the general formulation of our IRF-prior; for completeness a concise reference to the DS prior is also included. Section 3 provides a description of the two main experiments we undertake with the IRF-prior: we introduce the DSGE model used for experiments and we detail our estimation strategy, the microprior used as benchmark and the dataset used. We then report and discuss our results in Sections 3.5 and 3.6. In Section 4, we assess the sensitivity of the DS and our IRF-prior to parameters that need to be calibrated; conclusions follow.

2. Introducing macropriors

2.1. Impulse response priors

To introduce our IRF-based prior let us start from the definition of impulse responses. The reduced form of a (linearized) DSGE model can be written in state space form (An and Schorfheide, 2007) as

$$X_t = F(\psi)X_{t-1} + G(\psi)e_t, \quad e_t \sim N(0,\Sigma_e), \tag{1}$$

$$Y_t = HX_t, \tag{2}$$

where the transition Eq. (1) corresponds to the solution (e.g. the reduced form) of the model. X_t is the n-dimensional vector of states at time t; e_t denotes the l-dimension vector of i.i.d. normally distributed (and uncorrelated) structural shocks, whose standard deviations are collected in the matrix Σ_e. Both the F and the G matrix in the transition equation are functions of ψ, the vector of structural parameters (e.g. parameters in the utility functions of the agents or parameters related to technology). We gather the entries of the Σ_e matrix, σ_e, and the structural parameters ψ in a vector

$$\xi \equiv [\psi', \sigma_e'],$$

which will be referred to as the vector of deep parameters.

The selection matrix H links states X_t to the observable variables Y_t. As common in the literature, we assume that the number of observed variables is equal to the number of shocks, so that Y_t is an l-dimensional vector.[1]

[1] This assumption is not important to define impulse responses, but it plays a relevant role in the estimation process, as it is not possible to run the Kalman filter on a stochastically singular system.