Economics Letters 172 (2018) 134–137

Contents lists available at ScienceDirect

Economics Letters

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

Econometrics with system priors

Michal Andrle [a,*], Miroslav Plašil [b]

[a] International Monetary Fund, Research Department, 700 19th Street, N.W., WA, D.C. 20431, USA
[b] Czech National Bank, Financial Stability Division, Na Příkopě 28, 115 03 Prague 1, Czech Republic

HIGHLIGHTS

- System priors allow researchers to incorporate economically meaningful priors.
- Any economically meaningful feature of the model can be used to form a system prior.
- System priors are easy to implement within existing Bayesian toolkits.
- System priors help prevent unintended consequences of marginal independent priors.

ARTICLE INFO

Article history:
Received 15 May 2018
Received in revised form 25 August 2018
Accepted 28 August 2018
Available online xxxx

JEL classification:
C11
C18
C22
C51

Keywords:
Bayesian analysis
System priors
Time series

ABSTRACT

This paper introduces "system priors" into Bayesian econometrics and provides an illustrative application. Unlike priors on individual parameters, system priors offer and efficient way of formulating economically-meaningful priors about model properties. Illustrative example restricts output gap dynamics to business-cycle frequencies.

1. Introduction

This paper introduces "system priors", a simple and tractable way of employing economically meaningful a priori judgment about model properties in statistical inference. Often, there are situations where the a priori beliefs of analysts relate to features of the full model, features that are non-linear functions of individual coefficients — often with no closed-form solution. There are also situations where the individual model parameters are hard to interpret and to elicit priors about, whereas the model system properties are easy to interpret and can be readily communicated to researchers, policymakers, and interested audience. These are all situations where system priors offer a solution to the problem.

Examples of system priors include views on steady-states of the model variables or their growth rates (Villani, 2009; Del Negro and Schorfheide, 2008), beliefs on plausible shape and magnitude of impulse-response (Kocięcki, 2010), or expert judgment about policy scenarios, say sacrifice ratio after permanent disinflation shock (Andrle and Benes, 2013). Prior views related to spectral characteristics of the model, such as filter frequency-response function (Andrle and Brůha, 2017), coherence between variables or duration of business cycles (this paper and Jarociński and Lenza, 2016) are also natural candidates for system priors.

The contribution of the paper is twofold. In comparison to other literature, our approach aims at maximum generality and simplification to equip researchers with easily accessible tool addressing a wide range of potential applications. Second, we provide a concise and standalone exposition of the approach. While similar approaches have been used in the literature, they are often buried inside the papers or in their technical appendices. The common principles that underline ad-hoc approaches may thus go unnoticed. To illustrate our approach to system priors, we use an example of a stationary second-order auto-regressive – AR(2) – process (for a cyclical component of output) and incorporate a prior belief that most of its variance comes from business-cycle frequencies.

* Corresponding author.
E-mail addresses: mandrle@imf.org (M. Andrle), miroslav.plasil@cnb.cz (M. Plašil).

https://doi.org/10.1016/j.econlet.2018.08.038
0165-1765/

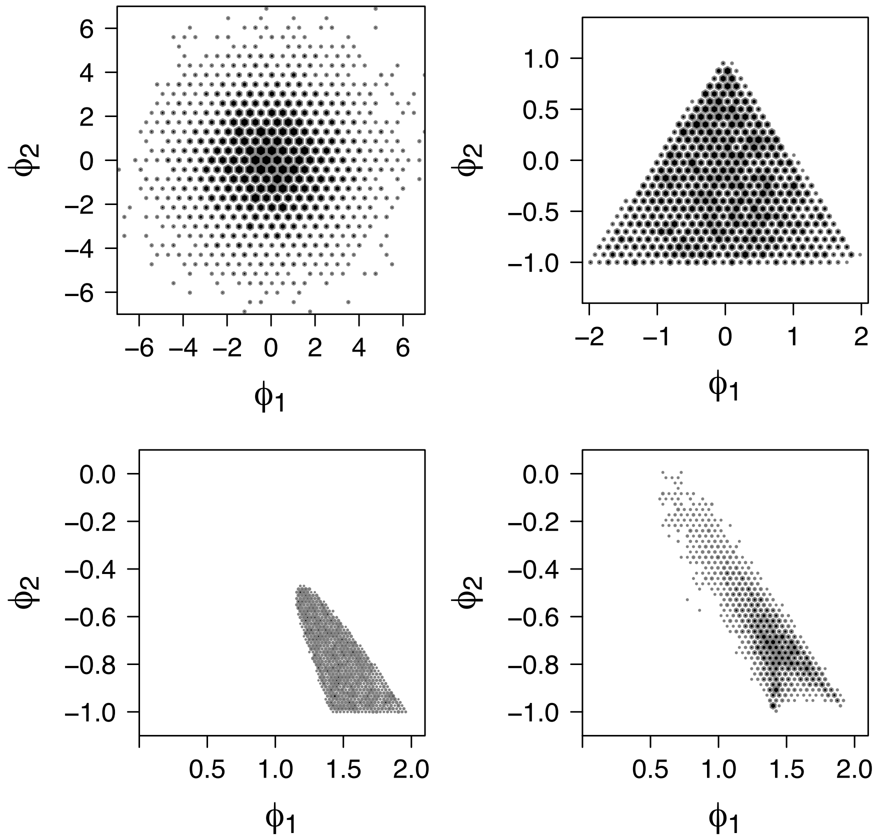

Fig. 1. Parameter regions for different priors. *Note:* Hexagonal binning density plot. Upper left panel: normally distributed independent marginal priors for ϕ_1 and ϕ_2 ; upper right: identical priors restricted to the stationarity region; bottom left: stationarity + at least 60% of the variability comes from business-cycle frequencies; bottom right: stationarity + the share of business-cycle frequencies given by Y .

2. System priors

The estimation of models with system priors closely follows the general principles of Bayesian inference. The difference rests in the form of the prior distribution formulation. We assume that the joint prior beliefs about a $(k \times 1)$ vector of individual parameters, θ, of a model M are expressed using independent marginal probability distributions, i.e., as: $p_m(\theta) = p_m(\theta_1) \times \cdots \times p_m(\theta_k)$. We further assume that given the observed data, Y, it is possible to evaluate the likelihood function of the model, $L(Y \mid \theta; M)$ for different parameter values. Applying Bayes' law, the posterior distribution of the parameters is proportional to the product of the likelihood and the prior distribution:

$$p(\theta \mid Y; M) \propto L(Y \mid \theta; M) \times p_m(\theta). \qquad (1)$$

Now let us incorporate a priori views about the model's system features. Let us define a feature of interest, r, that a prior view will be formulated about. The feature, r, is a function of the individual parameters θ given the model M: $r = h(\theta; M)$. We assume that the feature can be evaluated for different parameter values. As in the case of the individual parameters, the prior beliefs about the values of feature r can be summarized by a functional form that best captures analysts' views. We will call it the "system prior" and denote it as $p_s(r;\ h, M) \equiv p_s(h(\theta);\ M)$. Putting together the effects of the marginal prior, the system prior, and the likelihood function, the posterior distribution of the parameters emerges as

$$p(\theta \mid Y;\ M) \propto L(Y \mid \theta;\ M) \times [p_s(h(\theta); M) \times p_m(\theta)]. \qquad (2)$$

The form of the posterior kernel in (2) is intuitive, read from the right to left. For a given value of parameter θ, the posterior distribution is based on a two-step updating process. In the *first step* the marginal prior, p_m, is updated with the system priors, p_s, resulting in a composite prior distribution. Updating with system priors can also be understood as combining priors on parameters with artificial likelihood function that arises from a stochastic model with the structure $r = h(\theta; M)$. As the system priors operate on functions of parameters, the composite prior $p_c = [p_s(h(\theta);\ M) \times p_m(\theta)]$ implies restrictions on individual coefficients but generally not in a unique or invertible way. In the *second step*, the composite prior beliefs are updated with the information contained in the data using a likelihood function of the model. Estimation with system priors thus represents a two-layer approach: beliefs about the parameters are complemented with beliefs about the model properties and then combined with information in the data.

To analyze the implications of the composite joint prior distribution in greater detail, computations analogous to posterior sampling in (2) are needed, with evaluation of the conventional likelihood function switched off.[1] In practical applications, the advantage of system priors is that existing Bayesian computations and computer code can stay almost unchanged when system priors are used. Metropolis–Hastings or Sequential Monte Carlo (SMC) methods are easily applicable.[2]

1 In some applications, an option would be to entirely discard information arising from the conventional likelihood and combine parameter priors with any preferred loss function.

2 The codes R and Matlab/Iris toolbox and Dynare for DSGE and VAR models are available upon request or at http://michalandrle.weebly.com/systempriors. Also, the European Central Bank's toolbox YADA for New-Area Wide Model (NAWM) features implementation of system priors following this paper, http://www.texlips.net/yada/.