

System Priors in DYNARE: A Quick Hack

Michal Andrle

August 2017

I. INTRODUCTION

This short note describes implementing a bare-bone version of system priors for the DYNARE Toolbox for Matlab. The implementation follows Andrle and Benes (2013) and Andrle and Plasil (2016). More refined implementations are available for the IRIS Toolbox and YADA Toolbox for Matlab.

Dynare toolbox already seem to feature system priors in principle but narrowly focused on pre-sample moments or relative impulse-response priors, which the users specify in the Dynare *.mod file and in Dynare language. Dynare calls these ‘endogenous priors’.

Implementation of system priors is specified in the most general way, when the analyst is responsible for writing her or his Matlab file that parameterizes and evaluates the system priors using the current model.

II. SYSTEM PRIORS IN BRIEF

For a thorough and step-by-step explanation of system priors, see Andrle and Plasil (2006), using the same example of an AR(2) model used in the example Dynare file.

System priors generalize the marginal independent priors commonly used for DSGE or VAR models. System priors allow to specify priors about any computable feature of the model as a system (IRFs, variances, scenario results, spectral properties, etc.)

In the AR(2) model example below the system prior about the model is that the share of variance at business cycle frequencies should be ‘around’ seventy percent. Since not all combinations of parameters are consistent with this prior, the model parameters are a priori nudged towards a subset of coefficients.

More formally, if $p_m(\theta)$ are the marginal priors for the parameters of the model, and $L(Y|\theta)$ is the likelihood function of the model, system priors introduce another term (or penalty)

$p_s(h(\theta))$. Here, $z = h(\theta)$ is a mapping of parameters into the system features of interest, z , and $p_s(\cdot)$ is a formulation of the prior distribution.

The posterior distribution is thus a result of two Bayesian updates. First, the marginal priors, $p_m(\theta)$ are updated by $p_s(h(\theta))$, resulting into a composite prior. This composite prior is further updated by the likelihood of the data, $L(Y|\theta)$, i.e.:

$$p(\theta|Y) \propto L(Y|\theta) \times p_s(h(\theta)) \times p_m(\theta). \quad (1)$$

III. IMPLEMENTATION

The example implemenation modifies the file `dsge_likelihood.m` in the Matlab directory of Dynare. Here a call to a function `modelname_spriors.m` is carried out if such a file exists. In this file the user can implement the system prior.

The file `ar2_spriors.m` is the system prior for the AR(2) model, `ar2.mod`. The signature of the function must be consistent with the signature in the caller function. Most of the relevant information are passed, e.g. the model structure, Dynare options, etc. Adjustments are easy. The system prior function needs to return a scalar value of the system prior evaluated at the current draw/vector of parameters.



The AR(2) example function checks if the AR(2) is stationary and then evaluates what portion of the current parameterization lies at business-cycle frequencies, using the spectrum of the AR(2) process.

This Dynare implemenation computes the value of the posterior distribution with system priors. In order to understand just the implied prior, the likelihood function must be ignored and only the composite prior analyzed. To analyze the composite prior a proper Bayesian updating step is needed with the likelihood function switched off. The current implementation does not allow for this, in order to stay non-intrusive to Dynare, but setting the number of observations in the Dynare file to a trivial number that won't affect the computations can bypass that.

A. AR(2) Example

The model is in the file `ar2.mod`. The initial value of the the two parameters, $\alpha_1 = \alpha_2 = 0.5$ and the data for the model are generated with these parameters. The marginal independent priors for both of the coefficients is $\alpha_i N(0.4, 0.2)$, very wide. Prior for the variance of the shock is not affected by the system priors and so is not discussed here.

The system prior example for Dynare test is that the share of variance of the model at business cycle frequencies is distributed as $N(0.7, 0.1)$ and stationary, for simplicity ignoring the share cannot be negative.

Figure 1. Composite Prior Draws**Figure 2. Marginal Priors and Marginal Composite Priors**

When the model is estimated with all the 500 observations, the likelihood dominates the system prior. When the number of observations is set to a trivial number, say 20, the estimator focuses only at the combinations of the auto-regressive parameters depicted in the Fig. 1.

Clearly, the joint composite prior is materially different from the joint prior implied by the two independent normal distributions.

IV. CONCLUSION

Have fun and use system priors.