Sensitivity Analysis Toolbox for Dynare*

Marco Ratto
European Commission, Joint Research Centre
TP581
21027 Ispra (VA) Italy
Marco.Ratto@ec.europa.eu †

October 18, 2026

Abstract

The Sensitivity Analysis Toolbox for Dynare is a set of MATLAB routines for the analysis of DSGE models with global sensitivity analysis. The routines are thought to be used within the Dynare 6 environment.

Keywords: Stability Mapping , Reduced form solution, DSGE models, Monte Carlo filtering, Global Sensitivity Analysis.

*Copyright © 2012-2024 Dynare Team. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license can be found at: https://www.gnu.org/licenses/fdl.txt

†The author gratefully thanks Christophe Planas, Kenneth Judd, Michel Juillard, Alessandro Rossi, Frank Schorfheide and the participants to the Courses on Global Sensitivity Analysis for Macroeconomic Models (Ispra, 2006-2007-2008-2010) for interesting discussions and helpful suggestions.

1 Introduction

The Sensitivity Analysis Toolbox for Dynare is a collection of MATLAB routines implemented to answer the following questions: (i) Which is the domain of structural coefficients assuring the stability and determinacy of a DSGE model? (ii) Which parameters mostly drive the fit of, e.g., GDP and which the fit of inflation? Is there any conflict between the optimal fit of one observed series versus another one? (iii) How to represent in a direct, albeit approximated, form the relationship between structural parameters and the reduced form of a rational expectations model?

The discussion of the methodologies and their application is described in Ratto (2008).

2 Use of the Toolbox

The Dynare parser now recognizes sensitivity analysis commands. The syntax is based on a single command:

```
sensitivity(option1=<opt1_val>,option2=<opt2_val>,...)
```

with a list of options described in the next section.

In order to work properly, the sensitivity analysis Toolbox does not need a Dynare estimation environment to be set up. Rather, `sensitivity` is the only command to run to make a sensitivity analysis on a DSGE model[1].

[1]Of course, when the user needs to perform the mapping of the fit with a posterior sample, a Bayesian estimation has to be performed beforehand

3 List of options

3.1 Sampling options

option name	default	description
`Nsam`	2048	Size of MC sample
`ilptau`	1	1 = use LP_τ quasi-Monte Carlo 0 = use LHS Monte Carlo
`pprior`	1	1 = sample from prior distributions 0 = sample from multivariate normal $N(\hat{\theta}, \Sigma)$, $\hat{\theta}$ is posterior mode $\Sigma = H^{-1}$, H is Hessian at the mode
`prior_range`	1	1 = sample *uniformly* from prior ranges 0 = sample from prior distributions:
`morris`	0	0 = no Morris sampling for screening 1 = Morris sampling for screening
`morris_nliv`	6	number of levels in Morris design
`morris_ntra`	20	number of trajectories in Morris design
`ppost`	0	0 = don't use Metropolis posterior sample 1 = use Metropolis posterior sample: this overrides any other sampling option!
`neighborhood_width`	[]	δ (real number> 0) uniform sample in the neighborhood of the posterior mode $\hat{\theta}$ interval width: $\hat{\theta}(1 \pm \delta)$

3.2 Stability mapping

option name	default	description
`stab`	1	1 = perform stability mapping 0 = no stability mapping is performed
`load_stab`	0	0 = generate a new sample 1 = load a previously created sample
`pvalue_corr`	0.001	critical p-value for correlations ρ in filtered samples: plot couples of parameters with p-value<`pvalue_corr`
`pvalue_ks`	0.001	critical p-value for Smirnov statistics d: plot parameters with p-value<`pvalue_ks`
`lik_init`	1	1 = the model is stationary (unit roots are 'explosive') 3 = the model has unit roots (unit roots are 'stable')

3.3 Reduced form mapping

The mapping of the reduced form solution forces the use of samples from prior ranges or prior distributions, i.e.:

```
options_.opt_gsa.pprior=1;
options_.opt_gsa.ppost=0;
```

It uses 250 samples to optimize smoothing parameters and 1000 samples to compute the fit. The rest of the sample is used for out-of-sample validation.

option name	default	description
`redform`	0	0 = don't prepare MC sample of reduced form matrices 1 = prepare MC sample of reduced form matrices
`load_redform`	0	0 = estimate the mapping of reduced form model 1 = load previously estimated mapping
`logtrans_redform`	0	0 = use raw entries 1 = use log-transformed entries
`threshold_redform`	[]	[] = don't filter MC entries of reduced form coefficients [`max max`] = analyse filtered entries within the range [`max max`]
`ksstat_redform`	0.001	critical p-value for Smirnov statistics d when `threshold_redform` is active plot parameters with p-value<`ksstat_redform`
`alpha2_redform`	0	critical p-value for correlation ρ when `threshold_redform` is active plot couples of parameters with p-value<`alpha2_redform`
`namendo`	()	list of endogenous variables
	:	jolly character to indicate ALL endogenous
`namlagendo`	()	list of lagged endogenous variables: analyse entries [`namendo`×`namlagendo`]
	:	jolly character to indicate ALL endogenous
`namexo`	()	list of exogenous variables: analyse entries [`namendo`×`namexo`]
	:	jolly character to indicate ALL exogenous

One can also load a previously estimated mapping with a new MC sample, to look at the forecast for the new MC sample.

3.4 Mapping the fit

The RMSE analysis can be performed with different types of sampling options:

1. when `pprior=1` and `ppost=0`, the Toolbox analyses the RMSE's for the MC sample obtained by sampling parameters from their prior distributions (or prior ranges): this analysis provides some hints about what parameter drives the fit of which observed series, prior to the full estimation;

2. when `pprior=0` and `ppost=0`, the Toolbox analyses the RMSE's for a multivariate normal MC sample, with covariance matrix based on the inverse Hessian at the optimum: this analysis is useful when ML estimation is done (i.e. no Bayesian estimation);

3. when `ppost=1` the Toolbox analyses the RMSE's for the posterior sample obtained by Dynare's Metropolis procedure.

The use of cases 2. and 3. require an estimation step beforehand! To facilitate the sensitivity analysis after estimation, the `sensitivity` command also allows to indicate some options of `estimation`. These are:

- `datafile`
- `diffuse_filter`
- `mode_file`
- `first_obs`
- `lik_init`
- `nobs`
- `prefilter`
- `presample`
- `loglinear`

option name	default	description
`rmse`	0	0 = no RMSE analysis 1 = do RMSE analysis
`load_rmse`	0	0 = make a new RMSE analysis 1 = load previous RMSE analysis
`lik_only`	0	0 = compute RMSE's for all observed series 1 = compute only likelihood and posterior
`var_rmse`	varobs	list of observed series to be considered
`pfilt_rmse`	0.1	filtering threshold for RMSE's: default it to filter the best 10% for each observed series
`istart_rmse`	1	start computing RMSE's from `istart_rmse`: use 2 to avoid big initial error
`alpha_rmse`	0.001	p-value for Smirnov statistics d: plot parameters with p-value<`alpha_rmse`
`alpha2_rmse`	0	p-value for correlation ρ plot couples of parameters with p-value<`alpha2_rmse`

3.5 Screening analysis

The screening analysis does not require any additional options with respect to those listed in the 'Sampling options': `morris`, `morris_nliv`, `morris_ntra`. The Toolbox performs all the analyses required and displays results.

3.6 Identification analysis

Setting the option `identification=1`, an identification analysis based on theoretical moments is performed. Sensitivity plots are provided that allow to infer which parameters are most likely to be less identifiable.

option name	default	description
`identification`	0	0 = no identification analysis 1 = performs identification analysis: this forces `redform`=0 and default `morris`=1
`morris`	1	1 = Screening analysis (Type II error) 2 = Analytic derivatives (Iskrev, 2010, 2011)
`morris_nliv`	6	number of levels in Morris design
`morris_ntra`	20	number of trajectories in Morris design

For example, the following commands in the Dynare model file

```
sensitivity(identification=1, morris=2);
```

trigger the identification analysis using Iskrev (2010, 2011), jointly with the mapping of the acceptable region.

4 Directory structure

Sensitivity analysis results are saved on the hard-disk of the computer. The Toolbox uses a dedicated folder called `GSA`, located in

`<Dynare_file>\GSA`,

where `<Dynare_file>.mod` is the name of the Dynare model file.

4.1 Binary data files

A set of binary data files is saved in the `GSA` folder:

`<Dynare_file>_prior.mat`: this file stores information about the analyses performed sampling from the prior ranges, i.e. `pprior=1` and `ppost=0`;

`<Dynare_file>_mc.mat`: this file stores information about the analyses performed sampling from multivariate normal, i.e. `pprior=0` and `ppost=0`;

`<Dynare_file>_post.mat`: this file stores information about analyses performed using the Metropolis posterior sample, i.e. `ppost=1`.

`<Dynare_file>_prior_*.mat`: these files store the filtered and smoothed variables for the prior MC sample, generated when doing RMSE analysis (`pprior=1` and `ppost=0`);

`<Dynare_file>_mc_*.mat`: these files store the filtered and smoothed variables for the multivariate normal MC sample, generated when doing RMSE analysis (`pprior=0` and `ppost=0`).

4.2 Stability analysis

Figure files `<Dynare_file>_prior_*.fig` store results for the stability mapping from prior MC samples:

`<Dynare_file>_prior_stab_SA_*.fig`: plots of the Smirnov test analyses confronting the CDF of the sample fulfilling Blanchard-Kahn conditions with the CDF of the rest of the sample;

`<Dynare_file>_prior_stab_indet_SA_*.fig`: plots of the Smirnov test analyses confronting the CDF of the sample producing indeterminacy with the CDF of the original prior sample;

<Dynare_file>_prior_stab_unst_SA_*.fig: plots of the Smirnov test analyses confronting the CDF of the sample producing unstable (explosive roots) behaviour with the CDF of the original prior sample;

<Dynare_file>_prior_stable_corr_*.fig: plots of bivariate projections of the sample fulfilling Blanchard-Kahn conditions;

<Dynare_file>_prior_indeterm_corr_*.fig: plots of bivariate projections of the sample producing indeterminacy;

<Dynare_file>_prior_unstable_corr_*.fig: plots of bivariate projections of the sample producing instability;

<Dynare_file>_prior_unacceptable_corr_*.fig: plots of bivariate projections of the sample producing unacceptable solutions, i.e. either instability or indeterminacy or the solution could not be found (e.g. the steady state solution could not be found by the solver).

Similar conventions apply for <Dynare_file>_mc_*.fig files, obtained when samples from multivariate normal are used.

4.3 RMSE analysis

Figure files <Dynare_file>_rmse_*.fig store results for the RMSE analysis.

<Dynare_file>_rmse_prior*.fig: save results for the analysis using prior MC samples;

<Dynare_file>_rmse_mc*.fig: save results for the analysis using multivariate normal MC samples;

<Dynare_file>_rmse_post*.fig: save results for the analysis using Metropolis posterior samples.

The following types of figures are saved (we show prior sample to fix ideas, but the same conventions are used for multivariate normal and posterior):

<Dynare_file>_rmse_prior_*.fig: for each parameter, plots the CDF's corresponding to the best 10% RMES's of each observed series;

<Dynare_file>_rmse_prior_dens_*.fig: for each parameter, plots the pdf's corresponding to the best 10% RMES's of each observed series;

<Dynare_file>_rmse_prior_<name of observedseries>_corr_*.fig: for each observed series plots the bi-dimensional projections of samples with the best 10% RMSE's, when the correlation is significant;

`<Dynare_file>_rmse_prior_lnlik*.fig`: for each observed series, plots *in red* the CDF of the log-likelihood corresponding to the best 10% RMSE's, *in green* the CDF of the rest of the sample and *in blue* the CDF of the full sample; this allows to see the presence of some idiosyncratic behaviour;

`<Dynare_file>_rmse_prior_lnpost*.fig`: for each observed series, plots *in red* the CDF of the log-posterior corresponding to the best 10% RMSE's, *in green* the CDF of the rest of the sample and *in blue* the CDF of the full sample; this allows to see idiosyncratic behaviour;

`<Dynare_file>_rmse_prior_lnprior*.fig`: for each observed series, plots *in red* the CDF of the log-prior corresponding to the best 10% RMSE's, *in green* the CDF of the rest of the sample and *in blue* the CDF of the full sample; this allows to see idiosyncratic behaviour;

`<Dynare_file>_rmse_prior_lik_SA_*.fig`: when `lik_only=1`, this shows the Smirnov tests for the filtering of the best 10% log-likelihood values;

`<Dynare_file>_rmse_prior_post_SA_*.fig`: when `lik_only=1`, this shows the Smirnov test for the filtering of the best 10% log-posterior values.

4.4 Reduced form mapping

In the case of the mapping of the reduced form solution, synthetic figures are saved in the `\GSA` folder:

`<Dynare_file>_redform_<endo name>_vs_lags_*.fig`: shows bar charts of the sensitivity indices for the *ten most important* parameters driving the reduced form coefficients of the selected endogenous variables (`namendo`) versus lagged endogenous variables (`namlagendo`); suffix `log` indicates the results for log-transformed entries;

`<Dynare_file>_redform_<endo name>_vs_shocks_*.fig`: shows bar charts of the sensitivity indices for the *ten most important* parameters driving the reduced form coefficients of the selected endogenous variables (`namendo`) versus exogenous variables (`namexo`); suffix `log` indicates the results for log-transformed entries;

`<Dynare_file>_redform_GSA(_log).fig`: shows bar chart of all sensitivity indices for each parameter: this allows to notice parameters that have a minor effect for *any* of the reduced form coefficients,

Detailed results of the analyses are shown in the subfolder `\GSA\redform_stab`, where the detailed results of the estimation of the single functional relationships between parameters θ and reduced form coefficient are stored in separate directories named as:

`<namendo>_vs_<namlagendo>`: for the entries of the transition matrix;

`<namendo>_vs_<namexo>`: for entries of the matrix of the shocks.

Moreover, analyses for log-transformed entries are denoted with the following suffixes (y denotes the generic reduced form coefficient):

`log`: $y^* = \log(y)$;

`minuslog`: $y^* = \log(-y)$;

`logsquared`: $y^* = \log(y^2)$ for symmetric fat tails;

`logskew`: $y^* = \log(|y + \lambda|)$ for asymmetric fat tails.

The optimal type of transformation is automatically selected without the need of any user's intervention.

4.5 Screening analysis

The results of the screening analysis with Morris sampling design are stored in the subfolder `\GSA\SCREEN`. The data file `<Dynare_file>_prior` stores all the information of the analysis (Morris sample, reduced form coefficients, etc.).

Screening analysis merely concerns reduced form coefficients. Similar synthetic bar charts as for the reduced form analysis with MC samples are saved:

`<Dynare_file>_redform_<endo name>_vs_lags_*.fig`: shows bar charts of the elementary effect tests for the *ten most important* parameters driving the reduced form coefficients of the selected endogenous variables (`namendo`) versus lagged endogenous variables (`namlagendo`);

`<Dynare_file>_redform_<endo name>_vs_shocks_*.fig`: shows bar charts of the elementary effect tests for the *ten most important* parameters driving the reduced form coefficients of the selected endogenous variables (`namendo`) versus exogenous variables (`namexo`);

`<Dynare_file>_redform_screen.fig`: shows bar chart of all elementary effect tests for each parameter: this allows to identify parameters that have a minor effect for *any* of the reduced form coefficients.

References

Nikolay Iskrev. Local identification in DSGE models. *Journal of Monetary Economics*, 57(2):189–202, 2010. doi:10.1016/j.jmoneco.2009.12.007.

Nikolay Iskrev. Evaluating the strenght of identification in DSGE models. an a priori approach. mimeo, 2011.

Marco Ratto. Analysing DSGE models with global sensitivity analysis. *Computational Economics*, 31:115–139, 2008. doi:10.1007/s10614-007-9110-6.