VAR EXERCISE

In this exercise you will estimate and analyze a four-variable non-structural, VAR, in the four variables real GDP (logged), unemployment rate (as a fraction, not a per cent), employment (logged), and employment-to-population ratio (not logged). All the data are quarterly and are for dates from 1948:I to 2017:III. Use 5 lags (not `rfvar3`)'s default of 6 lags). An R data file with the data already logged, quarterly averaged, etc. is available on the course web site. Raw data in .csv files is also available, as is a text file with the transformed data as an unlabeled matrix. R and Matlab code that will be useful is also available. Doing the exercise in Matlab may require more coding effort. The R code is organized so that the directory it is in can be installed as an R package that will integrate with R's help system.

(1) Estimate the model with the default Minnesota prior parameter values used in the `rfvar3` program. This program is available in both R and Matlab versions. The default prior uses the "cointegration" dummy observation and "individual unit root" dummy observations, but not the rest of the Minnesota prior. Calculate and plot the triangularized impulse responses for the posterior modal parameter values. (Use `impulsdtrf`) and (in R) `plotir`.)

(2) Plot the residuals from the estimated model as time series. Comment on whether you can see the "Great Moderation", during which disturbances remained relatively small over an extended period.

(3) Find the eigenvalue decomposition of the system matrix and use it (and the output of `rfvar3`) to determine the covariance matrix of the components of the transformed, stacked data vector that are stationary and farther than $1/T$ from 1. Check whether any of them are multiple standard deviations away from steady state.

(4) Calculate forecasts of the model variables (`fcast`) will be useful here) using the posterior modal parameter values and the sample initial conditions. Plot each variable's forecast on the same graph with a plot of the actual data. Comment on whether these forecasts agree with your eigenvalue/eigenvector analysis on whether problematic long-run predictability is implied by the estimates.

(5) Calculate and plot forecasts based on the last 5 observations in the sample, over 10 years into the future.

Date: November 16, 2017.

(6) Make 1000 draws from the joint posterior of the parameters and the covariance matrix. (`postdraw`) will be helpful here.) Then for each draw, form a forecast for the 10 years after the end of the sample, adding shocks for each future period drawn from the distribution of shocks implied by the drawn value of the covariance matrix of residuals.

(7) Using your collection of 1000 future simulated histories of these four variables, display the histogram of 10-year average annual growth rates for log GDP and log employment, and the histogram of 10-year average levels of the unemployment rate and the employment-population ratio.