REDUCED FORM VAR FORECASTING EXERCISE

This exercise tries to determine the dates and variables in which the big surprises occurred during 2007-2009 and to see how forecasts from a time series model evolved over this period.

Repeat the calculations asked for below for the full sample and for samples ending in 2008:6, 2008:9, 2008:11, and 2009:1. In part (b), the calculation of standardized residuals, the results may not change a lot between time periods. Note any substantial changes (i.e., situations where new data change the model's notion of what in the past was surprising), but there is no need to produce charts for part (b) except for the full sample, except to the extent there are big changes.

(a) Estimate a reduced form monthly VAR using industrial production, the CPI for all urban consumers, the 90 day treasury bill rate, and the 10 year constant maturity government bond rate. These are all available for download at FRED, the St. Louis Fed data web site. Use the log of CPI and the log of industrial production, and multiply the interest rates by .01 so that their residual variance is the same order of magnitude as that of the other two variables. Use the default `lambda=5`, `mu=1` settings from `rfvar3`. (If you use non-R, non-Matlab/Octave software, try to mimic these settings for the prior.) Use the longest time span you can without encountering missing data.

(b) Plot a histogram of standardized residuals (i.e., residuals divided by estimated standard deviations) for each variable. Also a q-q plot for residuals of each variable. (A q-q plot plots the sample cdf against the Gaussian cdf. In R, it is produced by the `qqnorm()` function.) Also a time series plot of each variable's standardized residuals. Does 2007-2009 stand out as a time of unusually large shocks? Do you see instances of the whip-saw pattern often generated by single outliers in an otherwise smooth series?

(c) Construct forecasts, with 90% error bands for each variable, conditional on data up through. On the same plot, show the actual time path of the variable. (For the full sample, there might be a month or two of actual data for some series in the forecast period, even though there will not be any for some series.)

(d) For the truncated-sample calculations, did actual time paths run outside of 90% error bands repeatedly, at just one or two dates, or not at all?

Date: November 18, 2009.