TAKEHOME EXAM

On the course web site are two data files, `examdata.RData` and `examdata.csv`. Both contain data, starting in January 1971, for the LIBOR-TBill spread and for the log of US industrial production. LIBOR is the London Interbank Offer (or Overnight?) Rate. It is usually close to the US 3-month TBill rate, but when banks start to worry about counterparty risk when lending to other banks, they can diverge. This spread is often referred to as the "TED spread", with TED standing for "Treasury-EuroDollar".

(1) Estimate a bivariate VAR for these two variables using the data through January 2007. You should use 13 lags (this is just over one year, and allows for catching seasonal effects if there are any). It is OK to use `rfvar3.R` pr `rfvar3.m` from my website, with the default settings other than lag length. (This implements a sum-of-coefficients and joint-persistence prior via dummy observations. This is not a full proper prior. It uses 3 dummy observations, and the "residuals" in the 3 dummy observations appear at the end of the matrix of residuals `vout$u` returned by `rfvar3`.)

(2) Calculate the impulse response functions and plot them. Check whether they imply that the spread has substantial predictive value for industrial production. Explain how you reached your conclusion. (Here `impulsdtrf.R` could be useful. For the plots, it is handy to use
`plot(ts(t(resp[j, , ]),`
for example, to plot the responses of variable j.)

(3) Form forecasts (`fcast.R` may be useful) for industrial production from data through January, through July, and through December of 2007. You need not re-estimate the VAR coefficients each time. Assess the extent to which the model anticipated the decline in production.

(4) Assess, using whatever approach you think sensible, whether the residuals in the model are normally distributed. You need not do formal tests or odds ratio calculations. Convincing graphical evidence is enough.

(5) Assess, using whatever approach you think sensible, whether the residuals have a constant covariance matrix over time. Again, formal tests are not necessary and good eyeball-checks are OK.

(6) Form forecasts for the entire sample starting from initial conditions based on the estimated coefficients. Does the model imply unreasonable long-term forecastability? If you used the default `rfvar3.R` prior, this is likely enough to prevent this problem from arising, but it's worth checking.

©2012 by Christopher A. Sims. This document is licensed under the
Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported License

(7) Explain how to expand the model to allow for either non-normal residuals or a time-varying covariance matrix of disturbances. The latter could be by a hidden Markov chain approach. In either case, give as detailed a description as time allows of how to implement Bayesian estimation of your expanded model. (Total time on the exam should not need to exceed six hours at the max, and four hours if you are well prepared). You can provide draft R or matlab code, but that is not essential. If you are really fast at coding, you might actually estimate your extended model, but that likely would take unreasonably long.