Econ. 556b Spring 1999 Chris Sims

VAR Testing Exercise

Obtain from Citibase (reachable from Yale computers or from a dialup to a Yale PPP account at

`http://statlab.stat.yale.edu/yale_only/Citibase-ext.html`)

US quarterly data on the three-month treasury bill rate (r), real GDP (y) (chain-linked, if possible), GDP deflator (p) (also preferably chain-linked), civilian unemployment rate (u), and M1 (m). Probably M1 is available only back to about 1959, so start the sample at the earliest date for which you can get M1. In all cases use seasonally adjusted data if there is a choice. Take logarithms of y, p, and m. Divide r and u (which are in per cent units to start with) by 100 to make them of the same scale as the other variables.

1. Estimate by OLS a bivariate VAR using just u and p, plus a constant term. Plot the four impulse responses over a 16-quarter horizon, using a Choleski decomposition with u first — i.e. arranged so that there is no contemporaneous response of u to p shocks. Assess whether u surprises affect p as the Phillips curve suggests. What about the effect of p surprises on u?
2. Test the hypothesis that u does not help predict p, i.e. that p is Granger-causally prior. Use a classical asymptotic likelihood ratio test, a Schwarz criterion test, and a test that uses the difference in log determinants of variance matrices of parameter estimates as well as the likelihood ratio. (This is the test statistic from which we derived the Schwarz criterion in lectures.) Also use a chi-squared test based on the estimated coefficients and their covariance matrix.
3. Estimate a VAR for the full vector of five variables, again including a constant. Plot the responses of p to all five innovations, orthogonalizing with u first in the ordering and p last. Assess whether these impulse responses suggest a large role for u in predicting p.
4. Test the hypothesis that u does not help predict p in this system, using the same set of test criteria as in 2.

Matlab `m`-files `rfvar.m` and `impulsdt.m`, which compute OLS VAR estimates and construct the impulse responses from the estimates, are available via the course web page. For this exercise you might also consider using RATS, a program available at Statlab that computes VAR estimates and impulse responses more automatically.