

Estimation, Solution and Analysis of Equilibrium Monetary Models

November 2-5, 2004

Assignment 2: Tutorial on Estimation and Analysis of a VAR

This is a tutorial that takes you through the estimation and analysis of the vector autoregression (VAR) used in Altig, Christiano, Eichenbaum and Linde ('Firm-Specific Capital, Nominal Rigidities and the Business Cycle') (ACEL). The questions allow you to reproduce the results reported in the paper and in addition, they allow you to assess the robustness of the results to changes in sample period, choice of data and other features of the analysis. The code also allows the user to study the dynamic equilibrium model in ACEL. This will be explored in a later assignment.

To answer the questions, execute `main.m`. Different questions are answered by choosing different settings for the user-controlled parameters in the first part of that program. At the moment, the code requires MATLAB 6.5 (not 7), and the user should set `vver = 0` at the beginning.

In the 'benchmark model' the sample period is 1959:1 - 2001:4, the measure of money is MZM, the price of investment goods is the investment good price deflator, labor productivity is measured by GDP divided by non-farm business hours, and the measure of population is the non-institutional population 16 years and older (P16). In addition, there are four lags in the VAR and the VAR includes a constant term. Moreover, the data, Y_t , in the VAR are given by:

$$\underbrace{Y_t}_{10 \times 1} = \begin{pmatrix} \Delta \ln(\text{relative price of investment}_t) \\ \Delta \ln(GDP_t/\text{Hours}_t) \\ \Delta \ln(GDP \text{ deflator}_t) \\ \text{capacity utilization}_t \\ \ln(\text{Hours}_t) \\ \ln(GDP_t/\text{Hours}_t) - \ln(W_t/P_t) \\ \ln(C_t/GDP_t) \\ \ln(I_t/GDP_t) \\ \text{Federal Funds Rate}_t \\ \ln(GDP \text{ deflator}_t) + \ln(GDP_t) - \ln(MZM_t) \end{pmatrix}.$$

1. Estimate the benchmark VAR by executing `main.m`, at the command prompt in MATLAB. Print the impulse responses to the policy, neutral and embodied technology shocks (these are graphed if `imp = 1`). Print the graphs of the historical decomposition of the data in terms of the individual identified shocks, as well as all three shocks simultaneously (these are graphed if

$decomp = 1$). Instead of printing these graphs, you may simply want to verify that they coincide with the graphs reported in ACEL.

- a Consider the lag length of the VAR. One strategy for picking lag length is to minimize one of the lag length selection criteria: Akaike, Hannan-Quinn or Schwartz (see Bierens 2004 on the web site for a discussion of these criteria). What lag length do these criteria suggest choosing?
 - b Consider the multivariate Portmanteau (Q) statistics (see the excerpt from the Stata Technical Bulletin). This is a statistic for testing the null hypothesis that the first n autocorrelations of the fitted disturbances in a VAR are zero. This is used as a specification test, since the hypothesis that the lag length of a VAR is q corresponds to the hypothesis that the fitted disturbances in a VAR(q) are white noise. The Stata Technical Bulletin (available on the course web site) indicates that the Q statistic is, under the null hypothesis, a realization from a chi-square distribution with a number of degrees of freedom that is a function of q and n , as well as the number of variables in the VAR.
 - (i) Do the Q statistics associated with $q = 1, \dots, 8$ show evidence of serial correlation in the residuals if the chi-square sampling theory is used?
 - (ii) What if the null distribution of the Q statistic is instead obtained by the bootstrap?
2. Prove that the estimated VAR fit in ACEL has the following property. A disturbance in any of the non-technology shocks (i.e., any element of e_t other than the first two) has no effect on the level of labor productivity and the level of the price of investment goods, in the long run.
 3. Plot the estimated monetary policy shocks, as well as their standard deviation over time (for the latter, use a centered rolling window of 7 observations to compute the standard deviation). Determine the standard deviation of the monetary shocks for the whole sample by looking at the intercept in the impulse response function of the interest rate to a monetary policy shock (see your answer to 1 above). (The monetary policy shock is in units of percentage points, so that multiplication by 100 converts to basis points). Does the overall estimate of the standard deviation of a monetary policy shock seem high in light of what you know about monetary policy in central

banks? Based on examination of the estimate of the policy shocks, is there a subperiod that plays a particularly large role in determining the overall standard deviation?

4. Conventional macroeconomic analyses use a measure of population that corresponds to the non-institutional population aged 16 and over. Francis and Ramey (2004, course web site) have recently argued for a different measure of population.
 - a. How does the time series plot of per capita hours based on the Francis and Ramey population data compare with the time series of the data used in the benchmark data set? For example, do the two series have the same trend? (Hint: the required time series plots appear automatically if you run `main.m` with $FR = 1$ or 2 .)
 - b. How do their growth rates compare?
 - c. Estimate the VAR using data constructed with the Francis and Ramey population data.
 - (i) Note that the VAR now has an explosive root. Determine if this is because of the apparent trend in per capita hours based on the Francis-Ramey population data by removing a time trend from the per capita hours series ($FR = 2$). Do the results support the idea that the explosive root is due to the trend in per capita hours?
 - (ii) Compare the impulse response functions based on the Francis-Ramey population measure with what you found in 1 above. In particular, consider the response of hours to a neutral technology shock.
 - (iii) Have a look at the Francis and Ramey paper, to see how their population measure is constructed. What do you think the best way is to measure population for the purpose of fitting a general equilibrium macroeconomic model?
5. Estimate the benchmark VAR over different sample periods. How do the impulse response functions compare? Are the estimated VARs always covariance stationary? How might you interpret the fact that sometimes the VAR has an explosive root when fit over some subperiods? Other robustness checks can be performed too. For example, the benchmark data set uses non-farm business hours as its hours worked measure. One could instead use

business hours (see the parameter, `hours`, at the beginning of `main.m`). The population data has suspicious spikes in them (see the previous question). Check whether these influence the results by working instead with a smooth HP filter of the population data.

6. Consider the Variance Decompositions. Five types are computed in `main.m`, and which one looks at makes a difference. They can be categorized according to two different criteria. One is based on the length of the data set: whether it is based on the population of data implied by the VAR, or on a sample the same length as the data set. The second criterion has to do with the aspect of the data used in the criterion. For example, the conventional variance decomposition is concerned with the forecast error variance implied by the VAR. In practice, this measure is evaluated in population. Another focusses on different frequency components of the data: say the Hodrick-Prescott filtered component, or the band-pass filtered component.
 - a. Consider the historical decomposition of the data, from question 1 above. One measure of variance decomposition takes the ratio of the variance of the thin lines in that graph to the thick lines, after these lines have been filtered (thick line: raw - detrended - data; thin line - what the data would have been had only the indicated estimated shocks been operative). For example, they could be HP filtered or Band Pass filtered. We call these the *In-sample Variance Decomposition, HP Filtered Data* and the *In-sample Variance Decomposition, BP Filtered Data*. Consider the tables of variance decompositions (the Scientific Word table, `table.tex`, generated by running `main.m` with `vardecomp = 1`.) The fourth and fifth tables show the results of applying the in-sample variance decompositions. The tables also show the mean value (in parentheses) of this measure of decomposition and the standard deviation (in square brackets), across `ndraws` bootstrap simulations of the estimated VAR. Note how high the fraction of variance is attributed to the monetary policy shock: Based on the BP filtered data the monetary policy shock accounts for 75 percent of the variation of the data. Note that the variances do not add up across shocks. Can you explain this? Note that the in-sample point estimate is substantially higher than the corresponding mean in repeated bootstrap samples. Can you explain this?

- b. The program, `main.m`, computes two other types of variance decomposition. The program reports the decomposition of forecast error variance 1, 4, 8, 12 and 30 quarters out. Why is the variance due to monetary policy shocks zero for some variables at the 1 quarter horizon? The program also reports the population decomposition of variance in the BP filtered data and the HP filtered data. Note that for the monetary policy shock, policy shocks contribute relatively more to variance in HP and BP filtered data than they do to forecast error variance. How can you interpret this difference? The reverse is true for the neutral technology shock. How do you interpret this?
- c. Recompute the variance decompositions with $nlags = 6$. Is there still a big discrepancy between the estimated in-sample variance decomposition and the others? If not, then what should we infer from the fact that there is a discrepancy when $nlags = 4$, but there isn't when $nlags = 6$? (Hint: in the bootstrap simulations the disturbances are by construction serially uncorrelated, while they are whatever they are in the in-sample computations. Perhaps this evidence suggests that there is evidence of serial correlation in the $nlags = 4$ VAR.)