Washington University
Department of Economics

Spring 2010
James Morley

Economics 4151

Homework #4
Due on Tuesday 4/26

In this assignment, you will use maximum likelihood estimation (MLE) to find estimates of parameters for a few ARMA models and an ADL forecasting model for U.S. real GDP growth. You will also conduct hypothesis tests and construct confidence intervals for some of these parameters.

1. Estimation in EViews (15 points)

Download the latest vintage of postwar quarterly U.S. real GDP from FRED. Also, download postwar quarterly "real change in private inventories" from FRED (it is also under GDP and components). Import the data into EViews.

i. Based on the notion that the long-run growth rate of U.S. real GDP is relatively stable, take natural logarithms of the raw data and multiply by 100 so that movements in the transformed series, denoted y_t, have the interpretation of percentage point movements. Run an ADF unit root test, including a constant and trend in the test regression. Consider both SIC and AIC lag selection. What are the test statistics vs. 5% critical values? Do you reject the unit root null? Why did you include a constant and a time trend? Repeat the test for the first differences of 100 times ln(real GDP), which we will denote as Δy_t. This time only include a constant in the regression. Again, consider both SIC and AIC lag selection. What are the test statistics vs. 5% critical values? Do you reject the unit root null? Why did you include and constant, but no trend in the test regression? When reporting results, include plots of y_t and Δy_t. (1 page)

ii. In terms of change in inventories, divide real change in inventories by real GDP (level, not logs) and multiply by 100 to get change in inventories as a percentage of GDP. Run an ADF unit root test for the change in inventory ratio, including a constant in the test regression. Consider both SIC and AIC lag selection. What are the test statistics vs. 5% critical values? Do you reject the unit root null? Why did you include a constant? When reporting results, include a plot of $\Delta H_t^* = 100 \times \dfrac{\Delta H_t}{Y_t}$, where ΔH_t denotes the real change in private inventories and Y_t denotes real GDP. (1 page)

iii. Still in EViews, estimate the following four models for Δy_t: an AR(1) model, an AR(2) model, an ARMA(2,1) model, and an ADL(1,1) model with a constant, lagged Δy_t, and lagged ΔH_t^* on the right-hand-side (no contemporaneous change in inventories for the forecasting model). For example, to estimate the ARMA(2,1) model, you can select "Estimate Equation…" from the "Quick" menu and type "dlrgdp c ar(1) ar(2) ma(1)" in the box, where "dlrgdp" is the name for Δy_t. Note that estimating the univariate ARMA models in this way (instead of using, say, "dlrgdp(-1 to -2)" for the AR terms) produces an estimate for c that is actually an estimate of the unconditional mean $\mu = c/(1-\phi_1-\phi_2)$, not of the intercept. This has no impact on the other estimates in the model. On the other hand, to estimate the ADL model with inventories, you need to use "dlrgdp c dlrgdp(-1) dhstar(-1)" and not "ar(1)" for the lagged output growth term (they are not equivalent in a multivariate setting). Report your estimates for the four models in a concise manner. The key information to report in each case is point estimates, standard errors, log likelihood, and sample period. Note that the sample periods will not be the same in every case. How would you rank the forecasting models? Are inventories helpful for forecasting? (1 page)

2. **Estimation in GAUSS** (10 points)

The previous estimation in EViews corresponded to conditional MLE (under the assumption of Gaussian errors). In this exercise, you will estimate the first three models (i.e., the univariate models) in GAUSS using exact MLE (under the assumption of Gaussian errors). I have put GAUSS code on the class website that gives an idea of how to do this.

i. Adjust the GAUSS code for your data file and the model that you want to estimate. Estimate the three models and report your estimates, standard errors, and log likelihood, in a concise manner. Are the exact MLE estimates similar to those for conditional MLE? (1 page)

3. **Hypothesis Tests and Confidence Intervals** (35 points)

Background: It is possible to use a *t*-test to determine if a parameter is significantly different from zero or an *F*-test to test multiple parameters. However, in time series, this "Wald test" approach (i.e., testing based on estimates under the alternative only) has the problem that it generally relies on numerical derivatives to find estimates of asymptotic approximation of the variance-covariance matrix for the maximum likelihood estimates. For basic models, it is possible to solve analytically for the derivatives. However, it is often easier to simply estimate the model under the null and the alternative and then construct a Likelihood Ratio (LR) test. This approach to testing does not rely on numerical derivatives.

In this exercise, you will consider both *t*-tests and *LR*-tests. You will also construct confidence intervals based on these tests. Finally, you will consider the bootstrap distribution of the *t*-test and you will construct bootstrap percentile confidence intervals. Again, in order to help, I have put related GAUSS code on the class website.

i. Test the AR(1) vs. the AR(2) model. The null hypothesis is $H_0 : \phi_2 = 0$. In particular, based on your estimates in question 2, construct a *t*-test for this hypothesis. What is the test statistic? What is the 95% critical value (consider a two tailed test)? Report the asymptotic critical value (based on N(0,1), not the finite sample value based on a t-dist with T-1 degrees of freedom). Also, report a p-value based on the asymptotic distribution. Now, based on your estimates in question 2, construct an LR test for this hypothesis. What is the test statistic? What is the 95% critical value? Again, report the asymptotic critical value (in this case based on a $\chi^2(1)$ distribution) and the p-value. How well do your results accord across the two tests? (1 page)

ii. Construct an asymptotic 95% confidence interval for ϕ_2 based on the *t*-test. Then, construct an asymptotic 95% confidence interval for ϕ_2 by inverting an LR test. In particular, compare models that impose values for ϕ_2 with the model in which it is freely estimated. Let the imposed value range from -0.9 to 0.9 in increments of 0.05 and re-estimate the other parameters. This produces a "likelihood profile" for ϕ_2. For each null model, construct the LR test based on your estimates in exercise 2 for the alternative model. Plot the LR test statistics against the null ϕ_2. There should be a "V" shape, with the base of the V occurring at the null ϕ_2 that is closest to your estimated $\hat{\phi}_2$ and the value of the LR statistic being close to zero (if one of your null ϕ_2's is exactly equal to $\hat{\phi}_2$, the LR stat will be exactly zero). In the same plot, plot a line equal to the 95% critical value for the LR test (assume it has an asymptotic $\chi^2(1)$ distribution). Define your inverted LR confidence interval as the outermost points of the set of null hypotheses that cannot be rejected. How well does this inverted LR confidence interval accord with the standard one? (1 page)

iii. Now let's consider the bootstrap approach to finding a better approximation of the finite sample distribution for the *t*-test in part i. Estimate the model under the null hypothesis. These estimates will form the basis for your bootstrap data generating process (BDGP). Draw repeated samples from your BDGP and record the *t*-statistic for $H_0 : \phi_2 = 0$ for each sample. Consider 999 bootstrap experiments (see MacKinnon, 2006, in the reading package on why 999, instead of 1000). Given the bootstrap distribution for the *t*-statistic, find the 95% critical value and find the bootstrap p-value of the *t*-statistic you found using the real data. (1/2 page)

iv. Finally, consider the bootstrap percentile approach to constructing confidence interval. In particular, use the estimates of the model under the alternative (i.e., the model that includes ϕ_2) to determine a BDGP. Draw repeated samples from this BDGP and record the estimates for ϕ_2 for each sample. Consider 999 bootstrap experiments. Given the bootstrap distribution for $\hat{\phi}_2$, determine the 95% confidence interval based on the percentiles of the bootstrap distribution. Remember to "flip" the quantiles. Explain why you should flip the quantiles and why this approach requires that the bootstrap distribution is close to being pivotal. (1/2 page)

v. Repeat i.-iv. for testing the MA(1) parameter θ in the ARMA(2,1) model. Again, let the imposed value range from -0.9 to 0.9 in increments of 0.05 and re-estimate the other parameters for ii. In this case, you should find very different results for the *t*-based confidence interval than for the inverted LR confidence interval. Why do you think the results are so different? (same page suggestions as before)