Problem Set 3

Question 1

Consider the AR(1) process

$$y_t = \rho y_{t-1} + \epsilon_t,$$

where ϵ_t is normally distributed with mean 0 and variance σ^2. Show that the maximum likelihood estimates of ρ and σ conditional on the first observation being non-stochastic can also be obtained as the OLS estimates of the process. What are the asymptotic approximations to the standard errors of the maximum likelihood estimates for ρ and σ?

Question 2

Consider the process

$$y_t = 0.9 y_{t-1} + 0.01 \epsilon_t$$

where ϵ_t is drawn from a standard normal process. Using a random number generator such as "randn" in Matlab, generate 100 observations of this process. You may assume that the initial point in the sample is non-stochastic and coincides with the steady state for the process. For replication purposes, before invoking the random number generator, set the "seed" for the random number generator.

Given your first Monte Carlo sample, compute the conditional maximum likelihood estimates of the autoregressive coefficient and the standard deviation for the innovation of the autoregressive process that you used in generating the sample. Following Handout 6, compute bootstrap standard errors for your estimates.

What is your bootstrap estimate of the bias for the original estimates?

Question 3

Using the same process considered in Question 2, reset the seed to 1 and draw 1000 samples containing 100 observations, each starting from the steady state. For each sample, form the conditional maximum likelihood estimates for the autoregressive coefficient and for the standard deviation of the innovation. Use those estimates to form bootstrap 90% confidence intervals, as described in Handout 6. What is the effective coverage of your confidence intervals (i.e. the frequency with which the confidence intervals include the parameters used in the data-generating process)?

Question 4

Using the same process of Question 2, construct a new sample containing 100 observations. This time, however, to avoid dependence on the initial point, start from the steady state, draw 200 observations, then discard the first 100 (remember to set the seed for your random number generator for ease of replication). What are the unconditional maximum likelihood estimates for the autoregressive coefficient and the standard deviation of the innovation? What are the conditional maximum likelihood estimates?

If you are using Matlab, I would advise you to use the command "fminunc" to compute the unconditional maximum likelihood estimates. Remember the numerical minimizer might elicit to evaluate your likelihood function for values of the parameters implying nonstationarity of the AR(1) process. Use a penalty function to resolve this problem.

Question 5

Download the NIPA Table 1.1.5 "Gross Domestic Product" from the website of the Bureau of Economic Analysis. Consider the series in lines 2, 7, 14, and 21, respectively Personal Consumption Expenditures, Gross Private Domestic Investment, Net Exports of Goods and Services, Government Consumption Expenditures. Normalize each of these series by line 1 Gross Domestic Product. Plot each ratio and decide whether or not to detrend, then choose and estimate an AR process. Approximate MA components by increasing the order of the AR process. Explain your choices.

For this question, you may work in MATLAB, or use your preferred packaged econometric software, such as EVIEWS. In both cases, attach the relevant replication code to your answers.