An Example of Dealing with Non-Stationarity[1]

Earlier, we failed miserably at estimating the effect of monetary policy on the real economy. Among the many flaws of our specification was that we may have included non-stationary variables. These notes simply apply the recent material on trends and unit roots to these time series, showing how to render the transformed timer series stationary.

We will use housing starts as an example. I begin by adjusting for population growth by creating a per-capita version of the variable"

gen pchs=ln(hs)-ln(pop);

The first choice we must make is whether or not we are concerned that the variable may include both a trend and a unit root(s). In some cases, there may be established results or related work that convinces us that we need only test for either a trend or a unit root. Here, however, we will consider the possibility of both.

The first step is to choose the number of lags we should include in our Augmented Dickey Fuller tests. Stata's *dfgls* command is a good way to do this:

dfgls *pchs*

This command includes a trend unless we specifically include the *notrend* option. Stata wil use a linear trend by default. This shows another case for logging variables. A constant trend growth rate is a linear trend when we work with logs. Stata will choose the maximum number of lags to include, or we can impose a value using the *maxlag*(#) option.

Below the output, *dfgls* reports the optimal lag length. The result depends on different information criteria, here both 1 and 16 are suggested. This matters. For 16 lags, we reject the null hypothesis of a unit root with 95% confidence (the test stat is less than the critical value). For 1 lag, however, we fail to reject. If different information criteria yield different results we may 1) make a case for using on of the information criteria, probably by referencing the literature that tests them in different circumstances, or 2) error on the side of differencing. Although over-differencing is problem, non-stationarity is worse.

[1] These are undergraduate lecture notes. They do not represent academic work. Expect typos, sloppy formatting, and occasional (possibly stupefying) errors.

The main table reports that, for any lag length, we fail to reject the null hypothesis of a unit root because the Dickey Fuller test statistic is not less than the corresponding critical value. Because this result is invariant to the choice of lag length, we need not worry, in this case, about the proper lag length.

Stata's *dfuller* command can provide more detail:

dfuller *pchs*, lags(4) trend regress

Notice that the results are a little different. This is because, the *dfgls* automatically tests for a unit root around de-trended data. To test for a trend itself, we should use *dfuller*. We must include the *trend* option if we want to allow for a linear trend. We must also include the *regress* option if we want to see more than the test statistic and critical values.

In this example, we fail to reject both a unit root and the null of no trend. We may thus proceed as if our time series has no trend but is at least I(1), meaning it has at least one unit root. If we use the log of GDP in our specifications, then we will violate Gauss-Markov.

A variable may have more than one unit root. We must thus check if the first difference of logged GDP also has a unit root.

dfgls *d.pchs*

Each time we difference, we lose an observation. Because Dickey Fuller tests require that we reject a null hypothesis, they are known for finding unit roots that are nit really there. We can find ourselves in a trap where by continually differencing, we continually increase our standard errors, and we thus continually find unit roots. WE thus cannot ignore intuition. I know housing starts are not I(65), even if this process yields such a result.