

Statistics 434: Bullet Points for Day 8

ARIMA(p,d,q) In Full

After considering the “stylized facts” suggested by HW3, we complete our description of the most widely used class of univariate time series models, the ARIMA(p,d,q) models. Our main task is to understand how the models reflect structural properties of time series. Now in addition to worrying about stationarity and ergodicity there is a new kid on the block *identifiability*.

- Return HW2. Class discussion of the in-coming HW3
 1. First encounter with the notion of a “stylized fact”
 2. Consideration of individual results
- The MA(q) model
- The ARMA(p,q) model
 1. The Φ and Θ polynomials
 2. Stationarity
 - (a) The general root criterion
 - (b) The stationarity triangle for the AR(2) model
 3. Wold Decomposition for the ARMA(p,q) model
 4. Issues of Identifiability
- ARIMA(p,d,q)
 1. Differencing to achieve stationarity — part of “Box-Jenkins Methodology”. Surprising how often it works.
 2. Since $R_t = \log(P_t/P_{t-1}) = \Delta \log(P_t)$ we’ve usually already differenced.
- More on Ergodicity, Stationarity, and “Practical” assumptions
 1. Two problems that we typically *assume away*. But, please, let’s not forget that we did this. There are times when it can be very big trouble. These are usually *interesting* times.
 2. A variation on the “super dull” non-ergodic stationary process which shows that at a certain bizarre theoretical level no fixed amount of testing can resolve the issue of ergodicity or stationarity.

WHAT CAN I ASSUME?

Modelers must make assumptions, and modelers who have any where to turn inevitably turn and ask “Is it OK if I assume that ...?” The honest answers to this are pretty lame. Most honest is “I don’t know.” Less honest is “It depends...” The second answer is less honest because it just tries to turn the question back to the questioner. We’ll dig into this one shovel full at a time. It is possible to get past the absent minded patter — but its not easy.