

Statistics 434 : Bullet Points for Day 9 Simulating and Fitting ARIMA(p,d,q) Models

The main formal task is to pick up the computation tools for studying the ARIMA(p,d,q) models. In particular, we consider the tools for simulation and for fitting using the method of maximum likelihood. We also start to consider when one model might be preferred to another based on the ACF or PACF plots; more model selection tools are on the way. Finally, we do inventory of the big picture issues that are dogging our heels.

- Simulation Using `arma.sim()`
 1. ACF and PACF diagnostics
 2. Examples in simulation
- Fitting ARIMA models using S-Plus `arma.mle()`
 1. Model specification — with examples
 2. Attention to the “theta sign” convention for MA models
 3. Example of estimations
- Discussion of HW4
- How and what we can learn from simulations
 1. Behavior of methods and estimators — with the knowledge that the model is right
 2. Understanding of the wide ranging impacts of luck — in a controlled theoretical context. This is an underestimate of “luck” in the real world, yet it can still be very informative — though mostly with a cautionary spin.
 3. Human Frailty. Our intuition about chance is not very good — so calculation and simulation are valuable guides — even if we are forced to use “toy” models.
- *Big Picture* Issue Inventory
 1. Stationarity — Important and (Almost) Untestable
 2. Ergodicity — Important and Provably Untestable
 3. Stylized Facts — Humbly seat-of-the-pants stuff, but still the most honest stuff we’ve got
 4. Normality — and collective Schizophrenia (e.g. our LB test and our MLE estimates vs our JB and WS tests.)
 5. Parsimony — the principle everyone believes, yet for which there is no honestly compelling “theory.”

QUOTE OF THE DAY:

“IT ISN’T WHAT WE DON’T KNOW THAT GIVES US TROUBLE, IT’S WHAT WE KNOW THAT AIN’T SO.” — WILL ROGERS