Statistics 434: Bullet Points for Day 3
AR(1) Estimation — Point Estimates and Their Distributions

The AR(1) model will serve as our "model for a model." That is, what we can see and say about AR(1) provides an outline of what we would hope to see or say about any of the other models that we will consider. A further benefit of the AR(1) model is that it captures a variety of intuitive behaviors, yet is simple enough to permit us to work out explicit formulas for key features like the "long term distribution," and the maximum likelihood estimates for the parameters.

- Features of the AR(1) Model
  - "Long term distribution" vs conditional distributions
  - Stationarity by "starting with the long-run distribution".
  - Stationarity by "Burn-in" — the more typical method
- Estimation of the parameters ρ and σ^2
  - Review of maximum likelihood estimation
  - Derivation of the *approximate* MLE for ρ and σ^2
  - Distribution of $\hat{\rho}$ — or at least the *traditional approximation*
- Simulation and Distribution (as needed in Homework No. 1)
  - Simulation Tools — rnorm, rt, etc.
  - Comparison Tools — qqplot, qqnorm, also abline.
  - Step 1 (for free!): Writing an S-function for $\hat{\rho}$
  - Generating one AR(1) sample and calculate $\hat{\rho}$
  - Replicating this 1000 times
  - Looking at 1000 observations of $\hat{\rho}$
- Some philosophical — and practical — additions
  - A discrete example where you can guess the long-term distribution.
  - Why stationarity is about the least we can assume and still hope to learn from the past.

We Weave a Web

Pure sweet Gaussian noise is the simplest model for the returns on an asset, and the AR(1) model is its simplest alternative. More precisely, it is the simplest alternative model that makes intuitive sense and that contains the noise model as a special case. Finally, even in its extreme simplicity the AR(1) model brings us face to face some crucial theoretical constructs — such as *stationarity* and the "long-term" (or *ergodic*) distribution.