Econ 706 Prelim

June 2011

All questions are weighted equally. Good luck!

Consider Bayesian analysis of a simple Markovian dynamic model:

$$y_t = \phi y_{t-1} + \varepsilon_t$$

$$\varepsilon_t \sim N(0, \sigma^2)$$

$$|\phi| < 1.$$

(1) Derive the natural conjugate prior (NCP) for $\phi|\sigma$. What is the posterior mean of $\phi|\sigma$ and how is it related to the prior mean, the MLE, the prior precision, and the MLE precision?

From this point onward, use the NCP for $\phi|\sigma$ (and, when relevant, for $\sigma|\phi$) unless explicitly instructed otherwise.

(2) What is the key benefit of using NCPs? What is the key cost? Maintaining use of NCPs, how might you attempt to represent prior ignorance regarding $\phi|\sigma$? What complication arises? If you could use non-NCPs, how might you attempt to represent prior ignorance regarding $\phi|\sigma$? What complication arises?

(3) Why might you be interested in posterior inference not on $\phi|\sigma$, but rather on ϕ? Design a Markov chain Monte Carlo algorithm (in this case, a Gibbs sampler) to sample from the joint posterior distribution of (ϕ, σ). Be precise. How would you assess convergence to steady state of your Gibbs sampler? How would you use the Gibbs draws from the joint posterior to sample from the marginal posterior of ϕ?

(4) Is the Gibbs sequence for ϕ likely to be serially uncorrelated? If not, what is the likely form of the serial correlation? If you know the spectral density at frequency zero of the Gibbs sequence of ϕ, $f_\phi(0)$, how would you use it to assess the accuracy of your estimate of the posterior mean of ϕ based on the Gibbs sequence?

(5) If instead $f_\phi(0)$ is unknown, how would you estimate it consistently using an autoregressive model-based estimator? What determines the bandwidth, what conditions must it satisfy, and how might you select it in practice?