Econ 226, Spring 2015

Midterm Exam

DIRECTIONS: No books or notes of any kind are allowed. Answer all questions on separate paper. 100 points are possible

1.) (50 points total) A pth-order Gaussian vector autoregression (denoted VAR(p)) can be written as

$$\underset{(n\times 1)}{\mathbf{y}_t} = \underset{(n\times k)}{\mathbf{\Pi}'}\underset{(k\times 1)}{\mathbf{x}_{t-1}} + \underset{(n\times 1)}{\boldsymbol{\varepsilon}_t}$$

$$\boldsymbol{\varepsilon}_t \sim N(\mathbf{0}, \mathbf{\Omega})$$

where n is the number of different variables. The Normal-Wishart distribution turns out to be the natural conjugate. For the Normal-Wishart prior, the prior for $\mathbf{\Omega}^{-1}$ is distributed Wishart with N degrees of freedom and scale matrix $\mathbf{\Lambda}$ (denoted $\mathbf{\Omega}^{-1} \sim W(N, \mathbf{\Lambda})$) while $\boldsymbol{\pi} = vec(\mathbf{\Pi})$ has a distribution conditional on $\mathbf{\Omega}$ that is $N(\mathbf{m}, \mathbf{\Omega} \otimes \mathbf{M})$.

a.) (10 points) What are the values of $\mathbf{x}_{t-1}$ and k?

b.) (10 points) What is the dimension of $\mathbf{\Omega}$? Describe in words the meaning of the first element of $\boldsymbol{\varepsilon}_t$ and the (1,1) element of $\mathbf{\Omega}$.

c.) (10 points) Describe in words the meaning of the (1,1) element of $\mathbf{M}$. Indicate what it would correspond to if your prior came from observation of an earlier data set.

d.) (10 points) According to the Minnesota prior (developed by Doan, Litterman, and Sims) what values should you use for the first k elements of $\mathbf{m}$? Write these in the form of an equation, $\mathbf{m}_1 = ...$, where $\mathbf{m}_1$ denotes the first k elements of $\mathbf{m}$.

e.) (10 points) Name one advantage of using the natural-conjugate prior in this example. Name one disadvantage.

Econ 226, Spring 2015

2.) (50 points total) Let $\mathbf{Y}$ denote an observed vector of data and $\boldsymbol{\theta}$ a vector of parameters. Suppose you are able to write a computer procedure that could calculate the value of the prior density $p(\boldsymbol{\theta})$ and of the likelihood $p(\mathbf{Y}|\boldsymbol{\theta})$ but are not able to calculate any magnitudes analytically.

a.) (10 points) You know using only the information given that the posterior density $p(\boldsymbol{\theta}|\mathbf{Y})$ is proportional to some function $h(\boldsymbol{\theta})$. Write down the formula for $h(\boldsymbol{\theta})$.

b.) (10 points) Importance sampling involves generating draws from some proposal density $g(\boldsymbol{\theta})$. Suppose you generated M realizations, denoted $\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}, ..., \boldsymbol{\theta}^{(M)}$ where $\boldsymbol{\theta}^{(m)}$ has the density $g(\boldsymbol{\theta})$. Write down an expression you could use to estimate the posterior mean $E(\boldsymbol{\theta}|\mathbf{Y})$.

c.) (10 points) What is required of the proposal density $g(.)$ in order for the procedure you described in (b) to give a good estimate of $E(\boldsymbol{\theta}|\mathbf{Y})$? Can you suggest any practical guidelines for chosing $g(.)$?

d.) (10 points) Would you throw out any of your generated $\boldsymbol{\theta}^{(m)}$ as "burn-in" draws? Why or why not?

e.) (10 points) Suggest two approaches you might use to convince yourself you obtained the correct answer by following the procedure you described in part (b).