Statistics 5314: Homework 5

For each homework assignment, turn in at the beginning of class on the indicated due date. Late assignments will only be accepted with special permission. Write each problem up *very* neatly (LaTeX is preferred). Show all of your work.

Problem 1

For this problem we will implement the Gibbs sampling algorithm. Consider the model $x_i \sim N(\mu, \phi = 1/\sigma^2)$. Denote $\boldsymbol{X} = \{x_1, \ldots, x_N\}$.

Part a

Under the prior distribution $p(\mu) \propto 1$, find the conditional distribution $p(\mu|\boldsymbol{X}, \phi)$. Hint: $p(\mu|\boldsymbol{X}, \phi)$ will be normally distributed. Your job is to find the parameters (which will depend on $\boldsymbol{X}$ and ϕ) that define this distribution.

Part b

Under the prior distribution $p(\phi) \propto 1/\phi$, find the conditional distribution $p(\phi|\boldsymbol{X}, \mu)$. Hint:

$$p(\phi|\boldsymbol{X}, \mu) \propto \phi^{\alpha-1}e^{-\beta\phi}$$

which is the kernel of a Gamma distribution, with parameters (α, β). Your job is to find α and β.

Part c

Under $\mu = 8$ and $\sigma = 2$, simulate $N = 1,000$ points. Under these points, implement the full conditionals that you derived to implement the Gibbs sampler. Show a plot (2-dimensional) of the iterations for the joint samples. Also, show trace and autocorrelation plots for the margins (μ and ϕ).

Problem 2

Consider the log-$\chi^2/2$ distribution defined by $\nu = \log(\kappa)/2$, where $\kappa \sim \chi^2_1$. This distribution can be very closely approximated by discrete mixture of normals, with known parameters, i.e.,

$$p(\nu) = \sum_{j=1}^{J} q_j N(b_j, w_j)$$

where,

q_j :	0.0073	0.0001	0.1056	0.2575	0.3400	0.2455	0.0440
b_j :	-5.7002	-4.9186	-2.6216	-1.1793	-0.3255	0.2624	0.7537
w_j :	1.4490	1.2949	0.6534	0.3157	0.1600	0.0851	0.0418

See Kim, Shephard and Chib, 1998, *Review of Economic Studies* for details on this. More Components can refine this approximation, but we will limit ourselves to this small mixture model.

To sample from this model, sample component j with probability q_j, then sample from the corresponding normal $N(b_j, w_j)$. Repeat this 1,000 times to obtain 1,000 samples. Plot a histogram of the results.

Problem 3

Sample $\kappa_i \sim \chi^2_1$ for $i = 1, \dots, 1,000$. Transform each sample by $\nu = \log(\kappa)/2$.

Part a

Plot a histogram of the results.

Part b

plot a Q-Q plot comparing these samples to those obtained in problem 2.

Problem 4

Consider the Stochastic Volatility model discussed in class. Under reparameterization $y_t = \log(r_t^2)/2$, we obtained the HMM:

$$\begin{aligned} y_t &= \mu + x_t + \nu_t \\ x_t &= \theta x_{t-1} + \epsilon_t, \end{aligned}$$

where $\nu_t = \log(\kappa_t)/2$ and $\kappa_t \sim \chi_1^2$, and $\epsilon_t \sim N(0, \psi^2)$.

Part a

List all the unknown parameters in this model, which would be required for simulation.

Part b

Choose 3 configurations of parameters and simulate the time series for y_t. Your simulation should use the approximation to the log-$\chi^2/2$ as given in problem 2. Show plots of your results.

Part c

Given a time series for y_t, you might want to infer all the unknown parameters in the model (especially the latent time series for x_t). Consider the model as specified under the approximation to the log-$\chi^2/2$ distribution, and explain how Gibbs sampling can be used for inferring all the unknown parameters. You need to write down, for each unknown parameter, its full conditional distribution (when possible) and detail the steps in the algorithm. You do not need to code this up.