Bayesian Econometrics

Problem Set

This problem set consists of 5 questions, and is due on 30 April.

1. (5 points) Suppose $x_1 = 1.11, x_2 = 0.53, x_3 = 11.14, x_4 = 0.49, x_5 = 2.47$ is an observed sample from the $\mathcal{G}(1, \lambda)$ distribution. Consider Bayesian inference for the parameter λ, using the Gamma prior $\lambda \sim \mathcal{G}(1, 2)$.

   (a) Find the posterior distribution of λ.

   (b) Find the posterior mean.

2. (5 points) Suppose we wish to obtain a random sample from the distribution

$$f(x, y) \propto x\mathrm{e}^{-x-2y-10xy}, \qquad x, y > 0.$$

   (a) Derive the conditional densities $f(x \,|\, y)$ and $f(y \,|\, x)$.

   (b) Write a Gibbs sampler to obtain a random sample of size $50,000$ from $f(x, y)$. Estimate the expected values $\mathbb{E}\mathbf{X}$ and $\mathbb{E}\mathbf{XY}$ from the Gibbs output.

3. (10 points) Consider the linear regression model

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \qquad \boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \sigma^2\mathbf{I}_T)$$

   with the natural conjugate prior $(\boldsymbol{\beta}, \sigma^2) \sim \mathcal{NIG}(\boldsymbol{\beta}_0, \mathbf{V}_{\boldsymbol{\beta}}, \nu_0, S_0)$. Derive a formula for the marginal likelihood

$$p(\mathbf{y}) = \int p(\mathbf{y} \,|\, \boldsymbol{\beta}, \sigma^2)p(\boldsymbol{\beta}, \sigma^2)\mathrm{d}(\boldsymbol{\beta}, \sigma^2).$$

4. (10 points) Consider the following AR(p) model:

$$y_t = \beta_0 + y_{t-1}\beta_1 + \cdots + y_{t-p}\beta_p + \varepsilon_t, \qquad \varepsilon_t \sim \mathcal{N}(0, \sigma^2),$$

   and assume the natural conjugate prior $(\boldsymbol{\beta}, \sigma^2) \sim \mathcal{NIG}(\mathbf{0}, 10\mathbf{I}_{p+1}, 3, 10)$. Using the US PCE dataset, compute the marginal likelihoods of the AR(p) models with $p = 1, 2, 3$ and 4. Which lag length should we use?

   [Hint: always use the first 4 observations as the initial conditions y_0, y_{-1}, y_{-2} and y_{-3} so that $\mathbf{y}$ remain the same for all models. For numerical stability, compute the marginal likelihood in log.]

5. (20 points) This question concerns an approximate method to estimate the following stochastic volatility model:

$$y_t = \mathrm{e}^{\frac{1}{2}h_t}\varepsilon_t, \qquad \varepsilon_t \sim \mathcal{N}(0, 1), \tag{1}$$

$$h_t = h_{t-1} + \zeta_t, \qquad \zeta_t \sim \mathcal{N}(0, \sigma_h^2) \tag{2}$$

for $t = 1, \dots, T$, where the initial condition h_0 is treated as an unknown parameter. Since h_t enters the measurement equation in (1) nonlinearly, the joint conditional distribution of $\mathbf{h} = (h_1, \dots, h_T)'$ given the data and other parameters is no longer Gaussian. Suppose we transform the data as $y_t^* \equiv \log y_t^2$. Then, (1) can be rewritten as

$$y_t^* = h_t + \varepsilon_t^*,$$

where $\varepsilon_t^* \equiv \log \varepsilon_t^2$ has a log-χ_1^2 distribution. So we have transformed a nonlinear Gaussian state space model into a linear *non-Gaussian* state space model.

(In practice, we use $y_t^* \equiv \log(y_t^2 + 10^{-4})$ to avoid numerical issues.)

(a) Suppose we approximate the log-χ_1^2 distribution using the $\mathcal{N}(-1.27, 4.94)$ distribution. Under this approximate model and the priors $h_0 \sim \mathcal{N}(0, V_{h_0})$ and $\sigma_h^2 \sim \mathcal{IG}(\nu_h, S_h)$, derive a 3-block Gibbs sampler to sample from $p(\mathbf{h} \,|\, \mathbf{y}, \sigma_h^2, h_0)$, $p(\sigma_h^2 \,|\, \mathbf{y}, \mathbf{h}, h_0)$ and $p(h_0 \,|\, \mathbf{y}, \mathbf{h}, \sigma_h^2)$.

(b) Estimate the approximate model using the YEN/USD exchange rate data in `YENUSD.csv` with $V_{h_0} = 100, \nu_0 = 3$ and $S_0 = 0.4$. Report the posterior mean of σ_h^2 and plot the posterior means of $\mathrm{e}^{\frac{1}{2}h_t}$.