Advanced Bayesian Econometrics

Homework Assignment #2

- Due Date: November 18th, 2026
- Instructor: Hedibert F. Lopes
- Teaching Assistant: Ricardo Pedroso

Homework 2 is to be done individually and submitted electronically via email to the TA, Ricardo Pedroso.

Model

Let $x_1, \ldots, x_n$ be iid $\mathrm{N}(\theta, 1)$, so that

$$p(\mathbf{x} \mid \theta) \propto \exp\Big\{ -\frac{1}{2} \sum_{i=1}^{n} (x_i - \theta)^2 \Big\}.$$

Consider two representations of the same prior for θ:

(A) Double q-exponential prior:

$$p(\theta) = \frac{1}{4} \left(1 + \frac{|\theta|}{2} \right)^{-2}, \qquad \theta \in \mathbb{R}.$$

(B) Hierarchical (scale-mixture) prior:

$$\theta \mid \tau^2 \sim \mathrm{N}(0, \tau^2), \qquad \tau^2 \mid \lambda \sim \mathrm{Exp}(\lambda^2/2), \qquad \lambda \sim \mathrm{Exp}(2),$$

where $\mathrm{Exp}(r)$ denotes the exponential distribution with rate r (mean $1/r$).

Questions

0. Data. Simulate $n = 20$ observations $x_1, \ldots, x_{20}$ iid $\mathrm{N}(\theta, 1)$ with $\theta = 0$. Use this same dataset in all questions below.

1. SIR. Under prior (A), use sampling importance resampling (SIR) to approximate the posterior $p(\theta \mid \mathbf{x})$. Report the posterior mean, the posterior standard deviation, and $\Pr(\theta > 1 \mid \mathbf{x})$. Justify your choice of proposal density and report the effective sample size.

2. Gibbs sampler. Under prior (B), derive the full conditional distributions and implement a Gibbs sampler for $(\theta, \tau^2, \lambda)$. Report the same summaries as in Question 1, together with trace plots and autocorrelation functions for θ. Compare the results with those from Question 1.

3. Extra. Prove that priors (A) and (B) are exactly the same, i.e., that integrating τ^2 and λ out of (B) yields the density in (A).

Hint: first show that $\theta \mid \lambda$ follows a Laplace distribution with rate λ.