Trying different Metropolis-Hastings samplers for gamma parameters

Generate $n = 30$ observations y_i from the gamma distribution with, say, parameters $\alpha^* = 13$ and $\beta^* = 1$. This is your data.

Now try to estimate the parameters α and β using the data $y = (y_1, \ldots, y_n)$. We use the true likelihood

$$[Y_i \mid \alpha, \beta] \overset{\text{i.i.d.}}{\sim} \text{Gam}(\alpha, \beta), \qquad i = 1, \ldots, n,$$

and the flat prior

$$\alpha \sim \text{Exp}(\lambda), \qquad \beta \sim \text{Exp}(\lambda),$$

where $\lambda = 1/1000$.

- Do you recognize the full conditional of α or β in the posterior?
- First try M–H sampling using the parameters

  $$\phi = \log \alpha, \qquad \psi = \log \beta$$

  You can, e.g., use independent random walk proposals from normal distributions. Select the standard deviations so that the acceptance rate becomes reasonable (10–40 %). Produce autocorrealtion plots for the parameters.
- Next try the reparametrization

  $$\phi = \log \mu, \qquad \psi = \log \sigma,$$

  where μ and σ are the mean and standard error of the $\text{Gam}(\alpha, \beta)$ distribution. Again, write a Metropolis–Hastings algorithm, where the proposal consists of independent normal random walks for ϕ and ψ Select the standard deviations so that the acceptance rate becomes reasonable. Produce autocorrelation plots for the parameters. (If you come up with a better reparametrization, use that instead.)
- Last, implement a Metropolis–Hastings algorithm, where you take advantage of conditional conjugacy. E.g., you can update one of the parameters with a random walk proposal on the logarithmic scale, and the other one from its full conditional distribution. Then you accept or reject the proposed pair. Alternatively, you can implement an independence sampler, where one of the parameters is drawn from a step function approximation to its marginal posterior, and the other one from its full conditional. Again, tune the sampler, and produce autocorrelation plots.