Three different Metropolis-Hastings samplers for gamma parameters

Generate $n = 50$ observations y_i from the gamma distribution with, say, parameters $\alpha^* = 2.1$ and $\beta^* = 11$. This is your data.

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

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

and the flat prior

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

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 (α, β) which are constrained to be positive. You can, e.g., use independent random walk proposals from normal distributions. Reject immediately those proposed values which are not both positive. Try to select the standard deviations of the proposal distributions so that the acceptance rate becomes reasonable (10–40 %). Produce autocorrelation plots for the parameters.

- Next try the reparametrization

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

  where $\mu = \alpha/\beta$ is the mean of the $\mathrm{Gam}(\alpha, \beta)$ distribution. These parameters do not have any restrictions. 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.

- Last, implement a Metropolis–Hastings algorithm, where you take advantage of conditional conjugacy. Update one of the parameters with a random walk proposal on the logarithmic scale, and the second by drawing a value from its full conditional distribution, conditionally on the value of the first. Then you accept or reject the proposed pair. Again, tune the sampler and produce autocorrelation plots.