

Stat 544 – Spring 2005 Homework assignment 3

Due on Friday, February 25 in TA's office by 5:00 pm

Problem 1

You have not been able to clean your computer from the virus that affected it (recall that you cannot produce samples other than uniform random variables). This time we want to sample from a Beta distribution, but instead of using the method that you used to do homework 1, you should use rejection sampling to obtain 1000 draws from a Beta(2,2) distribution. In addition to your R program you should provide a terse description on how your program works.

Problem 2

Use the rejection method to obtain a sample of size 1000 from the Gamma($\alpha, 1$) distribution for some $\alpha > 1$ of your choice. As your instrumental density you should use an exponential distribution with mean equal to α .

- Specify the constant M.
- For your chosen α plot the target density and the envelope in one display.
- Make a histogram using your draws.
- Make a plot of the trial ratio, i.e. the average number of points needed to produce one acceptable value from your target distribution.

Problem 3

Recall the bioassay example on page 88 of the textbook. We will draw new samples from (α, β) and from $LD50 = \alpha/\beta$.

- Obtain a sample of size 2000 of (α, β) using a normal approximation to their joint posterior density. Draw the following figures:
 - Scatterplot of your sample.
 - Histogram of your draws from the approximate posterior density of LD50.
- Using your sample from part (a) obtain a SIR (sampling importance resampling) sample of size 500 of (α, β) . Draw the following figures:
 - Histogram of importance weights
 - Scatterplot of your sample.

Problem 4

What used to be called the National Bureau of Standards in Washington, DC, conducts extremely high precision measurement of physical constants, such as the actual weight of so-called *check-weights* that are supposed to serve as reference standards (like the official kg). The file nb10.dat contains 100 weighings of a block of metal called *NB10*, which was supposed to weigh exactly 10g (the units are micrograms below 10g). Those weighings were made under conditions as close to IID as possible.

Here we wish to estimate the variance of the NB10 check-weight measurements. Assume that the mean $\mu = 404.59$ is known and

$$Y_i | \sigma^2 \sim N(\mu, \sigma^2) \quad i = 1, \dots, n$$

Further, we assume a conjugate prior for σ^2 :

$$\sigma^2 \sim \text{Inv-}\chi^2(\nu_p, \sigma_p^2)$$

with $\nu_p = 10$ and $\sigma_p^2 = 40$.

- Construct a 95% credible set for σ^2 . To do this you should obtain a sample of size 1000 from the posterior distribution of σ^2 using the random-walk Metropolis-Hastings algorithm.
- Use WinBUGS to obtain 1000 draws from the posterior distribution of σ^2 . Report the 95% credible set for σ^2 .
- Compare both credible sets.

Problem 5

Observations y_1, y_2, \dots, y_n are independently distributed as Poisson random variables given parameters $\theta_1, \theta_2, \dots, \theta_n$ so that

$$p(y_i | \theta_i) = (y_i!)^{-1} \theta_i^{y_i} \exp(-\theta_i).$$

The prior distribution for the vector $(\theta_1, \theta_2, \dots, \theta_n)$ is constructed hierarchically (or in steps). First, given a hyperparameter ϕ , the θ_i are assumed to be identically and independently distributed a priori as exponential random variables, so that

$$p(\theta_i | \phi) = \phi \exp(-\phi \theta_i), \quad \theta_i > 0.$$

Second, the hyperparameter ϕ gets a noninformative uniform prior distribution, so that

$$p(\phi) \propto 1, \quad \phi \geq 0.$$

- Write down the joint posterior distribution $p(\theta, \phi | y)$.
- Provided that $n\bar{y} > 1$, prove that the posterior distribution of the random variable $z = (1 + \phi)^{-1}$ has a beta form, so that
$$p(z | y) \propto z^{n\bar{y}-2} (1-z)^n.$$
- Show that the posterior means of the θ_i , $E(\theta_i | y)$, are shrunk by a factor equal to $(n\bar{y} - 1)/(n\bar{y} + n)$ relative to the conditional means $E(\theta_i | y, z)$.
- What happens if $n\bar{y} \leq 1$?