

## MATLAB Computer Session 1: Basics of Bayesian Computation

Exercises 1 and 4 are taken from the book: Gary Koop, Dale Poirier and Justin Tobias (2007), *Bayesian Econometric Methods*, Cambridge University Press. In particular, these are exercises 11.2 and 11.7, from this book. There is MATLAB code associated with *Bayesian Econometric Methods*, available at: <http://web.ics.purdue.edu/~jltobias/bem.html>. However, you do not need to download this code. I provide the code for all the exercises in the website associated with this course: [http://personal.strath.ac.uk/gary.koop/BoK\\_course.html](http://personal.strath.ac.uk/gary.koop/BoK_course.html)

### MATLAB Exercises:

#### 1. *Drawing from Standard Distributions.:*

Simulation-based inference using algorithms such as the Gibbs sampler requires the researcher to be able to draw from standard distributions. In this exercise we discuss how MATLAB can be used to obtain draws from a variety of standard continuous distributions. Specifically, we obtain draws from the Uniform, Normal, Student-t, Beta, Exponential and Chi-squared distributions (see the Appendix of Bayesian Econometric Methods for definitions of these distributions). Using the Matlab program for this exercise (Ex1.m), obtain sets of 10, 100 and 100,000 draws from the Uniform, standard Normal, Student-t(3) (denoted  $t(0, 1, 3)$  in the notation of the Appendix to the book), Beta(3,2), Exponential with mean and  $\chi^2(3)$  distributions. For each sample size calculate the mean and standard deviation and compare these quantities to the known means and standard deviations from each distribution.

#### 2. *Monte Carlo integration:*

If the posterior density  $p(\theta|y)$  takes the a familiar form (e.g. a Normal or Student-t or Gamma or other distribution for which computer algorithms exist to take random draws) then we can obtain  $R$  iid draws of the parameters, which we denote  $\theta^{(r)}$ ,  $r = 1, \dots, R$ . Usually, quantities of interest to the researcher are functions of the model parameters. Let us call such a function  $g(\theta)$ . The researcher would then often be interested in calculating:

$$E(g(\theta) | y) = \int g(\theta) p(\theta|y) d\theta$$

Monte Carlo integration allows us to calculate integrals of this form. The weak law of large numbers implies that

$$E(g(\theta) | y) \simeq \frac{\sum_{r=1}^R g(\theta^{(r)})}{R}$$

This means that the posterior mean of  $g(\theta)$  can be calculated by drawing from the posterior and then averaging functions of the posterior draws. Exercise: Suppose  $p(\theta|y) \sim N(1, 4)$  and the quantity of interest is  $g(\theta) = \theta^2$ . Use Monte Carlo integration to calculate  $E(\theta^2)$ . Code for this question is in Ex2.m. Note: in this case, you know the correct answer is  $E(\theta^2) = 5$  (since the definition of variance tells you that  $var(\theta) = E(\theta^2) - [E(\theta)]^2$  and, in this exercise,  $var(\theta) = 4$  and  $E(\theta) = 1$ ), so you would not need to have done Monte Carlo integration. Optional exercise: modify Ex2.m to calculate the posterior mean for a more complicated quantity of interest for which analytical results are not so easily available (e.g. calculate  $\Pr(\theta^2 > 2)$  or  $E(\ln(\theta))$  or some other choice for  $g(\theta)$ ).

#### 3. *Analysis of the Normal model using Monte Carlo integration:*

Assume that we observe  $y = (y_1, \dots, y_T)$  which come from a Normal density with unknown mean  $\theta$  but with known variance  $\sigma^2$ . The likelihood function of this model has only one unknown parameter,  $\theta$ . We have to specify a prior on  $\theta$  and we assume that a normal density is a reasonable choice. Consequently

we define  $\theta \sim N(\mu, \tau^2)$  where  $\mu$  is the prior mean and  $\tau^2$  is the prior variance. Analytical results are available for this example. The Likelihood and prior are:

$$p(y|\theta) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(\frac{-\sum (y_t - \theta)^2}{2\sigma^2}\right)$$

$$p(\theta) = \frac{1}{\sqrt{2\pi\tau^2}} \exp\left(\frac{-(\theta - \mu)^2}{\tau^2}\right)$$

so that the posterior of the parameters is (see Koop et al. (2007, Exercise 2.3)):

$$\theta|y \sim N(\xi(\sigma^2\mu + \tau^2\bar{y}), \xi\sigma^2\tau^2) \tag{1}$$

$$\xi = (\sigma^2 + \tau^2)^{-1}$$

where  $\bar{y}$  is the sample mean. Accordingly, to obtain draws from the posterior density, we have to draw from normal density with the parameters specified in (1), i.e. the posterior mean  $\xi(\sigma^2\mu + \tau^2\bar{y})$  and the posterior variance  $\xi\sigma^2\tau^2$ . The MATLAB code Ex2.m provided, generates  $y$  from a known Normal distribution of your choice, and takes samples from the posterior using the Monte Carlo integration. You are asked to explore 3 things:

- (a) Change the values of the prior hyperparameters  $\mu, \tau^2$  and see what happens to the posterior mean. Begin with uninformative values ( $\mu = 0, \tau^2 = 100,000$ ). Then provide a prior belief for  $\mu$  which is far away from the true mean we just used to generate  $y$ . What happens? Try to set a very tight prior variance  $\tau^2 = 0.01$ . What happens in this case?
- (b) Keep the prior hyperparameters constant and change the number of samples  $T$  of the generated variable  $y$ . Report what changes for  $T = 10, 100, 1000, 10000$ . Compare with the MLE estimate of the parameter  $\theta$ .
- (c) Compare the histograms of the prior and resulting posterior, for different choices of  $T$ , and the prior hyperparameters  $\mu, \tau^2$ .

#### 4. Gibbs Sampling from the Bivariate Normal:

The purpose of this question is to learn about the properties of the Gibbs sampler in a very simple case. Assume that you have a model which yields a bivariate Normal posterior,

$$\begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix} \sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}\right),$$

where  $|\rho| < 1$  is the (known) posterior correlation between  $\theta_1$  and  $\theta_2$ .

- (a) Write a program which uses Monte Carlo integration to calculate the posterior means and standard deviations of  $\theta_1$  and  $\theta_2$ .
- (b) Write a program which uses Gibbs sampling to calculate the posterior means and standard deviations of  $\theta_1$  and  $\theta_2$ .
- (c) Set  $\rho = 0$  and compare the programs from parts a) and b) for a given number of replications (e.g.  $R = 100$ ) and compare the accuracy of the two algorithms.
- (d) Repeat part (c) of this question for  $\rho = .5, .9, .99$  and  $.999$ . Discuss how the degree of correlation between  $\theta_1$  and  $\theta_2$  affects the performance of the Gibbs sampler. Make graphs of the Monte Carlo and Gibbs sampler replications of  $\theta_1$  (i.e. make a graph with x-axis being replication number and y-axis being  $\theta_1$ ). What can the graphs you have made tell you about the properties of Monte Carlo and Gibbs sampling algorithms.
- (d) Repeat parts (c) and (d) using more replications (e.g.  $R = 50,000$ ) and discuss how Gibbs sampling accuracy improves with number of replications.