UNIVERSITY OF OSLO

Faculty of Mathematics and Natural Sciences

Examination in: STK4021/9021 — Applied Bayesian Analysis and Numerical Methods

Day of examination: Tuesday, December 16th, 2014

Examination hours: 09.00 – 13.00

This problem set consists of 2 pages.

Appendices: List of distributions

Permitted aids: One A4 sheet with notes, approved calculator

Please make sure that your copy of the problem set is complete before you attempt to answer anything.

Problem 1

Suppose y has an exponential sampling model with a gamma prior distribution. That is,

$$Y|\theta \sim \text{Exp}(\theta),$$
$$\theta|\alpha, \beta \sim \Gamma(\alpha, \beta).$$

(a) Show that the gamma prior distribution is conjugate. Show that the equivalent prior specification for the mean, $\varphi = 1/\theta$, is inverse-gamma. That is, derive the latter density function.

(b) The length of life of a light bulb manufactured by a certain process follows the model above where the prior parameters are chosen such that the coefficient of variation is 0.5. (The coefficient of variation is defined as the standard deviation divided by the mean). A random sample of light bulbs is to be tested and the lifetime of each obtained. If the coefficient of variation of the posterior distribution of θ is to be reduced to 0.1, how many light bulbs need to be tested? If the coefficient of variation refers to φ instead of θ, how would your answer change?

(c) Suppose we observe that $y \geq 100$, but do not observe the exact value of y. What is the posterior distribution $p(\theta|y \geq 100)$ as a function of α and β? Write down the posterior mean and variance of θ. Repeat this calculation for the case when we observe $y = 100$. Explain why the posterior variance of θ is higher for $p(\theta|y = 100)$ even though more information has been observed.

(Continued on page 2.)

Problem 2

Suppose that $y_1, \ldots, y_n$ are independent samples from a Cauchy distribution with unknown location θ and known scale 1,

$$Y_i|\theta \sim \text{Cauchy}(\theta, 1), \qquad i = 1, \ldots, n.$$

(a) Assume that the prior distribution for θ is uniform on $[0, 1]$, $\theta \sim \mathcal{U}([0, 1])$, and that we observe $(y_1, y_2) = (0, 1)$. Determine the derivative and the second derivative of the log posterior density. Show that the posterior mode of θ is equal to 0.5. Construct the normal approximation for the posterior distribution based on the posterior mode and the second derivative of the log posterior density at the mode.

(b) Assume again that $n = 2$ and use the improper prior $p(\theta) \propto 1$. Show that the posterior distribution is proper. Under what conditions will the posterior density be unimodal?

Problem 3

Suppose y has a binomial sampling model with a beta prior distribution for the unknown success probability. That is,

$$\begin{aligned} Y|\theta &\sim \text{Binomial}(n, \theta), \\ \theta|\alpha, \beta &\sim \text{Beta}(\alpha, \beta), \end{aligned}$$

where the number of trials n is given.

(a) Find $p(y|\alpha, \beta)$, the marginal distribution for y, for $y = 0, \ldots, n$ (unconditional on θ). This discrete distribution is known as the beta-binomial distribution. Show that if the beta-binomial probability is constant as a function of y, that is, if every y has the same probability, then the prior distribution has to have $\alpha = \beta = 1$.

(b) Find the maximum likelihood estimate (MLE), the posterior distribution, and the Bayes rule under the squared error loss.

(c) Calculate the risk for both the MLE and the Bayes rule. For $\frac{\alpha}{\alpha+\beta} = 0.5$, find the region where the Bayes rule has smaller risk than the MLE.

THE END