University of Regina
Department of Mathematics & Statistics
Final Examination (PART I)
200810

Statistics 352
Advanced Mathematical Statistics

Instructor: Michael Kozdron

This final exam consists of two parts.

Part I of this exam is due at 9:00 am on Tuesday, April 15, 2008, in Education Building 311 (ED 311).

Part II of this exam will begin at 9:00 am on Tuesday, April 15, 2008, in Education Building 311 (ED 311).

Read all of the following information before starting Part I of the exam.

Please read all instructions carefully, and check your answers. Show all work neatly and in order, and clearly indicate your final answers. Answers must be justified whenever possible in order to earn full credit.

You may use standard notation; however, any new notations or abbreviations that you introduce must be clearly defined.

You may consult any notes or books or websites that you wish, provided that proper acknowledgements are included. However, you may not discuss Part I of this exam with anyone before 9:00 am on April 15, 2008. This includes students, professors, and colleagues.

*Part I of this exam has **2** problems and is worth a total of **42** points.*

*Part II of this exam has **6** problems and is worth a total of **78** points.*

1. (*24 points*) Suppose that Y_1, Y_2, Y_3, Y_4 are a random sample with $Y_i \sim \mathcal{N}(1, \theta)$ for each $i = 1, 2, 3, 4$ where $\theta > 0$ is an unknown parameter. (Note that the variance of this normal distribution is θ.) Suppose further that your prior beliefs about θ are given by

$$g(\theta) = \theta^{-2} e^{-1/\theta}, \quad \theta > 0.$$

(a) If the four observations produce data

$$y = \{1.08, 1.13, 1.25, 1.34\},$$

determine an expression for the posterior density function

$$f(\theta \,|\, y = \{1.08, 1.13, 1.25, 1.34\})$$

by numerically approximating the normalizing constant.

(b) Construct an equal-tailed 90% Bayesian credible interval (i.e., a Bayesian confidence interval) for θ based on the data $y = \{1.08, 1.13, 1.25, 1.34\}$.

(c) Use the envelope method in R to simulate several hundred random variables having the posterior density $f(\theta \,|\, y = \{1.08, 1.13, 1.25, 1.34\})$ you found in **(a)**. Use the prior as your envelope.

Hint: If a random variable X has density function

$$g(x) = x^{-2} e^{-1/x}, \quad x > 0,$$

then $1/X$ has a known distribution.

2. (*18 points*) The goal of this problem is to have you simulate random variables (X, Y) from a multivariate normal distribution. Specifically, suppose that (X, Y) is multivariate normal with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Lambda}$ where

$$\boldsymbol{\mu} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \quad \text{and} \quad \boldsymbol{\Lambda} = \begin{bmatrix} 4 & 5 \\ 5 & 9 \end{bmatrix}.$$

Recall that in this notation $\text{Var}(X) = 4$, $\text{Var}(Y) = 9$, and $\text{Cov}(X, Y) = 5$.

(a) Using facts from STAT 351, identify the marginal distribution for X.

(b) Using facts from STAT 351, identify the conditional distribution for $Y \,|\, X = x$.

(c) Using facts from STAT 351, identify the conditional distribution for $X \,|\, Y = y$.

(d) Use Gibbs sampling in R to simulate 10000 random variables from this $\text{MVN}(\boldsymbol{\mu}, \boldsymbol{\Lambda})$ distribution. Include your R code as well as a plot (not a histogram) of your pairs (X, Y). (Use the R command `plot` for this.) Note that since the marginal distribution for X is known, there is no need to burn-in the Gibbs sampler. In particular, every pair of points that you generate will have the required distribution.

University of Regina
Department of Mathematics & Statistics
Final Examination (PART II)
200810
(April 15, 2008)

Statistics 352
Advanced Mathematical Statistics

Name: ______________________________ **Student Number:** ______________

Instructor: Michael Kozdron **Time:** 3 hours

This final exam consists of two parts.

Part I of this exam is due at 9:00 am on Tuesday, April 15, 2008, in Education Building 311 (ED 311).

Part II of this exam will begin at 9:00 am on Tuesday, April 15, 2008, in Education Building 311 (ED 311).

Read all of the following information before starting Part II of the exam.

*You have **3** hours to complete Part II of this exam. Please read all instructions carefully, and check your answers. Show all work neatly and in order, and clearly indicate your final answers. Answers must be justified whenever possible in order to earn full credit.* ***Unless otherwise specified, no credit will be given for unsupported answers, even if your final answer is correct.*** *Points will be deducted for incoherent, incorrect, and/or irrelevant statements.*

You may use standard notation; however, any new notations or abbreviations that you introduce must be clearly defined. Calculators are permitted; however, you must still show all your work. You are also permitted to have ONE 8.5×11 *page of handwritten notes (double-sided) for your personal use. Other than these exceptions, no other aids are allowed.*

The order of the test questions is essentially random; they are not intentionally written easiest-to-hardest.

*Part I of this exam has **2** problems and is worth a total of **42** points.*
*Part II of this exam has **6** problems and is worth a total of **78** points.*

DO NOT WRITE BELOW THIS LINE

Problem 3	______	Problem 4	______	Problem 5	______
Problem 6	______	Problem 7	______	Problem 8	______
		Part I Total	______	Part II Total	______
				TOTAL	______

3. (*10 points*) Suppose that Y is a binomial random variable with parameters $n = 4$ and θ where $0 < \theta < 1$ is unknown. Assume that θ can take on only two values and that the prior beliefs about θ are given by

$$g(0.3) = \frac{2}{3} \quad \text{and} \quad g(0.4) = \frac{1}{3}.$$

Compute the posterior density for θ if the data $y = 2$ is observed.

4. (*10 points*) At a certain factory, machines I and II are all producing springs of the same length. Machine I and II produce 1% and 4% defective springs, respectively. Of the total production of springs in the factory, machine I prduces 40% and machine II produces 60%. If one spring is selected at random and it is found to be defective, determine the probability that it produced by machine II.

5. (*14 points*) Suppose that a random variable X has a $\Gamma(a, 1/b)$ distribution with $a > 0$ and $b > 0$ so that the density function of X is

$$f_X(x) = \frac{b^a}{\Gamma(a)} x^{a-1} e^{-bx}, \quad x > 0.$$

(a) Determine the density function of the random variable $1/X$. A random variable with this density function is said to have an *inverse Gamma distribution with parameters a and b.*

Suppose that Y is a $\mathcal{N}(0, \theta)$ random variable where $\theta > 0$ an unknown parameter. Note that θ is the variance of this normal distribution and so the likelihood function is

$$f(y|\theta) = \frac{1}{\sqrt{2\pi\theta}} \exp\left\{-\frac{y^2}{2\theta}\right\}, \quad -\infty < y < \infty.$$

(b) If the prior distribution for θ is an inverse Gamma distribution with parameters a and b, determine an exact, closed-form expression for the posterior density $f(\theta|y)$ and use this to prove that the inverse Gamma distribution is the natural conjugate prior for this likelihood function.

6. (*14 points*) Consider a random variable Y having the Weibull distribution with unknown parameter $\theta > 0$ so that the likelihood function is

$$f(y|\theta) = \frac{3y^2}{\theta^3} \exp\left\{-\frac{y^3}{\theta^3}\right\}, \quad y > 0.$$

(a) Compute $I(\theta)$, the Fisher information for this Weibull density function.

(b) Determine an expression (up to a constant of proportionality) for the density function of the Jeffreys prior for θ. *Note that the Jeffreys prior for θ is improper.*

(c) Even though the Jeffreys prior is improper, it can still be used to produce a proper posterior density. Write down an exact, closed-form expression for the posterior density $f(\theta|y)$ based on a Weibull(θ) likelihood and the Jeffreys prior.

7. (*16 points*) Engineers are testing three types of filaments for lightbulbs. To test the performance of the filaments, for each type of filament, they plug in 50 bulbs of the same wattage and record how long each one stays lit before burning out. The data values are identified as $Y_{i,j}$, $i = 1, 2, 3$ and $j = 1, 2, \ldots, 50$, for the burning time of the jth bulb of filament type i. (It is assumed that the $Y_{i,j}$ are independent.)

The engineers will fit a *hierarchical model* to estimate the average burning time of each type of filament. They assume that the burning times for the ith type of filament are a random sample from an $\text{Exp}(1/\lambda_i)$ distribution with $\lambda_i > 0$.

They further assume that each λ_i is drawn from a $\Gamma(\alpha, 1/\beta)$ distribution with $\alpha > 0$ and $\beta > 0$. Finally, they put $\Gamma(0.5, 2)$ priors on both α and β.

Hence, this hierarchical model may be summarized as follows:

- $Y_{i,j} \sim \text{Exp}(1/\lambda_i)$, $i = 1, 2, 3$; $j = 1, 2, \ldots, 50$, so that

$$f(y_{i,j} \,|\, \lambda_i) = \lambda_i e^{-\lambda_i y_{i,j}}, \quad y_{i,j} > 0,$$

- $\lambda_i \sim \Gamma(\alpha, 1/\beta)$, $i = 1, 2, 3$, so that

$$f(\lambda_i \,|\, \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \lambda_i^{\alpha-1} e^{-\beta\lambda_i}, \quad \lambda_i > 0,$$

- $\alpha \sim \Gamma(0.5, 2)$ so that

$$g(\alpha) = \frac{1}{\sqrt{2\pi\alpha}} e^{-\alpha/2}, \quad \alpha > 0.$$

- $\beta \sim \Gamma(0.5, 2)$ so that

$$g(\beta) = \frac{1}{\sqrt{2\pi\beta}} e^{-\beta/2}, \quad \beta > 0.$$

(a) Write down an expression (up to a constant of proportionality) for the likelihood function

$$f(y_{1,1}, \ldots, y_{3,50} \,|\, \lambda_1, \lambda_2, \lambda_3).$$

(b) Write down an expression (up to a constant of proportionality) for the (joint) prior density function

$$g(\lambda_1, \lambda_2, \lambda_3, \alpha, \beta).$$

(c) Write down an expression (up to a constant of proportionality) for the (joint) posterior density

$$f(\lambda_1, \lambda_2, \lambda_3, \alpha, \beta \,|\, y_{1,1}, \ldots, y_{3,50}).$$

(d) Derive an expression (up to a constant of proportionality) for

$$f(\alpha \,|\, \lambda_1, \lambda_2, \lambda_3, \beta, y_{1,1}, \ldots, y_{3,50}),$$

the (full) posterior conditional density of α. If this distribution belongs to a standard parametric family, identify that family.

8. (*14 points*) Suppose that Y has a $\mathcal{N}(\theta, 1)$ distribution where $-\infty < \theta < \infty$ is an unknown parameter. You think that your prior beliefs about θ can be described by a normal distribution with mean 2 and variance 1 so that

$$g(\theta) = \frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{(\theta-2)^2}{2}\right\}, \quad -\infty < \theta < \infty.$$

Assume that the single data value $y = 1$ was observed.

(a) Determine an exact, closed-form expression for $f(\theta|y = 1)$, the posterior density of θ given $y = 1$.

(b) Based on your answer to **(a)**, determine a 95% Bayesian credible interval for θ. *Recall that if* $Z \sim \mathcal{N}(0, 1)$, *then* $P\{-1.96 \le Z \le 1.96\} = 0.95$.