1. Suppose that $Y_1, Y_2, ..., Y_n$ is an iid sample from

$$f_Y(y;\theta) = \begin{cases} \theta e^{-\theta y}, & y > 0 \\ 0, & \text{otherwise.} \end{cases}$$

In addition, suppose that $\theta$ varies according to a gamma$(a, b)$ prior distribution; i.e.,

$$g(\theta) = \begin{cases} \dfrac{1}{\Gamma(a)b^a}\theta^{a-1}e^{-\theta/b}, & \theta > 0 \\ 0, & \text{otherwise,} \end{cases}$$

where $a$ and $b$ are both known constants larger than zero.
(a) Show that the posterior distribution $g(\theta|\boldsymbol{y})$ is gamma with shape parameter $n + a$ and scale parameter $b/(1 + bu)$, where $u = \sum_{i=1}^n y_i$.
(b) Suppose that $n = 5$ and $\boldsymbol{y} = (0.13, 0.21, 0.33, 0.05, 0.62)'$. Find the posterior mean, median, and mode if $a = 2$ and $b = 2$.

2. Suppose that $Y_1, Y_2, ..., Y_n$ is an iid sample from a $\mathcal{U}(0, \theta)$ distribution, where $\theta > 0$. In turn, the parameter $\theta$ is best regarded as a random variable with a Pareto$(\alpha, \beta)$ distribution, that is,

$$g(\theta) = \begin{cases} \dfrac{\beta\alpha^\beta}{\theta^{\beta+1}}, & \theta > \alpha \\ 0, & \text{otherwise,} \end{cases}$$

where $\alpha > 0$ and $\beta > 0$ are known.
(a) Under the $\mathcal{U}(0, \theta)$ model, we proved in STAT 512 that $Y_{(n)}$ is a sufficient statistic for $\theta$. Find the distribution of $Y_{(n)}$.
(b) Find the posterior mean $\widehat{\theta}_B = E(\theta|\boldsymbol{Y} = \boldsymbol{y})$.
(c) Can $\widehat{\theta}_B$ be written as a linear combination of the prior mean and the MLE $\widehat{\theta} = Y_{(n)}$? If so, prove it. If not, show that this can not be done.

3. Coach Spurrier has asked for your help in assessing whether or not his team is being penalized too often. Through six games this season, here are the penalty counts (the total number of penalties per game):

| Game | Opponent | Number of Penalties |
|---|---|---|
| 1 | East Carolina | 8 |
| 2 | Georgia | 5 |
| 3 | Navy | 5 |
| 4 | Vanderbilt | 9 |
| 5 | Auburn | 7 |
| 6 | Kentucky | 7 |

For purposes of modeling, suppose that the number of USC penalties per game has a Poisson distribution with mean $\theta$. My inside sources tell me that Coach Spurrier is a

hardcore Bayesian, so we better do a Bayesian analysis.
(a) Pick a gamma($\alpha, \beta$) prior that is reasonable for this problem. To do this, a sensible course of action would be to do the following. Find the sample mean and sample variance of the number of penalties in 12 USC's regular season games last year; see, e.g.,

`http://rivals.yahoo.com/ncaa/football/teams/ssi?y=2010`

Then, you can set

$$\begin{aligned}\text{sample mean} &= \alpha\beta \\ \text{sample variance} &= \alpha\beta^2\end{aligned}$$

and solve for $\alpha$ and $\beta$ (this has a method of moments-type feel to it). If you wanted to choose your prior distribution in another way, that is fine, but your prior selection better be defendable (we don't want Coach Spurrier to be unhappy).
(b) Find the posterior distribution and the posterior mean.
(c) Compute a 90 percent credible interval for $\theta$. Interpret the interval.
(d) Is it likely that $\theta$ is larger than 7 penalties per game? To answer this question, perform a Bayesian hypothesis test of

$$H_0 : \theta > 7$$
versus
$$H_a : \theta \leq 7.$$

4. Insurance payments data in actuarial industries are typically highly positively skewed and distributed with a larger upper tail. A reasonable model for this type of data is the Weibull distribution. Specifically, suppose that $Y_1, Y_2, ..., Y_n$ denote $n$ payments, modeled as iid random variables with common Weibull pdf

$$f_Y(y;\theta) = \begin{cases} \dfrac{m}{\theta} y^{m-1} e^{-y^m/\theta}, & y > 0 \\ 0, & \text{otherwise} \end{cases}$$

where $m > 0$ is known and $\theta$ is unknown. In turn, suppose that $\theta \sim \text{IG}(\alpha, \beta)$, that is, $\theta$ has (prior) pdf

$$g(\theta) = \begin{cases} \dfrac{1}{\Gamma(\alpha)\beta^\alpha} \dfrac{1}{\theta^{\alpha+1}} e^{-1/\beta\theta}, & \theta > 0 \\ 0, & \text{otherwise.} \end{cases}$$

(a) Prove that the inverted gamma IG($\alpha, \beta$) prior is a conjugate prior.
(b) Suppose that $m = 2$, $\alpha = 0.5$, and $\beta = 2$. Here are $n = 10$ insurance payments (measured in \$10,000s):

`0.2697 0.0719 0.4698 0.8192 3.9700 0.2681 0.2415 2.8351 0.0885 0.1180`

Compute a 95 percent credible interval for $\theta$. Interpret the interval.