Sta 114 = Mth 136 : Homework 2

When Morris DeGroot wrote the first edition of our text, different authors used a variety of different (mostly Greek) letters for the commonly occuring objects in statistical analysis; more recently, most authors have converged on the same combinations which differ somewhat from those in our text. In the homework exercises (and lectures) I will most often use the currently fashionable letters, including:

Sample Space	Ω	$= \{\omega\}$
Parameter Space	Θ	$= \{\theta\}$
Outcome Space	$\mathcal{X}$	$= \{\mathbf{x}\}$
Prior pdf/pf	$\pi(\theta)$	
Likelihood	$f_n(\mathbf{x} \mid \theta)$	$= \prod_{j=1}^n f(x_j \mid \theta)$
or	$\ell(\theta)$	$= c f_n(\mathbf{x} \mid \theta)$, for any $c > 0$
Marginal pdf/pf	$m(\mathbf{x})$	$= \int f(\mathbf{x} \mid \theta)\, \pi(\theta)\, d\theta$
Posterior pdf/pf	$\pi(\theta \mid \mathbf{x})$	$= f(\mathbf{x} \mid \theta)\, \pi(\theta)/m(\mathbf{x})$

Also, we will parametrize commonly-occuring probability distributions as on the class "PDF sheet" (see course syllabus, near 1st Midterm, or click here).

1. Let the lifetimes $\{X_j\}$ of several CFC lightbulbs be iid with the $\mathsf{Ex}(\theta)$ distribution, and suppose we observe four lifetimes (in hours)

$$\mathbf{x} = \{\ 2083,\ 4500,\ 5185,\ 10545\ \}$$

From past experience the experimentor believes that the prior distribution for θ is $\mathsf{Ga}(2, 8000)$, so

$$\pi(\theta) = \frac{\lambda^\alpha}{\Gamma(\alpha)} \theta^{\alpha-1}\, e^{-\lambda\theta}, \qquad \theta > 0$$

with $\alpha = 2$ and $\lambda = 8000$. Evaluate

$$\Pr[X_5 > 3000 \mid \mathbf{x}].$$

2. The proportion θ of defective items in a large manufacturing lot is known to be either 0.10 or 0.25, with prior pf

$$\pi(0.10) = 0.80 \qquad \pi(0.25) = 0.20$$

In a random sample of eight items from the lot, exactly one was defective. Determine the posterior pf $\pi(\theta \mid \mathbf{x})$ of θ.

3. The parameter θ has a Gamma prior distribution $\theta \sim \mathsf{Ga}(\alpha, \lambda)$ with mean $\mathsf{E}[\theta] = 12$ and variance $\mathsf{V}[\theta] = 48$. Find the parameters α, λ and give the pdf $\pi(\theta)$ correctly for every $\theta \in \mathbb{R}$ (simplify— no Greek letters in your answer! Okay, maybe θ as a dummy variable...)

4. The proportion θ of defective items in a large manufacturing lot is unknown, with uniform prior distribution on the unit interval $[0, 1]$. When a random sample of ten items it taken, exactly one was found to be defective. Determine the posterior distribution of θ— give both its pdf $\pi(\theta \mid \mathbf{x})$ and its name, with the value(s) of any parameter(s).

5. Consider the same manufacturing situation as in Problem 4, with the same prior distribution, but a different sampling scheme— now we continue sampling sequentially until we discover the first defective item. If we sample exactly nine working items before discovering the first defective one, what is the posterior distribution for θ?

6. Let $\{X_j\} \overset{\text{iid}}{\sim} \mathsf{Un}(0, \theta)$ be independent uniformly-distributed random variables on the interval $[0, \theta]$, for unknown $\theta > 0$ with Pareto $\mathsf{Pa}(1, 1)$ prior distribution with pdf

$$\pi(\theta) = \begin{cases} \theta^{-2} & \theta > 1 \\ 0 & \theta \leq 1. \end{cases}$$

   If we observe a sample of size five to be:

$$\mathbf{x} = \{\ 6.9,\ 1.5,\ 5.6,\ 6.5,\ 3.4\ \}$$

   find the posterior pdf for θ.

7. The proportion θ of defective items in a large manufacturing lot is unknown, with a beta $\mathsf{Be}(\alpha = 1, \beta = 99)$ prior distribution.

   (a) About what fraction of items are defective, on average?

   (b) If a random sample of 100 items reveals two defective ones, what is the posterior distribution for θ?

   (c) After observing the random sample of (7b) above, what fraction of *this lot* do you expect are defective?

8. Suppose that a random sample of 100 observations is to be taken from a normal $\mathsf{No}(\theta, \sigma^2)$ distribution with unknown mean θ and standard

deviation $\sigma = 10$, and that $\theta \sim \mathsf{No}(\mu, \tau^2)$ has a normal prior distribution. Show that, no matter how large the prior standard deviation τ might be, the standard deviation for the *posterior* distribution is less than one.