Economics 520, Fall 2008
Homework 9
Due Tuesday, November 25 at beginning of class

1. R exercise: Let (X_i, Y_i) be the age and unemployment duration for person i, for $i = 1, \ldots, N$. Assume that age has a normal distribution with mean 40 and standard deviation 10. Given age X_i, Y_i, the duration of unemployment in weeks is assumed to have an exponential distribution with mean $\exp(\theta \cdot X_i)$. Assume that the observations for different individuals are independent. The observations are (25,10), (43,29), (40,25), (30,9), (55,40), (30,24), (24,18), (24,20), (42,30), (45,43).

   (a) Show that the log likelihood function is concave.

   (b) In R, write a function that takes θ as input and calculates the log likelihood (based on the data values given in the previous part of this question). The function should only have one input, θ. Use the `curve` procedure in R to plot the log likelihood function for $\theta \in [0, .4]$, and discuss the shape of the log likelihood function.

   (c) In R, numerically solve for the maximum of the log likelihood function. You can use the procedure `optimize`, which does numerical minimization/maximization on functions with a scalar input. It has syntax:

   ```
   optimize(f= , lower= , upper= , maximum=TRUE)
   ```

   where `f` is the function being optimized, `lower` is the lower bound of the values of the input, and `upper` is the upper bound of the values of the inputs. The term `maximum=TRUE` tells it to perform maximization instead of minimization.

2. Suppose that conditional on τ the random variable X has a normal distribution with mean zero and variance $1/\tau$. The prior distribution for τ is Gamma with parameters α and β.

   (a) Find the parameters of the posterior distribution of τ.

   (b) What is the prior mean of τ?

   (c) What is the posterior mean of τ?

3. Show that the priors in the following cases are conjugate priors:

   (a) $X_1, \ldots, X_n$ is a random sample from the Binomial(p, k) distribution with probability p and size k. Assume that k is known. The prior for p is a Beta distribution with parameters α and β.

   (b) $X_1, \ldots, X_n$ is a random sample from the uniform distribution on $(0, \theta)$. The prior for θ is
   $$f(\theta) = ba^b\theta^{-(b+1)} \cdot 1(\theta > a).$$

   (c) $X_1, \ldots, X_n$ is a random sample from the exponential distribution with density $f(x; \lambda) = \lambda \exp(-\lambda x)$ for $x > 0$. The prior for λ is a Gamma distribution with parameters α and γ.

4. Let Y_1, Y_2 be a pair of independent, identically distributed random variables with common density

$$f_Y(y|\theta) = \frac{1}{\theta}\exp(-y/\theta)$$

for $y > 0$ and zero elsewhere, and for some $\theta > 0$. Consider the estimator $W_{\lambda,\delta} = \lambda Y_1 + \delta Y_2$ for fixed values of λ and δ. For which values of λ and δ is $W_{\lambda,\delta}$ unbiased? Within the set of unbiased estimators $W_{\lambda,\delta}$, which is minimum variance?

5. X has a binomial distribution with parameters $N = 1$ and $p = 1/2$. Y, which is independent of X, has a normal distribution with mean μ and variance 1. Consider the estimator for μ of the form $W_1 = Y + 2X - 1$.

   (a) Is W_1 unbiased?

   (b) What is the variance of W_1?

   (c) Consider the estimator $W_2 = E[W_1|Y]$. Is W_2 unbiased? How does its variance compare to that of W_1?