STA 260s20 Assignment Nine: Bayesian Statistics[1]

The following homework problems are not to be handed in. They are preparation for the final exam. **Please try each question before looking at the solution**. Use the formula sheet.

1. Let $X_1, \ldots, X_n$ be a random sample from a Poisson distribution with parameter $\lambda > 0$. The prior on λ is Gamma(α, β). This makes the prior expected value $\frac{\alpha}{\beta}$.
   (a) Give the posterior density of λ, including the constant that makes it integrate to one.
   (b) Derive the posterior predictive distribution – actually, the posterior predictive probability mass function. Do you recognize it?
2. Let $X_1, \ldots, X_n$ be random sample from a binomial distribution with parameters 4 and θ, where θ is unknown. The prior distribution of θ is beta with parameters α and β.
   (a) Find the posterior density of θ, including the constant that makes it integrate to one.
   (b) For $n = 20$ observations and prior parameters $\alpha = \beta = 1$ (the uniform distribution), we obtain $\overline{x}_n = 2.3$.
      i. What is the posterior mean? The answer is a number.
      ii. What is the posterior mode? The answer is a number.
      iii. We need to know if the coin is biased.
         A. What is $P\left(\Theta = \frac{1}{2}|\mathbf{x}\right)$?
         B. Using R, find $P\left(\Theta < \frac{1}{2}|\mathbf{x}\right)$ and $P\left(\Theta > \frac{1}{2}|\mathbf{x}\right)$. What do you conclude?
      iv. Give a 95% posterior credible interval for Θ, with 2.5% in each tail.
      v. How about a 95% *prior* credible interval? Is such a thing possible?
3. Let $X_1, \ldots, X_n$ be a random sample from an exponential distribution with parameter λ. As in Question 1, let the prior distribution of λ be Gamma(α, β).
   (a) Find the posterior distribution. Show your work. The answer is one of the distributions on the formula sheet. Name the distribution and give formulas for its parameters.
   (b) Derive a formula for the posterior mode — that is, the value of λ for which the posterior density is greatest.
   (c) Imagine a universe in which there are true fixed parameter values, and suppose that the data really do come from an exponential distribution, with fixed true parameter λ_0. If we use the posterior expected value to estimate λ_0, is the estimator consistent? Answer Yes or No and show your work.
   (d) What happens to the posterior variance as $n \to \infty$? Show your work.
   (e) Is the posterior mode consistent? Answer Yes or No and show your work.

[1]Copyright information is at the end of the last page.

4. Let $X_1, \ldots, X_n$ be random sample from a normal distribution with mean μ and precision τ (the precision is one over the variance).

   (a) Suppose that the parameter μ is known, while τ is unknown. The prior on τ is Gamma(α, β). Give the posterior distribution of τ, including the parameters.

   (b) Suppose that τ is known, while this time μ is unknown. The prior on μ is standard normal. Find the posterior distribution of μ.

5. Suppose the prior is a finite mixture of prior distributions. That is, the parameter θ has prior density

$$\pi(\theta) = \sum_{j=1}^{k} a_j \, \pi_j(\theta)$$

   The constants $a_1, \ldots, a_j$ are called *mixing weights*; they are non-negative and they add up to one.

   Show that the posterior distribution is a mixture of the posterior distributions corresponding to $\pi_1(\theta), \ldots, \pi_k(\theta)$. What are the mixing weights of the posterior?

   This result can be useful if your model has a conjugate prior family, because you can represent virtually any prior opinion by a mixture of conjugate priors. For example, a bimodal prior might be just a mixture of two normals with different expected values. Thus, you can have essentially any prior you wish, and also the convenience of an exact posterior distribution.
