These problems are related to the material on Bayesian inference from week 9 and 10. These are for practice and exam preparation and will not be graded. The solutions will be posted on Friday.

1. Exercise 16.8 on page 807 in the textbook.

2. Two friends Bob and Carol decide to play 5 sets of tennis. Before they begin playing, Carol and Bob agree to a wager: he'll pay her \$40 dollars if she wins and she'll pay him \$60 if he wins (she's a better tennis player). Carol wins two of the first three sets.

   Assume that the prior distribution on the probability p that Carol will win any particular set is a Beta distribution with $\alpha = 1.2$ and $\beta = 0.8$, and that the outcomes of the 5 sets are independent.

   (a) Find the Bayes estimate for p given that Carol won 2 of the first 3 sets.

   (b) Find the Bayes estimate for the probability that Bob will win 3 out of 5 sets given that he has already lost 2 of the first three.

   (c) Calculate Carol's posterior expectation of the amount of money she will win (or lose).

3. A survey of 226 residents of California found that 83% supported the latest bond initiative.

   (a) Calculate a 95% Confidence Interval for p.

   (b) Assuming that we start with a p that has a Beta prior with $\alpha = \beta = 1/2$, find a 95% credible interval for p. (Hint: Approximate the Beta probabilities by a normal probability with the same mean and standard deviation.)

   (c) How do these two intervals compare? Which one is shorter?

4. We start with $\mu \sim \mathcal{N}(0, \tau^2)$ as a prior distribution on the normal parameter. We observe n independent observations $x_1, \ldots, x_n$ from a $\mathcal{N}(\mu, \sigma^2)$, and find that the posterior distribution is

$$\mu \mid \bar{x} \sim \mathcal{N}\left(\bar{x}\left(\frac{n\tau^2}{n\tau^2+\sigma^2}\right),\ \frac{\sigma^2\tau^2}{n\tau^2+\sigma^2}\right)$$

   (see Example 16.4 in the text). If we then observe an additional m observations, $y_1, \ldots, y_m$, independent normals with mean μ and variance σ^2, show that the Bayesian estimator computed by using the posterior $\mu \mid \bar{x}$ as the prior distribution on μ will be the same as the estimator which puts the x's and y's together into one sample of $n + m$ observations.

5. We are going to measure the LDL cholesterol levels of 150 patients on a special exercise regimen. The prior distribution on μ for these measurements is a normal with mean 120 with a standard deviation of 40.

   (a) If the average cholesterol of this group is $\bar{x} = 114.73$, then calculate the Bayesian estimate of μ if we assume that $\sigma^2 = 20$.

   (b) Calculate the posterior probability that μ is less than 120.