STAT 7630: Homework 3
(Due: Thursday, 09/26/2024)

Note: Show all your work for the necessary steps to receive full credit.

Please turn in the HW on paper, hand-written and/or typed. For computational problems, return only the relevant parts of the output with comments/annotations. Questions taken from the textbook are marked with BR for "Bayes Rules!". From the code you are using to answer the problems, turn in the relevant output, and the figures (if requested), preferably printed from the output. No need to turn in your code or long lists of generated samples.

Please disclose any us of AI in your solutions. Regardless of you use it or not, make sure you submit your own work, not copy from other source(s). Any suspicion of AI use will result in automatic 0 or substantial point loss in any question.

BR Chapter 5:
Do Exercises 5.5, 5.6, 5.12, and 5.19.
Hint for 5.12: For part (a), use "`filter(group == "control")`" to pick out the "control" subjects.

Additional Questions (AQ):
AQ1. The eBay selling prices for auctioned Palm M515 PDAs are assumed to follow a normal distribution with unknown μ and σ^2. We wish to perform inference on the mean selling price μ.

(a) Suppose we assume an IG(1100, 250000) prior for σ^2 and let the prior for $\mu \mid \sigma^2$ be

$$p(\mu \mid \sigma^2) \propto (\sigma^2)^{-1/2} \exp\left(-\frac{1}{2\sigma^2/s_0}(\mu - \delta)^2\right),$$

with $s_0 = 1$ and $\delta = 220$. If our sample data are: (212, 249, 250, 240, 210, 234, 195, 199, 222, 213, 233, 251), find a point estimate and a 95% credible interval for μ.

(b) Now suppose (perhaps unrealistically) that we had known the true population variance was $\sigma^2 = 228$. Assuming a conjugate prior for μ with $\delta = 220$ and $\tau^2 = 25$, find a point estimate and a 95% credible interval for the single unknown parameter μ.

(c) How (if at all) does the inference in part (b) differ from the inferences in part (a)? Explain your answer intuitively.

AQ2. A researcher is trying to estimate the mean number of accidents per month within 100 feet of the Gervais Street/Assembly Street intersection in Columbia. She assumes a Poisson(λ) model for the number of accidents Y per month, so that the density function for Y given λ is

$$p(y \mid \lambda) = \frac{\lambda^y e^{-\lambda}}{y!}, \quad y = 0, 1, 2, \ldots, \lambda \geq 0.$$

(a) She uses a standard exponential prior distribution for λ (i.e., an exponential distribution with mean 1, which is the same as a gamma distribution with shape 1 and rate 1). Derive the general form of her posterior distribution for λ given a random sample $y_1, \ldots, y_n$ from n months.

(b) If she gathers the following accident counts from 15 randomly selected months: $1, 0, 4, 1, 4, 2, 5, 3, 0, 3, 1, 2, 2, 4, 1,$
find the posterior mean and a 95% credible interval (get both a quantile-based interval and a highest posterior density (HPD) interval) for λ using the standard exponential prior, along with these data.