YOUR NAME
YOUR EMAIL
February 18, 2026

Homework 4

1. (from last HW 1+2, 4 points) Suppose that X is a discrete random variable with:

$$P(X = x) = \begin{cases} \frac{2}{3}\theta & x = 0 \\ \frac{1}{3}\theta & x = 1 \\ \frac{2}{3}(1-\theta) & x = 2 \\ \frac{1}{3}(1-\theta) & x = 3 \end{cases}$$

where $0 \leq \theta \leq 1$. Suppose we observe 10 independent observations from such a distribution: $(3, 0, 2, 1, 3, 2, 1, 0, 2, 1)$. (We'll return to this same distribution later)

- (2 points) Find a Bayesian posterior for the parameter value θ (Hint: Try uniform prior)
- (1 point) Plot the posterior density.
- (1 point) Report both the posterior mean and mode.

2. (5 points) Suppose that 100 items are sampled from a manufacturing process, and 3 are found to be defective. Assuming that each item status is independent, we will use a binomial model for this data.

   Consider using a Beta(α, β) prior for Θ, the unknown proportion of defective items.

   (a) (2 points) Use a Beta$(\alpha = 1, \beta = 1)$ prior, and find the posterior distribution. Plot the posterior distribution. What is the posterior mean?

   (b) (2 points) Use a Beta$(\alpha = 0.5, \beta = 5)$ as the prior, and find the posterior distribution. Plot the posterior distribution. What is the posterior mean?

   (c) (1 point) Comment on the differences between the two posterior distributions and the posterior means.

3. (6 points) Suppose that the waiting time X_i in a queue is modeled as an exponential random variable with unknown parameter λ.

   (a) (3 points) Assuming we can observe n independent observations $X_1, X_2, \ldots, X_n$, use a Gamma(λ, α) as a prior distribution for Λ, the random variable denoting the rate λ in the exponential model. **Derive the posterior distribution of $\Lambda|X$ under this model and prior.**

   (b) (1 point) Now suppose we observe an average waiting time of 5.1 minutes for 20 customers, and we use a Gamma prior that satisfies $E[\Lambda] = 10$ and $\text{Var}(\Lambda) = 20^2$. **What is the posterior distribution in this case? What about the posterior mean**?

   (c) (1 point) How does this compare if we get the same sample average, but there are only $n = 5$ total customers?

   (d) (1 point) Compare both posteriors to the shape of the likelihood function (create / draw a figure, or comment on the differences or similarities of the shapes.)