Homework 2

Due end-of-day Monday, February 17th, via Canvas

PUBH 8442: Bayes Decision Theory and Data Analysis

1. Show that the posterior median minimizes posterior risk under the absolute error loss $l(\theta, a) = |\theta - a|$.

2. Let $y \sim \text{Binomial}(n, \theta)$ with prior $\theta \sim \text{Uniform}(0, 1)$. Consider a point estimate for θ, based on y.

   (a) In a decision-theoretic context, what is the action space $\mathcal{A}$?

   Consider squared error loss $l(\theta, a) = (\theta - a)^2$ for parts (b-g).

   (b) What is the Bayes decision rule?

   (c) What is the frequentist risk function for the Bayes decision rule?

   (d) The maximum likelihood estiamte (MLE) is $\hat{\theta} = y/n$. What is the frequentist risk function for the MLE?

   (e) For which values of θ is the MLE superior to the Bayes rule?

   (f) What is the bias (if any) of the MLE?

   (g) What is the bias (if any) of the Bayes rule?

3. Prove that any Bayes rule that has constant frequentist risk (risk does not depend on θ) is minimax.

4. Let $y_1, \ldots, y_n$ be iid with mean μ and variance σ^2, where μ has prior with mean μ_0 and variance τ^2. Consider shrinkage estimators for μ of the form $d_B(\mathbf{y}) = B\mu_0 + (1 - B)\bar{y}$. Show that the rule minimizing Bayes risk under squared error loss has shrinkage factor

$$B = \frac{\sigma^2}{\sigma^2 + n\tau^2}.$$

   (This relates to slides 5-7 of "The Bias-Variance Tradeoff")

5. *Advertising a Sunscreen Pill*
   An advertising firm creates a commercial for a "sunscreen pill" that one can swallow to provide mild SPF protection throughout the day. A focus group of $n = 16$ individuals are randomly selected to view the commercial. Afterward, they are given the opportunity to purchase the product, and $y = 10$ of the individuals choose to do so. Because such a product is completely new, it is reasonable to assume a uniform prior for θ, the probability that a randomly selected individual will be inclined to buy the product after viewing the commercial.

   (a) Compute and interpret a 95% highest posterior density credible interval for θ based on $y = 10$.

   (b) Compute and interpret a 95% symmetric quantile credible interval for θ based on $y = 10$.

   (c) Which value(s) of y would result in equivalent highest posterior density and symmetric quantile intervals?

   (Note: Your solutions will likely require some computing. Feel free to use or modify any of the R code that is linked to slides from class.)