Midterm [20 pts]

March 18th, 2024 in class

PUBH 8442: Bayes Decision Theory and Data Analysis

Give your final answers in simplified, closed form wherever possible. However, partial credit will be awarded for incomplete solutions. Good luck!

1. *In-vitro fertilization* [11 pts]
   Let θ be the unknown success rate of a novel *in vitro* fertilization method. Let X be the number of failed attempts before the first success. Then X has a geometric distribution with probability θ:

   $$P(X = k \mid \theta) = (1-\theta)^k \theta$$

   for $k = 0, 1, 2, 3, \ldots$.

   (a) (3 points) Show that the Jeffreys prior for θ is

   $$p(\theta) \propto \theta^{-1}(1-\theta)^{-\frac{1}{2}}$$

   for $\theta \in (0, 1)$.

   The log-likelihood is

   $$L = \log P(X \mid \theta) = X \log(1-\theta) + \log \theta.$$

   Differentiating twice gives

   $$\frac{d^2 L}{d\theta^2} = -\frac{1}{\theta^2} - \frac{X}{(1-\theta)^2}.$$

   The Fisher information $I(\theta)$ is the negative expectation of the log-likelihood. Using $E(X \mid \theta) = (1-\theta)\theta^{-1}$,

   $$I(\theta) = \frac{1}{\theta^2} + \frac{1-\theta}{\theta(1-\theta)^2} = \theta^{-2}(1-\theta)^{-1}.$$

   The Jeffreys prior is the square root of $I(\theta)$:

   $$p(\theta) \propto \sqrt{I(\theta)} = \theta^{-1}(1-\theta)^{-\frac{1}{2}}.$$

   For parts (b–e), use the Jeffreys prior given in part (a).

   (b) (1 points) Is the prior proper?
   No, the prior is improper because its integral is not finite on $[0, 1]$.

(c) (2 points) For which values of X, if any, will the posterior be proper?

The posterior is

$$p(\theta \mid X) \propto (1-\theta)^{X-1/2}.$$

Thus the posterior is a proper Beta$\{1, X+1/2\}$ *distribution for all possible* X.

(d) (2 points) What is the Bayes estimator for θ, given X, under squared error loss?

The Bayes estimator under squared error loss is the conditional expectation. Using $\theta \mid X \sim$ *Beta*$\{1, X+1/2\}$*, the Bayes estimator is*

$$E(\theta \mid X) = \frac{1}{1+X+1/2} = \frac{2}{2X+3}.$$

(e) (3 points) Let X_2 and X_3 be distributed identically to X, and assume X,X_2, and X_3 are independent given θ. If $X = 1$, what is $P(X_2 + X_3 = 0 \mid X = 1)$, the probability that there are no failed attempts in the next two trials?

The posterior after observing $X = 1$ *is* $p(\theta \mid X = 1) =$ *Beta*$(1, 3/2)$*, and after observing* $X = 1$, $X_2 = 0$ *is Beta*$(2, 3/2)$ *distributed. Thus,*

$$\begin{aligned} P(X_2 + X_3 = 0 \mid X = 1) &= P(X_2 = 0 \mid X = 1) P(X_3 = 0 \mid X_2 = 0, X = 1) \\ &= \frac{1}{B(1,3/2)} \int_0^1 \theta \cdot (1-\theta)^{1/2} d\theta \times \frac{1}{B(2,3/2)} \int_0^1 \theta \cdot \theta (1-\theta)^{1/2} d\theta \\ &= \frac{B(3,3/2)}{B(1,3/2)} \\ &= 8/35, \end{aligned}$$

where $B(\cdot, \cdot)$ *denotes the Beta function.*

2. *Bayesian P-Value* [3 pts]
Consider $y_1, \ldots, y_n \overset{iid}{\sim} N(\theta, \sigma^2)$, with σ^2 known and a flat prior on θ: $p(\theta) = 1 \; \forall \; \theta$. Let $\mathbf{y} = (y_1, \ldots, y_n)$, and let $T(\cdot)$ give the sample mean: $T(\mathbf{y}) = \bar{y}$. What is the posterior predictive p-value,

$$P(T(\mathbf{Y}) \geq T(\mathbf{y}) \mid \mathbf{y})?$$

The posterior distribution of θ *given* $\mathbf{y}$ *is*

$$\theta \mid \mathbf{y} \sim N(\bar{y}, \sigma^2/n)$$

(see slide 9 of the "More on Interval Estimation" notes on 2/18). So, because $\bar{Y} \sim N(\theta, \sigma^2/n)$*, it follows from the marginal distribution of the normal-normal model that*

$$\bar{Y} \mid \mathbf{y} \sim N(\bar{y}, 2\sigma^2/n).$$

Thus, $P(\bar{Y} \geq \bar{y} \mid \mathbf{y}) = \frac{1}{2}$.

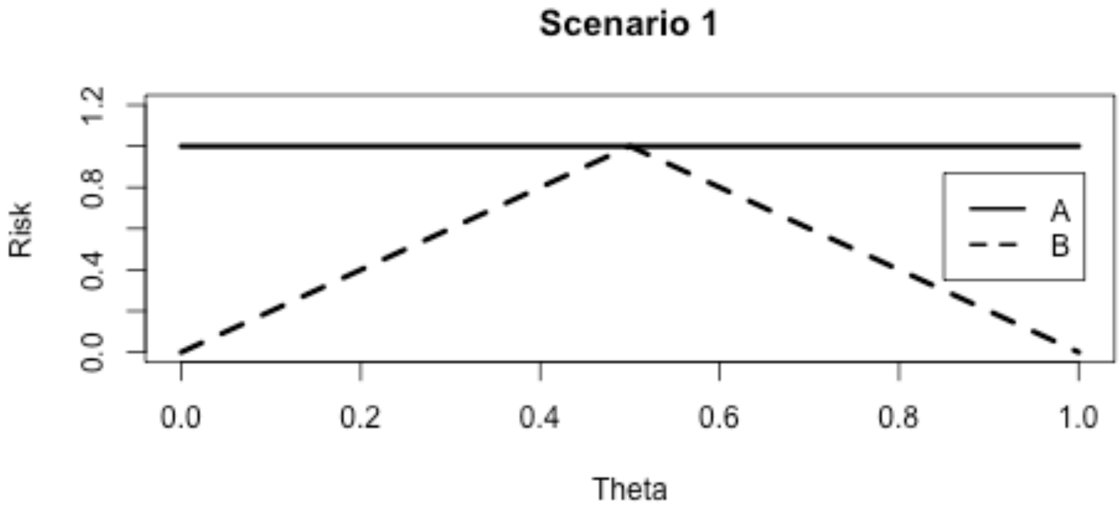

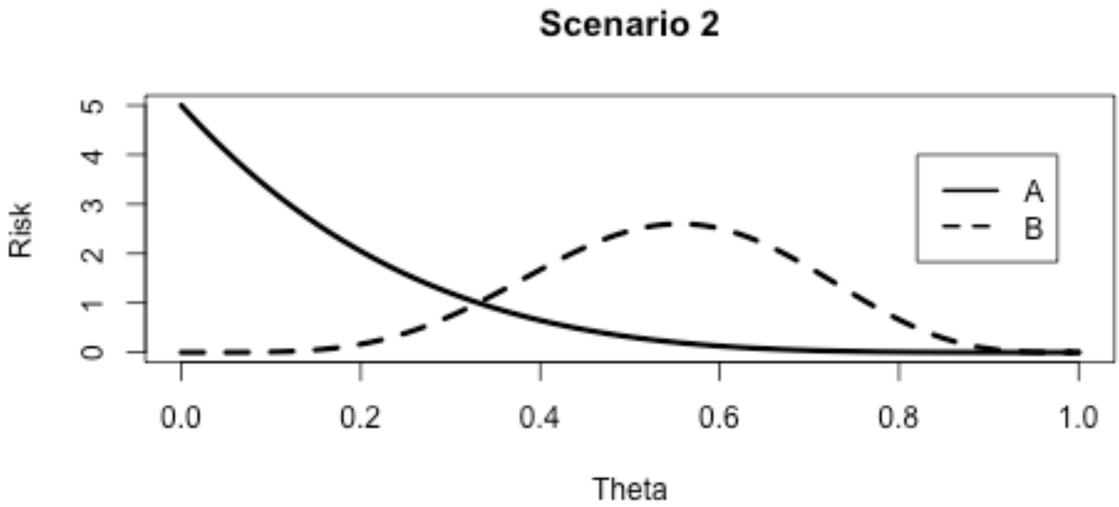

3. *Decision rules* [6 pts]
   The above figure shows the frequentist risk of two decision rules (A and B) as a function of an unknown parameter θ, under two different scenarios. For each scenario, A and B are the only two potential decisions (action space = {A,B}), θ can take any value between 0 and 1, and θ has a prior that is positive over the unit interval $p(\theta) > 0 \ \forall \ \theta \in [0, 1]$.

   For each of the following, answer "yes", "no", or "not enough informations":

   (a) Is A admissible in Scenario 1? *No*
   (b) Is A minimax in Scenario 1? *Yes*
   (c) Is A a Bayes rule in Scenario 1? *No*
   (d) Is B admissible in Scenario 1? *Yes*
   (e) Is B minimax in Scenario 1? *Yes*
   (f) Is B a Bayes rule in Scenario 1? *Yes*

(g) Is A admissible in Scenario 2? *Yes*

(h) Is A minimax in Scenario 2? *No*

(i) Is A a Bayes rule in Scenario 2? *Not enough information*

(j) Is B admissible in Scenario 2? *Yes*

(k) Is B minimax in Scenario 2? *Yes*

(l) Is B a Bayes rule in Scenario 2? *Not enough information*