4 The Reverend Thomas Bayes / Decisions, decisions (AJ Chapter 9)

1. Let $X_{1,}, \ldots, X_n \sim \text{Uniform}(0, 1/\theta)$ where $\theta > 0$ is some unknown parameter; $X = (X_1, \ldots, X_n)$. Taking the Bayesian route, take (θ, X) to have a joint distribution f. Take your prior distribution to be the distribution to be

$$f(\theta) = 1/\theta, \qquad\qquad f(\theta) = \int f(\theta, x)dx.$$

   (a) Is this a proper prior? i.e. is $\int f(\theta)d\theta = 1$? In words, what does our choice of prior say about our beliefs about θ?

   (b) Use the formula

$$f(\theta \mid x) \propto f(\theta) \prod_{i=1}^{n} f(X_i \mid \theta)$$

   to find the posterior distribution? Is it a proper distribution? i.e. is $\int f(\theta \mid x)d\theta = 1$?

   (c) What is the posterior mean? Does that seem reasonable?

2. There are two envelopes, one containing £$\theta$and the other containing £2θ. Your prior beliefs about θ are given by the function $f(\theta)$.You are allowed to pick one of the envelopes uniformly at random, and to see how much money is inside. Call the amount you see £X. Given θ, X is uniformly distributed on the set $\{\theta, 2\theta\}$.
The other other envelope therefore either has £$X/2$ in it, or £$2X$. You now have the choice of keeping your envelope or switching to the other envelope.

   (a) You start off reasoning that you are equally likely to have $X = \theta$ as to have $X = 2\theta$. If you switch you can expect to be better off by

$$\frac{1}{2} \cdot \frac{X}{2} + \frac{1}{2} \cdot 2X - X = \frac{X}{2} > 0?$$

   This suggests you should always switch. Under what kind of prior will that be the case? If you always switch, why not just pick the other envelope to start with and then not bother switching?

   (b) What kind of prior do you think you should use? When should you switch envelopes?

3. Let $X_1, \ldots, X_n \sim N(\theta, 1)$ and take your prior beliefs about θ be $N(a, b^2)$. You will observe $X_1, ..., X_n$ from the distribution $N(\theta, \sigma^2)$ where σ is known.

(a) Find the posterior distribution. Try to express your answer as simply as possible.

(b) Calculate the limit of the posterior distribution as $n \to \infty$ with a, b, σ fixed.

(c) Let $a = 0$, $b = 10$. Setting $\theta = 100$, $\sigma = 1$ and $n = 8$, sample $X_1, ..., X_n$.Calculate the posterior distribution. Find a 98% credible interval for θ by calculating the 1st and 99th quantiles of the posterior distribution using `qnorm`.

(d) Repeat part (c) but $b = 1/10$. Comment on the difference.

4. An experiment will produce a result X with distribution Binomial$(100, p)$. A scientist tells you that their prior beliefs is $p \sim$ Beta$(1000, 1000)$. You observe a sample $x = 2$.

   (a) Plot the prior distribution on the interval $[0, 1]$.

   (b) Calculate the posterior distribution and plot it. Compare with part (a).

   (c) Calculate a 95% posterior credible interval for p. [Hint: use `qbeta`] Does this seem to be compatible with your observation?

   (d) What do you think about the prior belief the scientist supplied? Should we begin again with a different prior?

5. Your prior beliefs are given by a positive probability density function $f : \mathbb{R} \to (0, \infty)$. You choose to use the absolute value loss function,

$$L(\theta, \delta) = |\theta - \delta|$$

   You will oberve a sample X that depends on θ in some way. You want to find an estimator $\delta = \delta(x)$ that minizes your expected loss, the Bayes risk,

$$r(f, \delta) = \int R(\theta, \delta) f(\theta)\, d\theta = \int \int L(\theta, \delta(x))\, dx\, f(\theta) d\theta$$

   Show that the median of the posterior distribution minimizes the posterior risk

$$r(\delta \mid x) = \int L(\theta, \delta(x)) f(\theta \mid x) d\theta,$$

   and that therefore the posterior median minimizes the Bayes risk. [Reminder: Under squared error loss/quadratic loss, the posterior mean minimizes the Bayes risk.]

6. Let $X_1, \ldots, X_n \sim$ Poisson(λ) where $\lambda$is unknown.

   (a) Calculate the Jeffreys prior for the Poisson(λ) distribution

$$\text{prior} \propto \sqrt{I(\theta)} \qquad I(\theta) = \text{Var}_{X \sim \text{Poisson}(\theta)} \left(\frac{\partial}{\partial \theta} \log f(X \mid \theta) \right) \qquad f(x \mid \theta) = e^{-\theta} \frac{\theta^x}{x!}$$

   Calculate the posterior distribution with respect to the Jeffreys prior.

(b) After taking expert advice you decide instead that your prior beliefs about λ are given by a Gamma distribution with mean 100 and variance 200.

i. Calculate the posterior distribution.

ii. Find the Bayes estimator with respect to squared error loss.

iii. Find the Bayes estimator with respect to absolute error loss when $n = 4$, and $X_1 = 57$, $X_2 = 104$, $X_3 = 97$, $X_4 = 120$. (i.e. the posterior median)

7. Let $X_1, \ldots X_n \sim N(\theta, \theta)$ and consider squared error loss/quadratic loss

$$L(\theta, \delta) = (\theta - \delta)^2$$

Consider two estimators, the sample mean and the sample variance:

$$\delta_1(x) = \overline{x} = \frac{1}{n} \sum_i x_i, \qquad \delta_2(x) = \frac{1}{n-1} \sum_i (x_i - \overline{x})^2.$$

(a) Calculate the risks of δ_1 and δ_2 as functions of θ.

(b) Give an example of a prior distribution under which $\delta_1$has lower Bayes risk, and one where δ_2 has lower Bayes risk.