HOMEWORK SOLUTIONS #2

1. **Computer Exercise**

$$\tau = \frac{\pi_0(C_{10} - C_{00})}{\pi_1(C_{01} - C_{11})} = \frac{70\pi_0}{75(1-\pi_0)}$$

From the text/notes, we know that

$$\begin{aligned}
P_j(\Gamma_1) &= 1 - \Phi\left(\frac{\log(\tau)}{d} + (-1)^j \frac{d}{2}\right) \\
\text{where} \quad d &= \frac{\mu_1 - \mu_0}{\sigma} = \frac{6}{2} = 3 \\
\text{thus} \quad P_D &= P_1(\Gamma_1) = 1 - \Phi\left(\frac{\log(\tau)}{3} - \frac{3}{2}\right) \\
P_F &= P_0(\Gamma_1) = 1 - \Phi\left(\frac{\log(\tau)}{3} + \frac{3}{2}\right) \\
R_1(\delta) &= C_{01} + P_D(C_{11} - C_{01}) \\
R_0(\delta) &= C_{00} + P_F(C_{10} - C_{00}) \\
\text{and} \quad V(\pi_0) &= \pi_0 R_0(\delta) + (1-\pi_0) R_1(\delta)
\end{aligned}$$

The desired figure is shown below. Matlab says the least favorable prior is $\pi_L = 0.101$ and the corresponding conditional risks are $R_1(\delta_{\pi_L}) = R_0(\delta_{\pi_L}) = 0.259$. Thus δ_{π_L} is a minimax (equalizer) decision rule.

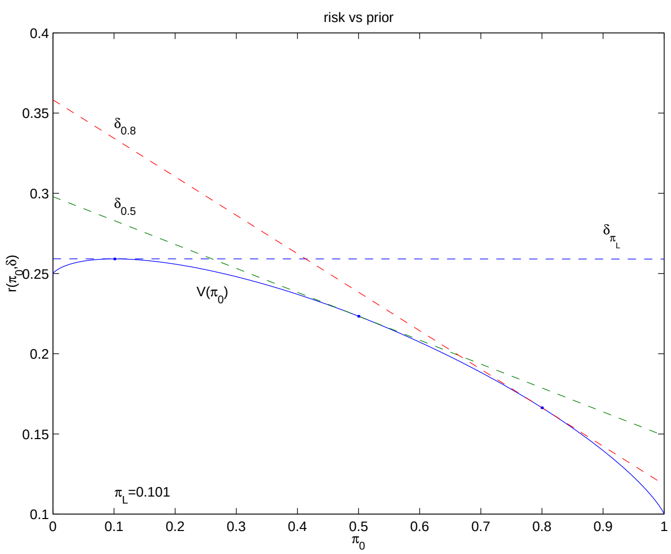

2. [Poor II.2] Recall that the likelihood ratio is given by

$$L(y) = \frac{3}{2(y+1)}, \quad 0 \le y \le 1.$$

(b) With uniform costs, the least-favorable prior will be interior to $(0,1)$, so we examine the conditional risks of Bayes rules for an equalizer condition. The critical region for the Bayes rule δ_{π_0} is given by

$$\Gamma_1 = \left\{ y \in [0,1] \,\middle|\, L(y) \ge \frac{\pi_0}{1-\pi_0} \right\} = [0, \tau'],$$

where

$$\tau' = \begin{cases} 1 & \text{if } 0 \le \pi_0 \le \frac{3}{7} \\ \frac{1}{2}\left(\frac{3}{\pi_0} - 5\right) & \text{if } \frac{3}{7} < \pi_0 < \frac{3}{5} \\ 0 & \text{if } \frac{3}{7} \le \pi_0 \le 1 \end{cases}.$$

Thus, the conditional risks are:

$$R_0(\delta_{\pi_0}) = \int_0^{\tau'} \frac{2}{3}(y+1)dy = \begin{cases} 1 & \text{if } 0 \le \pi_0 \le \frac{3}{7} \\ \frac{2\tau'}{3}\left(\frac{\tau'}{2} + 1\right) & \text{if } \frac{3}{7} < \pi_0 < \frac{3}{5} \\ 0 & \text{if } \frac{3}{7} \le \pi_0 \le 1 \end{cases},$$

and

$$R_1(\delta_{\pi_0}) = \int_{\tau'}^{1} dy = \begin{cases} 0 & \text{if } 0 \le \pi_0 \le \frac{3}{7} \\ 1-\tau' & \text{if } \frac{3}{7} < \pi_0 < \frac{3}{5} \\ 1 & \text{if } \frac{3}{7} \le \pi_0 \le 1 \end{cases}.$$

By inspection, a minimax threshold τ'_L is the solution to the equation

$$\frac{2\tau'_L}{3}\left(\frac{\tau'_L}{2} + 1\right) = 1 - \tau'_L,$$

which yields $\tau'_L = (\sqrt{37} - 5)/2$. The minimax risk is the value of the equalized conditional risk; i.e., $V(\pi_L) = 1 - \tau'_L$.

(c) Since Y is a continuous random variable, the Neyman-Pearson test is given by

$$\delta_{NP}(y) = \begin{cases} 1 & \text{if } \frac{3}{2(y+1)} \ge \eta \\ 0 & \text{if } \frac{3}{2(y+1)} < \eta \end{cases},$$

where η is chosen to give false-alarm probability α. Since $L(y)$ is monotone decreasing in y, the above test is equivalent to

$$\delta_{NP}(y) = \begin{cases} 1 & \text{if } y \le \eta' \\ 0 & \text{if } y > \eta' \end{cases},$$

where $\eta' = \frac{3}{2\eta} - 1$. Thus, the false-alarm probability is:

$$P_F(\delta_{NP}) = P_0(Y < \eta') = \int_0^{\eta'} \frac{2}{3}(y+1)dy = \begin{cases} 0 & \text{if } \eta' \le 0 \\ \frac{2\eta'}{3}\left(\frac{\eta'}{2} + 1\right) & \text{if } 0 < \eta' < 1 \\ 1 & \text{if } \eta' \ge 1 \end{cases}.$$

The threshold for $P_F(\delta_{NP}) = \alpha$ is the solution to

$$\frac{2\eta'}{3}\left(\frac{\eta'}{2} + 1\right) = \alpha,$$

which is $\eta' = \pm\sqrt{1+3\alpha}-1$. Recalling $y \in [0,1]$, then noting that $\eta' < 0$ would yield the trivial rule $\delta(y) = 0$, we discard the solution $\eta' = -\sqrt{1+3\alpha}-1$. So, the α-level Neyman-Pearson test is

$$\delta_{NP}(y) = \begin{cases} 1 & \text{if } y \le \sqrt{1+3\alpha}-1 \\ 0 & \text{if } y > \sqrt{1+3\alpha}-1 \end{cases}.$$

The detection probability is

$$P_D(\delta_{NP}) = \int_0^{\eta'} dy = \eta' = \sqrt{1+3\alpha}-1, \quad 0 < \alpha < 1.$$

3. [Poor II.6]

Recall that $p_0(y) = p_N(y+s)$ and $p_1(y) = p_N(y-s)$, which gave

$$L(y) = \frac{1+(y+s)^2}{1+(y-s)^2}.$$

With equal priors and uniform costs, the critical region for Bayes testing was $\Gamma_1 = \{L(y) \ge 1\} = \left\{1+(y+s)^2 \ge 1+(y-s)^2\right\} = \{2sy \ge -2sy\} = [0,\infty)$. Thus, the Bayes test was

$$\delta_B(y) = \begin{cases} 1 & \text{if } y \ge 0 \\ 0 & \text{if } y < 0 \end{cases}$$

and the minimum Bayes risk was

$$r(\delta_B) = \frac{1}{2}\int_0^\infty \frac{1}{\pi\left[1+(y+s)^2\right]}dy + \frac{1}{2}\int_{-\infty}^0 \frac{1}{\pi\left[1+(y-s)^2\right]}dy = \frac{1}{2} - \frac{\tan^{-1}(s)}{\pi}.$$

(b) Because of the symmetry of this problem with uniform costs, we can guess that 1/2 is the least-favorable prior. To confirm this, we recall from part (a) that this choice gives an equalizer rule:

$$R_0(\delta_{1/2}) = \int_0^\infty \frac{1}{\pi\left[1+(y+s)^2\right]}dy = \int_{-\infty}^0 \frac{1}{\pi\left[1+(y-s)^2\right]}dy = R_1(\delta_{1/2}).$$

(c) Consider the likelihood ratio test for choosing H_1

$$\begin{aligned} \frac{1+(y+s)^2}{1+(y-s)^2} &> \tau \\ 1+(y+s)^2 &> \tau + \tau(y-s)^2 \\ (1-\tau)y^2 - 2ys(1+\tau) &> \tau - 1 + (\tau-1)s^2 \end{aligned}$$

If $\tau < 1$ then we can complete the square of the LHS as follows (considering $\tau > 1$ is done similarly).

$$\left(y - \frac{s(1+\tau)}{1-\tau}\right)^2 > \tau' = \frac{4s^2\tau}{(1-\tau)^2} - 1$$

So we can determine the decision regions from this

$$\Gamma_1 = \left(-\infty, \frac{s(1+\tau)}{1-\tau} - \sqrt{\tau'}\right) \cup \left(\frac{s(1+\tau)}{1-\tau} + \sqrt{\tau'}, \infty\right)$$

The value of the threshold is found by solving

$$\alpha = \int_{\Gamma_1} \frac{1}{\pi}\frac{1}{1+(y-s)^2}dy$$

4. [Poor II.16] We have M hypotheses $H_0, H_1 \ldots, H_{M-1}$, where Y has distribution P_i and density p_i under hypothesis H_i. A decision rule δ is a partition of the observation set Γ into regions $\Gamma_0, \Gamma_1, \ldots, \Gamma_{M-1}$, where δ chooses hypothesis H_i when we observe $y \in \Gamma_i$. Equivalently, a decision rule can be viewed as a mapping from Γ to the set of decisions $\{0, 1, \ldots, M-1\}$, where $\delta(y)$ is the index of the hypothesis accepted when we observe $Y = y$.

   On assigning costs C_{ij} to the acceptance of H_i when H_j is true, for $0 \leq i, j \leq (M-1)$, we can define *conditional risks*, $R_j(\delta), j = 0, 1, \ldots, M-1$, for a decision rule δ, where $R_j(\delta)$ is the conditional expected cost given that H_j is true. We have

$$R_j(\delta) = \sum_{i=0}^{M-1} C_{ij} P_j(\Gamma_i).$$

   Assuming priors $\pi_j = P(H_j), j = 0, 1, \ldots, M-1$, we can define an overall average risk or *Bayes risk* as

$$r(\delta) = \sum_{j=0}^{M-1} \pi_j R_j(\delta).$$

   A *Bayes decision rule* will minimize the Bayes risk.

   We can write

$$r(\delta) = \sum_{j=0}^{M-1} \sum_{i=0}^{M-1} \pi_j C_{ij} P_j(\Gamma_i) = \sum_{i=0}^{M-1} \left[\sum_{j=0}^{M-1} \pi_j C_{ij} P_j(\Gamma_i) \right]$$

$$= \sum_{i=0}^{M-1} \left[\sum_{j=0}^{M-1} \pi_j C_{ij} \int_{\Gamma_i} p_j(y) dy \right] = \sum_{i=0}^{M-1} \int_{\Gamma_i} \left[\sum_{j=0}^{M-1} \pi_j C_{ij} p_j(y) \right] dy.$$

   Thus, by inspection, we see that the Bayes rule has decision regions given by

$$\Gamma_i = \left\{ y \in \Gamma \,\middle|\, \sum_{j=0}^{M-1} \pi_j C_{ij} p_j(y) = \min_{0 \leq k \leq M-1} \sum_{j=0}^{M-1} \pi_j C_{kj} p_j(y) \right\}.$$

Comments:

If you can compute the minimax threshold or worst-case prior analytically, you should do so. Several of you tried very hard to write out the risk in terms of π_0, which can be very complicated. Often, the easiest thing to do is to calculate the worst-case effective threshold (what we call τ') rather than π_L directly. Once you have the threshold, you can compute π_L. Also, if you discard solutions, be very explicit as to why you are doing that (*i.e.,* an equation may have two solutions where only one is valid).