STA 6360, Report 3.3

Carson Slater *Baylor University*

Problem

Suppose you observe $\mathbf{x} = (x_1, \ldots, x_n)$ with IID $\text{Poisson}(\theta)$ components and that you want to estimate θ under squared error loss. Assume the prior is a $\text{Gamma}(\alpha, \beta)$ density.

1.

Show that the Bayes decision rule for this problem is of the form $\delta^\pi(\mathbf{x}) = a + b\bar{x}$, where $a > 0$, $b \in (0, 1)$, and $\bar{x}$ is the sample mean.

The Bayes decision rule minimizes the expected loss with respect to the posterior distribution. For a Poisson likelihood and a Gamma prior, the posterior is also Gamma distributed with parameters $\alpha + \sum_{i=1}^n x_i$ and $\beta + n\theta$.

We can write the Bayes decision rule as a weighted average of the sample mean and some constant a. To determine a and b, we need to minimize the expected squared error loss with respect to the posterior. The expected squared error loss is given by $E\left[(\theta - \delta^\pi(\mathbf{x}))^2\right] = E\left[(\theta - (a + b\bar{x}))^2\right]$. So then the risk is minimized when

$$a + b\bar{x} = \frac{\alpha + \sum_{i=1}^n x_i}{\beta + n} = \frac{\alpha}{\beta(\beta + n)} + \frac{n}{\beta + n}\bar{x}.$$

Hence, we would then say that $a = \frac{\alpha}{\beta(\beta+n)} > 0$ and $b = \frac{n}{\beta+n} \in (0, 1)$.

2.

Compute and graph the risk functions for $\delta^\pi(\mathbf{x})$ and the MLE, $\bar{x}$.

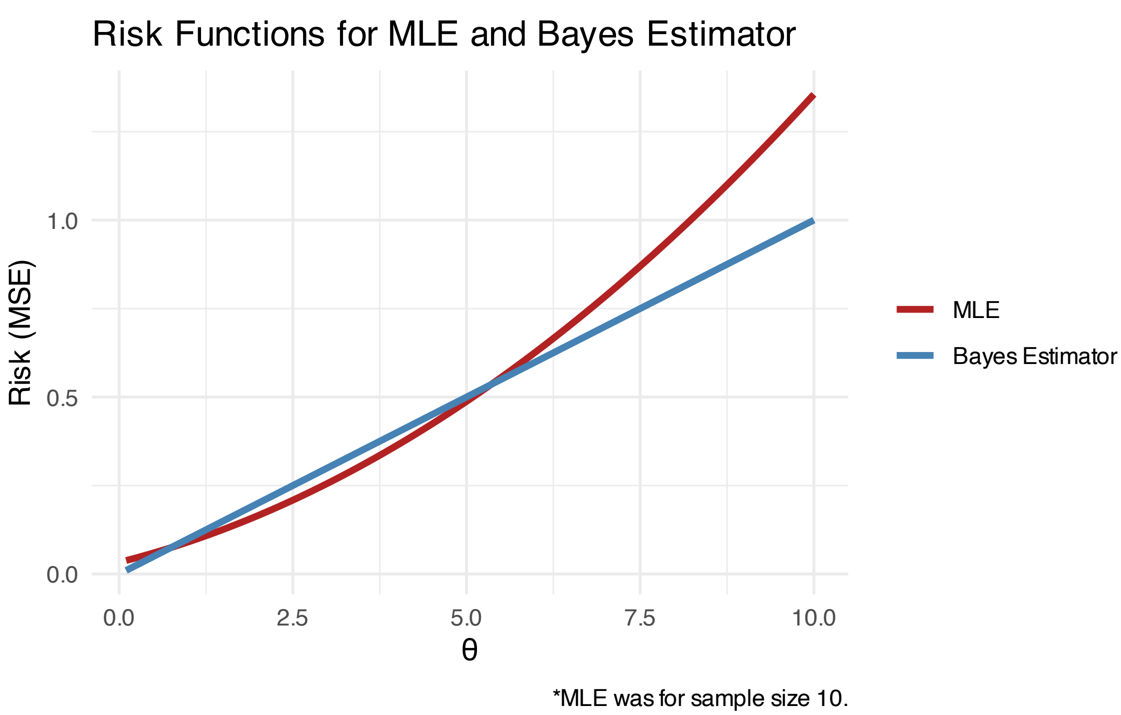

3.

Compute the integrated risk (Bayes risk) of $\delta^\pi(\mathbf{x})$ and show that it is decreasing in n, and goes to 0 as n gets large.

The Bayes risk of the estimator $\delta^\pi(\mathbf{x})$, also known as the integrated risk, is defined as

$$\mathbb{R}(\delta^\pi) = \int R(\theta, \delta^\pi)\pi(\theta)\, d\theta,$$

where $R(\theta, \delta^\pi)$ is the risk function, and $\pi(\theta)$ is the $\text{Gamma}(\alpha, \beta)$ prior density of θ. This becomes:

$$\mathbb{R}(\delta^\pi) = \int_0^\infty \left(\frac{\theta n}{(\beta+n)^2} + \frac{(\alpha - \beta\theta)^2}{(\beta+n)^2} \right) \cdot \frac{\beta^\alpha}{\Gamma(\alpha)} \theta^{\alpha-1} e^{-\beta\theta}\, d\theta.$$

We can separate this into two integrals:

$$\mathbb{R}(\delta^\pi) = \frac{n}{(\beta+n)^2} \int_0^\infty \theta \cdot \frac{\beta^\alpha}{\Gamma(\alpha)} \theta^{\alpha-1} e^{-\beta\theta}\, d\theta + \frac{1}{(\beta+n)^2} \int_0^\infty (\alpha - \beta\theta)^2 \cdot \frac{\beta^\alpha}{\Gamma(\alpha)} \theta^{\alpha-1} e^{-\beta\theta}\, d\theta.$$

Evaluating Each Integral

Part 1: $\int_0^\infty \theta \cdot \frac{\beta^\alpha}{\Gamma(\alpha)} \theta^{\alpha-1} e^{-\beta\theta}\, d\theta$

This is the mean of a $\text{Gamma}(\alpha+1, \beta)$ distribution, which is $\frac{\alpha}{\beta}$. Therefore,

$$\int_0^\infty \theta \cdot \frac{\beta^\alpha}{\Gamma(\alpha)} \theta^{\alpha-1} e^{-\beta\theta}\, d\theta = \frac{\alpha}{\beta}.$$

Part 2: $\int_0^\infty (\alpha - \beta\theta)^2 \cdot \frac{\beta^\alpha}{\Gamma(\alpha)} \theta^{\alpha-1} e^{-\beta\theta}\, d\theta$

Expanding $(\alpha - \beta\theta)^2$ and integrating term-by-term:

$$\int_0^\infty (\alpha - \beta\theta)^2 \cdot \frac{\beta^\alpha}{\Gamma(\alpha)} \theta^{\alpha-1} e^{-\beta\theta}\, d\theta = \alpha^2 - 2\alpha\beta \cdot \frac{\alpha}{\beta} + \beta^2 \cdot \frac{\alpha(\alpha+1)}{\beta^2}.$$

Simplifying, we find

$$= \alpha.$$

Substitute and Simplify

Substituting these results back, we get:

$$\mathbb{R}(\delta^\pi) = \frac{n}{(\beta+n)^2} \cdot \frac{\alpha}{\beta} + \frac{1}{(\beta+n)^2} \cdot \alpha.$$

Combining terms, we find

$$\mathbb{R}(\delta^\pi) = \frac{\alpha(\beta+n)}{\beta(\beta+n)^2} = \frac{\alpha}{\beta(\beta+n)}.$$

Conclusion: Behavior of the Bayes Risk as $n \to \infty$

The Bayes risk of $\delta^\pi(\mathbf{x})$ is

$$\mathbb{R}(\delta^\pi) = \frac{\alpha}{\beta(\beta+n)}.$$

This expression is clearly decreasing in n, as $\beta + n$ increases with n, causing $\mathbb{R}(\delta^\pi)$ to decrease. Moreover,

$$\lim_{n\to\infty} \mathbb{R}(\delta^\pi) = \lim_{n\to\infty} \frac{\alpha}{\beta(\beta+n)} = 0.$$

Thus, the Bayes risk of $\delta^\pi(\mathbf{x})$ is decreasing in n and approaches 0 as n grows large, reflecting the increasing accuracy of the estimator as the sample size increases.

4.

A researcher wants to design an experiment with a sample large enough that the Bayes risk after the experiment is 1/2 of that before the experiment. Find that sample size.

The initial Bayes risk, which we denote as $\mathbb{R}_0(\delta^\pi)$, occurs when we have no additional sample data, so effectively, $n = 0$:

$$\mathbb{R}_0(\delta^\pi) = \frac{\alpha}{\beta^2}.$$

With a sample size n, the Bayes risk becomes:

$$\mathbb{R}(\delta^\pi) = \frac{\alpha}{\beta(\beta+n)}.$$

The problem states that the researcher wants the Bayes risk after the experiment to be half of the initial Bayes risk. Therefore, we need:

$$\mathbb{R}(\delta^\pi) = \frac{1}{2}\mathbb{R}_0(\delta^\pi).$$

Substituting the expressions for $\mathbb{R}(\delta^\pi)$ and $\mathbb{R}_0(\delta^\pi)$, we get:

$$\frac{\alpha}{\beta(\beta+n)} = \frac{1}{2} \cdot \frac{\alpha}{\beta^2}.$$

Cancel α from both sides (assuming $\alpha > 0$) and simplify:

$$\frac{1}{\beta(\beta+n)} = \frac{1}{2\beta^2}.$$

Multiplying both sides by $2\beta^2(\beta+n)$, we get:

$$2\beta^2 = \beta(\beta+n).$$

Dividing by β (assuming $\beta > 0$):

$$2\beta = \beta + n.$$

Now, solving for n:

$$n = \beta.$$

The researcher should use a sample size of $n = \beta$ to ensure that the Bayes risk after the experiment is half of the Bayes risk before the experiment.