APMA 1650 – Homework 9

1. The reading on a voltmeter is uniformly distributed over the interval $[\theta, \theta+1]$, where θ is the true voltage of the circuit. Using your voltmeter, you take n consecutive voltage readings $Y_1, \dots, Y_n$ from a single circuit.

   (a) Show that the sample mean $\bar{Y}$ is a biased estimator for θ, and compute the bias of $\bar{Y}$.

   First we find the expected value of $\bar{Y}$. We know from class that $\mathbb{E}(\bar{Y} = \mu$, where μ is the population mean. Since the population has a uniform distribution, $\mu = (\theta + (\theta+1))/2 = \theta + 1/2$. Since this is not θ, our estimator is biased. The bias is $(\theta + 1/2) - \theta = 1/2$.

   (b) Find a function of $\bar{Y}$ which is an unbiased estimator of θ.

   To turn the expected value of $\bar{Y}$ into θ, all we have to do is subtract 1/2. Thus $\bar{Y} - 1/2$ is an unbiased esimator for θ.

   (c) Find the MSE of $\bar{Y}$ (the biased estimator) when $\bar{Y}$ is used as an estimator of θ.

   We have found the bias of $\bar{Y}$ so we only have to compute its variance. We know that $Var(\bar{Y}) = \sigma^2/n$, where σ^2 is the population variance. Since this is a uniform distribution,

$$\sigma^2 = \frac{((\theta+1)-\theta)^2}{12} = \frac{1}{12}$$

   Thus the MSE is $(1/2)^2 + 1/12n = 1/4 + 1/12n$.

2. Let $X \sim$ Binom $(n,\ p)$

   Consider the following estimator for p:

$$\hat{p}_1 = \frac{X+1}{n+2}$$

   The bias of $\hat{p}_1 = \frac{1-2p}{n+2}$

   The mean square error (MSE) of $\hat{p}_1$ is $\frac{(1-2p)^2 + np(1-p)}{(n+2)^2}$.

   The standard, unbiased estimator for p is

$$\hat{p} = \frac{X}{n}$$

   If the biased estimator ever has a lower MSE, it will likely be either on the extremes $(p = 0, p = 1)$ or in the middle $(p = 1/2)$. The $p = 1/2$ case looks promising, since in that case the bias is actually 0. The standard estimator $\hat{p}$ is unbiased, so its MSE is equal to its variance. Taking $p = 1/2$, we have:

$$MSE(\hat{p}) = Var\left(\frac{X}{n}\right) = \frac{np(1-p)}{n^2} = \frac{1}{4}\frac{n}{n^2}$$

For the biased estimator, taking $p = 1/2$, we have:

$$MSE(\hat{p}_1) = \frac{np(1-p)}{(n+2)^2} = \frac{1}{4}\frac{n}{(n+2)^2}$$

This is smaller than the MSE of the standard estimator $\hat{p}$.

3. Let $Y_1, Y_2, \ldots, Y_n$ be a random sample from a population with probability density function parameterized by θ given by

$$f_\theta(y) = \begin{cases} \theta y^{\theta-1} & 0 < y < 1 \\ 0 & \text{otherwise} \end{cases}$$

where $\theta > 0$ is the parameter of interest.

(a) Show that the sample mean $\bar{Y}$ is an unbiased estimator for $\frac{\theta}{\theta+1}$.

Since the expected value of $\bar{Y}$ is the population mean μ, first we find μ.

$$\begin{aligned} \mu &= \int_0^1 y\theta y^{\theta-1} dy \\ &= \int_0^1 \theta y^\theta dy \\ &= \theta \frac{y^{\theta+1}}{\theta+1}\Big|_1^0 \\ &= \frac{\theta}{\theta+1} \end{aligned}$$

(b) Show that the sample mean $\bar{Y}$ is a consistent estimator for $\frac{\theta}{\theta+1}$.

Since the estimator is unbiased, for consistency all we have to do is show that its variance goes to 0 as n goes to infinity. We use the Magic Variance Formula to find the population variance. Let Y be a sample from the population. Then

$$\begin{aligned} \mathbb{E}(Y^2) &= \int_0^1 y^2\theta y^{\theta-1} dy \\ &= \int_0^1 \theta y^{\theta+1} dy \\ &= \theta \frac{y^{\theta+2}}{\theta+2}\Big|_0^1 \\ &= \frac{\theta}{\theta+2} \end{aligned}$$

By the Magic Variance Formula, for the population variance we have

$$\begin{aligned}\sigma^2 &= \mathbb{E}(Y^2) - [\mathbb{E}(Y)]^2 \\ &= \frac{\theta}{\theta+2} - \left[\frac{\theta}{\theta+1}\right]^2\end{aligned}$$

To get the variance of the sample mean, we divide this by n. This gives us:

$$Var(\bar{Y}) = \frac{1}{n}\left(\frac{\theta}{\theta+2} - \left[\frac{\theta}{\theta+1}\right]^2\right)$$

Since everything inside the parentheses is constant, this goes to 0 as n goes to infinity. It is enough to note that since the population variance is finite, the sample variance must go to 0 as n goes to infinity.

4. Let $Y_1, Y_2, \ldots, Y_n$ be a random sample from a population with probability density function parameterized by θ given by

$$f_\theta(y) = \begin{cases} (\theta+1)y^\theta & 0 < y < 1 \\ 0 & \text{otherwise} \end{cases}$$

where $\theta > -1$ is the parameter of interest.

(a) Find an estimator for θ using the method of moments.

In the method of moments, we set the population mean $\mu = \bar{Y}$ and solve for θ.

$$\begin{aligned}\mu &= \int_0^1 y(\theta+1)y^\theta dy \\ &= \int_0^1 (\theta+1)y^{\theta+1}dy \\ &= (\theta+1)\frac{y^{\theta+2}}{\theta+2}\Big|_0^1 \\ &= \frac{\theta+1}{\theta+2}\end{aligned}$$

Now we set $\mu = \bar{Y}$ and solve for θ.

$$\begin{aligned}\frac{\theta+1}{\theta+2} &= \bar{Y} \\ \theta+1 &= \bar{Y}(\theta+2) = \bar{Y}\theta + 2\bar{Y} \\ \theta(1-\bar{Y}) &= 2\bar{Y} - 1 \\ \theta &= \frac{2\bar{Y}-1}{1-\bar{Y}}\end{aligned}$$

Thus the method of moments estimator is:

$$\hat{\theta} = \frac{2\bar{Y}-1}{1-\bar{Y}}$$

(b) Find the maximum likelihood estimator (MLE) for θ. Compare this to your answer from (a).

First we find the likelihood function. For this density function, we have:

$$\begin{aligned} L(Y_1, \ldots, Y_n | \theta) &= (\theta+1)Y_1^\theta \cdots (\theta+1)Y_n^\theta \\ &= (\theta+1)^n (Y_1 \cdots Y_n)^\theta \end{aligned}$$

We want to find the value of θ which maximizes this. To do so, we can take the derivative with respect to θ and set it equal to 0. Since we have a product of things involving θ, it is easier to take the log to turn the product into a sum and to then maximize the log likelihood function. This is the same thing we did for the geometric distribution. First we find the log likelihood function.

$$\begin{aligned} \log L(Y_1, \ldots, Y_n | \theta) &= \log \left[(\theta+1)^n (Y_1 \cdots Y_n)^\theta \right] \\ &= n \log(\theta+1) + \theta \log(Y_1 \cdots Y_n) \\ &= n \log(\theta+1) + \theta \sum_{i=1}^n \log(Y_i) \end{aligned}$$

To find the value of θ which maximizes this, we take the derivative with respect to θ and set it equal to 0.

$$\begin{aligned} \frac{d}{d\theta} \log L(Y_1, \ldots, Y_n | \theta) &= \frac{d}{d\theta} \left[n \log(\theta+1) + \theta \sum_{i=1}^n \log(Y_i) \right] \\ &= \frac{n}{\theta+1} + \sum_{i=1}^n \log(Y_i) \end{aligned}$$

Setting this equal to 0:

$$\begin{aligned} \frac{n}{\theta+1} &= -\sum_{i=1}^n \log(Y_i) \\ \theta + 1 &= -\frac{n}{\sum_{i=1}^n \log(Y_i)} \\ \theta &= -1 - \frac{n}{\sum_{i=1}^n \log(Y_i)} \end{aligned}$$