1. **Biased estimators**. Let $X \sim \text{Unif}[0, \theta]$, and suppose $n = 1$ for simplicity.
   (a) Show that the likelihood function is
   $$L(x; \theta) = \begin{cases} \frac{1}{\theta} & 0 \leq x \leq \theta \\ 0 & \text{otherwise} \end{cases}$$
   (b) Show that the MLE for θ is x.
   (c) Show that the method-of-moments estimator for θ is $2x$.
   (d) Which of these estimators is/are biased?
2. **Maximum likelihood estimators**. Let $X \sim \text{Binom}(2, p)$.
   (a) Show that the log-likelihood function is
   $$\ell(x_1, \ldots, x_n; p) = A_0 \log(1-p)^2 + A_1 \log(2p(1-p)) + A_2 \log p^2$$
   where A_k is the number of x_i that are equal to k. (So if our sample is $0, 2, 2, 2, 1, 0, 0, 2, 2$, then $A_0 = 3$, $A_1 = 1$, $A_2 = 5$).
   (b) Show that the MLE is $\frac{1}{2}\bar{X}$.
   (c) Is the MLE biased?
3. **Two-parameter estimation**. Let $X \sim \text{Unif}[\theta_1, \theta_2]$. Show that the MLEs of θ_1 and θ_2 are $\min x_i$ and $\max x_i$, respectively.

(1)

(a)

We use the PDF of a uniform random variable:

$$L(x;\theta) = f_\theta(x) = \begin{cases} \frac{1}{\theta} & 0 \leq x \leq \theta \\ 0 & \text{otherwise} \end{cases}$$

(b)

Let's rewrite the likelihood function to be more explicitly a function of θ:

$$L(x;\theta) = f_\theta(x) = \begin{cases} \frac{1}{\theta} & \theta \geq x \\ 0 & \text{otherwise} \end{cases}$$

Since $1/\theta$ is decreasing, this is maximized at the point $\theta = x$.

(c)

We compute the first population moment as $\mathbb{E}[X] = \theta/2$. For the method of moments, we want the first population moment to be estimated by the first sample moment:

$$\frac{\hat{\theta}}{2} = \frac{1}{n} \cdot \sum x_i = x$$

Solving, we get $\hat{\theta} = 2x$.

(d)

The expectation of the MLE is $\mathbb{E}[X] = \theta/2$, which is not θ. So this estimator is biased. Meanwhile, the expectation of the MoME is $\mathbb{E}[2X] = 2\theta/2 = \theta$, so this estimator is unbiased.

(2)

(a)

We find the log-likelihood function as the logarithm of the joint PMF:

$$\begin{aligned} \ell(x_1, \ldots, x_n; p) &= \log L(x_1, \ldots, x_n; p) \\ &= \log \prod_{i=1}^{n} \binom{2}{x_i} p^{x_i} (1-p)^{2-x_i} \\ &= \sum_{i=1}^{n} \log \left(\binom{2}{x_i} p^{x_i} (1-p)^{2-x_i} \right) \end{aligned}$$

If we break up the sum into the case where $x_i = 0$, $x_i = 1$, and $x_i = 2$, then we can separate out the terms as

$$A_0 \log(1-p)^2 + A_1 \log(2p(1-p)) + A_2 \log(p^2)$$

(b)

Before solving, we simplify with properties of logs:

$$\ell = 2A_0 \log(1-p) + A_1 \log 2 + A_1 \log p + A_1 \log(1-p) + 2A_2 \log(p)$$

Then, to maximize ℓ, we take the derivative with respect to p:

$$\begin{aligned}\frac{\partial \ell}{\partial p} &= \frac{-2A_0}{1-p} + 0 + \frac{A_1}{p} - \frac{A_1}{1-p} + \frac{2A_2}{p} \\ &= \frac{-2A_0 p + A_1(1-p) - A_1 p + 2A_2(1-p)}{p(1-p)} \\ &= \frac{(-2A_0 - 2A_1 - 2A_2)p + A_1 + 2A_2}{p(1-p)}\end{aligned}$$

Setting this equal to zero we get

$$(-2A_0 - 2A_1 - 2A_2)\hat{p} + A_1 + 2A_2 = 0$$

which means

$$\hat{p} = \frac{1}{2} \cdot \frac{A_1 + 2A_2}{A_0 + A_1 + A_2} = \frac{\sum X_i}{2n} = \frac{\bar{X}}{2}$$

(c)

We compute

$$\mathbb{E}[\hat{p}] = \mathbb{E}[\bar{X}/2] = \frac{1}{2}\mathbb{E}[\bar{X}] = \frac{1}{2}\bar{[}X_i] = \frac{1}{2} \cdot 2p = p$$

so it's an unbiased estimator.

(3)

Let's write down the likelihood function:

$$L(x_1, \ldots, x_n; \theta_1, \theta_2) = \begin{cases} \frac{1}{\theta_2 - \theta_1} & \theta_1 \le \text{all } x_i \le \theta_2 \\ 0 & \text{otherwise} \end{cases}$$

The likelihood function is zero if the θs are not outside of all the x_i, and it gets smaller the further apart θ_1 and θ_2 are. Therefore, it's maximized if $\theta_1 = \min x_i$ and $\theta_2 = \max x_i$, as required.