YOUR NAME
YOUR EMAIL
April 8, 2026

Homework 7

1. You are testing the lifespan of a new type of computer part. Let $X_1, X_2, \ldots, X_n$ be independent and identically distributed lifespans of a sample of n of these parts. A common model for the lifespan of an electrical component is the exponential distribution. We will parameterize this distribution by its unknown mean, θ. The probability density function (PDF) is:

$$f_X(x;\theta) = \frac{1}{\theta}e^{-x/\theta}, \quad x > 0$$

   (a) (1 point) Suppose an impatient technician decides to estimate the mean lifespan by only recording the lifespan of the very first computer part to fail. That is, let $Y = \min(X_1, X_2, \ldots, X_n)$, and we want to find an unbiased estimator of θ based only on Y (you may use the densities for order statistics derived last semester).

   (b) (1 point) Calculate the variance of your estimator in part 1a. Note that it may be helpful to compare the density of your estimate and compare to known distributions.

   (c) (1 point) Now derive the maximum likelihood estimator $\hat{\theta}$ for θ.

   (d) (1 point) Calculate the Bias of the MLE.

   (e) (1 point) Calculate the Variance of the MLE.

   (f) (1 point) Is the MLE a consistent estimator of θ? Why or why not.

   (g) (1 point) Which estimator do you prefer and why?

   (h) (1 point) Let $\tilde{\theta}$ be some other estimator with the property that $E[\tilde{\theta}] = E[\hat{\theta}]$ (i.e., another estimator with the same bias from part 1d). What is the absolute lowest possible variance for the estimate $\tilde{\theta}$? Does $\hat{\theta}$ achieve this lower bound?

2. Suppose you are helping an engineer analyze the magnitude of background static noise in a new wireless communication channel. You measure the noise amplitude over n independent trials, resulting in the sample $X_1, X_2, \ldots, X_n$.

   You model the noise amplitude using a Rayleigh distribution, which depends on the parameter θ, and has pdf:

$$f(x;\theta) = \frac{x}{\theta}\exp\left\{-\frac{x^2}{2\theta}\right\}, \quad x > 0.$$

   For your own reference, note that $E[X^2] = 2\theta$, and $\text{Var}(X^2) = 4\theta^2$.

   (a) (1 point) Write down the log-likelihood function, $\ell(\theta)$, for all n observations.

   (b) (1 point) Find the MLE of θ for this model.

   (c) (1 point) Calculate the bias and variance of the MLE. Is the MLE a consistent estimator?

   (d) (1 point) Calculate the Fisher Information $\mathcal{I}_n(\theta)$ for the entire sample (Note: you can use the expected value of the second derivative of the log-likelihood).

   (e) (1 point) Let $\tilde{\theta}$ be some other estimator with the property that $E[\tilde{\theta}] = E[\hat{\theta}]$ (i.e., another estimator with the same bias from part (c)). What is the absolute lowest possible variance for the estimate $\tilde{\theta}$? Does $\hat{\theta}$ achieve this lower bound?