

CSE 312: Foundations of Computing II
Quiz Section #9

1. (a) Suppose x_1, x_2, \dots, x_n are independent samples from a normal distribution whose mean is known to be zero, but whose variance is unknown. What is the maximum likelihood estimator for its variance?
(b) Suppose the mean is known to be μ but the variance is unknown. How does the maximum likelihood estimator for the variance differ from the maximum likelihood estimator when both mean and variance are unknown?
2. Let $f(x | \theta) = \theta x^{\theta-1}$ for $0 \leq x \leq 1$, where θ is any positive real number. Let x_1, x_2, \dots, x_n be independent samples from this distribution. Derive the maximum likelihood estimator $\hat{\theta}$.
3. You are given 100 independent samples x_1, x_2, \dots, x_{100} from $\text{Ber}(p)$, where p is unknown. These 100 samples sum to 30. You would like to estimate the distribution's parameter p . Give all answers to 3 significant digits.
 - (a) What is the maximum likelihood estimator \hat{p} of p ?
 - (b) Is \hat{p} an unbiased estimator of p ?
 - (c) Give your best approximation for the 95% confidence interval of p .
 - (d) Give your best approximation for the 90% confidence interval of p .
 - (e) Give three different reasons why your answers to (c) and (d) are only approximations.
 - (f) Explain why it makes sense that the interval in (d) is bigger (or smaller, depending on your answers) than the interval in (c).