1. Suppose $X_1, X_2, ..., X_n$ is an iid sample from a gamma distribution with known shape parameter $\alpha_0 > 0$ and unknown scale parameter $1/\theta$; i.e., the population probability density function is

$$f_X(x|\theta) = \frac{\theta^{\alpha_0}}{\Gamma(\alpha_0)} x^{\alpha_0 - 1} e^{-\theta x} I(x > 0),$$

where $\theta > 0$.
(a) Argue that $T = \sum_{i=1}^{n} X_i$ is a complete and sufficient statistic.
(b) Find the uniformly minimum variance unbiased estimator (UMVUE) of $\text{var}_\theta(X_1)$, the population variance.
(c) Carefully argue that X_1/T and T are independent statistics.

2. Suppose $X_1, X_2, ..., X_n$ are iid exponential with mean $\beta > 0$. Experimenter 1 gets to observe $X_1, X_2, ..., X_n$. Experimenter 2 does not. Instead, she observes $Y_1, Y_2, ..., Y_n$, where

$$Y_i = I(X_i \leq c) = \begin{cases} 0, & X_i > c \\ 1, & X_i \leq c, \end{cases}$$

$i = 1, 2, ..., n$, where $I(\cdot)$ is the indicator function and $c > 0$ is known. The second situation is common in survival applications involving "current status data;" i.e., one does not observe the time to event X, but only whether or not (1/0) the event has occurred up through time c.
(a) Find the maximum likelihood estimator (MLE) of β for both Experimenter 1 and Experimenter 2. For Experimenter 1, the MLE $\widehat{\beta}_X$ should depend on $\mathbf{X} = (X_1, X_2, ..., X_n)$. For Experimenter 2, the MLE $\widehat{\beta}_Y$ should depend on $\mathbf{Y} = (Y_1, Y_2, ..., Y_n)$ and c.
(b) Without appealing to the large-sample properties of MLEs, show that

$$\sqrt{n}(\widehat{\beta}_X - \beta) \xrightarrow{d} \mathcal{N}(0, \sigma^2_{\widehat{\beta}_X})$$
$$\sqrt{n}(\widehat{\beta}_Y - \beta) \xrightarrow{d} \mathcal{N}(0, \sigma^2_{\widehat{\beta}_Y})$$

and calculate $\sigma^2_{\widehat{\beta}_X}$ and $\sigma^2_{\widehat{\beta}_Y}$.
(c) Find consistent estimators of $\sigma^2_{\widehat{\beta}_X}$ and $\sigma^2_{\widehat{\beta}_Y}$. Use these to create large-sample statistics to test $H_0 : \beta = \beta_0$ versus $H_1 : \beta \neq \beta_0$ under both sampling strategies.

3. Suppose that $X_1, X_2, ..., X_n$ are iid Poisson with mean $\theta_1 > 0$. Suppose $Y_1, Y_2, ..., Y_n$ are iid Poisson with mean $\theta_2 > 0$. Suppose the samples are independent. In this problem, we will consider the following hypothesis test:

$$H_0 : \theta_1 = \theta_2$$
versus
$$H_1 : \theta_1 \neq \theta_2.$$

(a) Sketch a graph of the unrestricted and restricted parameter spaces (put on the same graph).
(b) Derive $\lambda(\mathbf{x}, \mathbf{y})$, the likelihood ratio test statistic to test H_0 versus H_1. Give a decision rule (i.e., when you would reject H_0) in terms of simple statistics and simplify this rule as much as possible. You do not need to find specific critical values for your test, but describe how they could be found. Do not appeal to any large-sample arguments on this part.

(c) Calculate $-2\ln\lambda(\mathbf{x},\mathbf{y})$ and describe how one could test H_0 versus H_1 using large-sample results.

4. Suppose that $X_1, X_2, ..., X_n$ is an iid sample from

$$f_X(x|\theta) = \frac{a_0}{\theta^{a_0}} x^{a_0-1} I(0 < x < \theta),$$

where $a_0 \geq 1$ is a known constant and $\theta > 0$ is an unknown parameter.
(a) Show that $Q = X_{(n)}/\theta$ is a pivotal quantity, where $X_{(n)}$ is the maximum order statistic. Use Q to construct a $1-\alpha$ confidence set for θ.
(b) Derive a $1-\alpha$ confidence set for θ by using the cumulative distribution function of $T = T(\mathbf{X}) = X_{(n)}$ as a pivot.
(c) Assume that θ has an inverse gamma prior distribution with parameters $c > 0$ and $d > 0$ (both known). Write a $1-\alpha$ credible set for θ.

5. If $X \sim \mathcal{N}(0,\sigma^2)$, then $Y = |X|$ follows a half-normal distribution; i.e., the probability density function of Y is

$$f_Y(y|\sigma^2) = \frac{\sqrt{2}}{\sigma\sqrt{\pi}} \exp\left(-\frac{y^2}{2\sigma^2}\right) I(y > 0).$$

This population level model might arise, for example, if X measures some type of zero-mean difference (e.g., predicted outcome from actual outcome) and we are interested in absolute differences. Suppose that $Y_1, Y_2, ..., Y_n$ is an iid sample from $f_Y(y|\sigma^2)$.
(a) Derive the uniformly most powerful (UMP) level α test of

$$\begin{gathered} H_0 : \sigma^2 \geq \sigma_0^2 \\ \text{versus} \\ H_1 : \sigma^2 < \sigma_0^2. \end{gathered}$$

Identify all critical values associated with your test.
(b) Derive an expression for the power function $\beta(\sigma^2)$ of your UMP test in part (a).
(c) Suppose $\phi^*(\mathbf{x})$ is another level α test of H_0 versus H_1 with power function $\beta^*(\sigma^2) = E_{\sigma^2}[\phi^*(\mathbf{X})]$. Plot $\beta(\sigma^2)$ and $\beta^*(\sigma^2)$ on the same graph and describe the relationship between them. Be specific.

6. Suppose $X_1, X_2, ..., X_n$ is an iid $\mathcal{N}(\mu, c^2\mu^2)$ sample, where c^2 is known. Let $\widetilde{\mu}$ and $\widehat{\mu}$ denote the method of moments and maximum likelihood estimators of μ, respectively.
(a) Show that

$$\widetilde{\mu} = \overline{X} \qquad \text{and} \qquad \widehat{\mu} = \frac{\sqrt{\overline{X}^2 + 4c^2 m_2'} - \overline{X}}{2c^2},$$

where $m_2' = n^{-1}\sum_{i=1}^n X_i^2$ is the second sample (uncentered) moment.
(b) Prove that both estimators $\widetilde{\mu}$ and $\widehat{\mu}$ are consistent estimators.
(c) Show that $\sqrt{n}(\widetilde{\mu} - \mu) \xrightarrow{d} \mathcal{N}(0, \sigma^2_{\widetilde{\mu}})$ and $\sqrt{n}(\widehat{\mu} - \mu) \xrightarrow{d} \mathcal{N}(0, \sigma^2_{\widehat{\mu}})$. Calculate $\sigma^2_{\widetilde{\mu}}$ and $\sigma^2_{\widehat{\mu}}$. Which estimator is more efficient?