GROUND RULES:

- This exam contains 10 questions. Each question is worth 10 points. This exam is worth 100 points.
- Print your name **at the top of this page in the upper right hand corner**.
- This is a closed-book and closed-notes exam. You may use a calculator if you wish, but SHOW ALL OF YOUR WORK AND EXPLAIN ALL OF YOUR REASONING!!!
- Summary information on the discrete and continuous distributions is provided.
- Any discussion or otherwise inappropriate communication between examinees, as well as the appearance of any unnecessary material, will be dealt with severely.
- You have 3 hours to complete this exam. GOOD LUCK!

HONOR PLEDGE FOR THIS EXAM:

After you have finished the exam, please read the following statement and sign your name below it.

> *I promise that I did not discuss any aspect of this exam with anyone other than the instructor, that I neither gave nor received any unauthorized assistance on this exam, and that the work presented herein is entirely my own.*

1. Suppose that Y follows a Pareto distribution with cumulative distribution function

$$F_Y(y) = \begin{cases} 0, & y < \beta \\ 1 - \left(\frac{\beta}{y}\right)^{\alpha}, & y \geq \beta, \end{cases}$$

where $\alpha > 0$ and $\beta > 0$. Note that I am giving you the cdf here (not the pdf).
(a) Find the probability density function of $U = 1/Y$.
(b) Find $E(U)$.

2. Waiting times in a hospital emergency room follow an exponential distribution with mean $\theta > 0$ (measured in hours). Suppose that an iid sample of n times is observed, denoted by $Y_1, Y_2, ..., Y_n$.
(a) Show that $\overline{Y}$ is a sufficient statistic for θ.
(b) Derive the exact sampling distribution of $\overline{Y}$; i.e., do not use a normal approximation.
Hint: Find the moment generating function of $\overline{Y}$.

3. Suppose that $Y_1, Y_2, ..., Y_n$ is an iid sample of Bernoulli(p) random variables. Recall that the Bernoulli(p) model is the same as the $b(n, p)$ model with $n = 1$.
(a) Prove that $\overline{Y}$ is the maximum likelihood estimator of p.
(b) Find the maximum likelihood estimator of $\tau(p) = \log[p/(1-p)]$, the *log-odds* of p.

4. Suppose that $Y_1, Y_2, ..., Y_n$ is an iid sample from $f_Y(y;\theta)$, where

$$f_Y(y;\theta) = \begin{cases} 2y/\theta^2, & 0 < y < \theta \\ 0, & \text{otherwise.} \end{cases}$$

(a) Let $Y_{(n)}$ denote the maximum order statistic. Derive the pdf of $Y_{(n)}$.
(b) Consider using $Y_{(n)}$ as a point estimator for θ. Find $\text{MSE}(Y_{(n)})$.

5. Suppose that we have two independent samples:

$$\begin{aligned} &\text{Sample 1}: \quad Y_{11}, Y_{12}, ..., Y_{1n_1} \sim \text{iid } \mathcal{N}(\mu_1, \sigma^2) \\ &\text{Sample 2}: \quad Y_{21}, Y_{22}, ..., Y_{2n_2} \sim \text{iid } \mathcal{N}(\mu_2, \sigma^2) \end{aligned}$$

Note that the variance in each population distribution is the same. Let S_1^2 and S_2^2 denote the sample variances from Sample 1 and Sample 2, respectively.

(a) Prove that

$$S_p^2 = \frac{(n_1 - 1)S_1^2 + (n_2 - 1)S_2^2}{n_1 + n_2 - 2}.$$

is an unbiased estimator of σ^2.

(b) Find $V(S_p^2)$.

6. Suppose that $Y_1, Y_2, ..., Y_n$ is an iid sample from a $\mathcal{U}(\theta, 1)$ distribution.
(a) Show that $\widehat{\theta} = 2\overline{Y} - 1$ is the method of moments (MOM) estimator of θ.
(b) Show that the standard error of $\widehat{\theta}$ is

$$\sigma_{\widehat{\theta}} = \frac{1-\theta}{\sqrt{3n}}.$$

(c) Find an unbiased estimator of $\sigma_{\widehat{\theta}}$. Prove that your estimator is unbiased.

7. Suppose that $Y_1, Y_2, ..., Y_n$ is an iid sample from a Weibull distribution with $m = 3$ and $\alpha = 1/\theta$ so that the common pdf

$$f_Y(y;\theta) = \begin{cases} 3\theta y^2 e^{-\theta y^3}, & y > 0 \\ 0, & \text{otherwise.} \end{cases}$$

Find the minimum variance unbiased estimator (MVUE) for $1/\theta$.

8. Suppose that $Y_1, Y_2, ..., Y_n$ is an iid sample from a gamma distribution with shape parameter $\alpha = 2$ and scale parameter $\beta > 0$.
(a) Argue that

$$Q_n = \frac{\overline{Y} - 2\beta}{\sqrt{2\beta^2/n}} \xrightarrow{d} \mathcal{N}(0, 1),$$

as $n \to \infty$. Therefore, Q_n is an approximate (i.e., large-sample) pivot.
(b) Use the result in part (a) to show that

$$\left(\frac{\overline{Y}}{2 + z_{\alpha/2}\sqrt{2/n}}, \frac{\overline{Y}}{2 - z_{\alpha/2}\sqrt{2/n}} \right)$$

is an approximate $100(1 - \alpha)$ percent confidence interval for β. As usual, $z_{\alpha/2}$ denotes the upper $\alpha/2$ quantile from a standard normal distribution.

9. I have 4 statistics T_1, T_2, T_3, and T_4. I know that

- T_1, T_2, T_3, and T_4 are (mutually) independent.
- $T_1 \sim \chi^2(4)$.
- $T_2 \sim \chi^2(5)$.
- $T_3 \sim \mathcal{N}(1, 4)$.
- T_4 has the same distribution as $T_1 + T_2$.

(a) Show that the moment generating function of $Z = (T_3 - 1)/2$ is $m_Z(t) = \exp(t^2/2)$. What is the distribution of Z?
(b) Find a statistic that has a $\chi^2(1)$ distribution.
(c) Find a statistic that has an $F(9, 4)$ distribution.
(d) Find a statistic that has a $t(18)$ distribution.

10. Suppose $Y_1, Y_2, ..., Y_n$ is an iid sample of Poisson(θ) random variables, where $\theta > 0$.
(a) Argue that $\overline{Y}$ is a consistent estimator for θ.
(b) Find a consistent estimator of $g(\theta) = \log \theta$.
(c) It is known that S^2 is also a consistent estimator of θ (you do not have to prove this). Which estimator would you rather use to estimate θ, $\overline{Y}$ or S^2? Explain.