1. Let $Y \sim N(\mu, 1)$ and consider testing the hypothesis $H : \mu = 0$. Consider an acceptance region of the form $A(0) = (Y : z_{\alpha(1-w)} < Y < z_{1-\alpha w})$.

   (a) Show that the type I error rate of such a test is α for $w \in [0, 1]$.

   (b) Obtain the power of the test, that is, $\Pr(Y \notin A_0|\mu)$ as a function of μ. Make a plot of this power function for $w = 1/2$ and $w = 1/4$. When would you use $w = 1/2$? When would you use $w = 1/4$?

2. Let $Y \sim P_\theta$ for some P_θ with density $p(y|\theta) = e^{-y/\theta}/\theta$ for $y \in R^+$ and $\theta \in R^+$.

   (a) Find the CDF of P_θ.

   (b) Suppose we want to test the *composite* hypothesis $H : \theta < \theta_0$, where θ_0 is some known number. Suppose we will reject H if $Y > b$ for some value b. Find the value of b that makes this a level-α test, that is, find the value of b as a function of θ_0 and α so that the probability of rejection is less than or equal to α for each θ less than θ_0. Plot the power (probability of rejection) as a function of θ.

   (c) Construct a confidence region for θ of the form $(c(y), \infty)$ such that $\Pr(\theta \in (c(Y), \infty)|\theta) = 1 - \alpha$. Prove directly that your region has $1 - \alpha$ coverage.

3. Suppose treatment A is assigned to a random selection of 4 of 8 experimental units, and treatment B is assigned to the 4 remaining experimental units.

   (a) Assuming the A and B outcomes are random samples from $N(\mu_A, \sigma^2)$ and $N(\mu_B, \sigma^2)$ populations, what is the smallest p-value that one could obtain when using the usual two-sample t-statistic (assuming equal variances) to test $H : \mu_A = \mu_B$?

(b) Suppose you will do a randomization test of H : "No treatment effect", using the same test statistic as in (a) for your test statistic. What is the smallest p-value you could obtain?

(c) The observed treatment assignments and measured responses are

$$(X_1, \ldots, X_8) = (B, A, B, B, A, A, B, A)$$
$$(Y_1, \ldots, Y_8) = (7.5, 1.2, 7.5, 8.7, 3.2, 5.1, 6.2, 1.7).$$

Graphically compare the null distribution for the normal-theory test in (a) to the null distribution for the randomization test in (b). Compute the p-value for each and discuss similarities and differences. Describe the differences in assumptions that the two testing procedures make.

4. Let $Y_1, \ldots, Y_n \sim$ i.i.d. $N(\mu, \sigma^2)$ and let $S^2 = \sum(Y_i - \bar{Y})^2/(n-1)$ be the sample variance. Recall that $S^2 \overset{d}{=} \sigma^2 X^2/(n-1)$ where X^2 has a χ^2_{n-1} distribution.

(a) Construct a level-α hypothesis test for evaluating $H : \sigma^2 = \sigma_0^2$.

(b) Construct a $1 - \alpha$ confidence interval for σ^2.