STAT 513 hw 2

1. Suppose $X_1, \dots, X_n$ is a random sample from the Normal(μ, σ^2) distribution, where μ is unknown but σ^2 is known, and it is of interest to test H_0: $\mu = \mu_0$ versus H_1: $\mu \neq \mu_0$ for some value μ_0. The R code below plots the power curve of the test

$$\text{Reject } H_0 \text{ iff } |\sqrt{n}(\bar{X}_n - \mu_0)/\sigma| > z_{\alpha/2}$$

for user-selected values of μ_0, n, σ, and α. For a sequence of values of μ, the code computes the probability that the null hypothesis will be rejected according to the above test. In addition, for each value of μ in the sequence, a simulation is run: 100 data sets with sample size n are generated from the Normal(μ, σ^2) distribution, and for each of the 100 data sets, it is recorded whether the null hypothesis was rejected. For each value of μ, the proportion of times the null hypothesis is rejected is recorded. This gets plotted as a dashed line.

```
mu.0 <- ???
n <- ???
sigma <- ???
alpha <- ???

mu.seq <- seq(mu.0 - 5,mu.0 + 5,length=50)

z_crit <- qnorm(1-alpha/2)
power.theoretical <- 1-(pnorm(z_crit-sqrt(n)*(mu.seq - mu.0)/sigma)
                    -pnorm(-z_crit-sqrt(n)*(mu.seq - mu.0)/sigma))

power.empirical <- numeric()
for(j in 1:length(mu.seq))
{
    reject <- numeric()
    for(s in 1:100)
    {
                    x <- rnorm(n,mu.seq[j],sigma)
                    x.bar <- mean(x)

                    reject[s] <- abs(sqrt(n)*(x.bar-mu.0)/sigma) > z_crit
    }
    power.empirical[j] <- mean(reject)
}

plot(mu.seq,power.theoretical,type="l",ylim=c(0,1),xlab="mu",ylab="power")
lines(mu.seq, power.empirical,lty=2)
abline(v=mu.0,lty=3)  # vert line at null value
abline(h=alpha,lty=3) # horiz line at size
```

(a) Put in $\mu_0 = 2$, $n = 5$, $\sigma = 2$, and $\alpha = 0.05$ and execute the code. Turn in the plot.

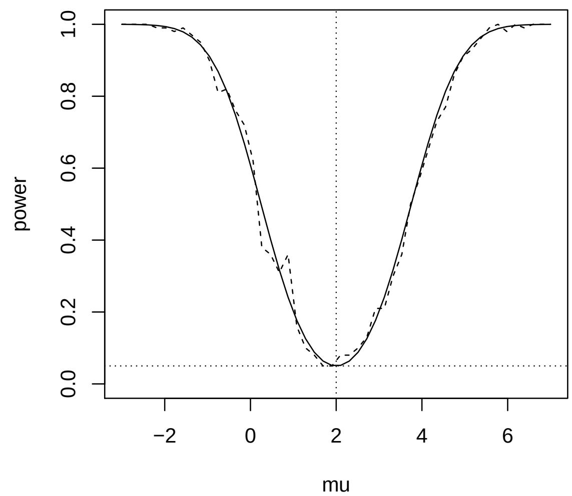

(b) Explain why the dashed line follows the solid line closely but not exactly.

The dashed line is the result of a randomized simulation.

(c) Interpret the height of the solid line at $\mu = 4$.

This is the probability, if the true mean is 4, that the test will reject H_0.

(d) Interpret the height of the solid line at $\mu = 2$.

The is the probability, if the true mean is 2, that the test will reject H_0. Since H_0: $\mu = 2$, this height represents the size of the test.

(e) Interpret the height of the dashed line at $\mu = 2$.

This is the proportion of times out of the 100 simulated data sets in which a Type I error was made.

(f) What would be the effect on the height of the solid line at $\mu = 4$ if

i. the sample size n were increased?

The height would increase.

ii. the standard deviation σ were increased?

The height would decrease.

iii. the size α of the test were increased?

The height would increase.

(g) What would be the effect on the height of the solid line at $\mu = 2$ if

i. the sample size n were increased?

The height would stay the same.

ii. the standard deviation σ were increased?

The height would stay the same.

iii. the size α of the test were increased?

The height would increase.

(h) What would be the effect on the dashed line of generating 500 data sets instead of only 100 data sets for the simulation at each value of μ?

The dashed line would more closely follow the solid line.

2. Suppose $X_1, \dots, X_n$ is a random sample from the Normal(μ, σ^2) distribution, where μ and σ^2 are unknown, and it is of interest to test H_0: $\mu = \mu_0$ versus H_1: $\mu \neq \mu_0$ for some value μ_0. Consider the test

$$\text{Reject } H_0 \text{ iff } |\sqrt{n}(\bar{X}_n - \mu_0)/S_n| > t_{n-1,\alpha/2}.$$

(a) Modify the code in Question 1 so that it displays the true power curve and a simulated power curve of this test. Run your modified code with $\mu_0 = 2$, $n = 5$, $\sigma = 2$, and $\alpha = 0.05$. Turn in your code and the resulting plot. *Hint: Refer to Lec 02. You will need to specify a noncentrality parameter for the t distribution.*

The code and the plot are:

```
mu.0 <- 2
mu.seq <- seq(mu.0 - 5,mu.0 + 5,length=50)
n <- 5
sigma <- 2
alpha <- 0.05

t_crit <- qt(1-alpha/2,n-1)
```

```
power.theoretical[j] <- 1-(pt(t_crit,n-1,ncp = sqrt(n)*(mu.seq[j] - mu.0)/sigma)
                     -pt(-t_crit,n-1,ncp = sqrt(n)*(mu.seq[j] - mu.0)/sigma))

power.empirical <- numeric()
for(j in 1:length(mu.seq))
{

    reject <- numeric()
    for(s in 1:100)
    {
            x <- rnorm(n,mu.seq[j],sigma)
            x.bar <- mean(x)

            reject[s] <- abs(sqrt(n)*(x.bar-mu.0)/sd(x) ) > t_crit
    }
    power.empirical[j] <- mean(reject)
}

plot(mu.seq,power.theoretical,type="l",ylim=c(0,1),xlab="mu",ylab="power")
lines(mu.seq, power.empirical,lty=2)
abline(v=mu.0,lty=3)  # vert line at null value
abline(h=alpha,lty=3) # horiz line at size
```

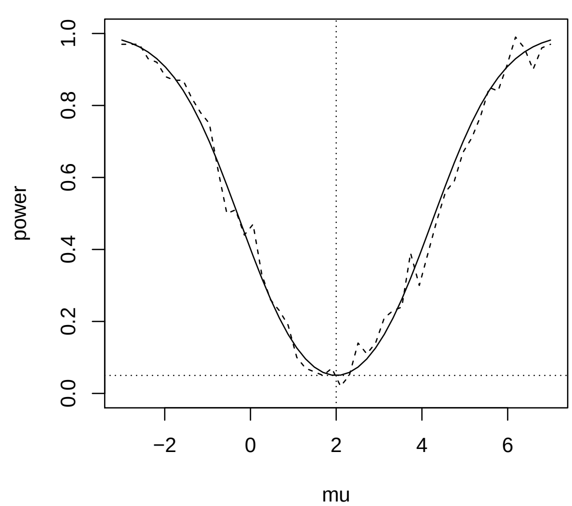

(b) Compare the height of the solid curve at $\mu = 4$ with that of the solid curve in Question 1 at $\mu = 4$. Comment on whether there is a difference and why/why not.

The power curve for the t-test is lower at $\mu = 4$ than that of the Z-test. This is the price we pay for having to estimate the variance.

(c) Compare the height of the solid curve at $\mu = 2$ with that of the solid curve in Question 1 at $\mu = 2$. Comment on whether there is a difference and why/why not.

The height is the same; both tests are calibrated to have size equal to 0.05.

3. Suppose the values

```
1.39 2.22 2.38 1.60 1.50
```

are a random sample from a distribution assumed to be Normal but for which the mean and variance are unknown.

(a) Give a 95% confidence interval for the mean.

We use $\bar{X}_5 \pm t_{4,.025} S_5/\sqrt{5}$, which we can compute with the code

```
x <- c(1.39, 2.22, 2.38, 1.60, 1.50)
c(mean(x) - qt(.975,4)*sd(x)/sqrt(5),mean(x) + qt(.975,4)*sd(x)/sqrt(5))
```

This gives $(1.2595, 2.3765)$.

(b) Test the hypotheses H_0: $\mu = 2.5$ versus H_1: $\mu \neq 2.5$ at the 0.05 significance level.

Since $2.5 \notin (1.2595, 2.3765)$, which is a 95% confidence interval, we reject H_0 at the 5% significance level.

(c) Give the p-values for testing the following hypotheses:

i. H_0: $\mu = 2.5$ versus H_1: $\mu \neq 2.5$

A sample which carries as much or more evidence against H_0 than the observed sample will have $|\sqrt{5}(\bar{X}_5 - 2.5)/S_5| >$ `abs(sqrt(5)*(mean(x) - 2.5)/sd(x))` $= 3.390392$. So the p-value is

$$\begin{aligned} P_{\mu=2.5}(|\sqrt{n}(\bar{X}_5 - 2.5)/S_5| > 3.390392) &= P(|T| > 3.390392), \quad T \sim t_4 \\ &= 2(1 - P(T < 3.390392)) \\ &= 2\texttt{*(1 - pt(3.390392,4))} \\ &= 0.02752004. \end{aligned}$$

ii. H_0: $\mu \leq 2.5$ versus H_1: $\mu > 2.5$

A sample which carries as much or more evidence against H_0 than the observed sample will have $\sqrt{5}(\bar{X}_5 - 2.5)/S_5 >$ `sqrt(5)*(mean(x) - 2.5)/sd(x)` $= -3.390392$. So the p-value is

$$\begin{aligned} P_{\mu=2.5}(\sqrt{n}(\bar{X}_5 - 2.5)/S_5 > -3.390392) &= P(T > -3.390392), \quad T \sim t_4 \\ &= P(T < 3.390392) \\ &= \texttt{pt(3.390392,4)} \\ &= 0.98624. \end{aligned}$$

iii. H_0: $\mu \geq 2.5$ versus H_1: $\mu < 2.5$

A sample which carries as much or more evidence against H_0 than the observed sample will have $\sqrt{5}(\bar{X}_5 - 2.5)/S_5 <$ `sqrt(5)*(mean(x) - 2.5)/sd(x)` $= -3.390392$. So the p-value is

$$\begin{aligned} P_{\mu=2.5}(\sqrt{n}(\bar{X}_5 - 2.5)/S_5 < -3.390392) &= P(T < -3.390392), \quad T \sim t_4 \\ &= \texttt{pt(-3.390392,4)} \\ &= 0.01376002. \end{aligned}$$

iv. State whether a 99% confidence interval for μ based on this sample would contain the value 2.5.

Yes, it would; the p-value for testing H_0: $\mu = 2.5$ versus H_1: $\mu \neq 2.5$ is greater that 0.01, so we would fail to reject the null hypotheses at the 0.01 significance level. Therefore the 99% confidence interval would contain the value 2.5.

4. Suppose the values

`1.39 2.22 2.38 1.60 1.50`

are a random sample from a distribution assumed to be Normal but for which the mean and variance are unknown.

(a) Report the estimated variance S_5^2.

The code

```
x <- c(1.39, 2.22, 2.38, 1.60, 1.50)
var(x)
```

gives $S_5^2 = 0.20232$.

(b) Give a 95% confidence interval for the unknown variance σ^2.

Use $(4S_5^2/\chi^2_{4,.025}, 4S_5^2/\chi^2_{4,.975})$, which we can compute with

```
x <- c(1.39, 2.22, 2.38, 1.60, 1.50)
c(4*var(x)/qchisq(.975,4),4*var(x)/qchisq(.025,4))
```

This gives $(0.07262489, 1.67062138)$.

(c) Based on your answer to part (b), what is your conclusion about the hypotheses H_0: $\sigma^2 = 1.5$ versus H_1: $\sigma^2 \neq 1.5$ at the 0.05 significance level?

We fail to reject H_0 at the 0.05 significance level because 1.5 is contained in the 95% confidence interval.

(d) Based on your answer to part (b), what is your conclusion about the hypotheses H_0: $\sigma^2 = 1.5$ versus H_1: $\sigma^2 \neq 1.5$ at the 0.01 significance level?

If we were to build a 99% confidence interval for σ^2 it would be wider than the 95% confidence interval; since the 95% confidence interval contains 1.5, so would the 99% confidence interval. Therefore we would fail to reject H_0: $\sigma^2 = 1.5$ at the 0.01 significance level because 1.5 would be contained in the 99% confidence interval.

(e) Give the p-values for testing the following hypotheses:

i. H_0: $\sigma^2 \leq 1.5$ versus H_1: $\sigma^2 > 1.5$

A sample which carries as much or more evidence against H_0 than the observed sample will have $4S_5^2/1.5 >$ `4*var(x)/1.5` $= 0.53952$. So the p-value is

$$\begin{aligned} P_{\sigma^2=1.5}(4S_5^2/1.5 > 0.53952) &= P(W > 0.53952), \quad W \sim \chi^2_4 \\ &= 1 - P(W < 0.53952) \\ &= \texttt{1 - pchisq(0.53952,4)} \\ &= 0.9695414. \end{aligned}$$

ii. H_0: $\sigma^2 \geq 1.5$ versus H_1: $\sigma^2 < 1.5$

A sample which carries as much or more evidence against H_0 than the observed sample will have $4S_5^2/1.5 <$ `4*var(x)/1.5` $= 0.53952$. So the p-value is

$$\begin{aligned} P_{\sigma^2=1.5}(4S_5^2/1.5 < 0.53952) &= P(W < 0.53952), \quad W \sim \chi^2_4 \\ &= \texttt{pchisq(0.53952,4)} \\ &= 0.03045859. \end{aligned}$$

iii. H_0: $\sigma^2 = 1.5$ versus H_1: $\sigma^2 \neq 1.5$

To compute the two-sided p-value, we compute the one-sided p-value in the direction suggested by the data and multiply it by 2. It is therefore

`2*0.03045859 = 0.06091718.`

5. You are importing a large number of small manufactured items and you want to know if more than 5% of them are defective. You randomly sample items one-by-one and determine whether each is defective or not. Let n be the number of items you sample and let Y be the number of defective items you discover out of the n items sampled. You decide to conclude that more that 5% are defective if $Y/n \geq 0.05 + 2\sqrt{(0.05)(0.95)/n}$.

   (a) State the relevant null and alternate hypotheses.

   If p is the true proportion of items that are defective, we are interested in

   $$H_0\text{: } p \leq 0.05 \text{ versus } H_1\text{: } p > 0.05.$$

   (b) Suppose that the true proportion of defective items is 0.04 and that you discover two defective items out of $n = 10$ sampled items. Do you make a correct decision, a Type I error, or a Type II error?

   For $n = 10$, $0.05 + 2\sqrt{(0.05)(0.95)/n} = 0.1878405$ and $Y/n = 0.2$, so we will reject H_0, but this will be a Type I error, since H_0: $p \leq 0.05$ is true.

   (c) Suppose that the true proportion of defective items is 0.11 and you decide to sample $n = 10$ items. What is the correct conclusion, and what is the probability that you will come to the correct conclusion?

   The correct conclusion is to reject H_0. The probability of this is

   $$\begin{aligned} P_{p=0.11}(Y/10 \geq 0.1878405) &= P_{p=0.11}(Y \geq 1.878405) \\ &= P_{p=0.11}(Y \geq 2) \\ &= 1 - P_{p=0.11}(Y \leq 1) \\ &= 1 \texttt{ - pbinom(1,10,.11)} \\ &= 0.3027908. \end{aligned}$$

   (d) If you sample $n = 100$ items and the true proportion of defective items is 0.05, what is the probability that you will make a Type I error?

If $n = 100$, $0.05 + 2\sqrt{(0.05)(0.95)/n} = 0.09358899$, and the probability of rejecting H_0 is

$$\begin{aligned} P_{p=0.05}(Y/100 \geq 0.09358899) &= P_{p=0.05}(Y \geq 9.358899) \\ &= P_{p=0.05}(Y \geq 10) \\ &= 1 - P_{p=0.05}(Y \leq 9) \\ &= 1 \texttt{ - pbinom(9,100,.05)} \\ &= 0.02818829. \end{aligned}$$