MS Module 7 Hypothesis testing – p values practice problems

(The attached PDF file has better formatting.)

Exercise 7.1: Proportions

The average incidence of resistance to a disease in a population is 25%. A study tests whether a drug raises the incidence of resistance: 5,000 subjects are given the drug and 1,300 show resistance to the disease. Let μ_0 = the percent of subjects who show resistance to the disease. The null hypothesis assumes μ = 25% (μ_0).

- The null hypothesis is H_0: the population mean $\mu = \mu_0$ = 25%
- The alternative hypothesis takes one of two forms:
  - The two-tailed alternative hypothesis is H_a: the population mean $\mu \neq \mu_0$ (25%)
  - The one-tailed (upper tailed) alternative hypothesis is H_a: the population mean $\mu > \mu_0$ (25%)

A. What is the incidence of resistance in the sample of 5,000 subjects?
B. What is the standard deviation of the incidence of resistance in the sample?
C. What is the z value used to test the null hypothesis?
D. What is the p value for the one-tailed alternative hypothesis?
E. What is the p value for the two-tailed alternative hypothesis?

Part A: The observed incidence of resistance in the sample is 1,300 / 5,000 = 26%.

Part B: The standard deviation of the incidence of resistance is $(p\,(1 - p) / N)^{½}$ =

$$(25\% \times (1 - 25\%) / 5{,}000)^{0.5} = 0.0061237$$

Question: Why do we use the assumed incidence of resistance in the population for the standard deviation? Why not use the observed incidence of resistance in the sample?

Answer: We are testing the null hypothesis to see the likelihood of observing the sample incidence; that is

> If the null hypothesis is true and the incidence of resistance in the sample is 25%, what is the probability of observing an incidence of 26% in the sample?

Part C: The z value to test the null hypothesis is (26% – 25%) / 0.0061237 = 1.633000

Part D: The p value for the one-tailed alternative hypothesis is $\Phi(-1.633) = 0.0512$.

Part E: The p value for the two-tailed alternative hypothesis is $2 \times \Phi(-1.633) = 0.1025$.

Answer: Why do we use the normal distribution for a percentage? Shouldn't we use a binomial distribution?

Answer: With 5,000 subjects of whom 1,300 are resistant, the central limit theorem says that the incidence of resistance is approximately normally distributed

Exercise 7.2: Power of a test

[This exercise helps you understand the reading in the text. You will not be asked to calculate powers on the final exam.]

Subjects in an untreated population have reaction times that are normally distributed with a mean of 80 and a standard deviation of 12.

To test if a new treatment reduces the reaction time,

- The null hypothesis is H_0: $\mu = 80$.
- The alternative hypothesis is H_a: $\mu < 80$.

The treatment does not change the standard deviation.

The sample has 100 observations and the significance level is 1%. Let $\bar{x}$ be the mean of the sample values.

A. What is the *z* value to test the null hypothesis as a function of $\bar{x}$?
B. What is the rejection region for the null hypothesis?
C. What is the upper bound of the rejection region?
D. If the true mean is 74, what is the power of the test?
E. How does the power of the test relate to the difference of the mean μ' from the μ_0 in the null hypothesis?
F. How does the power of the test relate to the standard deviation σ?
G. How does the power of the test relate to the significance level α?
H. How does the power of the test relate to the number of observations in the sample?

Part A: The *z* value is $(\bar{x} - \mu_0) / (\sigma / \sqrt{n}) = (\bar{x} - 80) / (12 / \sqrt{100}) = (\bar{x} - 80) / 1.2$

Part B: A one-sided lower tailed test with significance level α = 1% rejects the null hypothesis if $z \le \Phi(-0.01)$, where $z = (\bar{x} - 80) / (12 / \sqrt{100}) = (\bar{x} - 80) / 1.20$.

Part C: We solve for the upper bound of the rejection region:

$$(\bar{x} - 80) / 1.20 = -2.32635 \Rightarrow$$
$$\bar{x} = -2.32635 \times 1.2 + 80 = 77.2084$$

(A one-sided lower-tailed rejection region has no lower bound.)

Part D: The power of the test is the probability of rejecting the null hypothesis for a given value of the true μ.

For a true μ of 74, this probability is $\Phi($ (upper bound $- \mu'$) $/ (\sigma / \sqrt{n})$) =

$$\Phi(\ (77.2084 - 74) / 1.2\ ) = \Phi(2.67365) = 0.996248.$$

Part E: The power (the likelihood of rejecting the null hypothesis) increases as the as the mean μ' becomes farther away from μ_0, the mean in the null hypothesis.

Part F: As the standard deviation σ increases, the power of the test decreases, since random fluctuations are more likely to cause the sample mean to differ from the μ_0 in the null hypothesis and we are less likely to reject the null hypothesis.

Part G: As α decreases (the test becomes more stringent), the confidence interval around the μ_0 in the null hypothesis becomes wider, we are less likely to reject the null hypothesis and the power of the test decreases.

Part H: As the number of observations in the sample increases, the power of the test increases, since random fluctuations are smoothed.