Single Sample Inferences

Greg C Elvers

1

Single Samples

- A *single sample* implies that you have collected data from one group of people or objects
- You have not collected data from a comparison, or control, group
- Rather, you will compare your data to pre-existing data, perhaps from the census or other archive

2

Single Sample Inferences

- The basic questions that is asked is:
- Is the sample mean different from a pre-existing population mean?
  - E.g. Is the mean IQ of the students in this class different from the mean IQ of people in general?
  - E.g. Is the mean number of cigarettes smoked per hour higher in a sample of people with schizophrenia than in the population in general?

3

Steps To Follow

- When trying to answer these types of questions, there are several steps that you should follow:
  - Write the null and alternative hypotheses
    - Are they one or two tailed?
  - Specify the α level (usually .05)
  - Calculate the appropriate test statistic
  - Determine the critical value from a table
  - Decide whether to reject H_0 or not

4

Write the Hypotheses

- Lets consider the first example:
- The mean IQ of the people in a statistics class is 103. Is this value different from the population mean (100)?
- H_0: $\mu = 100$
  H_1: $\mu \neq 100$

1 vs 2 Tailed?

- The hypothesis does not state whether the sample mean should be larger (or smaller) than the population mean
- It only states that the sample mean should be different from the population mean
- Thus, this should be a two-tailed test

Specify the α Level

- The α level is the probability of making a Type-I error
- The α level specifies how willing we are to reject H_0 when in fact H_0 is true
- While α can take on any value between 0 and 1 inclusive, psychologists usually adopt an α level of either .05, .01, or .001
  - .05 is the most common
- $\alpha = .05$

Calculate the Appropriate Test Statistic

- First, you must decide what the appropriate test statistic is
- If the mean and standard deviation of the population are known, and the sampling distribution is normally distributed, then the appropriate test statistic is the *z-score* for the sampling distribution

Calculate the Appropriate Test Statistic

- IQs are normally distributed with a mean of 100, and a standard deviation of 15
- Thus, we are safe in using the z-score of the sampling distribution

z-scores of the Sampling Distribution

- The standard error of the mean is given by the population standard deviation (σ = 15) divided by the square root of the sample size (n = 225)
- $s_{\overline{X}}$ = 15 / √225 = 1

$$z = \frac{\overline{X} - \mu}{s_{\overline{X}}}$$

$$s_{\overline{X}} = \frac{\sigma}{\sqrt{n}}$$

z-scores of the Sampling Distribution

- The z-score is the difference of the sample and population means divided by the standard error of the mean
- z = (103 - 100) / 1
- z = 3

$$z = \frac{\overline{X} - \mu}{s_{\overline{X}}}$$

$$s_{\overline{X}} = \frac{\sigma}{\sqrt{n}}$$

Determine the Critical Value

- There are two ways of determining the critical value
- One way is used when calculating the statistic by hand
- The other way is used when calculating the statistic with statistical software such as SPSS, SAS, or BMDP

Determining the Critical Value by Hand (One-Tailed)

- When determining the critical value by hand, you want to determine the z-score beyond which is your α level
- The diagram to the right shows this for a one-tailed test

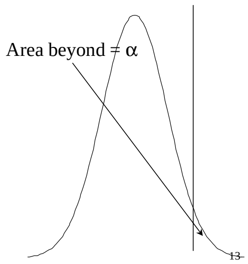

Determining the Critical Value by Hand (One-Tailed)

- In this example, find the z-score whose area above the z-score equals α (.05)
- Consult the table of areas under the unit normal curve
- A z-score of 1.65 has an area of .05 above it; 1.65 is our critical z

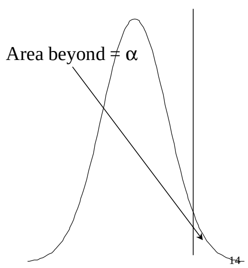

Determining the Critical Value by Hand (Two-Tailed)

- When determining the critical value by hand, you want to determine the z-score beyond which is your α level
- The diagram to the right shows this for a two-tailed test

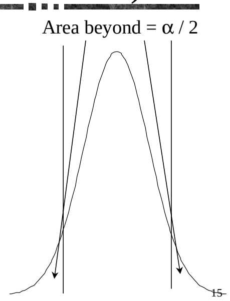

Determining the Critical Value by Hand

- In this example, find the z-score whose area above the z-score equals .5 X α (.025)
- Consult the table of areas under the unit normal curve
- A z-score of 1.96 has an area of .025 above it; 1.96 is our critical z

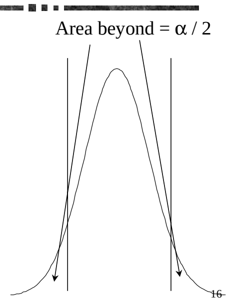

Decide Whether to Reject H_0

- When the absolute value of the observed z is larger than the critical z, you can reject H_0
  - That is, when | observed | > critical, the sample is different from the population
  - This is the 2-tailed rule; 1-tailed rule is slightly different
- Observed z = 3
- Critical z (two tailed) = 1.96
- Reject H_0
- The sample is probably different from the population

Deciding To Reject H_0 when Using the Computer

- When you use SPSS or similar software, the program will print the observed statistic and the probability that of observing a sample that large due to chance
- The probability is called the *p* value
- When the p value is less than or equal to α, you can reject H_0
- Thus, the sample is probably different from the population

Problem

- A professor gives a test in statistics. Based on the 81 students who took the test, the class average on the test is 75. From students who took the test in previous classes, the professor knows that the mean grade is 80 with a standard deviation of 27. Is the current class performing more poorly than the average?

Student's t Test

- When the population mean and / or standard deviation are not known, a different inferential statistical procedure should be used: *Student's t test*
- Student's t, or just t, test is, conceptually, very similar to the z-score test we have been using
- The t test is used to determine if a sample is different from the population

The t Test

- When the standard deviation of the population is not known, as is usually the case, we must estimate the standard deviation of the population
- We use the standard deviation of the sample to estimate the population standard deviation:

$$\sigma \approx s$$

Sample and Population Standard Deviations

- The sample standard deviation consistently underestimates the value of the population standard deviation
  - It is *biased*
- An unbiased estimate of the population standard deviation is given by:

$$\hat{s} = \sqrt{\frac{n}{n-1}s^2}$$

Sample and Population Standard Deviations

- Even the unbiased estimate of the population standard deviation will be inexact when the sample size is small (< 30)
- The smaller the sample size is, the less precise the unbiased estimate of the population standard deviation will be
- Because of this imprecision, it is inappropriate to use the normal distribution

The t Distributions

- William Gossett created a series of distributions known as the *t distributions*
- The t distributions are similar to the unit normal distributions, but account for the imprecision in the estimation of the population mean
- Because the imprecision in the estimate depends on sample size, there are multiple t distributions, depending on the *degrees of freedom*

Degrees of Freedom

- *Degrees of freedom* correspond to the number of scores that are free to take on any value after restrictions are placed on the set of data
- E.g., if the mean of 5 data points is 0, then how many data points can take on any value and still have the mean equal 0?

25

Degrees of Freedom

- 4 of the 5 numbers can take on any value
- But the fifth number must equal -1 times the sum of the other four for the mean to equal 0
- Thus n - 1 scores are free to vary
- In this case df = n - 1

1	X_1
2	X_2
3	X_3
4	X_4
5	$-(X_1 + X_2 + X_3 + X_4)$
Mean	0

26

The t Test

- The t test is used to decide if a sample is different from a population when the population standard deviation is unknown

$$t = \frac{\bar{X} - \mu_0}{s_{\bar{X}}}$$

$$s_{\bar{X}} = \frac{\hat{s}}{\sqrt{n}}$$

$$\hat{s} = \sqrt{\frac{n}{n-1} s^2}$$

27

Example

- On average, do people with schizophrenia smoke more cigarettes ($\bar{X}$ = 9 per day) than the population (μ_0 = 6 per day)
- Step 1: Write the hypotheses:
- H_0: $\mu \leq 6$
  H_1: $\mu > 6$

28

1 vs 2 Tailed? What is α?

- The hypothesis asks if people with schizophrenia smoke *more* cigarettes than average; thus we have a 1 tailed test
- We will adopt our standard α level, .05
  - α is the probability of making a Type - I error
  - α is the probability of rejecting H_0 when H_0 is true

29

Calculate the Appropriate Test Statistic

- Because we do not know the population standard deviation, we will estimate it from the sample standard deviation
- s = 5, n = 10

$$\hat{s} = \sqrt{\frac{n}{n-1}s^2} = \sqrt{\frac{10}{10-1}5^2} = 5.27$$

$$s_{\bar{x}} = \frac{\hat{s}}{\sqrt{n}} = \frac{5.27}{\sqrt{10}} = 1.67$$

30

Calculate the Appropriate Test Statistic

- Plug and chug the t test value
- Determine the degrees of freedom:
- df = n - 1 = 10 - 1 = 9

$$t_{obs} = \frac{\bar{X} - \mu}{s_{\bar{X}}} = \frac{9-6}{1.67} = 1.80$$

31

Determine the Critical Value

- To determine the critical t value, consult a table of critical t values
- Find the column that is labeled with your α level
  - Make sure you select the right number of tails (1 vs 2)
- Find the row that is labeled with your degrees of freedom
- The critical t value is at the intersection

32

Determine the Critical Value

- If the table does not contain the desired degrees of freedom, use the critical t value for the *next smallest* degrees of freedom
- With $\alpha = .05$, one tailed, and df = 9, the critical t value is 1.833

Decide Whether to Reject H_0

- If the observed t (the value you calculated) is larger than the critical t, then you can reject H_0
- Because our observed t (1.80) is not larger than the critical t (1.833), we fail to reject H_0 that people with schizophrenia smoke less than or equal to the population

Decide Whether to Reject H_0

- That is, there is no statistically reliable difference in the average number of cigarettes smoked by the population and by people with schizophrenia
- This does not claim that there is no difference, but rather that we failed to observe the difference if it did exist

Problem

- Do indoor cats weigh a different amount than outdoor cats which weigh an average of 11 pounds?
- $\overline{X}_{indoor} = 13$ pounds
- $s_{indoor} = 3.75$ pounds
- $n_{indoor} = 16$

Other Uses of t

- The t distributions can also be used to determine if a correlation coefficient is probably different from 0.
- The correlation between your scores on the first and second exam is r = .3786, n = 14
- Is this correlation probably different from 0?

Write Hypotheses; Specify α

- H_0: $\rho = 0$
  H_1: $\rho \neq 0$
- This is a two tailed hypothesis as it does not state whether the correlation should be positive or negative
- $\alpha = .05$

Calculate the Appropriate Statistic

- The formula for t given r and n is to the right:
- The degrees of freedom is the number of pairs of scores minus 2

$$t = \frac{r\sqrt{n-2}}{\sqrt{1-r^2}}$$

$$df = n-2$$

Calculate the Appropriate Statistic

- Plug and chug the t formula
- Calculate the degrees of freedom

$$t = \frac{r\sqrt{n-2}}{\sqrt{1-r^2}} = \frac{.3786\sqrt{14-2}}{\sqrt{1-.3786^2}} = 1.42$$

$$df = n-2 = 14-2 = 12$$

Determine the Critical t Value

- Consult a table to determine the critical t with $\alpha = .05$, two-tailed, and df = 12
- The critical t value is 2.179

Decide Whether to Reject H_0

- If the observed t (the calculated value) is larger than the critical t, we can reject H_0 that the correlation does not exist
- The observed t (1.42) is not larger than the critical t (2.179), so we fail to reject H_0
- This does not imply that a relation does not exist, but rather that we failed to observe it

Problem

- The correlation between how many cats you own and how introverted you are is r = 0.6 (made-up)
- The sample size was 102
- Is this correlation reliably different from $\rho = 0$?