

Example - GSS

The General Social Survey (GSS) conducted by the Census Bureau contains a standard 'core' of demographic, behavioral, and attitudinal questions, plus topics of special interest. Many of the core questions have remained unchanged since 1972 to facilitate time-trend studies as well as replication of earlier findings. Below is an excerpt from the 2010 data set. The variables are number of hours worked per week and highest educational attainment.

	degree	hrs1
1	BACHELOR	55
2	BACHELOR	45
3	JUNIOR COLLEGE	45
⋮		
1172	HIGH SCHOOL	40

Lecture 12 - Tests of Two Means

Sta102 / BME 102

Colin Rundel

October 8th, 2014


Sta102 / BME 102 (Colin Rundel)

Lec 12

October 8th, 2014 2 / 33

Difference of two means Confidence intervals for differences of means

Exploratory analysis



What can we say about the relationship between educational attainment and hours worked per week?

Sta102 / BME 102 (Colin Rundel)

Lec 12

October 8th, 2014 3 / 33

Difference of two means Confidence intervals for differences of means

Collapsing levels

- Say we are only interested the difference between the number of hours worked per week by college and non-college graduates.
- We can combine the levels of education into:
 - hs or lower ← less than high school or high school
 - coll or higher ← junior college, bachelor's, and graduate
- Here is how you can do this in R:

```
# create a new empty variable
gss$edu = NA

# if statements to determine levels of new variable
gss$edu[gss$degree == "LESS THAN HIGH SCHOOL" |
  gss$degree == "HIGH SCHOOL"] = "hs or lower"
gss$edu[gss$degree == "JUNIOR COLLEGE" |
  gss$degree == "BACHELOR" |
  gss$degree == "GRADUATE"] = "coll or higher"

# make sure new variable is categorical
gss$edu = as.factor(gss$edu)
```


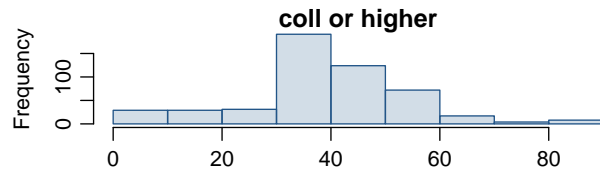
Sta102 / BME 102 (Colin Rundel)

Lec 12

October 8th, 2014 4 / 33

Exploratory analysis - another look

	\bar{x}	s	n
coll or higher	41.8	15.14	505
hs or lower	39.4	15.12	667



Parameter and point estimate

We want to construct a 95% confidence interval for the average difference between the number of hours worked per week by Americans with a college degree and those with a high school degree or lower. What are the parameter of interest and the point estimate?

- **Parameter of interest:** Average difference between the number of hours worked per week by *all* Americans with a college degree and those with a high school degree or lower.

$$\mu_c - \mu_{hs}$$

- **Point estimate:** Average difference between the number of hours worked per week by *sampled* Americans with a college degree and those with a high school degree or lower.

$$\bar{x}_c - \bar{x}_{hs}$$

Difference of Means and the CLT

We can think about our observations as being samples from two distributions D_x and D_y ,

$$X_1, X_2, \dots, X_{n_x} \sim D_x$$

$$Y_1, Y_2, \dots, Y_{n_y} \sim D_y.$$

We now want to know what the distribution of $\bar{x} - \bar{y}$ will be so that we can perform inference.

From our work with a single sample means, we know that the CLT tells us that both

$$\bar{x} \sim N(E(D_x), \text{Var}(D_x)/n_x),$$

$$\bar{y} \sim N(E(D_y), \text{Var}(D_y)/n_y),$$

Difference of Means and the CLT (cont.)

Proposition - the sum or difference of Normal RVs is always another Normal RV. (Not terribly hard to prove, but requires more probability theory than we've covered).

This proposition then tells us that $\bar{x} - \bar{y} \sim N(E(\bar{x} - \bar{y}), \text{Var}(\bar{x} - \bar{y}))$. Where,

$$\bar{x} - \bar{y} = \frac{1}{n_x} \sum_{i=1}^{n_x} x_i - \frac{1}{n_y} \sum_{i=1}^{n_y} y_i$$

Expected Value of the Difference

$$\begin{aligned}
 E(\bar{x} - \bar{y}) &= E\left(\frac{1}{n_x} \sum_{i=1}^{n_x} x_i - \frac{1}{n_y} \sum_{i=1}^{n_y} y_i\right) \\
 &= E\left(\frac{1}{n_x} \sum_{i=1}^{n_x} x_i\right) - E\left(\frac{1}{n_y} \sum_{i=1}^{n_y} y_i\right) \\
 &= \frac{1}{n_x} \sum_{i=1}^{n_x} E(x_i) - \frac{1}{n_y} \sum_{i=1}^{n_y} E(y_i) \\
 &= \frac{1}{n_x} \sum_{i=1}^{n_x} \mu_x - \frac{1}{n_y} \sum_{i=1}^{n_y} \mu_y \\
 &= \frac{n_x \mu_x}{n_x} - \frac{n_y \mu_y}{n_y} = \mu_x - \mu_y
 \end{aligned}$$

Variance of the Difference

$$\begin{aligned}
 \text{Var}(\bar{x} - \bar{y}) &= \text{Var}\left(\frac{1}{n_x} \sum_{i=1}^{n_x} x_i - \frac{1}{n_y} \sum_{i=1}^{n_y} y_i\right) \\
 &= \text{Var}\left(\frac{1}{n_x} \sum_{i=1}^{n_x} x_i\right) + \text{Var}\left(\frac{1}{n_y} \sum_{i=1}^{n_y} y_i\right) \\
 &= \frac{1}{n_x^2} \sum_{i=1}^{n_x} \text{Var}(x_i) + \frac{1}{n_y^2} \sum_{i=1}^{n_y} \text{Var}(y_i) \\
 &= \frac{1}{n_x^2} \sum_{i=1}^{n_x} \sigma_x^2 + \frac{1}{n_y^2} \sum_{i=1}^{n_y} \sigma_y^2 \\
 &= \frac{n_x \sigma_x^2}{n_x^2} + \frac{n_y \sigma_y^2}{n_y^2} = \frac{\sigma_x^2}{n_x} + \frac{\sigma_y^2}{n_y}
 \end{aligned}$$

Did I make any assumptions here?

Checking assumptions & conditions

1 Independence:

1 Independence within groups:

- Both the college graduates and those with HS degree or lower are sampled randomly.
- 505 < 10% of all college graduates and 667 < 10% of all students with a high school degree or lower.

We can assume that the number of hours worked per week by one college graduate in the sample is independent of another, and the number of hours worked per week by someone with a HS degree or lower in the sample is independent of another as well.

2 Independence between groups:

Since the sample is random, the college graduates in the sample are independent of those with a HS degree or lower.

3 Sample size / skew:

Both distributions look reasonably symmetric, and the sample sizes are at least 30, therefore we can assume that the sampling distribution of number of hours worked per week by college graduates and those with HS degree or lower are nearly normal. Hence the sampling distribution of the average difference will be nearly normal as well.

Confidence interval for difference between two means

- All confidence intervals have the same form:

$$\text{point estimate} \pm ME$$

- Always, $ME = \text{critical value} \times SE \text{ of point estimate}$
- In this case the point estimate is $\bar{x} - \bar{y}$
- Since the sample sizes are large enough, the critical value is z^*
- So the only new concept is the standard error of the difference between two means...

$$SE_{(\bar{x}-\bar{y})} = \sqrt{\frac{\sigma_x^2}{n_x} + \frac{\sigma_y^2}{n_y}} \approx \sqrt{\frac{s_x^2}{n_x} + \frac{s_y^2}{n_y}}$$

Let's put things in context

Calculate the standard error of the average difference between the number of hours worked per week by college graduates and those with a HS degree or lower.

	\bar{x}	s	n
college or higher	41.8	15.14	505
hs or lower	39.4	15.12	667

Confidence interval for the difference (cont.)

Estimate (using a 95% confidence interval) the average difference between the number of hours worked per week by Americans with a college degree and those with a high school degree or lower.

$$\bar{x}_c = 41.8 \quad \bar{x}_{hs} = 39.4 \quad SE_{(\bar{x}_c - \bar{x}_{hs})} = 0.89$$

Setting the hypotheses

What are the hypotheses for testing if there is a difference between the average number of hours worked per week by college graduates and those with a HS degree or lower?

$$H_0: \mu_c = \mu_{hs} \rightarrow \mu_c - \mu_{hs} = 0$$

There is no difference in the average number of hours worked per week by college graduates and those with a HS degree or lower. Any observed difference between the sample means is due to natural sampling variation (chance).

$$H_A: \mu_c \neq \mu_{hs} \rightarrow \mu_c - \mu_{hs} \neq 0$$

There is a difference in the average number of hours worked per week by college graduates and those with a HS degree or lower.

Calculating the test-statistic and the p-value

$$H_0: \mu_c - \mu_{hs} = 0$$

$$H_A: \mu_c - \mu_{hs} \neq 0$$

$$\bar{x}_c - \bar{x}_{hs} = 2.4, SE_{(\bar{x}_c - \bar{x}_{hs})} = 0.89$$



$$Z = \frac{(\bar{x}_c - \bar{x}_{hs}) - 0}{SE_{(\bar{x}_c - \bar{x}_{hs})}}$$

$$= \frac{2.4}{0.89} = 2.70$$

$$\text{upper tail} = 1 - 0.9965 = 0.0035$$

$$p\text{-value} = 2 \times 0.0035 = 0.007$$

Since the p-value is small, we reject H_0 . The data provide convincing evidence of a difference between the average number of hours worked per week by college graduates and those with a HS degree or lower.

Inference using difference of two means

- If $n_1 \geq 30$ and $n_2 \geq 30$, difference between the sample means follow a Z distribution with $SE \approx \sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}$.
- Conditions:
 - independence within groups (often verified by a random sample, and if sampling without replacement, $n < 10\%$ of population)
 - independence between groups
 - $n_1 \geq 30$ and $n_2 \geq 30$ and no extreme skew in either group
- Hypothesis testing:

$$Z = \frac{\text{point estimate} - \text{null value}}{SE}$$

- Confidence interval:


$$\text{point estimate} \pm z^* \times SE$$

Example - Diamonds

- Weights of diamonds are measured in carats.
- 1 carat = 100 points, 0.99 carats = 99 points, etc.
- The difference between the size of a 0.99 carat diamond and a 1 carat diamond is undetectable to the naked human eye, but the price of a 1 carat diamond tends to be much higher than the price of a 0.99 carat diamond.
- We are going to test to see if there is a difference between the average prices of 0.99 and 1 carat diamonds.
- In order to be able to compare equivalent units, we divide the prices of 0.99 carat diamonds by 99 and 1 carat diamonds by 100, and compare the average point prices.



Data



	0.99 carat pt99	1 carat pt100
\bar{x}	44.50	53.43
s	13.32	12.22
n	23	30

These data are a random sample from the diamonds data set in the ggplot2 R package.

Parameter and point estimate

- **Parameter of interest:** Average difference between the point prices of *all* 0.99 carat and 1 carat diamonds.

$$\mu_{pt99} - \mu_{pt100}$$

- **Point estimate:** Average difference between the point prices of *sampled* 0.99 carat and 1 carat diamonds.

$$\bar{x}_{pt99} - \bar{x}_{pt100}$$

Hypotheses and Conditions

Which of the following is the correct set of hypotheses for testing if the average point price of 1 carat diamonds (μ_{pt100}) is higher than the average point price of 0.99 carat diamonds (μ_{pt99})?

$$H_0: \mu_{pt99} = \mu_{pt100}$$

$$H_A: \mu_{pt99} < \mu_{pt100}$$

$$H_0: \bar{x}_{pt99} = \bar{x}_{pt100}$$

$$H_A: \bar{x}_{pt99} < \bar{x}_{pt100}$$

Is there any issue we need to be concerned with regard to assumptions?

In Context

	0.99 carat pt99	1 carat pt100
\bar{x}	44.50	53.43
s	13.32	12.22
n	23	30

What is the correct df for this hypothesis test?

Test statistic

Test statistic for inference on the difference of two small sample means

The test statistic for inference on the difference of two small sample means ($n_1 < 30$ and/or $n_2 < 30$) mean is the T statistic.

$$T_{df} = \frac{\text{point estimate} - \text{null value}}{SE}$$

where

$$SE \approx \sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}} \quad \text{and} \quad df = \min(n_1 - 1, n_2 - 1)$$

Note: The calculation of the df is actually much more complicated. For simplicity we'll use the above formula to estimate the true df when conducting the analysis by hand.

p-value

What is the correct p-value for the hypothesis test?

$$T = -2.508$$

	0.100	0.050	0.025	0.010	0.005
one tail	0.100	0.050	0.025	0.010	0.005
two tails	0.200	0.100	0.050	0.020	0.010
df 21	1.32	1.72	2.08	2.52	2.83
22	1.32	1.72	2.07	2.51	2.82
23	1.32	1.71	2.07	2.50	2.81
24	1.32	1.71	2.06	2.49	2.80
25	1.32	1.71	2.06	2.49	2.79

Synthesis

What is the conclusion of the hypothesis test? How (if at all) would this conclusion change your behavior if you went diamond shopping?

Confidence interval

Calculate the interval, and interpret it in context.

Critical value

What is the appropriate t^* for a confidence interval for the average difference between the point prices of 0.99 and 1 carat diamonds that would be equivalent to our hypothesis test?

		0.100	0.050	0.025	0.010	0.005
one tail		0.100	0.050	0.025	0.010	0.005
two tails		0.200	0.100	0.050	0.020	0.010
df	21	1.32	1.72	2.08	2.52	2.83
	22	1.32	1.72	2.07	2.51	2.82
	23	1.32	1.71	2.07	2.50	2.81
	24	1.32	1.71	2.06	2.49	2.80
	25	1.32	1.71	2.06	2.49	2.79

Inference using difference of two small sample means

- If $n_1 < 30$ and/or $n_2 < 30$, difference between the sample means follow a t distribution with $SE \approx \sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}$.
- Conditions:
 - independence within groups (often verified by a random sample, and if sampling without replacement, $n < 10\%$ of population)
 - independence between groups
 - $n_1 < 30$ and/or $n_2 < 30$ and no extreme skew in either group
- Hypothesis testing:

$$T_{df} = \frac{\text{point estimate} - \text{null value}}{SE}, \text{ where } df = \min(n_1 - 1, n_2 - 1)$$


- Confidence interval:

$$\text{point estimate} \pm t_{df}^* \times SE$$

Example - Reading and Writing

200 randomly selected high school students took a reading and writing test and their scores are shown below. At a first glance, does there appear to be a difference between the average reading and writing test score?

	id	read	write
	1	70	57
	2	86	44
	3	141	63
	4	172	47
	⋮	⋮	⋮
	200	137	65



Do you think reading and writing scores are independent?

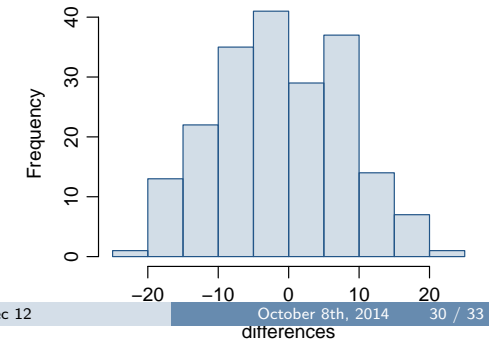
Analyzing paired data

- When two sets of observations have this special correspondence (not independent), they are said to be *paired*.
- To analyze paired data, it is often useful to look at the difference in outcomes of each pair of observations.

$$\text{diff} = \text{read} - \text{write}$$

- It is important that we always subtract using a consistent order.

	id	read	write	diff
	1	70	57	5
	2	86	44	11
	3	141	63	19
	4	172	47	-5
	⋮	⋮	⋮	⋮
	200	137	65	-2



Parameter and point estimate

- *Parameter of interest*: Average difference between the reading and writing scores of *all* high school students.

$$\mu_{diff}$$

- *Point estimate*: Average difference between the reading and writing scores of *sampled* high school students.

$$\bar{x}_{diff}$$

Setting the hypotheses

If in fact there was no difference between the scores on the reading and writing exams, what would you expect the average difference to be?

What are the hypotheses for testing if there is a difference between the average reading and writing scores?

H_0 : There is no difference between the average reading and writing score.

$$\mu_{diff} = 0$$

H_A : There is a difference between the average reading and writing score.

$$\mu_{diff} \neq 0$$

Nothing new here

- The analysis is no different than what we have done before.
- We have data from *one* sample: differences.
- We are testing to see if the average difference is different than 0.

	diff
\bar{x}	-0.545
s	8.89
n	200

$$Z = \frac{\bar{X} - \mu}{SE} = \frac{-0.545 - 0}{8.89/\sqrt{200}} = -0.877$$

$$\text{p-value} = 2 \times 0.19 = 0.38$$