

Lab 8

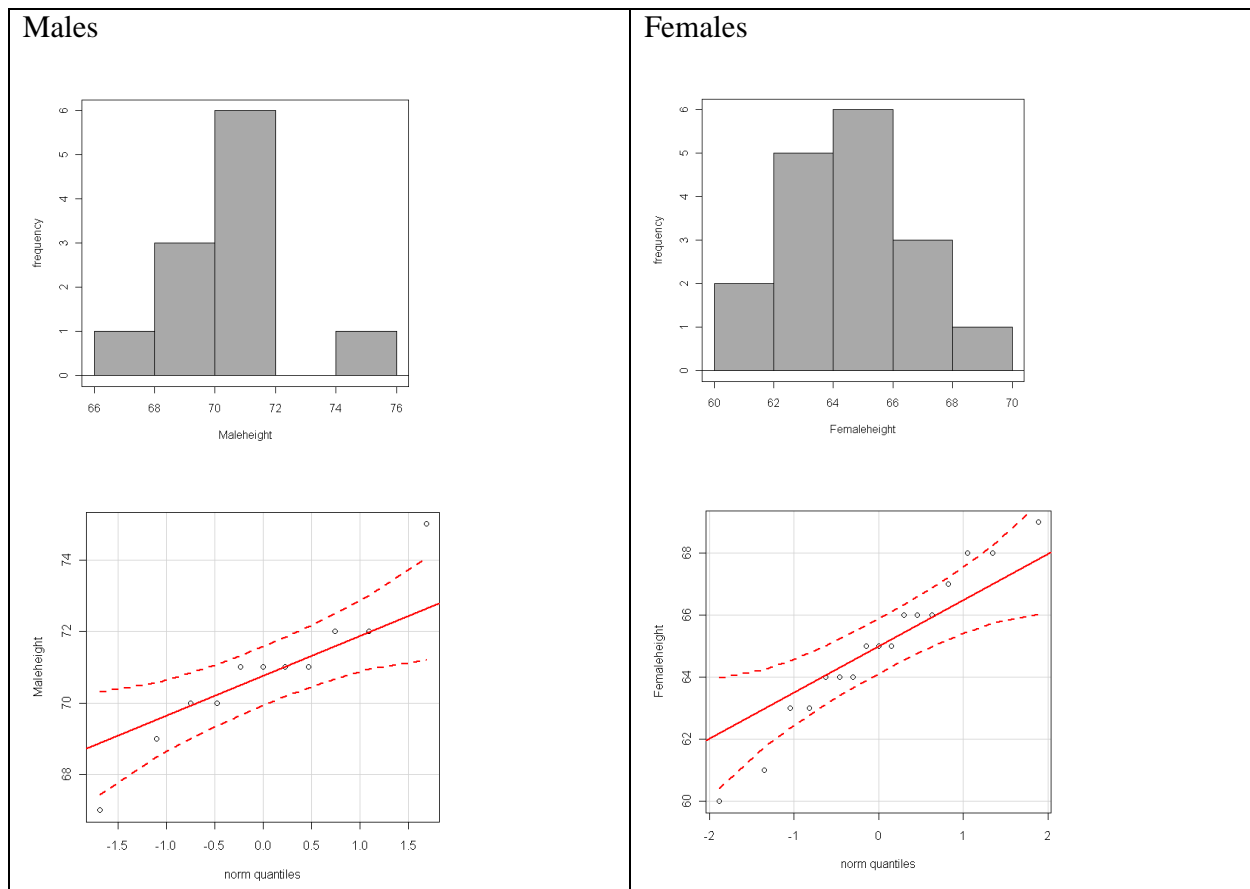
Goal: To gain experience with hypothesis tests for a mean, and to become familiar with using Rcmdr to run such hypothesis tests.

Part 1 – A One-Sample t -Test

We'll continue using our classroom data. For the purposes of this lab, we're going to treat our class as a random sample of Amherst College students. You may or may not agree with this, but we'll assume it is true. Let's focus on the heights of Amherst College students.

We'll consider the HEIGHT variable. According to the National Health Statistics Report number 10 (October 22, 2008), the average height of 20+ American males is 69.5 inches and the average height of 20+ American females is 64 inches.

Some summary statistics and graphs of the class data are given below.



We'll begin by testing whether the mean height of Amherst College males differs from the national average of 69.5 inches by hand.

To do that, compute the sample mean and standard deviation for HEIGHT by GENDER.

Go to: **Statistics → Summaries → Numerical Summaries**

Select HEIGHT as the variable.

Click the "Summarize by Groups" button and select GENDER as the grouping variable.

Q1] Give the sample size, sample mean, and standard deviation of HEIGHT for both genders.

The summary statistics are given below:

```
> numSummary(ClassData[, "Height"], groups=ClassData$Gender,
+   statistics=c("mean", "sd", "quantiles"),
+   quantiles=c(0, .25, .5, .75, 1))
      mean      sd 0% 25% 50% 75% 100%  n
Female 64.94118 2.410150 60 64 65 66.0 69 17
Male   70.81818 1.990888 67 70 71 71.5 75 11
```

Q2] Discuss whether the assumptions of the one-sample t -test are satisfied.

I think the assumptions are approximately satisfied:

- The heights are independent. One person's height doesn't affect another.
- We have sampled less than 10% of the total population
- The data are approximately normal, based on the histogram included.
- We don't have a random sample. We have a convenience sample. Using the t -test here may be inappropriate.

Q3] State the hypotheses for testing whether the mean height of Amherst College males is more than that of the national average.

Let μ be the mean height of Amherst College males. We are testing

$$H_0: \mu = 69.5$$

$$H_A: \mu > 69.5$$

Q4] Compute the test statistic value $\left(t_{n-1} = \frac{\bar{y} - \mu_0}{s/\sqrt{n}}\right)$.

The test statistic is: $t_{n-1} = \frac{\bar{y} - \mu_0}{s/\sqrt{n}} = \frac{70.8 - 69.5}{1.99/\sqrt{11}} = 2.17$

Q5] The p -value of this test can be found using the t -Table in your textbook as I described in class. You can also use Rcmdr.

Go to: **Distributions** → **Continuous distributions** → **t distribution** → **t probabilities**

Input the t -statistic value, degrees of freedom, and whether you want the upper or lower tail probability.

What is the p -value of your test?

The p -value of our test is 0.028.

Q6] What do you conclude from your hypothesis test?

This is a low p -value. We reject the null hypothesis and conclude that the mean height of Amherst College men is more than the national average.

Q7] Compute and interpret the 90% confidence interval for the mean height of Amherst College men.

```
> t.test(Maledata$Height, alternative='two.sided', mu=69.5,
conf.level=.90)
```

```
One Sample t-test
```

```
data: Maledata$Height
t = 2.196, df = 10, p-value = 0.0528
alternative hypothesis: true mean is not equal to 69.5
90 percent confidence interval:
 69.73021 71.90616
sample estimates:
mean of x
 70.81818
```

We are 90% confident that the true mean height of Amherst College men is between 69.7 and 71.8 inches.

A one-sample t -test using Rcmdr

Rcmdr can also do hypothesis tests for us, but it needs the raw data. Let's read in the classroom data we've used before. Note that it is still available on the course CMS site.

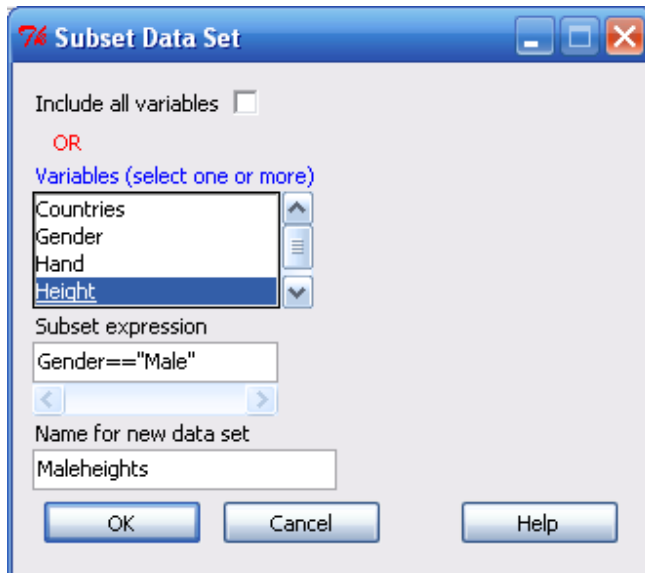
To open it in **Rcmdr**, follow the following steps:

1. Click Rcmdr Window, and select **Data > Import data > from Excel, Access, or dBase data set....**
2. In the window that opens, enter a **name** you would like to call your data set, e.g. *Class*. Click OK.

3. Find the data set file from wherever you saved it, click **Open**. In the window that opens, select the spreadsheet with the “cleaned” data and then click OK.

To run the test on men, we need to subset the data.

Go to **Data → Active data set → Subset active data set**



You can include all variables, or uncheck that box and only include HEIGHT.

Enter **Gender == “Male”** into the subset expression box (note I’ve used two equals signs here)

Give the new dataset a name.

Now, go to **Statistics → Means → Single-sample t test**

- Highlight the GENDER variable
- Select the “>” alternative hypothesis and set the null hypothesis value

To compute a confidence interval, rerun the “Single-sample t test” only using “≠” in the alternative hypothesis.

Q8] Compare the Rcmdr results to those you got by hand.

The Rcmdr results match our results by hand.

Part 2 – A Two-Sample t -Test

We use a two-sample t -test to compare the means from two different groups. We'll use it to compare the mean heights of male Amherst College students versus female Amherst College students.

Q9] Discuss whether the assumptions of the two-sample t -test are satisfied.

I think the assumptions are approximately satisfied:

- The heights are independent. One person's height doesn't affect another.
- We have sampled less than 10% of the total population
- The data are approximately normal for each group, based on the histograms included.
- We don't have a random sample. We have a convenience sample. Using the t -test here may be inappropriate.

Q10] State the hypotheses for testing whether the mean height of Amherst College males is more than that of Amherst College females.

Let μ_m be the mean height of Amherst College males and μ_f be the mean height of Amherst College females. We are testing

$$H_0: \mu_m = \mu_f$$

$$H_A: \mu_m > \mu_f$$

Q11] Compute the test statistic value. It is given here: $t_{df} = \frac{\bar{y}_1 - \bar{y}_2}{\sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}}$. The degrees of freedom

calculation is much more complicated now, though.

$$df = \frac{\left(\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}\right)^2}{\frac{1}{n_1 - 1} \left(\frac{s_1^2}{n_1}\right)^2 + \frac{1}{n_2 - 1} \left(\frac{s_2^2}{n_2}\right)^2} = \frac{\left(\frac{1.99^2}{11} + \frac{2.41^2}{17}\right)^2}{\frac{1}{10} \left(\frac{1.99^2}{11}\right)^2 + \frac{1}{16} \left(\frac{2.41^2}{17}\right)^2} = 24.3$$

$$t_{df} = \frac{\bar{y}_1 - \bar{y}_2}{\sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}} = \frac{70.818 - 64.941}{\sqrt{\frac{1.99^2}{11} + \frac{2.41^2}{17}}} = -36.67$$

Q12] The p -value of this test can be found using the t -Table in your textbook as I described in class. You can also use Rcmdr.

Go to: **Distributions** → **Continuous distributions** → **t distribution** → **t probabilities**

Input the t -statistic value, degrees of freedom, and whether you want the upper or lower tail probability.

What is the p -value of your test?

The p -value of our test is practically zero.

Q13] What do you conclude from your hypothesis test?

We conclude that we do have evidence to reject the null hypothesis. The mean height of Amherst College men is more than Amherst College women.

Q14] Compute and interpret the 95% confidence interval for difference in the mean heights of Amherst College men and women. $(\bar{y}_1 - \bar{y}_2) \pm t_{df} \sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}$.

$$\begin{aligned} (\bar{y}_1 - \bar{y}_2) \pm t_{df} \sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}} &= (70.818 - 64.941) \pm 2.064 \sqrt{\frac{1.99^2}{11} + \frac{2.41^2}{17}} = 5.877 \pm 1.729 \\ &= (4.148, 7.606) \end{aligned}$$

A two-sample t -test using Rcmdr

Rcmdr can also do hypothesis tests for us, but it needs the raw data. Let's read in the classroom data we've used before. Note that it is still available on the course CMS site.

Now, go to **Statistics** → **Means** → **Independent samples t test**

- Highlight the GENDER variable for "Groups" and HEIGHT as the "response"
- Select the appropriate alternative hypothesis and set the null hypothesis value
- Leave the test on "No" for "Assume equal variances"

To compute a confidence interval, rerun the test, only using " \neq " in the alternative hypothesis.

Q15] Compare the Rcmdr results to those you got by hand.

The Rcmdr interval is (-7.605, -4.149). This matches well with our hand calculations.

Part 3 – A Matched-Pairs t -Test

Some studies are interested in comparing the result of a treatment or study where two measurements are made on the same subject. Before/after studies are good examples of this. In such a situation, we do not have independent groups, so the usual two-sample t -test cannot be used.

A matched pairs t -test is used in this situation. The basic idea is that we take the difference between the two measurements, and focus on those. We can do a one-sample t -test on the *differences*.

Durability of shock absorbers. A manufacturer of automobile shock absorbers was interested in comparing the durability of its shocks with that of the shocks produced by its biggest competitor. To make the comparison, one of the manufacturer's and one of the competitor's shocks were randomly selected and installed on the rear wheels of each of six cars. After the cars had been driven 20,000 miles, the strength of each test shock was measured, coded, and recorded. Results of the examination are shown in the table.

Car Number	Manufacturer's Shock	Competitor's Shock	Difference
1	8.8	8.4	0.4
2	10.5	10.1	0.4
3	12.5	12.0	0.5
4	9.7	9.3	0.4
5	9.6	9.0	0.6
6	13.2	13.0	0.2

Enter these data into Rcmdr. **Go to Data → New data set**

Give the data a name. A data editor will open and you can input the data. Doubleclicking the “Var” at the top of the columns will bring up a window where you can each variable.

I named my dataset “Shocks” and gave the manufacturer data the name “OEM” and the competitor data the name “Comp”. We can compute the difference and add it to our dataset by going into the script window and typing:

```
Shocks$Difference <- Shocks$OEM - Shocks$Comp
```

You'll have to replace “Shocks” with your dataset name, and “OEM” and “Comp” with your variable names, too. Now, you've added column of differences to your dataset.

Q16] Explain why the data are collected as matched pairs.

The data are collected as matched pairs because both measurements are being made on the same individual sample unit.

Q17] Do the data present sufficient evidence to conclude that there is a difference in the mean strength of the two types of shocks after 20,000 miles of use? Use $\alpha = 0.05$.

By hand: Do a one-sample t -test using the differences

By Rcmdr: Go to **Statistics** → **Means** → **Paired t test**

```
> numSummary(Schocks[, "Difference"], statistics=c("mean", "sd",
"quantiles"),
+   quantiles=c(0, .25, .5, .75, 1))
      mean      sd  0%  25%  50%   75% 100% n
0.4166667 0.1329160 0.2 0.4 0.4 0.475 0.6 6
```

The matched pairs test is really nothing more than taking the *differences* between the two measurements and performing a one-sample t -test on them.

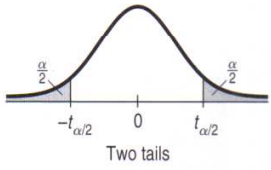
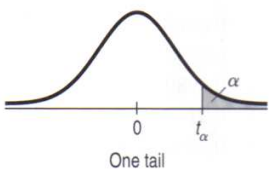
Let μ_d be the mean difference between the OEM shocks and the competitors. We are testing

$$H_0: \mu_d = 0$$
$$H_A: \mu_d \neq 0$$

The test statistic is: $t_{n-1} = \frac{\bar{y} - \mu_0}{s/\sqrt{n}} = \frac{0.417 - 0}{0.1329/\sqrt{6}} = 7.68$. With 5 degrees of freedom, our p -value is less than 0.005 (see the t -table). We reject the null hypothesis and conclude that the shocks are different. Rcmdr gives similar results.

Paired t -test

```
data: Schocks$OEM and Schocks$Comp
t = 7.6787, df = 5, p-value = 0.0005971
alternative hypothesis: true difference in means is not equal to
0
95 percent confidence interval:
 0.2771799 0.5561535
sample estimates:
mean of the differences
      0.4166667
```


		0.20	0.10	0.05	0.02	0.01	
		0.10	0.05	0.025	0.01	0.005	
Two-tail probability	One-tail probability						
Table T							
Values of t_α							df
							
Two tails							
							
One tail							
df							df
1		3.078	6.314	12.706	31.821	63.657	1
2		1.886	2.920	4.303	6.965	9.925	2
3		1.638	2.353	3.182	4.541	5.841	3
4		1.533	2.132	2.776	3.747	4.604	4
5		1.476	2.015	2.571	3.365	4.032	5
6		1.440	1.943	2.447	3.143	3.707	6
7		1.415	1.895	2.365	2.998	3.499	7
8		1.397	1.860	2.306	2.896	3.355	8
9		1.383	1.833	2.262	2.821	3.250	9
10		1.372	1.812	2.228	2.764	3.169	10
11		1.363	1.796	2.201	2.718	3.106	11
12		1.356	1.782	2.179	2.681	3.055	12
13		1.350	1.771	2.160	2.650	3.012	13
14		1.345	1.761	2.145	2.624	2.977	14
15		1.341	1.753	2.131	2.602	2.947	15
16		1.337	1.746	2.120	2.583	2.921	16
17		1.333	1.740	2.110	2.567	2.898	17
18		1.330	1.734	2.101	2.552	2.878	18
19		1.328	1.729	2.093	2.539	2.861	19
20		1.325	1.725	2.086	2.528	2.845	20
21		1.323	1.721	2.080	2.518	2.831	21
22		1.321	1.717	2.074	2.508	2.819	22
23		1.319	1.714	2.069	2.500	2.807	23
24		1.318	1.711	2.064	2.492	2.797	24
25		1.316	1.708	2.060	2.485	2.787	25
26		1.315	1.706	2.056	2.479	2.779	26
27		1.314	1.703	2.052	2.473	2.771	27
28		1.313	1.701	2.048	2.467	2.763	28
29		1.311	1.699	2.045	2.462	2.756	29
30		1.310	1.697	2.042	2.457	2.750	30
32		1.309	1.694	2.037	2.449	2.738	32
35		1.306	1.690	2.030	2.438	2.725	35
40		1.303	1.684	2.021	2.423	2.704	40
45		1.301	1.679	2.014	2.412	2.690	45
50		1.299	1.676	2.009	2.403	2.678	50
60		1.296	1.671	2.000	2.390	2.660	60
75		1.293	1.665	1.992	2.377	2.643	75
100		1.290	1.660	1.984	2.364	2.626	100
120		1.289	1.658	1.980	2.358	2.617	120
140		1.288	1.656	1.977	2.353	2.611	140
180		1.286	1.653	1.973	2.347	2.603	180
250		1.285	1.651	1.969	2.341	2.596	250
400		1.284	1.649	1.966	2.336	2.588	400
1000		1.282	1.646	1.962	2.330	2.581	1000
∞		1.282	1.645	1.960	2.326	2.576	∞
Confidence levels		80%	90%	95%	98%	99%	