Using R for Power Analysis
Presenters: Ding Yu and Jeremy Yagle
Spring 2013

R is a free software downloadable at http://www.r-project.org/

Notes:	Code and Output:

1. R Console Setup:

> prompts you for formula or function.

The result appears on the next line(s).

2. Comments begin with #

Anything in the line following a # is a comment.

```
> # This is a comment!
```

3. Installing a Package

To conduct a power or sample size analysis using R the pwr package must be installed and loaded.

1) Select "**Install Packages**" in the dropdown menu "**Packages**" at the top of the screen.
2) Select the country, and state that is nearest you.
3) Select the package "**pwr**" and press "**ok**".
4) Activate the package using the library command.

All the following functions are also explained in this link.
http://cran.r-project.org/web/packages/pwr/pwr.pdf

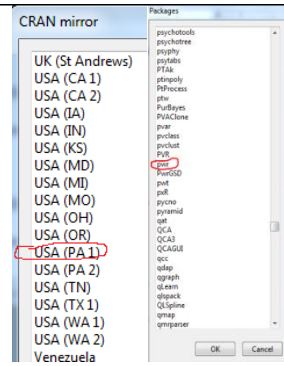

 # In step 3) you can also use the command:
>install.packages("pwr")

Step 4) type in
> library(pwr)

4. the conventional effect size (Cohen's Rules of Thumb)

cohen.ES(test = c("p", "t", "r", "anova", "chisq", "f2"), size = c("small", "medium", "large"))

Arguments:
Test: Choose one statistical test of interest

Size: Choose <u>one</u> effect size: small, medium, or large

5. Determine a medium effect size for a two-samples t-test according to Cohen's Rules of Thumb

Cohen suggests that d values of 0.2, 0.5, and 0.8 represent small, medium, and large effect sizes respectively.

```
>cohen.ES(test = c("t"), size = c("medium"))

Conventional effect size from Cohen (1982)
        test = t
        size = medium
   effect.size = 0.5
```

6. Determine a small effect size for multiple regression

```
>cohen.ES(test = c("f2"), size = c("small"))

Conventional effect size from Cohen (1982)

        test = f2
        size = small
   effect.size = 0.02
```

<table>
<tr>
<td>

7. power analysis function for t-tests
pwr.t.test(n = , d = , sig.level = , power = ,
type = c("two.sample", "one.sample", "paired"),
alternative = c("two.sided", "less","greater"))

<u>Arguments:</u>
d: Effect size
n: Number of observations in sample
sig.level: Significance level (default= .05)
power: Power of test (usually 0.8)
type: Choose one character string (default= two-sample) specifying the type of t-test.
alternative: choose one character string (default two.sided) specifying the alternative hypothesis.

</td>
<td></td>
</tr>
<tr>
<td>

8. Find the power for an independent-samples t-test given sample size

For an independent-sample two-sided t-test, with 30 subjects in each group, and a medium effect size. What is the power?

</td>
<td>

>pwr.t.test(n =30 , d = .5, sig.level =.05)

Two-sample t test power calculation

n = 30
d = 0.5
sig.level = 0.05
power = 0.4778965
alternative = two.sided

NOTE: n is number in *each* group

</td>
</tr>
<tr>
<td>

9. Find the required sample size for an independent-samples t-test and a given power

For an independent-sample two-sided t-test, 80% power, and a small effect size. What is the sample size?

</td>
<td>

>pwr.t.test(n =NULL , d = .2, sig.level =.05, power=.8)

Two-sample t test power calculation

n = 393.4057
d = 0.2
sig.level = 0.05
power = 0.8
alternative = two.sided

NOTE: n is number in *each* group

</td>
</tr>
<tr>
<td>

10. Displaying the sample size (n) only
Add "$n" to the end of the pwr.t.test command.

</td>
<td>

pwr.t.test(n =NULL , d = .2, sig.level =.05, power=.8)$n
[1] 393.4057

</td>
</tr>
<tr>
<td>

11. t-test (two samples with unequal n) example

Function: pwr.t2n.test

Example. Find the power for an independent-samples t-test at the 5% level of significance, medium effect size, with sample sizes sizes of 20 and 80.

</td>
<td>

>pwr.t2n.test(n1 =20 , n2=80 , d =.5 , sig.level = 0.05, power =NULL ,alternative = c("two.sided", "less","greater"))

t test power calculation

n1 = 20
n2 = 80
d = 0.5
sig.level = 0.05
power = 0.5081857
alternative = two.sided

</td>
</tr>
</table>

12. Balanced one way ANOVA

Function: pwr.anova.test

Arguments:

k: Number of groups
n: Number of observations per group
f: Effect size(Cohen suggests that f values of 0.1, 0.25, and 0.4 represent small, medium, and large effect sizes respectively).

sig.level: Significance level
power: Power of test

13. Balanced one way ANOVA Example

Example. Find the sample size for a one way ANOVA test, with 3 groups, a medium effect size, 80% power, and at the 5% level of significance.

So n= 53 per group (159 total)

> pwr.anova.test(k =3 , n =NULL , f =.25, sig.level =.05 , power =.8)

Balanced one-way analysis of variance power calculation

```
          k = 3
          n = 52.3966
          f = 0.25
  sig.level = 0.05
      power = 0.8
```

NOTE: n is number in each group

14. Chi Square

Function: pwr.chisq.test

Arguments:

w: Effect size
N: Total number of observations
df: Degrees of freedom.
Usually (number of rows-1)(number of columns -1)
sig.level: Significance level
power: Power of test

15. Chi Square Example

Example. Find the power for a Chi Square test for independence for two categorical variables (one with three levels and one with four levels), with 300 observation, a medium effect size, and at the 5% level of significance.

3 levels(rows)

4 levels (columns)

	Age			
Education	20-29	30-39	40-49	50 and older
G.E.D				
A.A				
B.S				

Note: df = (number of rows-1)(number of columns -1)
= (3-1) (4-1) = 2*3 = 6

> pwr.chisq.test(w =.3 , N =300 , df =6 , sig.level = 0.05, power =NULL)

(Cohen suggests that w values of 0.1, 0.3, and 0.5 represent small, medium, and large effect sizes respectively).

Chi squared power calculation

```
          w = 0.3
          N = 300
         df = 6
  sig.level = 0.05
      power = 0.9872113
```

NOTE: N is the number of observations

16. Multiple Regression Function: pwr.f2.test Arguments: **u**: Numerator degrees of freedom (the number of continuous predictors plus the number of dummy variables minus one) **v**: Denominator (error) degrees of freedom **f2**: Effect Size (Cohen suggests f2 values of 0.02, 0.15, and 0.35 represent small, medium, and large effect sizes). **sig.level**: Significance level **power**: Power of test ***Notes: every categorical variable with k levels has k-1 dummy variables because one level serves as the reference category*** ***The sample size(n) is the sum of u and v plus one.***	
17. Multiple Regression Example Example Find sample size for multiple regression test with two continuous predictor, * GRE score * Blood pressure, and two categorical predictors, Sex (Male or Female), Marital Status (married, single, or other), with a medium effect size, 80% power, and at the 5% level of significance. **u = 4** **i.e.** **2 continuous variables** **+ (2-1) dummy variables** **+ (3-1) dummy variables** **—1** **4** **Note: total sample size = N = 80+4+1=85**	**>pwr.f2.test(u = 4, v = NULL, f2 = .15, sig.level = 0.05, power = 0.8)** Multiple regression power calculation u = 4 **v = 79.44992** f2 = 0.15 sig.level = 0.05 power = 0.8
18. Example Generating a Table of sample sizes **Example:** Suppose that you are determining the required sample size for a two-sided upper-tailed independent samples t-test with 80% power. Generate a table showing the required sample size for each of the following 8 effect sizes: d = 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, and 0.8.	**>seq=c(0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8)** **>FindN=array(0,8)** **>for (i in 1:8) FindN[i]=pwr.t.test(d=seq[i],power=.8, sig.level=.05, type="two.sample", alternative="greater")$n** **>data.frame(d=seq ,N=ceiling(FindN))** **d N** **1 0.1 1238** **2 0.2 310** **3 0.3 139** **4 0.4 78** **5 0.5 51** **6 0.6 36** **7 0.7 26** **8 0.8 21**

19. Generating Power Curve
 	
Example:
Suppose that you are determining the required sample size for a two-sided upper-tailed independent samples t-test with 80% power.
 	
Generate power curves showing the required sample size for each of the following 8 effect sizes:
d = 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, and 0.8.
 	
On the graph output:
 	
x axis shows the range of effect sizes
 	
y axis shows required samples sizes
 	
Three power curves are generated (p=.6, .7, or .8). You can find the required sample size on the y axis respective to the desired effect size and power.
	**# Plot sample size curves for detecting correlations of
	# various sizes.
	
	library(pwr)
	
	# range of effect sizes
	d <- seq(.1,.8,.01)
	nd <- length(d)
	
	# power values
	p <- seq(.6,.8,.1)
	np <- length(p)
	
	# obtain sample sizes
	samsize <- array(numeric(nd*np), dim=c(nd,np))
	for (i in 1:np){
	 for (j in 1:nd){
	  result <- pwr.t.test(n = NULL, d = d[j],
	  sig.level = .05, power = p[i],
	  alternative = "two.sided")
	  samsize[j,i] <- ceiling(result$n)
	 }
	}
	
	# set up graph
	xrange <- range(d)
	yrange <- round(range(samsize))
	colors <- rainbow(length(p))
	plot(xrange, yrange, type="n",
	 xlab="Effect sizes (d)",
	 ylab="Sample Size (n)", ylim=c(0,1600))
	
	# add power curves
	for (i in 1:np){
	 lines(d, samsize[,i], type="l", lwd=2, col=colors[i])
	}
	
	# add annotation (grid lines, title, legend)
	abline(v=0, h=seq(0,yrange[2],50), lty=2, col="grey89")
	abline(h=0, v=seq(xrange[1],xrange[2],.02), lty=2,
	  col="grey89")
	title("Sample Size Estimation for t-test Studies\n
	 Sig=0.05 (Two-tailed)")
	legend("topright", title="Power", as.character(p),
	  fill=colors)**

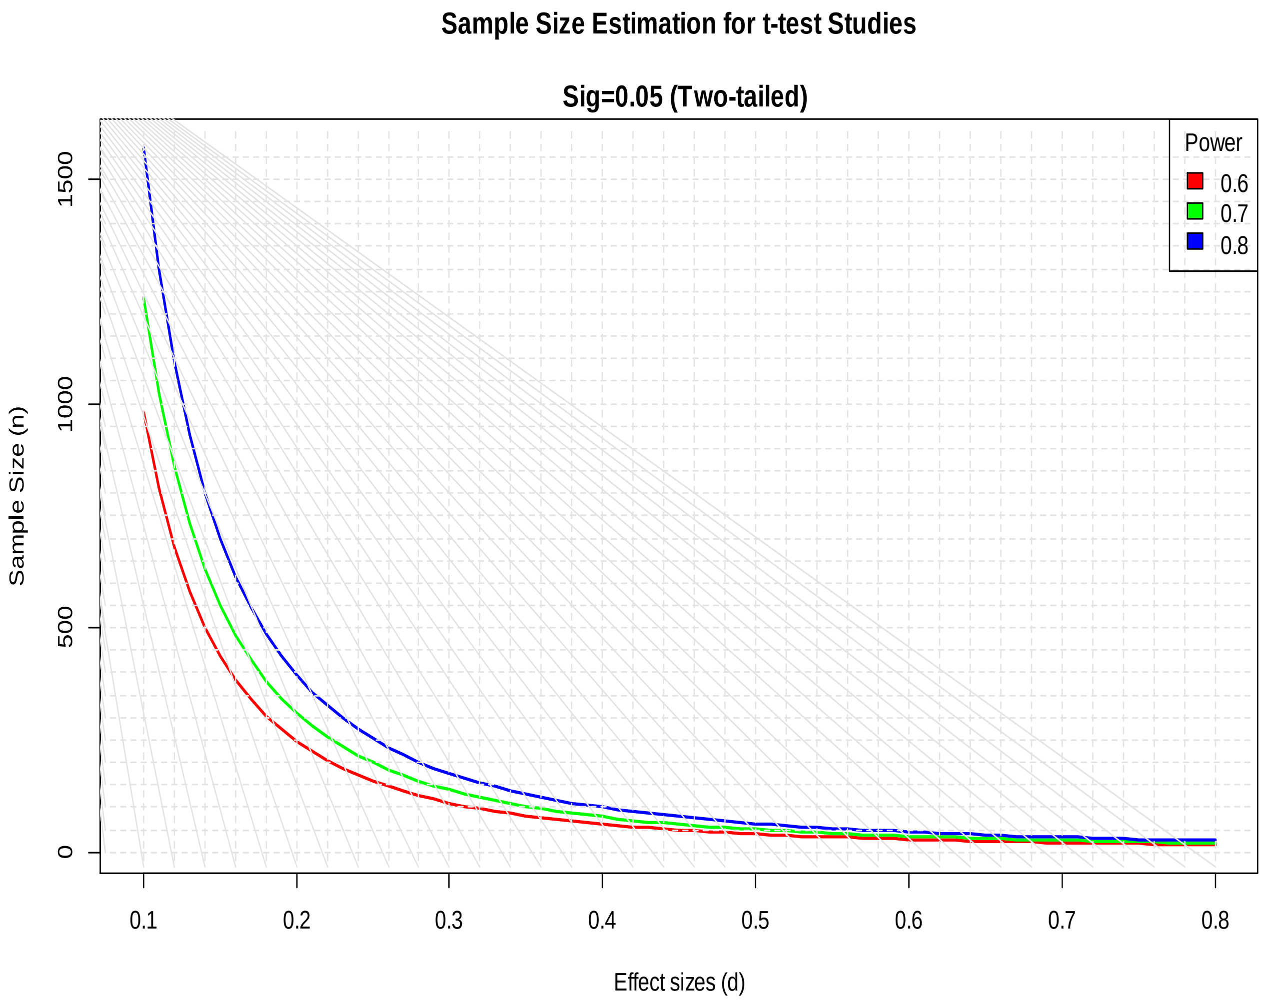

Sample Size Estimation for t-test Studies
Sig=0.05 (Two-tailed)
Power
0.6
0.7
0.8
Sample Size (n)
0
500
1000
1500
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
Effect sizes (d)