

An Introduction to R

2.4 Working with linear models

Dan Navarro (daniel.navarro@adelaide.edu.au)
School of Psychology, University of Adelaide
ua.edu.au/ccs/people/dan
DSTO R Workshop, 29-Apr-2015

Standard multiple regression

Background

- Regression usually involves two different things
 - Estimates regression coefficients from data
 - Runs hypothesis tests on coefficients and on the model fit as a whole

Background


- Regression usually involves two different things
 - Estimates regression coefficients from data
 - Runs hypothesis tests on coefficients and on the model fit as a whole
- R separates these into the distinct parts
 - `lm()` estimates the regression model
 - `summary()` runs the hypothesis tests
 - Various other functions to extract interesting things from the regression model

```
> mod <- lm( hormone ~ sad + happy + age, expt )
```



Model formula

```
> mod <- lm( hormone ~ sad + happy + age, expt )
```



Data frame

```
> mod <- lm( hormone ~ sad + happy + age, expt )
```



Variable to store the
regression model in

```
> mod <- lm( hormone ~ sad + happy + age, expt )  
> mod
```

Call:

```
lm(formula = hormone ~ sad + happy + age, data = expt)
```

Coefficients:

(Intercept)	sad	happy	age
47.1044	-10.3780	9.5906	-0.1422

Raw (unstandardised)
regression coefficients


```
> summary( mod )
```

```
> summary( mod )
```

```
Call:
```

```
lm(formula = hormone ~ sad + happy + age, data = expt)
```

```
Residuals:
```

Min	1Q	Median	3Q	Max
-5.5256	-2.5924	-0.5794	1.7483	7.4826

```
Coefficients:
```

	Estimate	Std. Error	t value	Pr(> t)	
(Intercept)	47.1044	15.1134	3.117	0.014299	*
sad	-10.3780	1.2013	-8.639	2.5e-05	***
happy	9.5906	1.8631	5.148	0.000877	***
age	-0.1422	0.4115	-0.346	0.738605	

```
---
```

```
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
Residual standard error: 4.278 on 8 degrees of freedom
```

```
Multiple R-squared: 0.9785, Adjusted R-squared: 0.9704
```

```
F-statistic: 121.2 on 3 and 8 DF,  p-value: 5.24e-07
```

```
> summary( mod )
```

```
Call:
lm(formula = hormone ~ sad + happy + age, data = expt)
```

Reminder of what the regression model is

```
Residuals:
```

```
      Min       1Q   Median       3Q      Max
-5.5256 -2.5924 -0.5794  1.7483  7.4826
```

```
Coefficients:
```

```
              Estimate Std. Error t value Pr(>|t|)
(Intercept)  47.1044    15.1134   3.117 0.014299 *
sad          -10.3780     1.2013  -8.639 2.5e-05 ***
happy         9.5906     1.8631   5.148 0.000877 ***
age          -0.1422     0.4115  -0.346 0.738605
```

```
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
Residual standard error: 4.278 on 8 degrees of freedom
Multiple R-squared: 0.9785, Adjusted R-squared: 0.9704
F-statistic: 121.2 on 3 and 8 DF, p-value: 5.24e-07
```

```
> summary( mod )
```

```
Call:
```

```
lm(formula = hormone ~ sad + happy + age, data = expt)
```

```
Residuals:
```

Min	1Q	Median	3Q	Max
-5.5256	-2.5924	-0.5794	1.7483	7.4826

Some information about the residuals

```
Coefficients:
```

	Estimate	Std. Error	t value	Pr(> t)	
(Intercept)	47.1044	15.1134	3.117	0.014299	*
sad	-10.3780	1.2013	-8.639	2.5e-05	***
happy	9.5906	1.8631	5.148	0.000877	***
age	-0.1422	0.4115	-0.346	0.738605	

```
---
```

```
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
Residual standard error: 4.278 on 8 degrees of freedom
```

```
Multiple R-squared: 0.9785, Adjusted R-squared: 0.9704
```

```
F-statistic: 121.2 on 3 and 8 DF,  p-value: 5.24e-07
```

```
> summary( mod )
```

```
Call:
```

```
lm(formula = hormone ~ sad + happy + age, data = expt)
```

```
Residuals:
```

```
      Min       1Q   Median       3Q      Max
-5.5256 -2.5924 -0.5794  1.7483  7.4826
```

Information about individual regression coefficients...

```
Coefficients:
```

	Estimate	Std. Error	t value	Pr(> t)	
(Intercept)	47.1044	15.1134	3.117	0.014299	*
sad	-10.3780	1.2013	-8.639	2.5e-05	***
happy	9.5906	1.8631	5.148	0.000877	***
age	-0.1422	0.4115	-0.346	0.738605	

```
---
```

```
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
Residual standard error: 4.278 on 8 degrees of freedom
```

```
Multiple R-squared: 0.9785, Adjusted R-squared: 0.9704
```

```
F-statistic: 121.2 on 3 and 8 DF,  p-value: 5.24e-07
```

```
> summary( mod )
```

```
Call:
```

```
lm(formula = hormone ~ sad + happy + age, data = expt)
```

```
Residuals:
```

```
      Min       1Q   Median       3Q      Max
-5.5256 -2.5924 -0.5794  1.7483  7.4826
```

The estimated (raw)
regression coefficient itself

```
Coefficients:
```

	Estimate	Std. Error	t value	Pr(> t)	
(Intercept)	47.1044	15.1134	3.117	0.014299	*
sad	-10.3780	1.2013	-8.639	2.5e-05	***
happy	9.5906	1.8631	5.148	0.000877	***
age	-0.1422	0.4115	-0.346	0.738605	

```
---  
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
Residual standard error: 4.278 on 8 degrees of freedom
```

```
Multiple R-squared: 0.9785, Adjusted R-squared: 0.9704
```

```
F-statistic: 121.2 on 3 and 8 DF,  p-value: 5.24e-07
```

```
> summary( mod )
```

```
Call:
```

```
lm(formula = hormone ~ sad + happy + age, data = expt)
```

```
Residuals:
```

```
      Min       1Q   Median       3Q      Max
-5.5256 -2.5924 -0.5794  1.7483  7.4826
```

The standard error associated with that estimate

```
Coefficients:
```

	Estimate	Std. Error	t value	Pr(> t)	
(Intercept)	47.1044	15.1134	3.117	0.014299	*
sad	-10.3780	1.2013	-8.639	2.5e-05	***
happy	9.5906	1.8631	5.148	0.000877	***
age	-0.1422	0.4115	-0.346	0.738605	

```
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
Residual standard error: 4.278 on 8 degrees of freedom
```

```
Multiple R-squared: 0.9785, Adjusted R-squared: 0.9704
```

```
F-statistic: 121.2 on 3 and 8 DF,  p-value: 5.24e-07
```

```
> summary( mod )
```

Call:

```
lm(formula = hormone ~ sad + happy + age, data = expt)
```

Residuals:

```
      Min       1Q   Median       3Q      Max
-5.5256 -2.5924 -0.5794  1.7483  7.4826
```

The corresponding t-statistic, and the associated degrees of freedom

Coefficients:

	Estimate	Std. Error	t value	Pr(> t)	
(Intercept)	47.1044	15.1134	3.117	0.014299	*
sad	-10.3780	1.2013	-8.639	2.5e-05	***
happy	9.5906	1.8631	5.148	0.000877	***
age	-0.1422	0.4115	-0.346	0.738605	

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 4.278 on 8 degrees of freedom

Multiple R-squared: 0.9785, Adjusted R-squared: 0.9704

F-statistic: 121.2 on 3 and 8 DF, p-value: 5.24e-07


```
> summary( mod )
```

```
Call:
```

```
lm(formula = hormone ~ sad + happy + age, data = expt)
```

```
Residuals:
```

```
      Min       1Q   Median       3Q      Max
-5.5256 -2.5924 -0.5794  1.7483  7.4826
```

And finally, the
corresponding p-value

```
Coefficients:
```

	Estimate	Std. Error	t value	Pr(> t)	
(Intercept)	47.1044	15.1134	3.117	0.014299	*
sad	-10.3780	1.2013	-8.639	2.5e-05	***
happy	9.5906	1.8631	5.148	0.000877	***
age	-0.1422	0.4115	-0.346	0.738605	

```
---
```

```
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
Residual standard error: 4.278 on 8 degrees of freedom
```

```
Multiple R-squared: 0.9785, Adjusted R-squared: 0.9704
```

```
F-statistic: 121.2 on 3 and 8 DF,  p-value: 5.24e-07
```

```
> summary( mod )
```

```
Call:
```

```
lm(formula = hormone ~ sad + happy + age, data = expt)
```

```
Residuals:
```

```
      Min       1Q   Median       3Q      Max
-5.5256 -2.5924 -0.5794  1.7483  7.4826
```

```
Coefficients:
```

```
              Estimate Std. Error t value Pr(>|t|)
(Intercept)  47.1044    15.1134   3.117 0.014299 *
sad          -10.3780     1.2013  -8.639 2.5e-05 ***
happy         9.5906     1.8631   5.148 0.000877 ***
age          -0.1422     0.4115  -0.346 0.738605
```

```
---
```

```
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
Residual standard error: 4.278 on 8 degrees of freedom
```

```
Multiple R-squared: 0.9785, Adjusted R-squared: 0.9704
```

```
F-statistic: 121.2 on 3 and 8 DF, p-value: 5.24e-07
```

Goodness of fit
information

```
> summary( mod )
```

```
Call:
```

```
lm(formula = hormone ~ sad + happy + age, data = expt)
```

```
Residuals:
```

```
      Min       1Q   Median       3Q      Max
-5.5256 -2.5924 -0.5794  1.7483  7.4826
```

```
Coefficients:
```

```
              Estimate Std. Error t value Pr(>|t|)
(Intercept)  47.1044    15.1134   3.117 0.014299 *
sad          -10.3780     1.2013  -8.639 2.5e-05 ***
happy         9.5906     1.8631   5.148 0.000877 ***
age          -0.1422     0.4115  -0.346 0.738605
```

```
---
```

```
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
Residual standard error: 4.278 on 8 degrees of freedom
```

```
Multiple R-squared: 0.9785, Adjusted R-squared: 0.9704
```

```
F-statistic: 121.2 on 3 and 8 DF,  p-value: 5.24e-07
```


F-test of the significance of
the model as a whole

Standardised coefficients (beta weights)

```
> library( lsr )  
> standardCoefs( mod )
```

	b	beta
sad	-10.3779780	-0.66936520
happy	9.5906022	0.39268976
age	-0.1422001	-0.01859143

Try it yourself (Exercise 2.4.1)




Regression diagnostics

Regression diagnostics

- Regression plots:
 - `plot(mod)`: several handy plots
- Not covered here:
 - `vif()` [`car` package]: variance inflation factor to check for multicollinearity
 - `ncvTest()` [`car` package] check linearity of the relationship between predictors and outcome.


Residuals at each fitted value

```
plot( mod, which=1 )
```



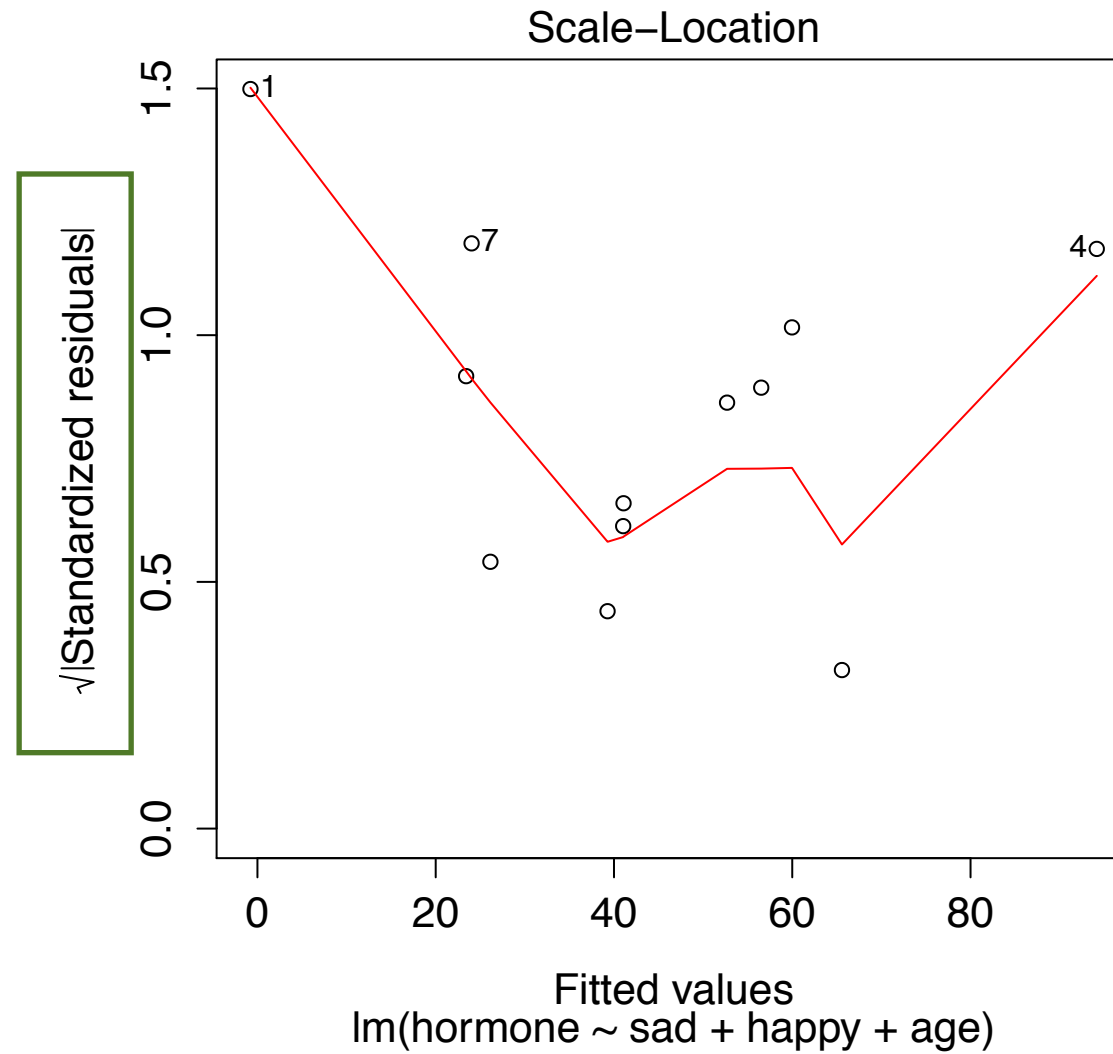
QQ plots for residuals

```
plot( mod, which=2 )
```




Residual “scale” by fitted value

```
plot( mod, which=3 )
```



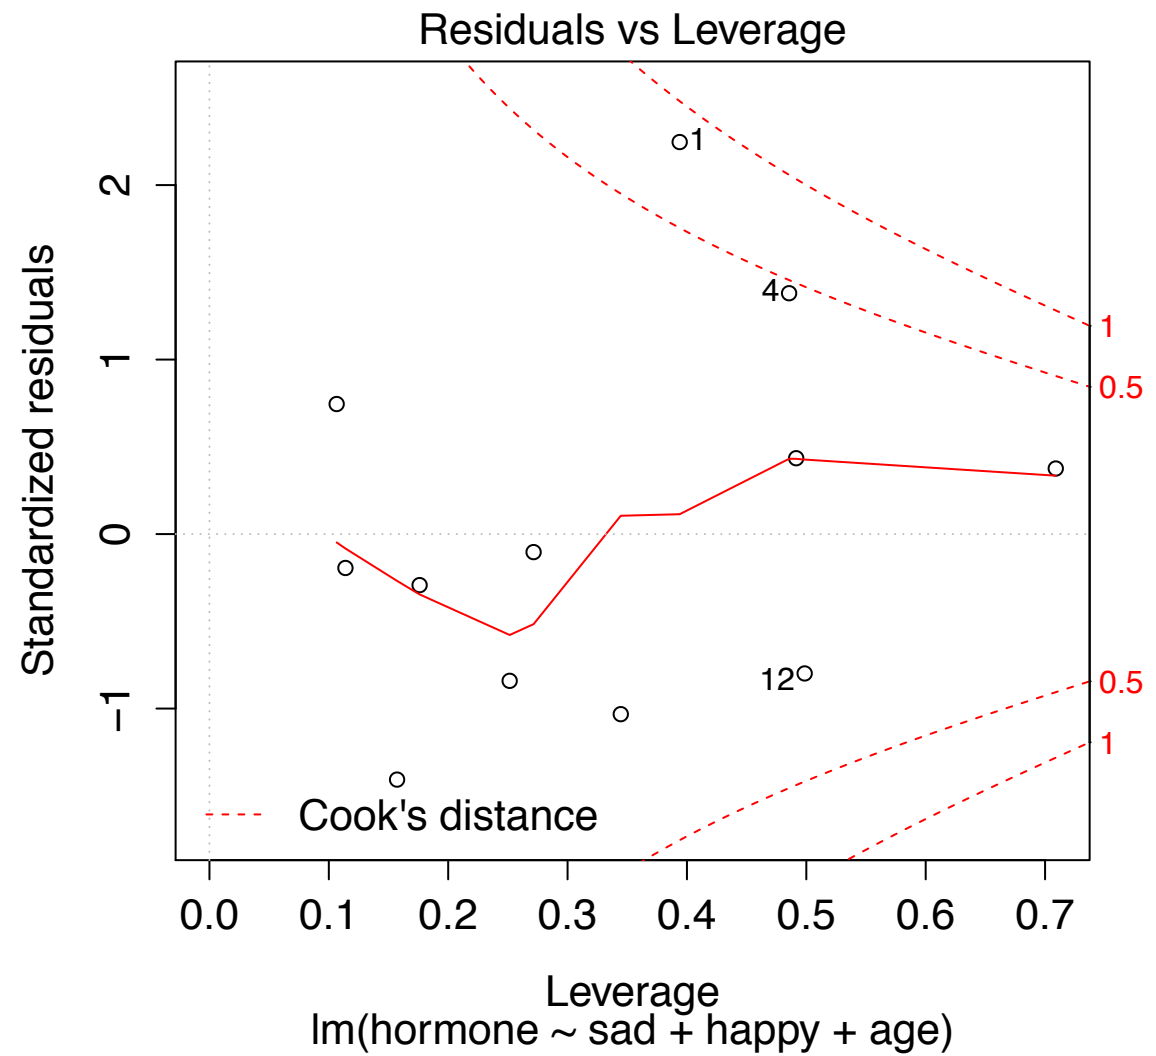
Cook's distance (influence) for all observations

```
plot( mod, which=4 )
```




Residuals by leverage

```
plot( mod, which=5 )
```



Try it yourself (Exercise 2.4.2)



One-way ANOVA

Background


- ANOVA isn't a statistical model
 - It is a hypothesis testing tool applied to a “regression” models (linear model) with all-categorical predictors
 - This is reflected in the way R implements ANOVA

Background

- ANOVA isn't a statistical model
 - It is a hypothesis testing tool applied to a “regression” models (linear model) with all-categorical predictors
 - This is reflected in the way R implements ANOVA
- Different functions for different jobs
 - Estimate the model using `aov()` or `lm()`
 - Run hypothesis tests using `anova()` or `summary()`. Or better yet, use `Anova()` [`car` package]
 - Run post hoc tests using `TukeysHSD()` or `posthocPairwiseT()` [`lsr` package]

One way ANOVA

```
> mod <- aov( hormone ~ treatment, data = expt )
```




Formula:

- hormone is the outcome variable (DV)
- treatment is the predictor (IV)

One way ANOVA


```
> mod <- aov( hormone ~ treatment, data = expt )
```



Tells R to look for these variables
inside the data frame `expt`

One way ANOVA

```
> mod <- aov( hormone ~ treatment, data = expt )
```



Store the output as a variable.

One way ANOVA

```
> mod <- aov( hormone ~ treatment, data = expt )  
> mod
```

Call:

```
aov(formula = hormone ~ treatment, data = expt)
```

Terms:

	treatment	Residuals
Sum of Squares	138.732	6662.177
Deg. of Freedom	2	9

Residual standard error: 27.20739

Estimated effects may be unbalanced

That's a bunch of ANOVA-relevant estimates, but not the ANOVA table itself...

One way ANOVA

```
> mod <- aov( hormone ~ treatment, data = expt )
```

```
> summary( mod )
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
treatment	2	139	69.4	0.094	0.911
Residuals	9	6662	740.2		

There's the ANOVA table
that we all know and love

One way ANOVA


```
> mod <- aov( hormone ~ treatment, data = expt )
```

```
> summary( mod )
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
treatment	2	139	69.4	0.094	0.911
Residuals	9	6662	740.2		

```
> bars( hormone ~ treatment, expt )
```

Reminder: that's what our data looked like, so $p=.911$ isn't a surprise!



Effect size using etaSquared()

```
> library(lsr)
> etaSquared( mod )
              eta.sq eta.sq.part
treatment 0.02039899 0.02039899
```

Effect size using eta-squared? Use the etaSquared() function...

Post hoc tests using Tukey's HSD

```
> TukeyHSD( mod )
```

```
Tukey multiple comparisons of means  
95% family-wise confidence level
```

```
Fit: aov(formula = hormone ~ treatment, data = expt)
```

```
$treatment
```

	diff	lwr	upr	p adj
drug1-control	2.175	-51.53908	55.88908	0.9929838
drug2-control	-5.875	-59.58908	47.83908	0.9501736
drug2-drug1	-8.050	-61.76408	45.66408	0.9089486

Post hoc tests using Tukey's HSD

What are we doing?

```
> TukeyHSD( mod )
```

```
Tukey multiple comparisons of means  
95% family-wise confidence level
```

```
Fit: aov(formula = hormone ~ treatment, data = expt)
```

```
$treatment
```

	diff	lwr	upr	p adj
drug1-control	2.175	-51.53908	55.88908	0.9929838
drug2-control	-5.875	-59.58908	47.83908	0.9501736
drug2-drug1	-8.050	-61.76408	45.66408	0.9089486

Post hoc tests using Tukey's HSD

```
> TukeyHSD( mod )
```

What data are we using?

```
Tukey multiple comparisons of means  
95% family-wise confidence level
```

```
Fit: aov(formula = hormone ~ treatment, data = expt)
```

```
$treatment
```

	diff	lwr	upr	p adj
drug1-control	2.175	-51.53908	55.88908	0.9929838
drug2-control	-5.875	-59.58908	47.83908	0.9501736
drug2-drug1	-8.050	-61.76408	45.66408	0.9089486

Post hoc tests using Tukey's HSD

```
> TukeyHSD( mod )
```

```
Tukey multiple comparisons of means  
95% family-wise confidence level
```

```
Fit: aov(formula = hormone ~ treatment, data = expt)
```

```
$treatment
```

	diff	lwr	upr	p adj
drug1-control	2.175	-51.53908	55.88908	0.9929838
drug2-control	-5.875	-59.58908	47.83908	0.9501736
drug2-drug1	-8.050	-61.76408	45.66408	0.9089486

All pairwise comparisons
between groups

Post hoc tests using Tukey's HSD

```
> TukeyHSD( mod )
```

```
Tukey multiple comparisons of means  
95% family-wise confidence level
```

```
Fit: aov(formula = hormone ~ treatment, data = expt)
```

```
$treatment
```

	diff	lwr	upr	p adj
drug1-control	2.175	-51.53908	55.88908	0.9929838
drug2-control	-5.875	-59.58908	47.83908	0.9501736
drug2-drug1	-8.050	-61.76408	45.66408	0.9089486

The observed difference
between sample means

Post hoc tests using Tukey's HSD

```
> TukeyHSD( mod )
```

```
Tukey multiple comparisons of means  
95% family-wise confidence level
```

```
Fit: aov(formula = hormone ~ treatment, data = expt)
```

```
$treatment
```

	diff	lwr	upr	p adj
drug1-control	2.175	-51.53908	55.88908	0.9929838
drug2-control	-5.875	-59.58908	47.83908	0.9501736
drug2-drug1	-8.050	-61.76408	45.66408	0.9089486

Simultaneous confidence
interval for the difference
between sample means

Post hoc tests using Tukey's HSD

```
> TukeyHSD( mod )
```

```
Tukey multiple comparisons of means  
95% family-wise confidence level
```


```
Fit: aov(formula = hormone ~ treatment, data = expt)
```

```
$treatment
```

	diff	lwr	upr	p adj
drug1-control	2.175	-51.53908	55.88908	0.9929838
drug2-control	-5.875	-59.58908	47.83908	0.9501736
drug2-drug1	-8.050	-61.76408	45.66408	0.9089486

p-value adjusted for
multiple comparisons via
Tukey's HSD method

Try it yourself (Exercise 2.4.3)



Factorial ANOVA

Factorial ANOVA is messier

- In real life, designs are almost always unbalanced
 - Sum of squares calculations aren't simple for unbalanced designs
 - Depends on whether you want Type I, II or III tests
 - It can also depend on the low level details of how contrasts are coded

Factorial ANOVA is messier

- In real life, designs are almost always unbalanced
 - Sum of squares calculations aren't simple for unbalanced designs
 - Depends on whether you want Type I, II or III tests
 - It can also depend on the low level details of how contrasts are coded
- Different defaults in different software!
 - SPSS default is Type III tests
 - R default in `anova()` and `summary()` is Type I tests
 - R default in `Anova()` is to use Type II tests
 - My opinion is that Type II tests are better.

Factorial ANOVA in R

```
> mod <- aov( hormone ~ treatment + gender, expt )
```

We still use the `aov()` function to specify the ANOVA model... the only difference is that we have two predictor variables

Factorial ANOVA in R

```
> mod <- aov( hormone ~ treatment + gender, expt )
```

```
> library( car )
```

```
> Anova( mod )
```

Anova Table (Type II tests)

Response: hormone

	Sum Sq	Df	F value	Pr(>F)
treatment	138.7	2	0.0873	0.9173
gender	305.0	1	0.3838	0.5528
Residuals	6357.2	8		

Here's our ANOVA table

Effect size for factorial ANOVA

```
> etaSquared( mod, type=2 )
```

	eta.sq	eta.sq.part
treatment	0.02039899	0.02135684
gender	0.04485001	0.04578395

eta-squared and partial eta-squared

Adding interaction terms

```
> mod <- aov( hormone ~ treatment * gender , expt )
```

The * here tells R to include the interaction terms

Adding interaction terms

```
> mod <- aov( hormone ~ treatment * gender , expt )
```

```
> Anova( mod )
```

Anova Table (Type II tests)

Response: hormone

	Sum Sq	Df	F value	Pr(>F)
treatment	138.7	2	0.0802	0.9239
gender	305.0	1	0.3527	0.5743
treatment:gender	1168.1	2	0.6753	0.5439
Residuals	5189.1	6		

Adding interaction terms

```
> mod <- aov( hormone ~ treatment * gender , expt )
```

```
> Anova( mod )
```

```
Anova Table (Type II tests)
```


```
Response: hormone
```

	Sum Sq	Df	F value	Pr(>F)
treatment	138.7	2	0.0802	0.9239
gender	305.0	1	0.3527	0.5743
treatment:gender	1168.1	2	0.6753	0.5439
Residuals	5189.1	6		

```
> etaSquared( mod )
```

	eta.sq	eta.sq.part
treatment	0.02039899	0.02603902
gender	0.04485001	0.05551763
treatment:gender	0.17174934	0.18373807

Try it yourself (Exercise 2.4.4)



Related tools [not covered here]

- Diagnostics
 - `ks.test()` runs a Kolmogorov-Smirnov test
 - `shapiro.test()` runs the Shapiro-Wilk test
 - `qqnorm()` draws a QQ plot
 - `leveneTest()` [`car` package] for the Levene test
 - `kruskal.test()` for nonparametric one way ANOVA
 - `oneway.test()` to allow unequal variances

Hierarchical regression

Hierarchical regression

```
> # first, create the two regression models...  
> stage.1.model <- lm( hormone ~ age + gender , expt )  
> stage.2.model <- lm( hormone ~ age + gender + happy + sad, expt )
```

Stage 1 model includes **age** and **gender**,
Stage 2 model adds **happy** and **sad**

Do “happy” and “sad” influence “hormone”, over
and above the effects of “age” and “gender”???

Hierarchical regression

```
> # first, create the two regression models...
> stage.1.model <- lm( hormone ~ age + gender , expt )
> stage.2.model <- lm( hormone ~ age + gender + happy + sad, expt )

> # next, run the F-test one against the other...
> anova( stage.1.model, stage.2.model )
```

Analysis of Variance Table

Model 1: hormone ~ age + gender

Model 2: hormone ~ age + gender + happy + sad

	Res.Df	RSS	Df	Sum of Sq	F	Pr(>F)
1	9	6049.2				
2	7	117.9	2	5931.2	176.05	1.034e-06 ***

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Hierarchical regression

```
> # first, create the two regression models...
> stage.1.model <- lm( hormone ~ age + gender , expt )
> stage.2.model <- lm( hormone ~ age + gender + happy + sad, expt )

> # next, run the F-test one against the other...
> anova( stage.1.model, stage.2.model )
```

Analysis of Variance Table

Model 1: hormone ~ age + gender


Model 2: hormone ~ age + gender + happy + sad

This line tests model 2 against
model 1: $F(2,7) = 176.05, p < .001$

	Res.Df	RSS	Df	Sum of Sq	F	Pr(>F)	
1	9	6049.2					
2	7	117.9	2	5931.2	176.05	1.034e-06	***

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Try it yourself (Exercise 2.4.5)



Analysis of Covariance

(no exercises: just watch!)

Linear models are pretty general...

If the predictors are then it is called:
all continuous	regression
all categorical	ANOVA
some continuous, some categorical	ANCOVA

So ANCOVA is pretty easy!

If the predictors are then it is called:
all continuous	regression
all categorical	ANOVA
some continuous, some categorical	ANCOVA

ANCOVA uses the same commands:

```
model <- lm( hormone ~ happy + gender + treatment, expt )
```

continuous



The diagram consists of three arrows pointing upwards from labels to terms in the R code. The first arrow points from 'continuous' to 'happy'. The second arrow points from 'two-level factor' to 'gender'. The third arrow points from 'three-level factor' to 'treatment'.

two-level
factor

three-level
factor

Anova() gives “omnibus” tests:

```
> Anova( model )
```

```
Anova Table (Type II tests)
```

```
Response: hormone
```

	Sum Sq	Df	F value	Pr(>F)	
happy	5003.6	1	25.8759	0.00142	**
gender	165.1	1	0.8541	0.38616	
treatment	19.7	2	0.0509	0.95076	
Residuals	1353.6	7			

```
---
```

```
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Use `summary()` if you want tests conducted separately for each contrast

```
> summary( model )
```

```
[blah blah blah]
```

```
Coefficients:
```

	Estimate	Std. Error	t value	Pr(> t)	
(Intercept)	-37.121	17.118	-2.169	0.06676	.
happy	21.186	4.165	5.087	0.00142	**
genderfemale	7.435	8.045	0.924	0.38616	
treatmentdrug1	-1.850	9.865	-0.188	0.85653	
treatmentdrug2	-3.121	9.848	-0.317	0.76055	

```
---  
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
Residual standard error: 13.91 on 7 degrees of freedom  
Multiple R-squared:  0.801, Adjusted R-squared:  0.6872  
F-statistic: 7.043 on 4 and 7 DF,  p-value: 0.01338
```

Want adjusted means?

Suppose we want to find the group means for **gender** and **treatment**, “adjusted” so that everyone has the same value for **happy**

Want adjusted means?

```
library( effects )  
  
Effect(  
  focal.predictors = c("gender","treatment"),  
  mod = model  
)
```

Suppose we want to find the group means for **gender** and **treatment**, “adjusted” so that everyone has the same value for **happy**

Want adjusted means?

```
library( effects )  
  
Effect(  
  focal.predictors = c("gender","treatment"),  
  mod = model  
)
```

```
gender*treatment effect  
      treatment  
gender  control  drug1  drug2  
male   41.53117  39.68085  38.41033  
female 48.96627  47.11595  45.84543
```

These are the expected group means
at the average level of “happy”

You can force it to use a different level as the baseline happiness

```
Effect(  
  focal.predictors = c("gender","treatment"),  
  mod = model,  
  given.values = c(happy=2)  
)
```

gender*treatment		effect	
	treatment		
gender	control	drug1	drug2
male	5.250333	3.400014	2.129499
female	12.685431	10.835112	9.564596

These are adjusted to a happiness
value of 2

```
> eff <- Effect(  
  focal.predictors = c("gender", "treatment"),  
  mod = model,  
  given.values = c(happy=2)  
)
```

You can get confidence intervals associated with your adjusted means.

First, save the output of the `Effect()` function to a variable

```

> eff <- Effect(
  focal.predictors = c("gender","treatment"),
  mod = model,
  given.values = c(happy=2)
)

> summary( eff )

```

```

gender*treatment effect
      treatment
gender  control  drug1  drug2
male   5.250333  3.400014  2.129499
female 12.685431 10.835112  9.564596

Lower 95 Percent Confidence Limits
      treatment
gender  control  drug1  drug2
male   -19.61115 -22.70879 -21.92520
female -12.98813 -16.13367 -15.26512

Upper 95 Percent Confidence Limits
      treatment
gender  control  drug1  drug2
male   30.11182  29.50882  26.18420
female 38.35899  37.80389  34.39431

```

Use the `summary()` function
to extract the information
you want!

Test equality of regression slopes by checking interaction terms

```
> mod1 <- lm( hormone ~ happy + gender + treatment, expt )  
> mod2 <- lm( hormone ~ happy + gender + treatment  
              + happy:gender + happy:treatment, expt )
```

mod2 is the same as our original model (i.e. mod1) with interaction terms added...

e.g. the happy:gender term allows the relationship between happiness and hormone to be different for males and females.

Test equality of regression slopes by checking interaction terms

```
> mod1 <- lm( hormone ~ happy + gender + treatment, expt )
```

```
> mod2 <- lm( hormone ~ happy + gender + treatment  
             + happy:gender + happy:treatment, expt )
```

```
> anova(mod1, mod2)
```

Analysis of Variance Table

Model 1: hormone ~ happy + gender + treatment

Model 2: hormone ~ happy + gender + treatment + happy:gender +
happy:treatment

	Res.Df	RSS	Df	Sum of Sq	F	Pr(>F)
1	7	1353.58				
2	4	640.88	3	712.7	1.4828	0.3466

Hierarchical regression to see if those
interaction terms are needed!

Repeated Measures ANOVA
(again, no exercises for this one!)

Where does this fit???

- Linear model (lm) / generalised linear model (glm) with univariate outcome
 - lm: Multiple regression ANOVA, ANCOVA, etc.
 - glm: Logistic regression, poisson regression, etc.
- Linear mixed effects model and generalised linear mixed effects model with univariate outcome:
 - RM-ANOVA (univariate), longitudinal data analysis, IRT models, crossed random effect models, etc
- Linear model with multivariate outcome
 - RM-ANOVA (multivariate), MANOVA, etc

There are many ways to do it in R

- Here are the ones I know of
 - `aov()` can (sort of) do univariate RM-ANOVA
 - `Anova()` [`car` package] has tools for multivariate tests that include MANOVA and multivariate RM-ANOVA
 - `nlme` package and `lme4` package have very powerful tools for handling mixed models of various forms
 - `ez` package is probably the best option if you want something that is familiar to an SPSS user

There are many ways to do it in R

- Here are the ones I know of
 - `aov()` can (sort of) do univariate RM-ANOVA
 - `Anova()` [car package] has tools for multivariate tests that include MANOVA and multivariate RM-ANOVA
 - nlme package and lme4 package have very powerful tools for handling mixed models of various forms
 - **ez** package is probably the best option if you want something that is familiar to an SPSS user


Hypothetical experiment

- Interested in the effect of different drugs on cognitive processing:
 - Measure working memory capacity (wm)
 - Measure median RT (rt) in a standard 2-AFC task

Hypothetical experiment

- Interested in the effect of different drugs on cognitive processing:
 - Measure working memory capacity (wm)
 - Measure median RT (rt) in a standard 2-AFC task
- Use a design with two within-subject factors:
 - Drug: Each person is tested under the influence of alcohol (alc), caffeine (caf) or no drug (nil)
 - Time: Each person is tested twice (t1 and t2)

Here's the working memory data



```
bars(  
  wm ~ drug + time,  
  dwr.long  
)
```

The data in “wide form”

```
> dwr.wide
```

	id	gender	wm_alc_t1	wm_caf_t1	wm_nil_t1	wm_alc_t2	wm_caf_t2	wm_nil_t2	rt_alc_t1
1	1	female	3.0	2.6	3.0	3.1	3.4	3.2	671
7	2	female	4.4	4.0	5.3	4.6	5.3	4.9	836
13	3	female	4.0	3.7	4.6	4.5	6.0	4.8	859
19	4	male	4.6	5.2	5.4	5.0	5.5	5.5	852
25	5	female	3.4	4.1	3.6	3.7	4.4	4.4	788
31	6	male	3.1	4.6	5.0	3.8	4.8	4.8	673
37	7	male	4.3	5.0	5.8	5.7	5.0	5.8	914
43	8	male	2.8	3.7	2.4	4.0	4.7	4.2	650
49	9	female	3.7	4.1	4.5	4.2	4.5	6.2	915
55	10	female	4.4	5.3	5.0	5.0	5.0	5.8	869

	rt_caf_t1	rt_nil_t1	rt_alc_t2	rt_caf_t2	rt_nil_t2
1	472	556	682	438	543
7	575	563	820	537	530
13	421	583	803	449	599
19	448	437	877	403	430
25	452	611	771	439	619
31	402	676	682	395	600
37	459	519	845	450	538
43	417	502	678	433	502
49	475	566	890	482	536
55	451	702	848	412	682

Each row corresponds to a specific subject

```
> dwr.wide
```

	id	gender	wm_alc_t1	wm_caf_t1	wm_nil_t1	wm_alc_t2	wm_caf_t2	wm_nil_t2	rt_alc_t1	
1	1	female	3.0	2.6	3.0	3.1	3.4	3.2	671	
7	2	female	4.4	4.0	5.3	4.6	5.3	4.9	836	
13	3	female	4.0	3.7	4.6	4.5	6.0	4.8	859	
19	4	male	4.6	5.2	5.4	5.0	5.5	5.5	852	
25	5	female	3.4	4.1	3.6	3.7	4.4	4.4	788	
31	6	male	3.1	4.6	5.0	3.8	4.8	4.8	673	
37	7	male	4.3	5.0	5.8	5.7	5.0	5.8	914	
43	8	male	2.8	3.7	2.4	4.0	4.7	4.2	650	
49	9	female	3.7	4.1	4.5	4.2	4.5	6.2	915	
55	10	female	4.4	5.3	5.0	5.0	5.0	5.8	869	
			rt_caf_t1	rt_nil_t1	rt_alc_t2	rt_caf_t2	rt_nil_t2			
1			472	556	682	438	543			
7			575	563	820	537	530			
13			421	583	803	449	599			
19			448	437	877	403	430			
25			452	611	771	439	619			
31			402	676	682	395	600			
37			459	519	845	450	538			
43			417	502	678	433	502			
49			475	566	890	482	536			
55			451	702	848	412	682			

Some variables are “id variables” or “between subjects” variables: things that aren’t measured repeatedly

```
> dwr.wide
```

	id	gender	wm_alc_t1	wm_caf_t1	wm_nil_t1	wm_alc_t2	wm_caf_t2	wm_nil_t2	rt_alc_t1
1	1	female	3.0	2.6	3.0	3.1	3.4	3.2	671
7	2	female	4.4	4.0	5.3	4.6	5.3	4.9	836
13	3	female	4.0	3.7	4.6	4.5	6.0	4.8	859
19	4	male	4.6	5.2	5.4	5.0	5.5	5.5	852
25	5	female	3.4	4.1	3.6	3.7	4.4	4.4	788
31	6	male	3.1	4.6	5.0	3.8	4.8	4.8	673
37	7	male	4.3	5.0	5.8	5.7	5.0	5.8	914
43	8	male	2.8	3.7	2.4	4.0	4.7	4.2	650
49	9	female	3.7	4.1	4.5	4.2	4.5	6.2	915
55	10	female	4.4	5.3	5.0	5.0	5.0	5.8	869

	rt_caf_t1	rt_nil_t1	rt_alc_t2	rt_caf_t2	rt_nil_t2
1	472	556	682	438	543
7	575	563	820	537	530
13	421	583	803	449	599
19	448	437	877	403	430
25	452	611	771	439	619
31	402	676	682	395	600
37	459	519	845	450	538
43	417	502	678	433	502
49	475	566	890	482	536
55	451	702	848	412	682

Other variables are “within subjects” or “repeated” variables.

```
> dwr.wide
```

	id	gender	wm_alc_t1	wm_caf_t1	wm_nil_t1	wm_alc_t2	wm_caf_t2	wm_nil_t2	rt_alc_t1
1	1	female	3.0	2.6	3.0	3.1	3.4	3.2	671
7	2	female	4.4	4.0	5.3	4.6	5.3	4.9	836
13	3	female	4.0	3.7	4.6	4.5	6.0	4.8	859
19	4	male	4.6	5.2	5.4	5.0	5.5	5.5	852
25	5	female	3.4	4.1	3.6	3.7	4.4	4.4	788
31	6	male	3.1	4.6	5.0	3.8	4.8	4.8	673
37	7	male	4.3	5.0	5.8	5.7	5.0	5.8	914
43	8	male	2.8	3.7	2.4	4.0	4.7	4.2	650
49	9	female	3.7	4.1	4.5	4.2	4.5	6.2	915
55	10	female	4.4	5.3	5.0	5.0	5.0	5.8	869

	rt_caf_t1	rt_nil_t1	rt_alc_t2	rt_caf_t2	rt_nil_t2
1	472	556	682	438	543
7	575	563	820	537	530
13	421	583	803	449	599
19	448	437	877	403	430
25	452	611	771	439	619
31	402	676	682	395	600
37	459	519	845	450	538
43	417	502	678	433	502
49	475	566	890	482	536
55	451	702	848	412	682

This is “working memory” (wm), measured under the influence of “alcohol” (alc), measured the “first time” (t1)

```
> dwr.wide
```

	id	gender	wm_alc_t1	wm_caf_t1	wm_nil_t1	wm_alc_t2	wm_caf_t2	wm_nil_t2	rt_alc_t1
1	1	female	3.0	2.6	3.0	3.1	3.4	3.2	671
7	2	female	4.4	4.0	5.3	4.6	5.3	4.9	836
13	3	female	4.0	3.7	4.6	4.5	6.0	4.8	859
19	4	male	4.6	5.2	5.4	5.0	5.5	5.5	852
25	5	female	3.4	4.1	3.6	3.7	4.4	4.4	788
31	6	male	3.1	4.6	5.0	3.8	4.8	4.8	673
37	7	male	4.3	5.0	5.8	5.7	5.0	5.8	914
43	8	male	2.8	3.7	2.4	4.0	4.7	4.2	650
49	9	female	3.7	4.1	4.5	4.2	4.5	6.2	915
55	10	female	4.4	5.3	5.0	5.0	5.0	5.8	869

	rt_caf_t1	rt_nil_t1	rt_alc_t2	rt_caf_t2	rt_nil_t2
1	472	556	682	438	543
7	575	563	820	537	530
13	421	583	803	449	599
19	448	437	877	403	430
25	452	611	771	439	619
31	402	676	682	395	600
37	459	519	845	450	538
43	417	502	678	433	502
49	475	566	890	482	536
55	451	702	848	412	682

There are three different levels of “drug”... alcohol, caffeine,
and no drug

```
> dwr.wide
```

	id	gender	wm_alc_t1	wm_caf_t1	wm_nil_t1	wm_alc_t2	wm_caf_t2	wm_nil_t2	rt_alc_t1
1	1	female	3.0	2.6	3.0	3.1	3.4	3.2	671
7	2	female	4.4	4.0	5.3	4.6	5.3	4.9	836
13	3	female	4.0	3.7	4.6	4.5	6.0	4.8	859
19	4	male	4.6	5.2	5.4	5.0	5.5	5.5	852
25	5	female	3.4	4.1	3.6	3.7	4.4	4.4	788
31	6	male	3.1	4.6	5.0	3.8	4.8	4.8	673
37	7	male	4.3	5.0	5.8	5.7	5.0	5.8	914
43	8	male	2.8	3.7	2.4	4.0	4.7	4.2	650
49	9	female	3.7	4.1	4.5	4.2	4.5	6.2	915
55	10	female	4.4	5.3	5.0	5.0	5.0	5.8	869

	rt_caf_t1	rt_nil_t1	rt_alc_t2	rt_caf_t2	rt_nil_t2
1	472	556	682	438	543
7	575	563	820	537	530
13	421	583	803	449	599
19	448	437	877	403	430
25	452	611	771	439	619
31	402	676	682	395	600
37	459	519	845	450	538
43	417	502	678	433	502
49	475	566	890	482	536
55	451	702	848	412	682

For all three “drugs”, we take measurements at the start of the session and the end of the session

```
> dwr.wide
```

	id	gender	wm_alc_t1	wm_caf_t1	wm_nil_t1	wm_alc_t2	wm_caf_t2	wm_nil_t2	rt_alc_t1
1	1	female	3.0	2.6	3.0	3.1	3.4	3.2	671
7	2	female	4.4	4.0	5.3	4.6	5.3	4.9	836
13	3	female	4.0	3.7	4.6	4.5	6.0	4.8	859
19	4	male	4.6	5.2	5.4	5.0	5.5	5.5	852
25	5	female	3.4	4.1	3.6	3.7	4.4	4.4	788
31	6	male	3.1	4.6	5.0	3.8	4.8	4.8	673
37	7	male	4.3	5.0	5.8	5.7	5.0	5.8	914
43	8	male	2.8	3.7	2.4	4.0	4.7	4.2	650
49	9	female	3.7	4.1	4.5	4.2	4.5	6.2	915
55	10	female	4.4	5.3	5.0	5.0	5.0	5.8	869

	rt_caf_t1	rt_nil_t1	rt_alc_t2	rt_caf_t2	rt_nil_t2
1	472	556	682	438	543
7	575	563	820	537	530
13	421	583	803	449	599
19	448	437	877	403	430
25	452	611	771	439	619
31	402	676	682	395	600
37	459	519	845	450	538
43	417	502	678	433	502
49	475	566	890	482	536
55	451	702	848	412	682

And we measure two things: working memory capacity, and median response time in a simple 2-AFC task

```
> dwr.wide
```

	id	gender	wm_alc_t1	wm_caf_t1	wm_nil_t1	wm_alc_t2	wm_caf_t2	wm_nil_t2	rt_alc_t1
1	1	female	3.0	2.6	3.0	3.1	3.4	3.2	671
7	2	female	4.4	4.0	5.3	4.6	5.3	4.9	836
13	3	female	4.0	3.7	4.6	4.5	6.0	4.8	859
19	4	male	4.6	5.2	5.4	5.0	5.5	5.5	852
25	5	female	3.4	4.1	3.6	3.7	4.4	4.4	788
31	6	male	3.1	4.6	5.0	3.8	4.8	4.8	673
37	7	male	4.3	5.0	5.8	5.7	5.0	5.8	914
43	8	male	2.8	3.7	2.4	4.0	4.7	4.2	650
49	9	female	3.7	4.1	4.5	4.2	4.5	6.2	915
55	10	female	4.4	5.3	5.0	5.0	5.0	5.8	869

	rt_caf_t1	rt_nil_t1	rt_alc_t2	rt_caf_t2	rt_nil_t2
1	472	556	682	438	543
7	575	563	820	537	530
13	421	583	803	449	599
19	448	437	877	403	430
25	452	611	771	439	619
31	402	676	682	395	600
37	459	519	845	450	538
43	417	502	678	433	502
49	475	566	890	482	536
55	451	702	848	412	682

The same data in “long form”

```
> dwr.long
```

	id	gender	rt	wm	drug	time
1	1	female	671	3.0	alc	t1
2	1	female	472	2.6	caf	t1
3	1	female	556	3.0	nil	t1
4	1	female	682	3.1	alc	t2
5	1	female	438	3.4	caf	t2
6	1	female	543	3.2	nil	t2
7	2	female	836	4.4	alc	t1
8	2	female	575	4.0	caf	t1
9	2	female	563	5.3	nil	t1
10	2	female	820	4.6	alc	t2
11	2	female	537	5.3	caf	t2
12	2	female	530	4.9	nil	t2

etc...

The same data in “long form”

```
> dwr.long
```

	id	gender	rt	wm	drug	time
1	1	female	671	3.0	alc	t1
2	1	female	472	2.6	caf	t1
3	1	female	556	3.0	nil	t1
4	1	female	682	3.1	alc	t2
5	1	female	438	3.4	caf	t2
6	1	female	543	3.2	nil	t2
7	2	female	836	4.4	alc	t1
8	2	female	575	4.0	caf	t1
9	2	female	563	5.3	nil	t1
10	2	female	820	4.6	alc	t2
11	2	female	537	5.3	caf	t2
12	2	female	530	4.9	nil	t2

Each subject contributes multiple rows... one for each of the 6 within subjects conditions (3 drugs x 2 times)

etc...

To do almost any repeated measures analysis in R, long form data are needed

```
> dwr.long
```

	id	gender	rt	wm	drug	time
1	1	female	671	3.0	alc	t1
2	1	female	472	2.6	caf	t1
3	1	female	556	3.0	nil	t1
4	1	female	682	3.1	alc	t2
5	1	female	438	3.4	caf	t2
6	1	female	543	3.2	nil	t2
7	2	female	836	4.4	alc	t1
8	2	female	575	4.0	caf	t1
9	2	female	563	5.3	nil	t1
10	2	female	820	4.6	alc	t2
11	2	female	537	5.3	caf	t2
12	2	female	530	4.9	nil	t2

etc...

I'll talk about how to switch between wide and long form data later!

ezANOVA

```
library(ez)

ezANOVA(
  data = dwr.long,
  dv = wm,
  wid = id,
  between = gender,
  within = .(drug,time)
)
```

The `ezANOVA()` function in the `ez` package can run “univariate” repeated measures ANOVA

ezANOVA

```
library(ez)

ezANOVA(
  data = dwr.long,
  dv = wm,
  wid = id,
  between = gender,
  within = .(drug,time)
)
```

This is the “long form” data frame
that contains all the variables

ezANOVA

```
library(ez)

ezANOVA(
  data = dwr.long,
  dv = wm,
  wid = id,
  between = gender,
  within = .(drug,time)
)
```

My dependent variable is “wm”

ezANOVA

```
library(ez)

ezANOVA(
  data = dwr.long,
  dv = wm,
  wid = id,
  between = gender,
  within = .(drug,time)
)
```

The id variable is called “id”

ezANOVA

```
library(ez)

ezANOVA(
  data = dwr.long,
  dv = wm,
  wid = id,
  between = gender,
  within = .(drug,time)
)
```

We have one between-subjects factor, “gender”

ezANOVA

```
library(ez)

ezANOVA(
  data = dwr.long,
  dv = wm,
  wid = id,
  between = gender,
  within = .(drug,time)
)
```

We have two within subjects factors, “drug” and “time”.

[Note the funny .() syntax]

The output:

```
$ANOVA
```

	Effect	DFn	DFd		F	p	p<.05	ges
2	gender	1	8	3.562678e-01	0.567090338			3.193780e-02
3	drug	2	16	8.948443e+00	0.002464231	*		1.107017e-01
5	time	1	8	2.180137e+01	0.001604195	*		1.202362e-01
4	gender:drug	2	16	3.998340e-01	0.676946405			5.531334e-03
6	gender:time	1	8	1.228218e-04	0.991429013			7.699468e-07
7	drug:time	2	16	1.005830e-01	0.904877327			1.227328e-03
8	gender:drug:time	2	16	1.382953e+00	0.279256780			1.661500e-02

degrees of freedom

The output:

```
$ANOVA
```

	Effect	DFn	DFd	F	p	p<.05	ges
2	gender	1	8	3.562678e-01	0.567090338		3.193780e-02
3	drug	2	16	8.948443e+00	0.002464231	*	1.107017e-01
5	time	1	8	2.180137e+01	0.001604195	*	1.202362e-01
4	gender:drug	2	16	3.998340e-01	0.676946405		5.531334e-03
6	gender:time	1	8	1.228218e-04	0.991429013		7.699468e-07
7	drug:time	2	16	1.005830e-01	0.904877327		1.227328e-03
8	gender:drug:time	2	16	1.382953e+00	0.279256780		1.661500e-02

F-statistic and p-value

The output:

\$ANOVA								
	Effect	DFn	DFd	F	p	p < .05	ges	
2	gender	1	8	3.562678e-01	0.567090338		3.193780e-02	
3	drug	2	16	8.948443e+00	0.002464231	*	1.107017e-01	
5	time	1	8	2.180137e+01	0.001604195	*	1.202362e-01	
4	gender:drug	2	16	3.998340e-01	0.676946405		5.531334e-03	
6	gender:time	1	8	1.228218e-04	0.991429013		7.699468e-07	
7	drug:time	2	16	1.005830e-01	0.904877327		1.227328e-03	
8	gender:drug:time	2	16	1.382953e+00	0.279256780		1.661500e-02	

generalized eta-squared
measure of effect size

The output:

```
$`Mauchly's Test for Sphericity`
```

	Effect	W	p	p<.05
3	drug	0.8289435	0.5186056	
4	gender:drug	0.8289435	0.5186056	
7	drug:time	0.6285687	0.1968950	
8	gender:drug:time	0.6285687	0.1968950	

Mauchly's test... kind of does what it
says on the label

The output:

```
$`Sphericity Corrections`
```

	Effect	GGe	p[GG]	p[GG]<.05	HFe	p[HF]	p[HF]<.05
3	drug	0.8539298	0.004324946	*	1.0624950	0.002464231	*
4	gender:drug	0.8539298	0.646539706		1.0624950	0.676946405	
7	drug:time	0.7291652	0.845365685		0.8503161	0.876048816	
8	gender:drug:time	0.7291652	0.279339317		0.8503161	0.279896289	

Greenhouse-Geisser
correction

Huynh-Feldt
correction

End of this section