

Multifactorial – ANOVA (more than one factor)

Part II: main *versus* interaction effects,
interaction plots and assessing
assumptions

Do exercise and sport affect weight loss?

Let's elaborate on this question further:

Main effects

- Are the differences in weight loss only due to exercise alone?
- Are the differences in weight loss only due to diet alone?

Interaction

- Does the effect of diet on weight loss depend on exercise? In other words, are the differences in weight loss attributable to some particular combinations of exercise and diet? (e.g., the biggest weight loss compared to any other combination of diet and exercise was observed when individuals both dieted and exercised).

Treatments

Main effects:

Diet - two treatments (yes/no).

Exercise - two treatments (yes/no).

Interaction:

Combination of diet and exercise treatments - four combinations:

- 1) No exercise but diet.
- 2) Exercise but no diet.
- 3) No exercise and no diet.
- 4) Exercise and diet.

Stating the 3 possible sets of statistical hypotheses in a two-factorial design:

Does *dieting* affect weight loss? DIET (main effect 1)

H_0 : There is no difference between diet treatments in mean weight loss (in the population).

H_A : There is a difference between diet treatments in mean weight loss (in the population).

Stating the 3 possible sets of statistical hypotheses in a two-factorial design:

Does exercising affect weight loss? EXERCISE (main effect 2)

H_0 : There is no difference between exercise treatments in mean weight loss (in the population).

H_A : There is a difference between exercise treatments in mean weight loss (in the population).

Stating the 3 possible sets of statistical hypotheses in a two-factorial design:

Are the differences in weight loss attributable to some combinations of exercise and diet? (interaction effect)

H_0 : The effect of diet on weight loss does not depend on exercise in the population (*or vice versa*).

H_A : The effect of diet on weight loss depends on exercise in the population (*or vice versa*).

ANOVA Table

Source of variation	Df	SS	Mean SS	F value	Prob
Diet	1	0.800	0.800	1.8089	0.1974
Exercise	1	28.800	28.800	65.1215	<0.0000001
Diet x Exercise	1	0.072	0.072	0.1628	0.6919
residuals	16	7.076	0.442		

H_0 : There is no difference between diet treatments in mean weight loss.

H_A : There is a difference between diet treatments in mean weight loss.

H_0 : There is no difference between exercise treatments in mean weight loss.

H_A : **There is a difference between exercise treatments in mean weight loss.**


H_0 : The effect of diet on weight loss does not depend on exercise (*or vice versa*).

H_A : The effect of diet on weight loss depends on exercise (*or vice versa*).

Conclusion: Only exercise affects weight loss!

Only exercise affects weight loss!
BUT HOW? Exercise increases weight loss ($P < 0.0000001$)

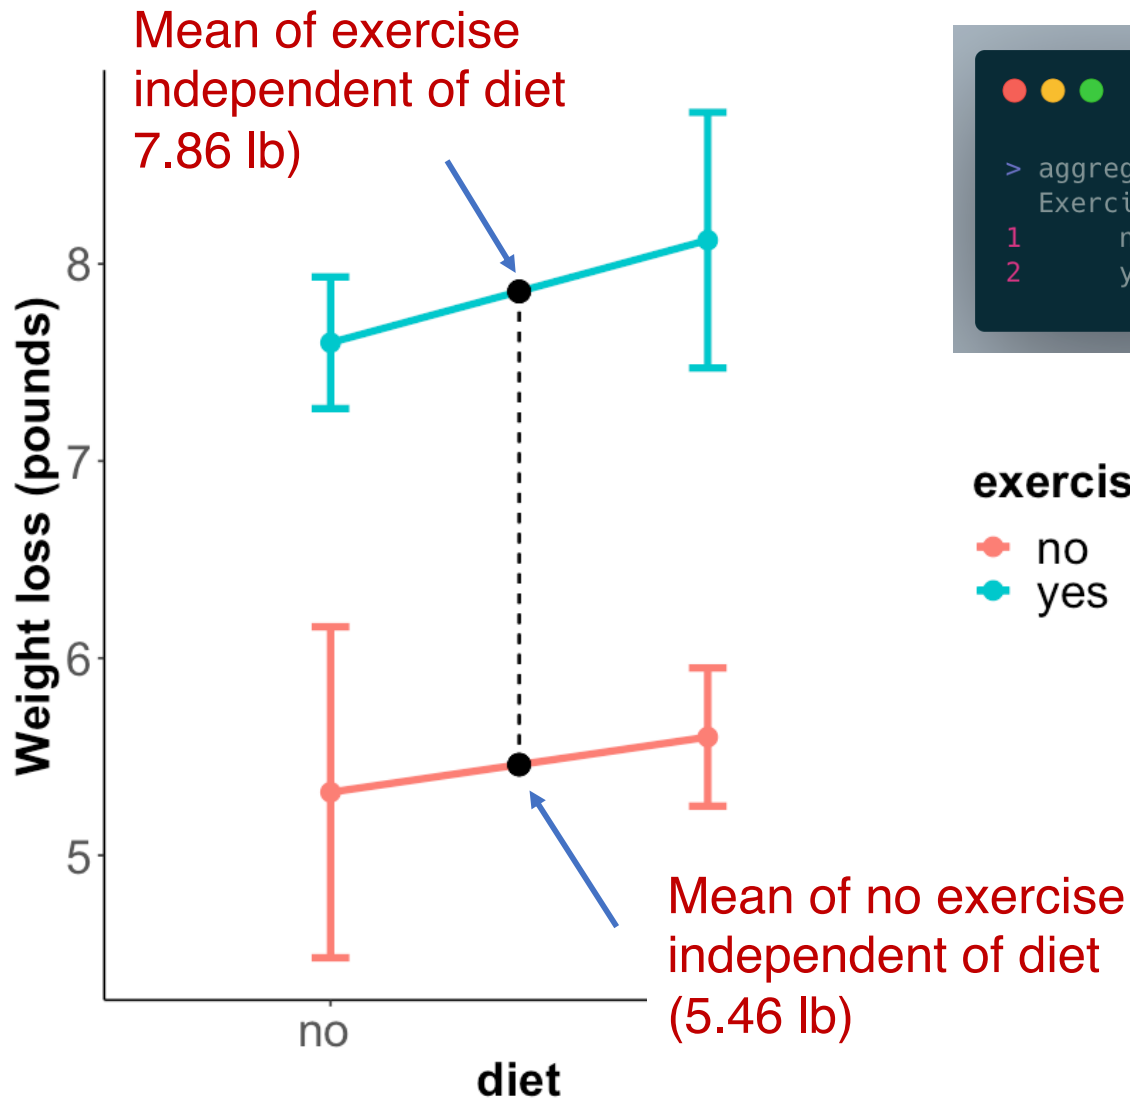
Interaction plot & 95% confidence intervals



Interaction plots are the best and most intuitive way to understand the outcomes (results) of complex factorial studies. We will use statistical significance tests in a later lecture to assess the significance of the patterns we detect visually.

Only exercise affected weight loss!
BUT HOW? Exercise increased weight loss ($P < 0.0000001$)

Understanding interaction plots

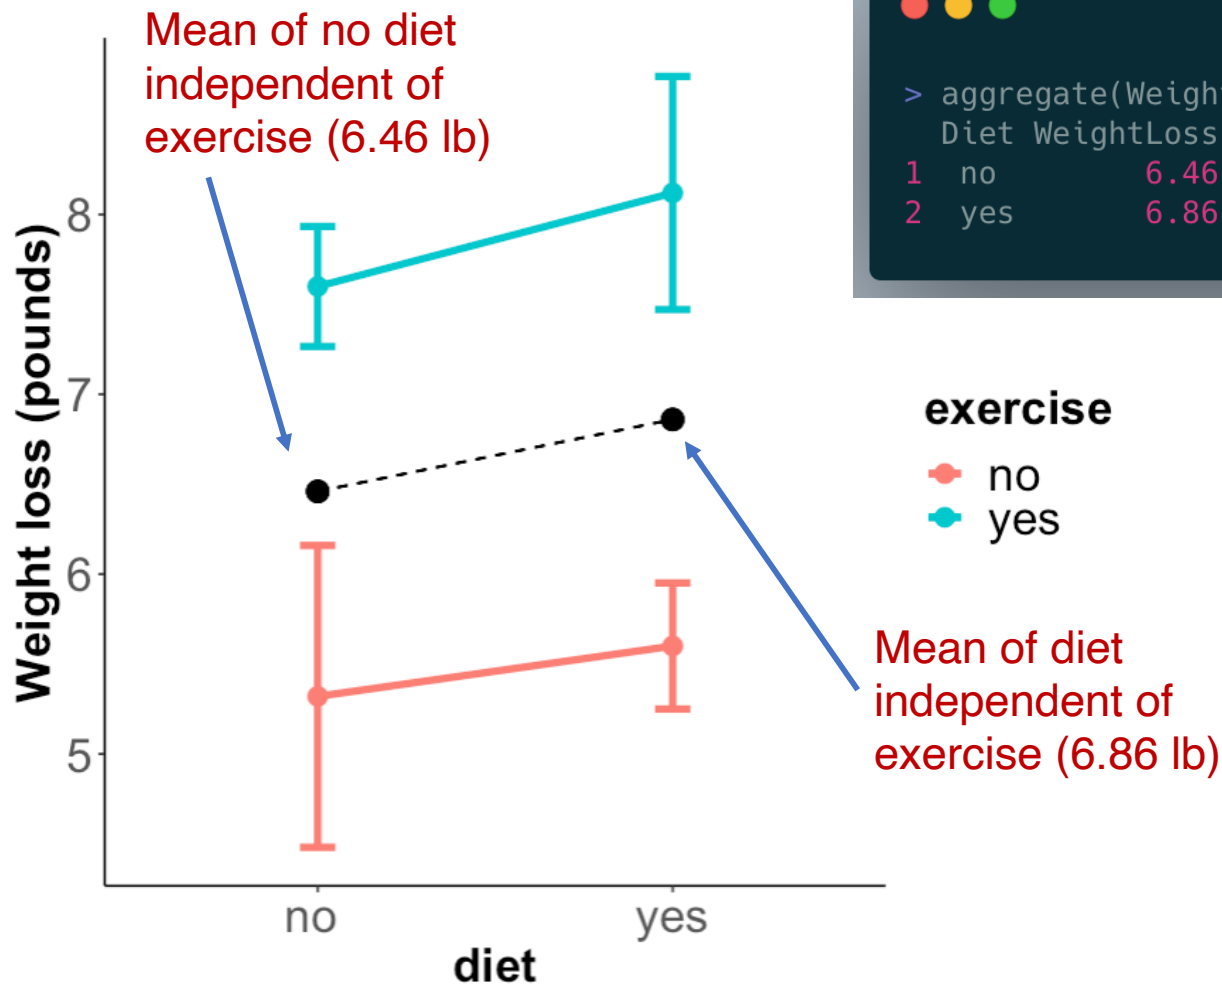


```
> aggregate(WeightLoss ~ Exercise, mean, data=diet.data)
Exercise WeightLoss
1      no          5.46
2      yes          7.86
```

NOTE: dots and dashed lines are not used in interaction plots; they are used here to facilitate understanding.

Diet did not affect weight loss! ($P=0.1974$), i.e., variation in mean (likely) due to sampling variation

Understanding interaction plots



```
> aggregate(WeightLoss ~ Diet, mean, data=diet.data)
Diet WeightLoss
1 no 6.46
2 yes 6.86
```



There are five different possible outcomes from a two-way factorial ANOVA:

CASE 1: Only one main effect is significant (either DIET or EXERCISE).



CASE 2: The two main effects are significant (both DIET AND EXERCISE) but not the interaction.

CASE 3: Only the interaction is significant.

CASE 4: One or both main factors are significant and the interaction as well.

CASE 5: No factor or interaction are significant (no need to cover this one; at least not graphically).

CASE 2: the two main effects are significant (DIET AND EXERCISE) but not the interaction.

```
anova(lm(WeightLoss~Diet*Exercise))
Analysis of Variance Table


Response: WeightLoss

      Df Sum Sq Mean Sq F value    Pr(>F)
Diet    1  6.9999   6.9999  69.8695 3.124e-07 ***
Exercise 1  7.6282   7.6282  76.1416 1.766e-07 ***
Diet:Exercise 1  0.0201   0.0201   0.2003  0.6605
Residuals 16  1.6030   0.1002


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

Note that I kept the “fictional study”, but I’ve created data for the different outcomes (cases).

CASE 2: The two main effects are significant (DIET AND EXERCISE) but not the interaction.



CASE 2: The two main effects are significant (DIET AND EXERCISE) but not the interaction.



CASE 3: Only the interaction is significant

```
anova(lm(WeightLoss~Diet*Exercise))
Analysis of Variance Table

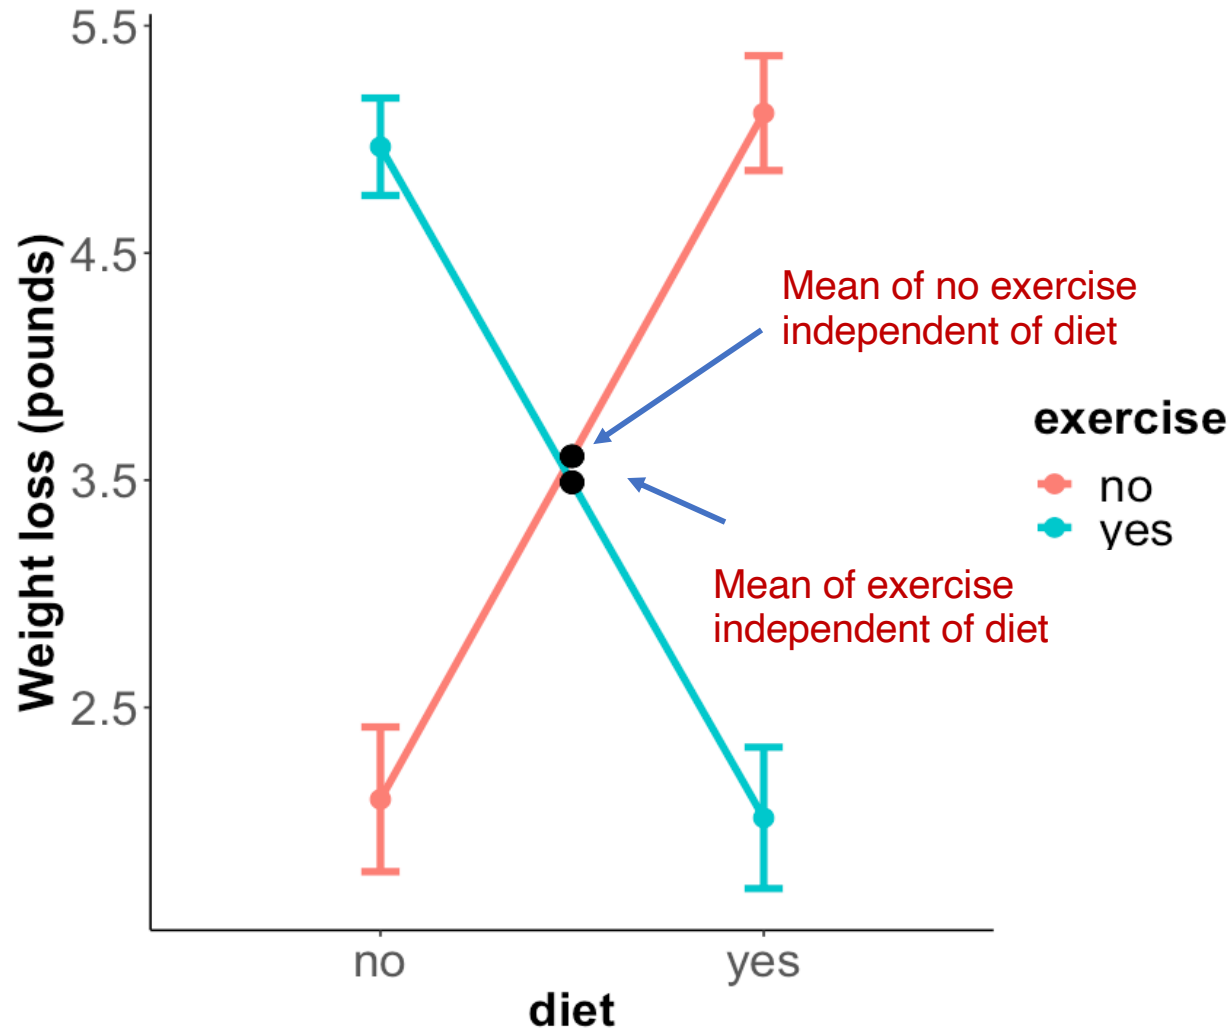
Response: WeightLoss

      Df Sum Sq Mean Sq  F value    Pr(>F)
Diet    1  0.006   0.006    0.0550    0.8175
Exercise 1  0.066   0.066    0.6581    0.4291
Diet:Exercise 1 44.600  44.600  445.1821 4.187e-13 ***
Residuals 16  1.603   0.100


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

Note that I kept the “fictional study”, but I’ve created different data set for the different possible outcomes (cases).

CASE 3: Only the interaction is significant, i.e., weight loss depends on the combinations of the levels of the main effects; greater when no diet and exercise OR when diet and no exercise.



CASE 3: Only the interaction is significant, i.e., weight loss depends on the combinations of the levels of the main effects; greater when no diet and exercise OR when diet and no exercise.



CASE 4: One or both main factors are significant and the interaction as well.

CASE 4.1: only interaction should be interpreted but not the main effect.



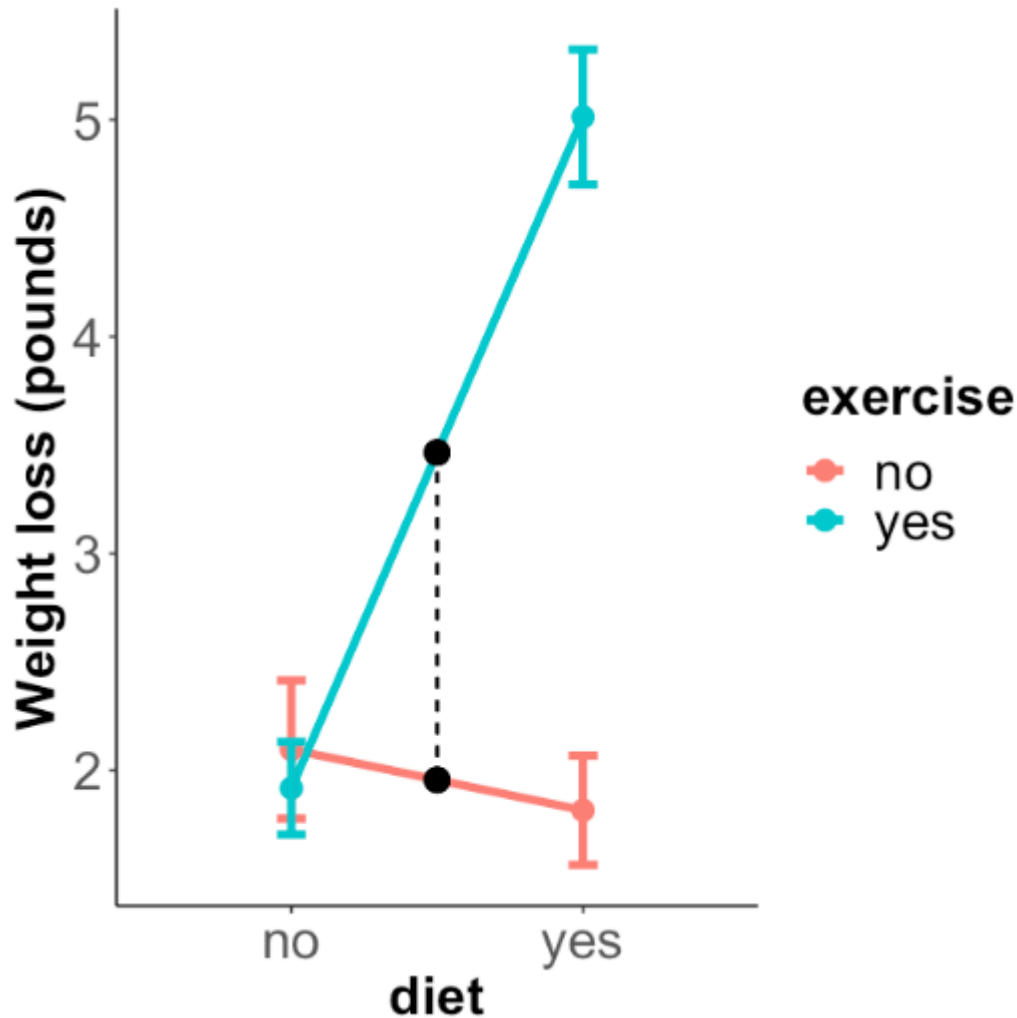
```
> anova(lm(WeightLoss~Diet*Exercise))
```

```
Analysis of Variance Table
```

```
Response: WeightLoss
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)	
Diet	1	9.9152	9.9152	98.969	2.952e-08	***
Exercise	1	11.4031	11.4031	113.820	1.108e-08	***
Diet:Exercise	1	14.2526	14.2526	142.263	2.246e-09	***
Residuals	16	1.6030	0.1002			

CASE 4: One or both main factors are significant and the interaction as well. **CASE 4.1:** only interaction should be interpreted but not the main effect.



A main effect says that there is a difference in weight loss between the exercise means, regardless of diet.

This may be technically true but only because of the big differences at diet.

It is not true that weight loss differ for the no diet case.

So, to say that there is weight loss regardless of diet (main effect) is not accurate!

CASE 4: One or both main factors are significant and the interaction as well.

CASE 4.2: the interaction & main effect can be interpreted.

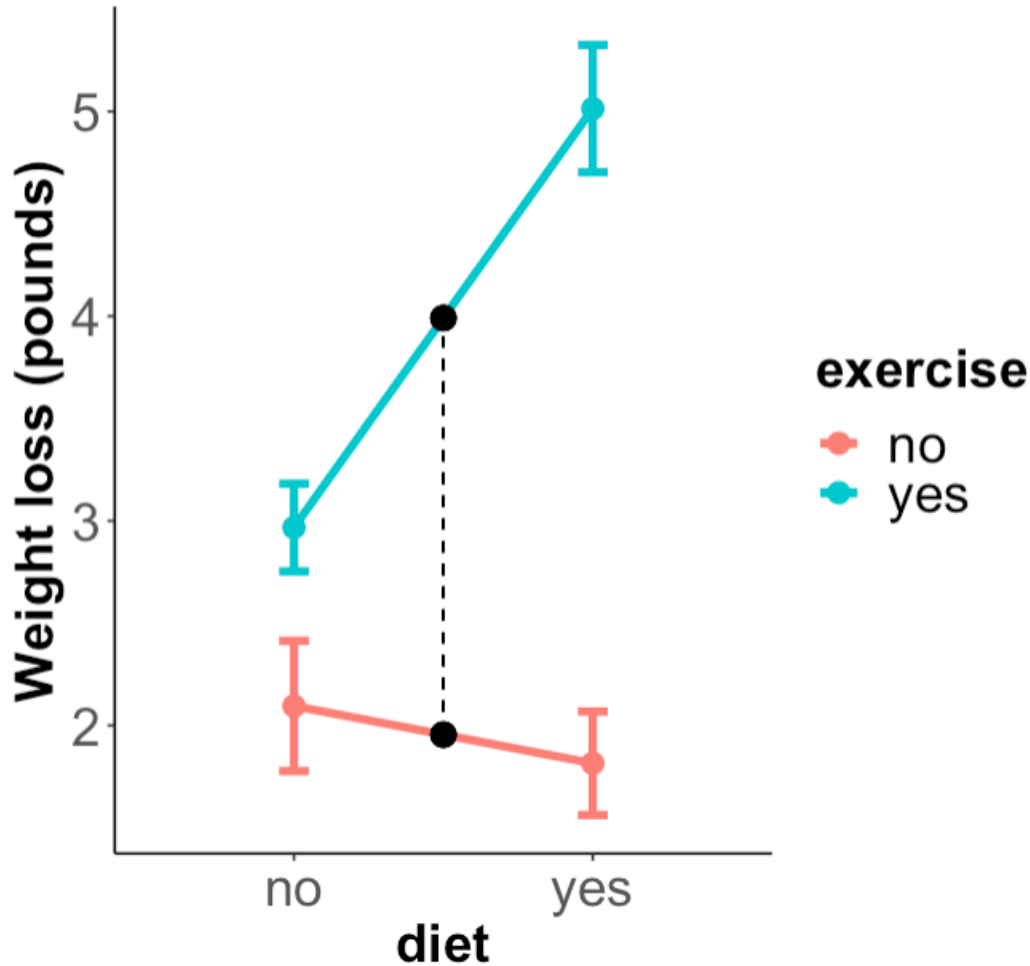


```
> anova(lm(WeightLoss~Diet*Exercise))  
Analysis of Variance Table
```

```
Response: WeightLoss
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)	
Diet	1	3.9002	3.9002	38.931	1.183e-05	***
Exercise	1	20.7096	20.7096	206.714	1.440e-10	***
Diet:Exercise	1	6.7669	6.7669	67.544	3.902e-07	***
Residuals	16	1.6030	0.1002			

CASE 4: One or both main factors are significant and the interaction as well. **CASE 4.2:** the interaction & main effect can be interpreted.




A main effect says that there is a difference in weight loss between the exercise means, regardless of diet.

This is the case because the weight loss when individuals exercised is consistently greater than no exercise regardless of the diet.

And, individuals that exercised and dieted loss even more weight than individuals than only dieted.

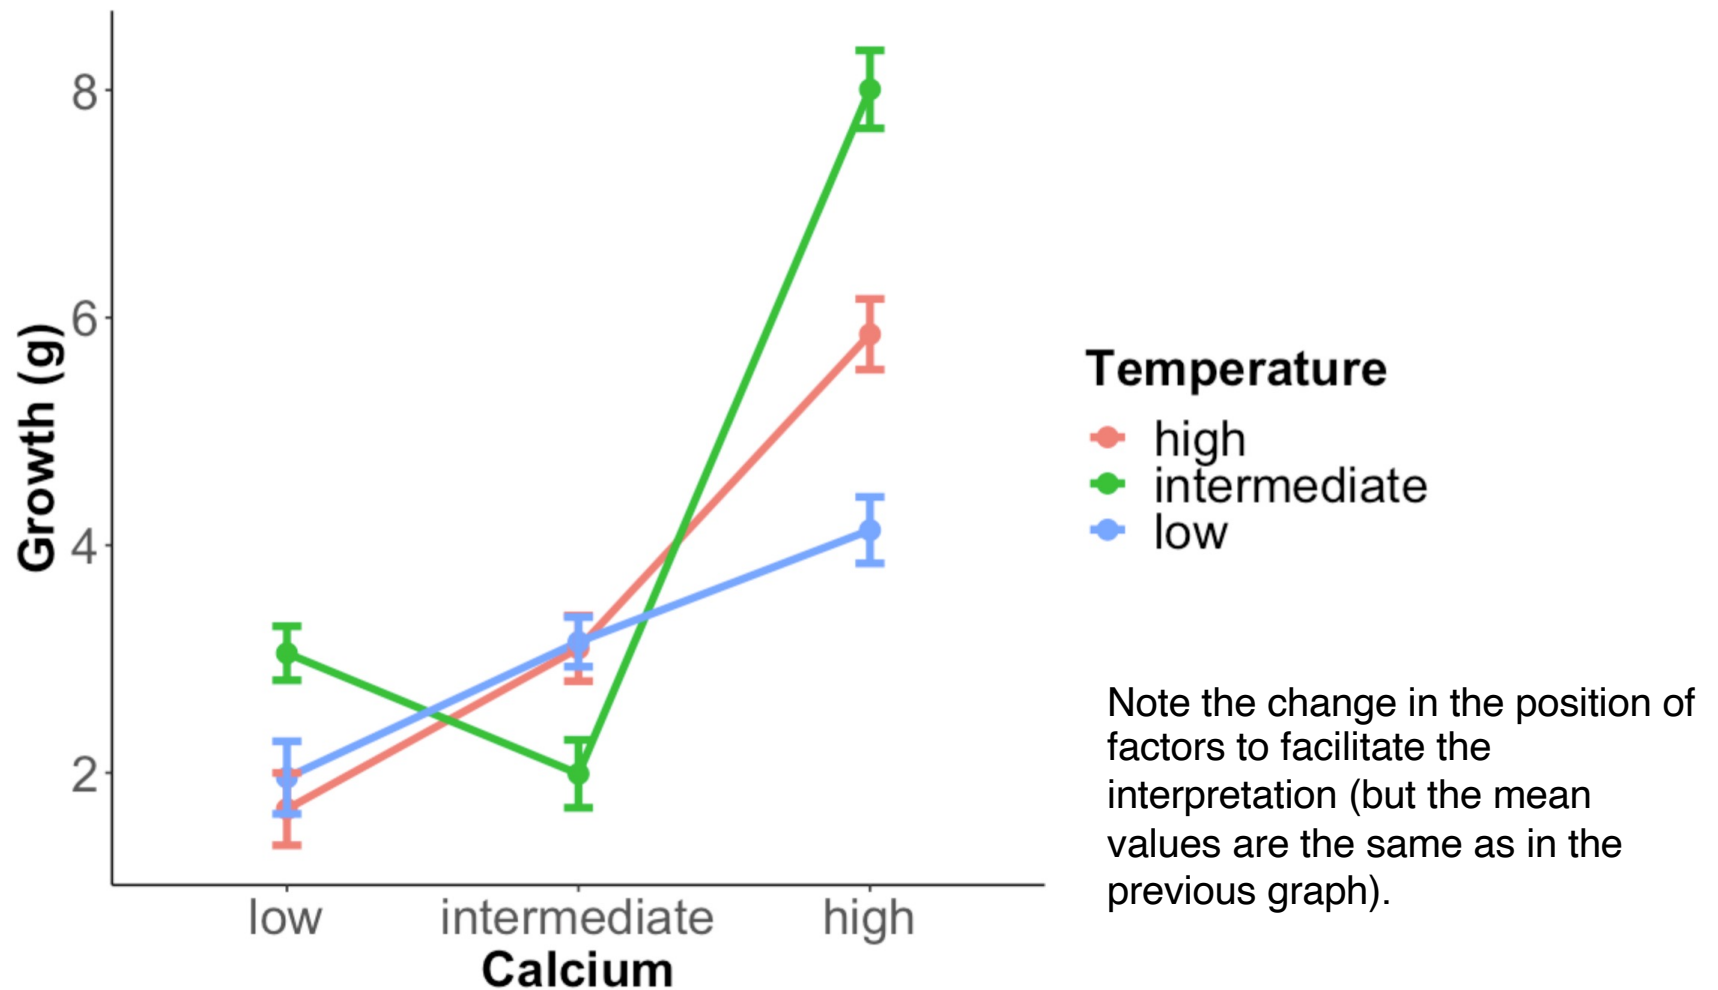
Should we interpret only the **interaction** or also the differences in **calcium** independent of temperature?

YES; one could interpret only the interaction because one level of calcium (high) is greater in average than all other levels. In this case one could state that calcium can affect growth independent of temperature; even given the fact the greatest growth is at an intermediate temperature and at a high level of calcium.



Should we interpret only the **interaction** or also the differences in **temperature** independent of calcium?

NO; one should not interpret the interaction because no level of temperature is greater or smaller in average regardless of calcium levels. In this case, one could not state that temperature can affect growth independent of calcium.





Multi-factorial ANOVA

Assumptions (the same as for the one-way ANOVA):

- 1) Each of the samples (observations within groups) is a random sample from its population (LATER IN THE COURSE).
- 2) The variable (e.g., weight loss) is normally distributed in each combination of treatment (e.g., no diet and exercise) population.
- 3) The variances are equal among all populations from which the treatments were sampled (otherwise the F values change in ways that may not measure difference among means).

Assessing the normality assumption


- ANOVAs are not very sensitive to lack of normality (i.e., they are robust against normality).
- Simulation studies, using a variety of non-normal distributions, have shown that the false positive rates (Type I error rates) in ANOVA are not strongly affected by the violation of the normality assumption (Harwell et al. 1992, Lix et al. 1996).

Assessing the normality assumption – some traditional tests

Test	Advantages	Disadvantages
Chi-Square test	<ul style="list-style-type: none">• appropriate for any level of measurement• ties may be problematic	<ul style="list-style-type: none">• grouping of observations required (frequencies per group must be > 5)• unsuitable for small samples• statistic based on squares
Kolmogorov-Smirnov test	<ul style="list-style-type: none">• suitable for small samples• ties are no problem• omnibus test	<ul style="list-style-type: none">• no categorical data• low power if prerequisites are not met
Lilliefors test	<ul style="list-style-type: none">• higher power than KS test	<ul style="list-style-type: none">• no categorical data
Anderson-Darling test	<ul style="list-style-type: none">• high power when testing for normal distribution• more precise than KS test (especially in the outer parts of the distribution)	<ul style="list-style-type: none">• no categorical data• statistic based on squares
Shapiro-Wilk test	<ul style="list-style-type: none">• highest power among all tests for normality	<ul style="list-style-type: none">• test for normality only• computer required due to complicated procedure
Cramér-von-Mises test	<ul style="list-style-type: none">• higher power than KS test	<ul style="list-style-type: none">• statistic based on squares• no categorical data

Assessing the normality assumption: The Quantile-Quantile normal plot (Q-Q normal plot)

The normal Q-Q plot is a graphical technique for determining if multiple data sets come from populations with a common distribution (here, if they all come from normally distributed populations regardless of their means and variances).



Tutorial 3: Factorial ANOVA

Factorial Analysis of Variance



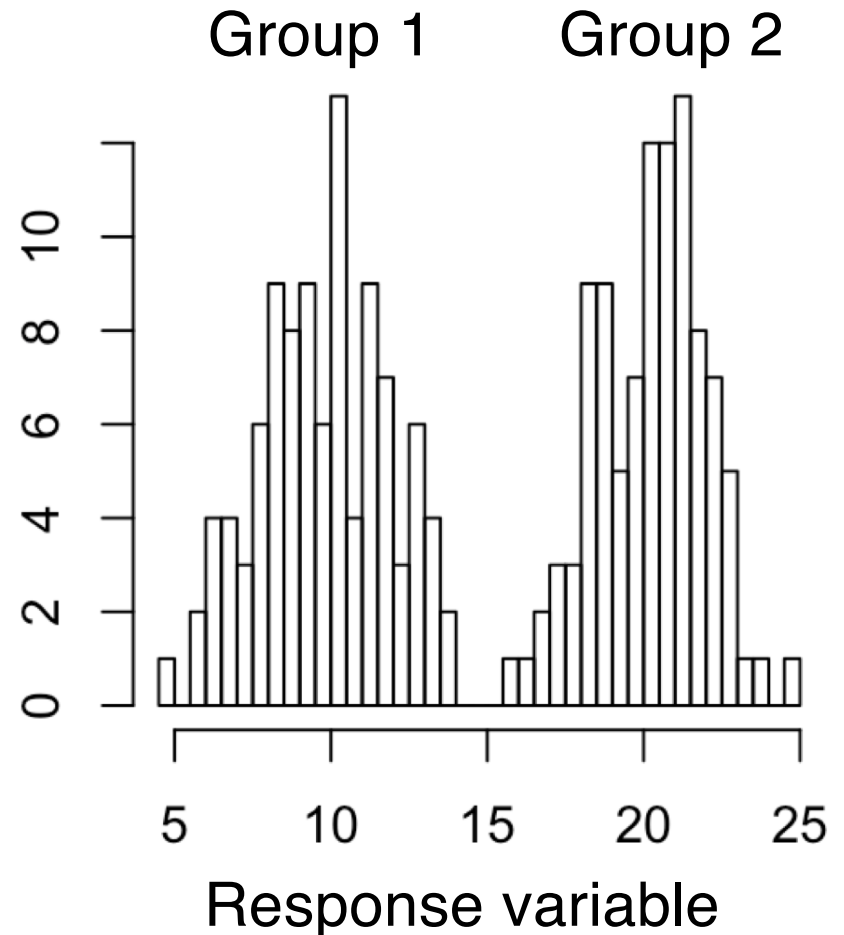
Assessing the normality assumption in linear models: The Quantile-Quantile normal (Q-Q normal plot)

In ANOVAs, it is not the response variable (Weight loss) as a whole that is required to be normal, but rather the response within groups.



Usual interpretation of the normality assumption in ANOVAs -
“Data have to be normal”

```
n <- 100  
Group1 <- rnorm(n, 10, 2)  
Group2 <- rnorm(n, 20, 2)  
hist(c(Group1, Group2), breaks=30)
```

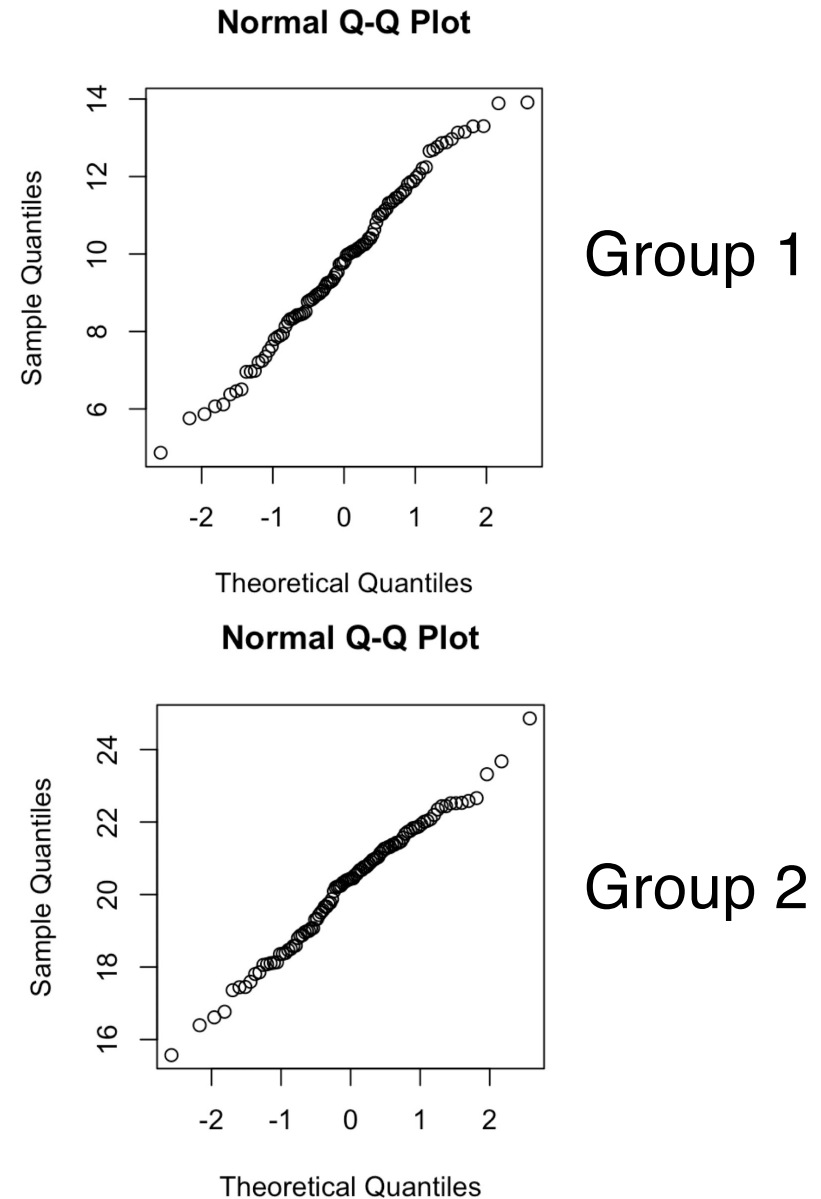

Response variable not normal across groups, but normal within groups (the correct assumption).



Assessing the normality assumption in linear models:
The Quantile-Quantile normal (Q-Q normal plot)




Assessing the normality assumption in linear models: The Quantile-Quantile normal (Q-Q normal plot)




Assessing the normality assumption in linear models: The Quantile-Quantile normal (Q-Q normal plot)

If there are too many groups and too many factors (e.g., multi-factorial ANOVA), it becomes impossible to analyze all Q-Q plots for all combinations of levels across factors!

Ex. 2 factors with 3 levels for factor 1 and 4 levels for factor 2 = 12 groups!



Group 1





Group 2

Assessing the normality assumption in linear models: The Quantile-Quantile normal plot of residuals (Q-Q normal residual plot)

ANOVA is a linear multiple regression model in which the response variable is continuous, and predictors are categorical.

$$Y = \text{Factor}(G1, G2) + \text{residuals}$$

So, instead of plotting all groups, we assess the residuals across all groups, i.e., variation not accounted by group mean differences.




Assessing the normality assumption in linear models: The Quantile-Quantile normal plot of residuals (Q-Q normal residual plot)

You will practice the application of Q-Q normal residual plots for two-factorials ANOVAs in tutorial 3.

$$Weight_{Loss} = Diet + Exercise + Diet \times Exercise + residuals$$

Tutorial 3: Factorial ANOVA

Factorial Analysis of Variance



Assessing the equality of variance (homoscedasticity) assumption

The two main assessments for testing the null hypothesis that multiple samples come from populations with equal variances are:

Levene's test and Barlett test (more sensitive to non-normality than Levene's).

H₀: The samples come from populations with the same variance.

H_A: At least two samples come from populations with different variances.

Assessing the equality of variance (homoscedasticity) assumption

```
n <- 100
Group1 <- rnorm(n, 10, 2)
Group2 <- rnorm(n, 20, 2)
Factor <- c(rep(1, n), rep(2, n))
```

```
> var(Group1)
[1] 3.911981
> var(Group2)
[1] 4.022584
```

The two samples come from populations with the same variances (they only vary in mean values).

```
Levene's Test for Homogeneity of Variance (center = median)
      Df F value Pr(>F)
group  1   0.428 0.5137
      198
```

Conclusion?

Assessing the equality of variance (homoscedasticity) assumption

```
n <- 100
Group1 <- rnorm(n, 10, 2)
Group2 <- rnorm(n, 20, 3)
Factor <- c(rep(1, n), rep(2, n))
```

```
> var(Group1)
[1] 4.11724
> var(Group2)
[1] 7.693817
```

The two samples come from populations with different variances (and they also vary in their means).

```
> leveneTest(c(Group1, Group2) ~ as.factor(Factor))
Levene's Test for Homogeneity of Variance (center = median)
      Df F value Pr(>F)
group  1  6.3814 0.01232 *
      198
```

Conclusion?

Let's contrast the Levene's and ANOVAs hypotheses

Levene's:

H₀: The samples come from populations with the same variance.

H_A: At least two samples come from populations with different variances.

ANOVA:

H₀: The samples come from the same population.

H_A: At least two samples come from different populations.

- If they have the same variances and same means, then we can state under the null hypothesis that they come from the same population. Remember, we should test for differences in variance (Levene's) before conducting an ANOVA.



A more complex (and real) biological data

Regional and strain-specific gene expression mapping in the adult mouse brain

11038–11043 | PNAS | September 26, 2000 | vol. 97 | no. 20

Rickard Sandberg*[†], Rie Yasuda*[‡], Daniel G. Pankratz*, Todd A. Carter*, Jo A. Del Rio[§], Lisa Wodicka[§], Mark Mayford[‡], David J. Lockhart[§], and Carrolee Barlow*[¶]

To determine the genetic causes and molecular mechanisms responsible for neurobehavioral differences in mice, we used highly parallel gene expression profiling to detect genes that are differentially expressed between the 129SvEv and C57BL/6 mouse strains at baseline and in response to seizure. In addition, we identified genes that are differentially expressed in specific brain regions. We found that approximately 1% of expressed genes are differentially expressed between strains in at least one region of the brain and that the gene expression response to seizure is significantly different between the two inbred strains. The results lead to the identification of differences in gene expression that may account for distinct phenotypes in inbred strains and the unique functions of specific brain regions.

Gene expression is standardized in relation to seizure versus base line

Case 1 - What are the significant effects?
a gene for which only strain is significant (i.e., they differ in gene expression levels)

gene aa119706.at - Only strain is significant (i.e., strains differ from one another in their mean gene expression levels, but these differences are independent of the brain region)


```
Response: aa119706.at
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)	
strain	1	45850	45850	15.5796	0.001938	**
brain.region	5	7434	1487	0.5052	0.767145	
strain:brain.region	5	2291	458	0.1557	0.974152	
Residuals	12	35315	2943			

```
---
```

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


gene aa217379.s.at - Only strain is significant



Black lines represent means of each strain independent of brain region and their differences are significant.

Note that the B6 and X129 lines (orange and green) do look parallel

gene aa217379.s.at - Only strain is significant





Black lines represent means of each brain region independent of strain and their differences are not significant.

Note that the brain region lines (6 colors) do look parallel

Which factor to plot where?

It depends on how differences facilitate interpretation.



Case 2 - What are the significant effects?
a gene for which only the brain region is significant


gene AA166452.at - Only brain region is significant (i.e., regions differ from one another in their mean gene expression levels. but these differences are independent of the strain)

```
Response: AA166452.at
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)	
strain	1	176	176	0.0150	0.9046	
brain.region	5	1435582	287116	24.4269	6.67e-06	***
strain:brain.region	5	21824	4365	0.3713	0.8587	
Residuals	12	141049	11754			

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


AA166452.at - Only brain region is significant



Black lines represent means of each strain independent of brain region and their differences are not significant.

Note that the B6 and X129 lines (orange and green) do look parallel

AA166452.at - Only brain region is significant





Black lines represent means of each brain region independent of strain and their differences are significant – BUT, which brain regions differ from one another? There are 15 possible pairwise comparisons (tests).

Note that the brain region lines (6 colors) do look parallel

Which factor to plot where?

It depends on how differences facilitate interpretation



Case 3 - What are the significant effects?
a gene for which only the interaction is significant

gene aa051500.at - Only the interaction between brain regions and strain is significant (i.e., differences in mean gene expression levels of brain regions depend on strain, or vice-versa)


```
Response: aa051500.at
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)	
strain	1	852.0	852.04	2.4399	0.144256	
brain.region	5	2212.9	442.58	1.2674	0.339231	
strain:brain.region	5	9087.2	1817.44	5.2045	0.009038	**
Residuals	12	4190.5	349.21			

```
---
```

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

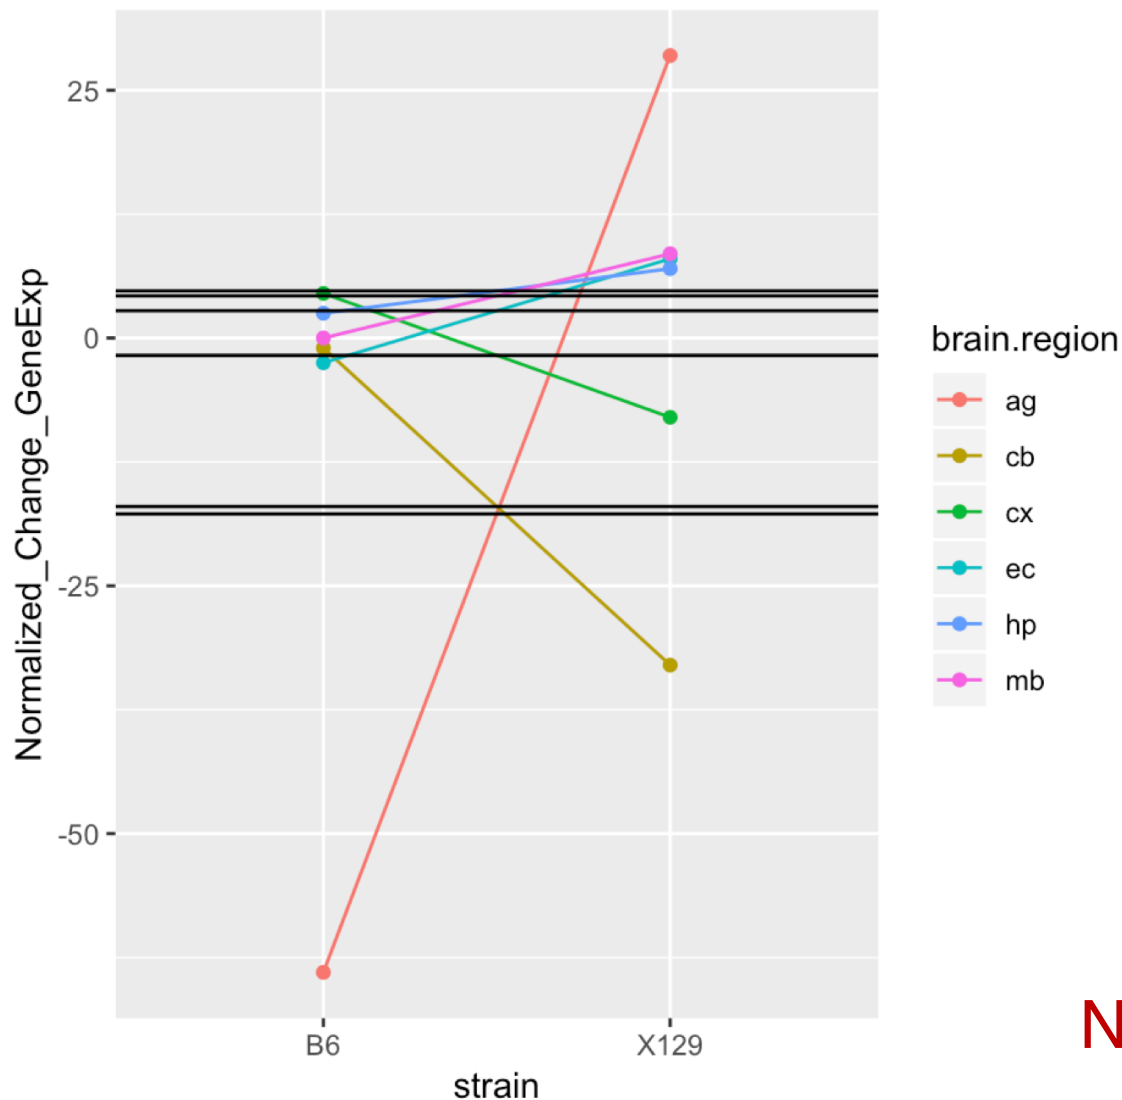
aa051500.at - Only the interaction between strain and brain region is significant



Black lines represent means of each strain independent of brain region and their differences are not significant – but some genes are more different than the other depending on the brain region

Note that the lines for strain are NOT parallel

aa051500.at - Only the interaction between strain and brain region is significant



Black lines represent means of each brain region independent of strain and their differences are not significant – but some genes are more different than the other depending on the strain. There are 66 possible pairwise comparisons (tests).

Note that the lines for brain region are NOT parallel

Case 4 - What are the significant effects? a gene for which at least one main factor and the interaction is significant

gene AA107725.f.at - The mean gene expression levels in brain regions vary, and the mean differences depend on the strain


```
Response: AA107725.f.at
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)	
strain	1	715.0	715.0	1.6751	0.2199350	
brain.region	5	12941.7	2588.3	6.0635	0.0050251	**
strain:brain.region	5	19124.7	3824.9	8.9603	0.0009664	***
Residuals	12	5122.5	426.9			

```
---
```

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


AA107725.f.at – brain region differs in gene expression
and interaction is significant
(only interpret interactions seem to be interpretable)



Black lines represent means of each strain independent of brain region and their differences are not significant – but some genes are more different than the other depending on the brain region

Note that the lines are NOT parallel

AA107725.f.at – brain region differs in gene expression
and interaction is significant
(only interpret interactions seem to be interpretable)



Black lines represent means of each brain region independent of strain and their differences are significant – and some genes are more different than the other depending on the strain

Note that the lines are NOT parallel

A word on balanced designs

The ANOVAs performed here (and in tutorial 3) are based on equal number of observations per combination of groups.

In the fictional diet example, there are 5 individuals in each of the 4 combinations of diet (yes/no) and exercise (yes/no).

In the gene expression study, there are 2 individuals in each of the 12 combinations of strain (2 strains) and brain region (6 regions).

For balanced designs, we say that the design is fully orthogonal because there is no variation that is shared between factors (a concept we will see in a few lectures; under ANCOVA).

For fully orthogonal designs, we use what is called a Type I Sum-of-Squares (Type I SS). When factors are not fully orthogonal, then we use the Type III SS (Sum-of-Squares). We will learn about Type III in the ANCOVA module).

Lecture 6: Which effects are significant?

which pairwise means to compare?

