Factorial ANOVA: main effects & interactions

Week 9

Concepts

- 2-way factorial experimental designs
- marginal means
- main effects & interaction effects
- ANOVA output table for a 2-way factorial design
- interpreting main effects vs interaction effects
- *next week: assumptions, more complex factorial designs, & post-hoc tests*

Factorial ANOVA

When to use:

- when you have 2 or more independent variables (factors)
- when you want to test
  - the **main effect** of each factor
  - the **interaction effect** between factors

Factorial ANOVA

Notation:

- 2x2 factorial design: *2 independent variables, each with 2 levels*
- 2x3 factorial design: *2 independent variables, one with 2 levels, one with 3 levels*
- 3x4 factorial design: *2 independent variables, one with 3 levels, one with 4 levels*
- etc.

2x2 Design

<table>
<tr><td colspan="2" rowspan="2">2x2 Design</td><td colspan="2">IV 1</td></tr>
<tr><td>IV1: Level 1</td><td>IV1: Level 2</td></tr>
<tr><td rowspan="2">IV 2</td><td>IV2: Level 1</td><td>dv</td><td>dv</td></tr>
<tr><td>IV2: Level 2</td><td>dv</td><td>dv</td></tr>
</table>

<table>
<tr><td colspan="2" rowspan="2">2x2 Design</td><td colspan="2">Time of Day</td></tr>
<tr><td>Morning</td><td>Afternoon</td></tr>
<tr><td rowspan="2">Caffeine</td><td>Some Caffeine</td><td>dv</td><td>dv</td></tr>
<tr><td>No Caffeine</td><td>dv</td><td>dv</td></tr>
</table>

2x3 Design

2x3 Design		IV 1	
		IV1: Level 1	IV1: Level 2
IV 2	IV2: Level 1	dv	dv
	IV2: Level 2	dv	dv
	IV2: Level 3	dv	dv

2x3 Design		Time of Day	
		Morning	Afternoon
Caffeine	1 coffee	dv	dv
	2 coffees	dv	dv
	3 coffees	dv	dv

Fully Factorial

- An experimental design is *fully factorial* when the levels of each factor are *fully crossed* with the levels of each other factor
- e.g. in a 2x2 design with factors A (A1,A2) and B (B1,B2), there are 4 conditions
  - A1B1, A1B2, A2B1, A2B2
  - this design is *fully factorial* if there are data for each of these 4 conditions

Fully Factorial

- An experimental design is *fully factorial* when the levels of each factor are *fully crossed* with the levels of each other factor
- in a 2x3 design, (A1,A2), (B1,B2,B3) there are 6 conditions
  - A1B1, A1B2, A1B3, A2B1, A2B2, A2B3
  - this design is *fully factorial* if there are data for each of these 6 conditions

Example Data

- sample data file from Navarro text: clinicaltrial.Rdata

```r
library(tidyverse)
load(url("https://www.gribblelab.o
```

```r
(clin.trial <- tibble(clin.trial))
```

```
# A tibble: 18 × 3
   drug     therapy    mood.gain
   <fct>    <fct>          <dbl>
 1 placebo  no.therapy       0.5
 2 placebo  no.therapy       0.3
 3 placebo  no.therapy       0.1
 4 anxifree no.therapy       0.6
 5 anxifree no.therapy       0.4
 6 anxifree no.therapy       0.2
 7 joyzepam no.therapy       1.4
 8 joyzepam no.therapy       1.7
 9 joyzepam no.therapy       1.3
10 placebo  CBT              0.6
11 placebo  CBT              0.9
12 placebo  CBT              0.3
13 anxifree CBT              1.1
14 anxifree CBT              0.8
15 anxifree CBT              1.2
16 joyzepam CBT              1.8
17 joyzepam CBT              1.3
18 joyzepam CBT              1.4
```

- DV is `mood.gain`
- 2 IVs: `drug` (3 levels) and `therapy` (2 levels)

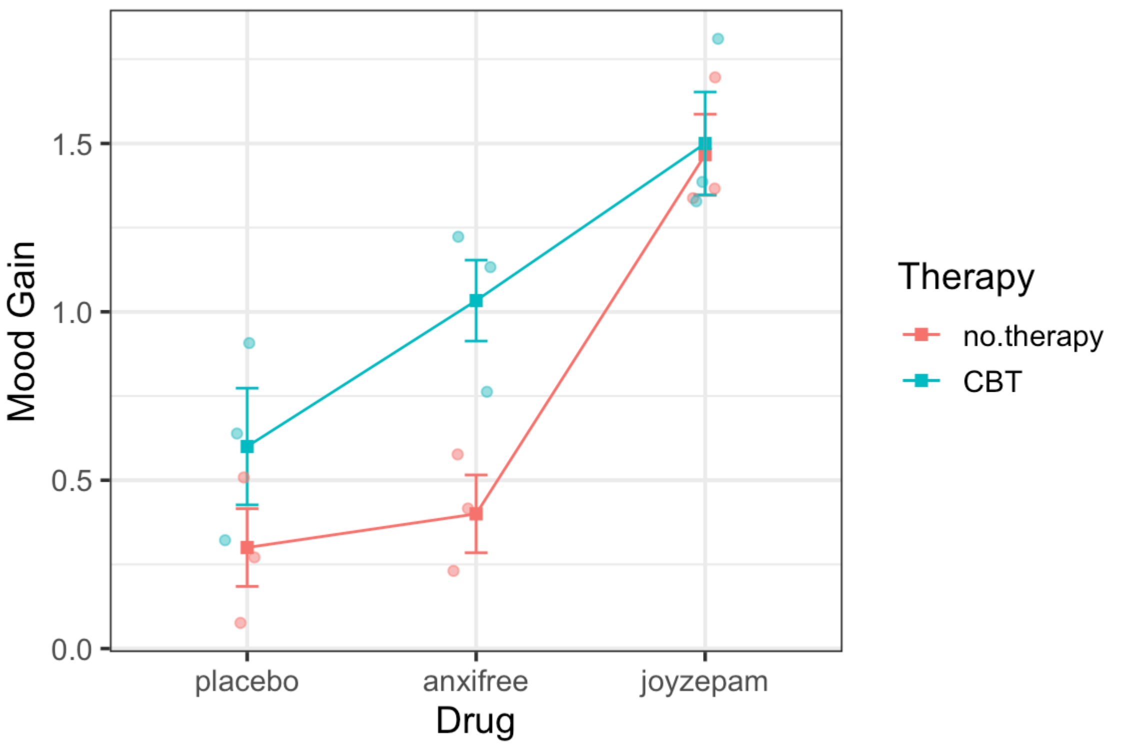

Marginal Means

```
1 clin.trial
```

```
# A tibble: 18 × 3
   drug     therapy    mood.gain
   <fct>    <fct>          <dbl>
 1 placebo  no.therapy       0.5
 2 placebo  no.therapy       0.3
 3 placebo  no.therapy       0.1
 4 anxifree no.therapy       0.6
 5 anxifree no.therapy       0.4
 6 anxifree no.therapy       0.2
 7 joyzepam no.therapy       1.4
 8 joyzepam no.therapy       1.7
 9 joyzepam no.therapy       1.3
10 placebo  CBT              0.6
11 placebo  CBT              0.9
12 placebo  CBT              0.3
13 anxifree CBT              1.1
14 anxifree CBT              0.8
15 anxifree CBT              1.2
16 joyzepam CBT              1.8
17 joyzepam CBT              1.3
18 joyzepam CBT              1.4
```

- **marginal means** are the average of the DV for each level of a factor, **across all levels of the other factor**
  - (ignoring the other factor)
- marginal means for `drug`:
- `placebo`: (0.5 + 0.3 + 0.1 + 0.6 + 0.9 + 0.3) / 6 = 0.45
- `anxifree`: (0.6 + 0.4 + 0.2 + 1.1 + 0.8 + 1.2) / 6 = 0.72
- `joyzepam`: (1.4 + 1.7 + 1.3 + 1.8 + 1.3 + 1.4) / 6 = 1.48
- marginal means for `therapy`:
- `no.therapy`: (0.5 + 0.3 + 0.1 + 0.6 + 0.4 + 0.2 + 1.4 + 1.7 + 1.3) / 9 = 0.72
- `CBT`: (0.6 + 0.9 + 0.3 + 1.1 + 0.8 + 1.2 + 1.8 + 1.3 + 1.4) / 9 = 1.04

Marginal Means

```
clin.trial
```

```
# A tibble: 18 × 3
   drug     therapy    mood.gain
   <fct>    <fct>          <dbl>
 1 placebo  no.therapy       0.5
 2 placebo  no.therapy       0.3
 3 placebo  no.therapy       0.1
 4 anxifree no.therapy       0.6
 5 anxifree no.therapy       0.4
 6 anxifree no.therapy       0.2
 7 joyzepam no.therapy       1.4
 8 joyzepam no.therapy       1.7
 9 joyzepam no.therapy       1.3
10 placebo  CBT              0.6
11 placebo  CBT              0.9
12 placebo  CBT              0.3
13 anxifree CBT              1.1
14 anxifree CBT              0.8
15 anxifree CBT              1.2
16 joyzepam CBT              1.8
17 joyzepam CBT              1.3
18 joyzepam CBT              1.4
```

- marginal means for `drug`:

```
clin.trial %>%
  group_by(drug) %>%  # (we are ignoring therapy)
  summarise(mean(mood.gain))
```

```
# A tibble: 3 × 2
  drug     `mean(mood.gain)`
  <fct>                <dbl>
1 placebo              0.45
2 anxifree             0.717
3 joyzepam             1.48
```

- marginal means for `therapy`:

```
clin.trial %>%
  group_by(therapy) %>%  # (we are ignoring drug)
  summarise(mean(mood.gain))
```

```
# A tibble: 2 × 2
  therapy    `mean(mood.gain)`
  <fct>                  <dbl>
1 no.therapy             0.722
2 CBT                    1.04
```

Marginal Means

```r
clin.trial
```

```
# A tibble: 18 × 3
   drug     therapy    mood.gain
   <fct>    <fct>          <dbl>
 1 placebo  no.therapy       0.5
 2 placebo  no.therapy       0.3
 3 placebo  no.therapy       0.1
 4 anxifree no.therapy       0.6
 5 anxifree no.therapy       0.4
 6 anxifree no.therapy       0.2
 7 joyzepam no.therapy       1.4
 8 joyzepam no.therapy       1.7
 9 joyzepam no.therapy       1.3
10 placebo  CBT              0.6
11 placebo  CBT              0.9
12 placebo  CBT              0.3
13 anxifree CBT              1.1
14 anxifree CBT              0.8
15 anxifree CBT              1.2
16 joyzepam CBT              1.8
17 joyzepam CBT              1.3
18 joyzepam CBT              1.4
```

- marginal means for `drug`:

```r
aggregate(mood.gain ~ drug, data=clin.trial, FUN=mean)
```

```
      drug mood.gain
1  placebo 0.4500000
2 anxifree 0.7166667
3 joyzepam 1.4833333
```

- marginal means for `therapy`:

```r
aggregate(mood.gain ~ therapy, data=clin.trial, FUN=mean)
```

```
     therapy mood.gain
1 no.therapy 0.7222222
2        CBT 1.0444444
```

- is there an effect of `drug` on `mood.gain`?
- is there an effect of `therapy` on `mood.gain`?
- let's use ANOVA to test these effects

Testing Main Effects separately

```
clin.trial
```

```
# A tibble: 18 × 3
   drug     therapy    mood.gain
   <fct>    <fct>          <dbl>
 1 placebo  no.therapy       0.5
 2 placebo  no.therapy       0.3
 3 placebo  no.therapy       0.1
 4 anxifree no.therapy       0.6
 5 anxifree no.therapy       0.4
 6 anxifree no.therapy       0.2
 7 joyzepam no.therapy       1.4
 8 joyzepam no.therapy       1.7
 9 joyzepam no.therapy       1.3
10 placebo  CBT              0.6
11 placebo  CBT              0.9
12 placebo  CBT              0.3
13 anxifree CBT              1.1
14 anxifree CBT              0.8
15 anxifree CBT              1.2
16 joyzepam CBT              1.8
17 joyzepam CBT              1.3
18 joyzepam CBT              1.4
```

- we could fit a one-way ANOVA to test the **main effect** of the `drug` factor, ignoring `therapy`:

```
mod.drug    <- lm(mood.gain ~ drug    , data=clin.trial)
summary(aov(mod.drug    ))
```

```
            Df  Sum Sq Mean Sq F value   Pr(>F)
drug         2   3.453  1.7267   18.61 8.65e-05
Residuals   15   1.392  0.0928
```

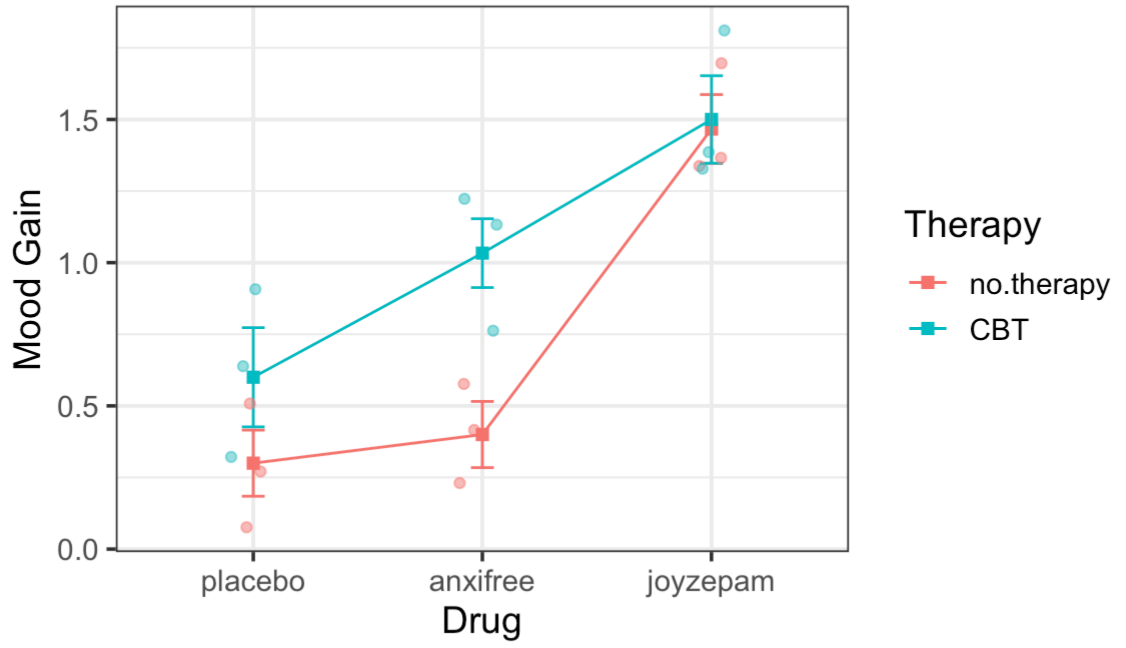

Testing Main Effects separately

```r
clin.trial
```

```
# A tibble: 18 × 3
   drug     therapy    mood.gain
   <fct>    <fct>          <dbl>
 1 placebo  no.therapy       0.5
 2 placebo  no.therapy       0.3
 3 placebo  no.therapy       0.1
 4 anxifree no.therapy       0.6
 5 anxifree no.therapy       0.4
 6 anxifree no.therapy       0.2
 7 joyzepam no.therapy       1.4
 8 joyzepam no.therapy       1.7
 9 joyzepam no.therapy       1.3
10 placebo  CBT              0.6
11 placebo  CBT              0.9
12 placebo  CBT              0.3
13 anxifree CBT              1.1
14 anxifree CBT              0.8
15 anxifree CBT              1.2
16 joyzepam CBT              1.8
17 joyzepam CBT              1.3
18 joyzepam CBT              1.4
```

- we could then fit a second one-way ANOVA to test the **main effect** of the `therapy` factor, ignoring `drug`:

```r
mod.therapy <- lm(mood.gain ~ therapy, data=clin.trial)
summary(aov(mod.therapy))
```

```
            Df Sum Sq Mean Sq F value Pr(>F)
therapy      1  0.467  0.4672   1.708   0.21
Residuals   16  4.378  0.2736
```

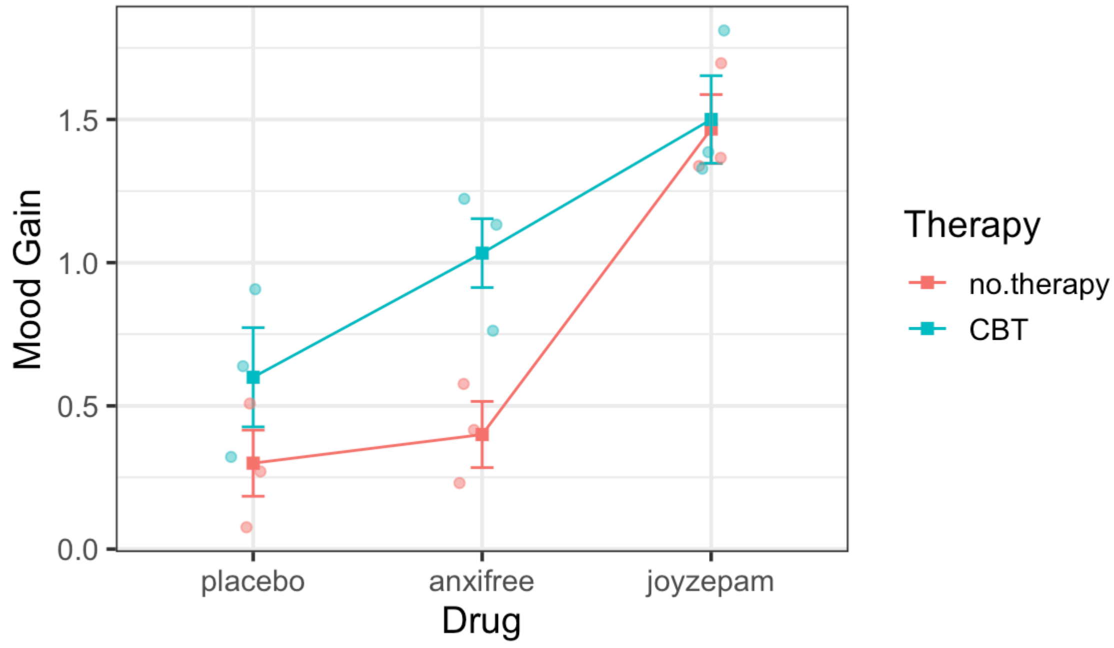

Testing Main Effects together

- **main effect** of `drug`:

```r
mod.drug <- lm(mood.gain ~ drug, data
summary(aov(mod.drug))
```

```
          Df Sum Sq Mean Sq F value   Pr(>F)
drug       2  3.453  1.7267   18.61 8.65e-05
Residuals 15  1.392  0.0928
```

- **main effect** of `therapy`:

```r
mod.therapy <- lm(mood.gain ~ therapy
summary(aov(mod.therapy))
```

```
          Df Sum Sq Mean Sq F value Pr(>F)
therapy    1  0.467  0.4672   1.708   0.21
Residuals 16  4.378  0.2736
```

- we can do better than this!
- if we put *both factors* in the model, we can test both **main effects** together
- each **main effect** accounts for a different portion of the total variance in the DV
- each **main effect** takes a bite out of the *SSres*

Testing Main Effects together

- **main effect** of `drug`:

```
mod.drug <- lm(mood.gain ~ drug, data
summary(aov(mod.drug))
```

```
            Df Sum Sq Mean Sq F value   Pr(>F)
drug         2  3.453  1.7267   18.61 8.65e-05
Residuals   15  1.392  0.0928
```

- **main effect** of `therapy`:

```
mod.therapy <- lm(mood.gain ~ therapy
summary(aov(mod.therapy))
```

```
            Df Sum Sq Mean Sq F value Pr(>F)
therapy      1  0.467  0.4672   1.708   0.21
Residuals   16  4.378  0.2736
```

- two-way ANOVA to test `drug` and `therapy` together
- both **main effects** are in the model

```
mod.both <- lm(mood.gain ~ drug + therapy, data=clin.tri
summary(aov(mod.both))
```

```
            Df Sum Sq Mean Sq F value   Pr(>F)
drug         2  3.453  1.7267  26.149 1.87e-05
therapy      1  0.467  0.4672   7.076   0.0187
Residuals   14  0.924  0.0660
```

- df, Sum Sq, Mean Sq for `drug` and `therapy` are the same as in the one-way ANOVAs
- **But** the F-statistics are larger, and the p-values are smaller
- why? the **Residuals** row is different (smaller!)

Testing Main Effects together

- SSdrug + SSres = SStot
- 3.453 + 1.392 = 4.85

```
          Df Sum Sq Mean Sq F value   Pr(>F)
drug       2  3.453  1.7267   18.61 8.65e-05
Residuals 15  1.392  0.0928
```

- SStherapy + SSres = SStot
- 0.467 + 4.378 = 4.85

```
          Df Sum Sq Mean Sq F value Pr(>F)
therapy    1  0.467  0.4672   1.708   0.21
Residuals 16  4.378  0.2736
```

- SSdrug + SStherapy + SSres = SStot
- 3.453 + 0.467 + 0.924 = 4.85

```
          Df Sum Sq Mean Sq F value   Pr(>F)
drug       2  3.453  1.7267  26.149 1.87e-05
therapy    1  0.467  0.4672   7.076   0.0187
Residuals 14  0.924  0.0660
```

- when we put both factors in the model, the SS for the Residuals is smaller
- **each factor in the model reduces the SSRes**
- each factor accounts for some variability in the DV
- **the more factors we add to the model, the smaller the SSRes**
- smaller SSres means smaller MSres
- smaller MSres means larger F-statistic
- larger F-statistic means smaller p-value
- smaller p-values means more confidence that the factor has an effect

The Interaction Term

- SSdrug + SSres = SStot
- 3.453 + 1.392 = 4.85

```
          Df Sum Sq Mean Sq F value   Pr(>F)
drug       2  3.453  1.7267   18.61 8.65e-05
Residuals 15  1.392  0.0928
```

- SStherapy + SSres = SStot
- 0.467 + 4.378 = 4.85

```
          Df Sum Sq Mean Sq F value Pr(>F)
therapy    1  0.467  0.4672   1.708   0.21
Residuals 16  4.378  0.2736
```

- SSdrug + SStherapy + SSres = SStot
- 3.453 + 0.467 + 0.924 = 4.85

```
          Df Sum Sq Mean Sq F value   Pr(>F)
drug       2  3.453  1.7267  26.149 1.87e-05
therapy    1  0.467  0.4672   7.076   0.0187
Residuals 14  0.924  0.0660
```

- **but we can do even better**
- we can take another bite out of SSRes
- we can add a third effect to the model:
  - an **interaction** term
  - the interaction between `drug` and `therapy`
- `mood.gain ~ drug + therapy + drug:therapy`
- or shorthand: `mood.gain ~ drug * therapy`

  `*` means include all main effects and interactions

```r
mod.full <- lm(mood.gain ~ drug * therapy, data=clin.tri
summary(aov(mod.full))
```

```
             Df Sum Sq Mean Sq F value   Pr(>F)
drug          2  3.453  1.7267  31.714 1.62e-05
therapy       1  0.467  0.4672   8.582   0.0126
drug:therapy  2  0.271  0.1356   2.490   0.1246
Residuals    12  0.653  0.0544
```

Main Effects & Interaction

- SSdrug + SSres = SStot
- 3.453 + 1.392 = 4.85

```
          Df Sum Sq Mean Sq F value   Pr(>F)
drug       2  3.453  1.7267   18.61 8.65e-05
Residuals 15  1.392  0.0928
```

- SStherapy + SSres = SStot
- 0.467 + 4.378 = 4.85

```
          Df Sum Sq Mean Sq F value Pr(>F)
therapy    1  0.467  0.4672   1.708   0.21
Residuals 16  4.378  0.2736
```

- SSdrug + SStherapy + SSres = SStot
- 3.453 + 0.467 + 0.924 = 4.85

```
          Df Sum Sq Mean Sq F value   Pr(>F)
drug       2  3.453  1.7267  26.149 1.87e-05
therapy    1  0.467  0.4672   7.076   0.0187
Residuals 14  0.924  0.0660
```

- SSdrug + SStherapy + SSinteraction + SSres = SStot
- 3.453 + 0.467 + 0.271 + 0.653 = 4.85

```r
mod.full <- lm(mood.gain ~ drug * therapy, data=clin.tri
summary(aov(mod.full))
```

```
             Df Sum Sq Mean Sq F value   Pr(>F)
drug          2  3.453  1.7267  31.714 1.62e-05
therapy       1  0.467  0.4672   8.582   0.0126
drug:therapy  2  0.271  0.1356   2.490   0.1246
Residuals    12  0.653  0.0544
```

- SSres is even smaller!
- reduces unexplained variance in the DV
- denominator of F-statistic is smaller
- F-statistics for our main effects are bigger
- p-values for our main effects are smaller

What is the Interaction?

- an **interaction** occurs when the effect of one factor on the DV depends on the level of the other factor
- the difference in means between the levels of one factor is different for each level of the other factor
- the interaction is a difference of differences

What is the Interaction?

- the effect of avocado on the deliciousness of a sandwich depends upon whether bacon was present or absent

What is the Interaction?

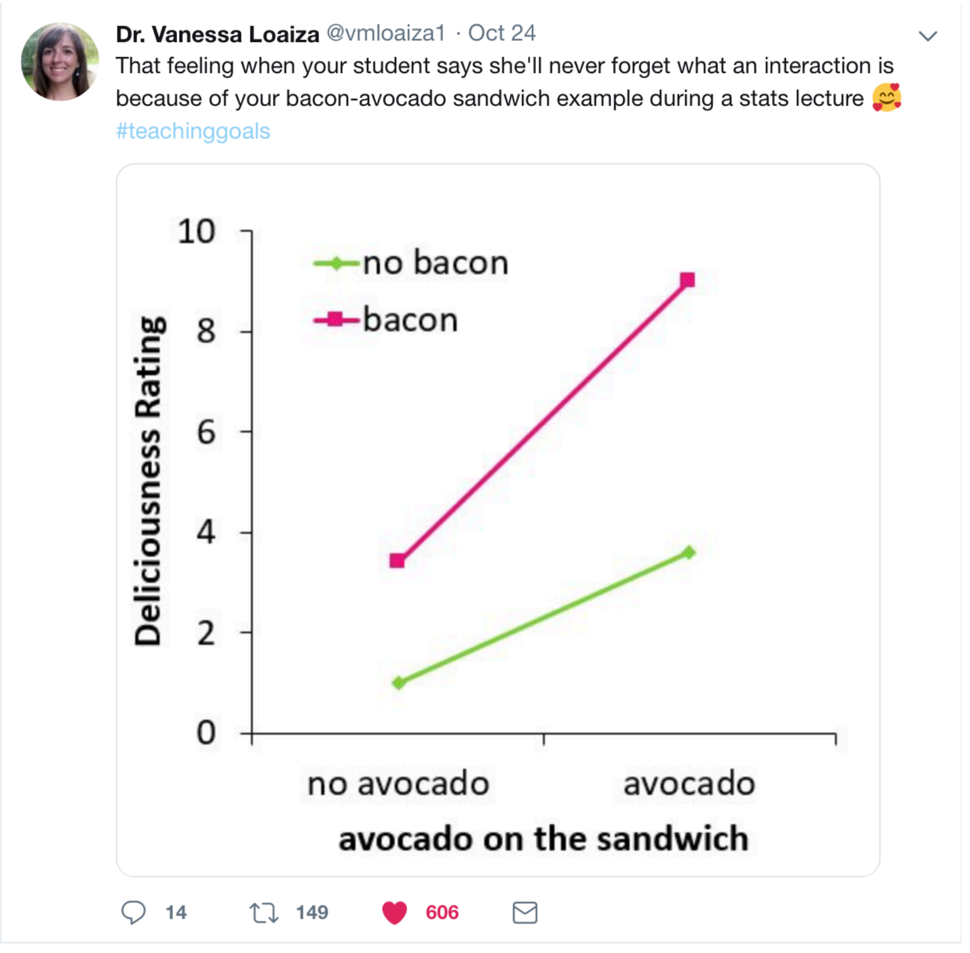

- **bacon absent:** avocado increases deliciousness by 3 (4 - 1)
- **bacon present:** avocado increases deliciousness by 6 (9 - 3)
- (4-1) is different than (9-3)
- a difference in the differences

What is the Interaction?

- a bigger effect of avocado when bacon is present compared to when bacon is absent
- the difference between avocado present vs absent is different for each level of bacon
- a difference in the difference

Example 1

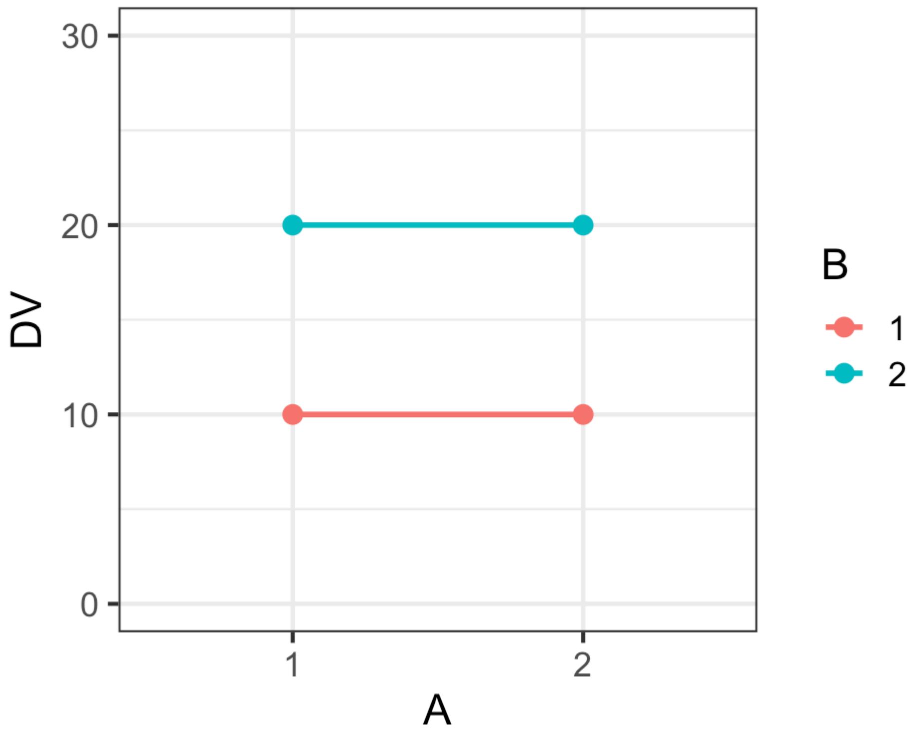

Example 1

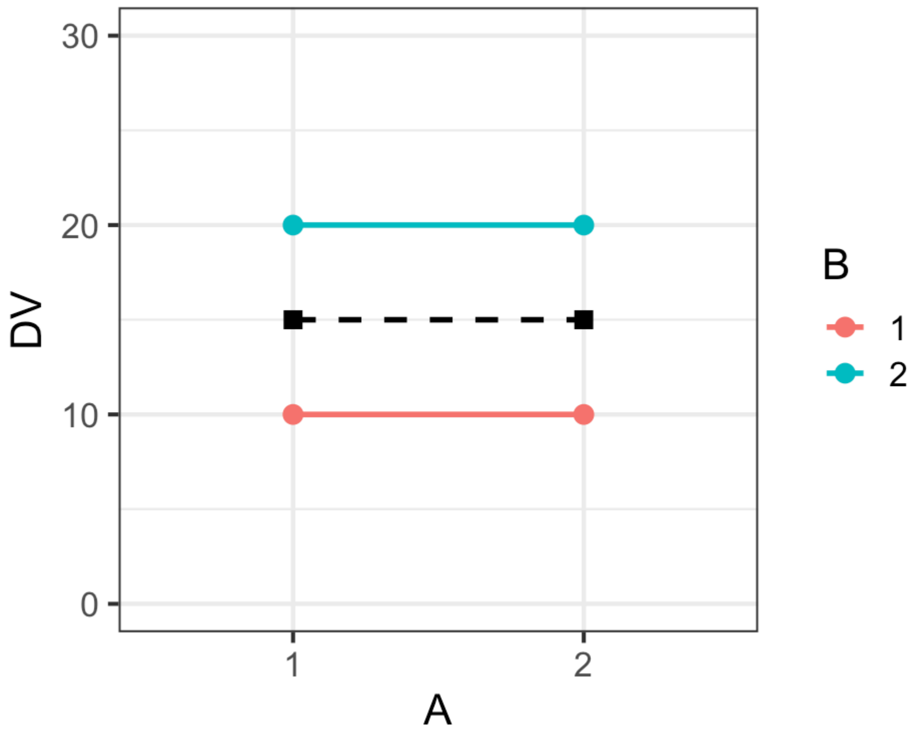

- main effect of A: no
- A1 - A2 = 0
- (10+20)/2 - (10+20)/2 = 0
- 15 - 15 = 0

Example 1

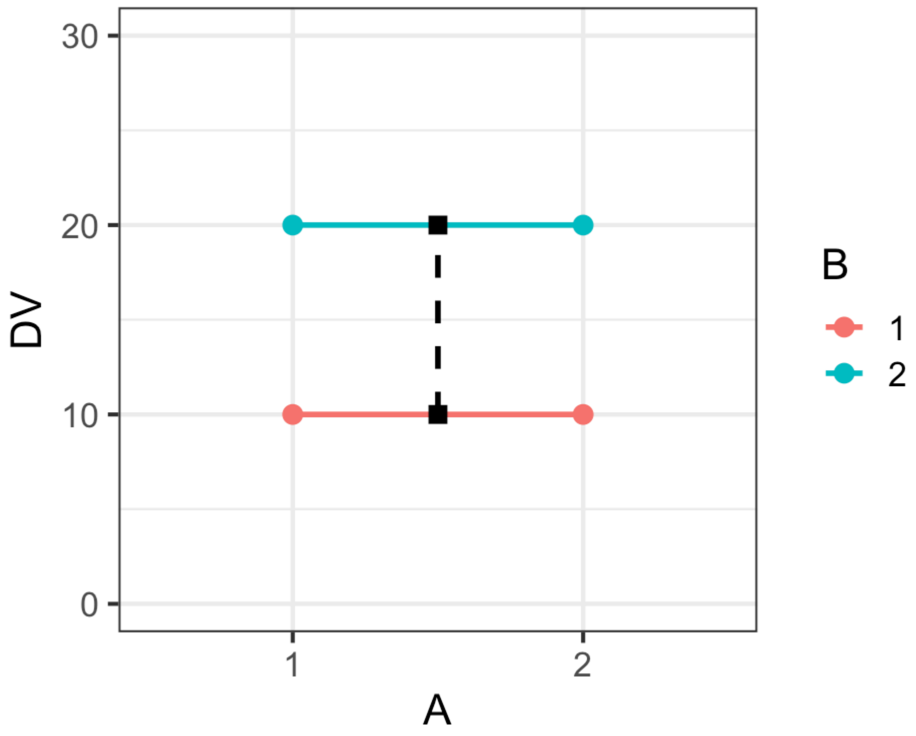

- main effect of B: **yes**
- B1 - B2 ≠ 0
- (20+20)/2 - (10+10)/2 ≠ 0
- 20 - 10 ≠ 0
- 10 ≠ 0

Example 1

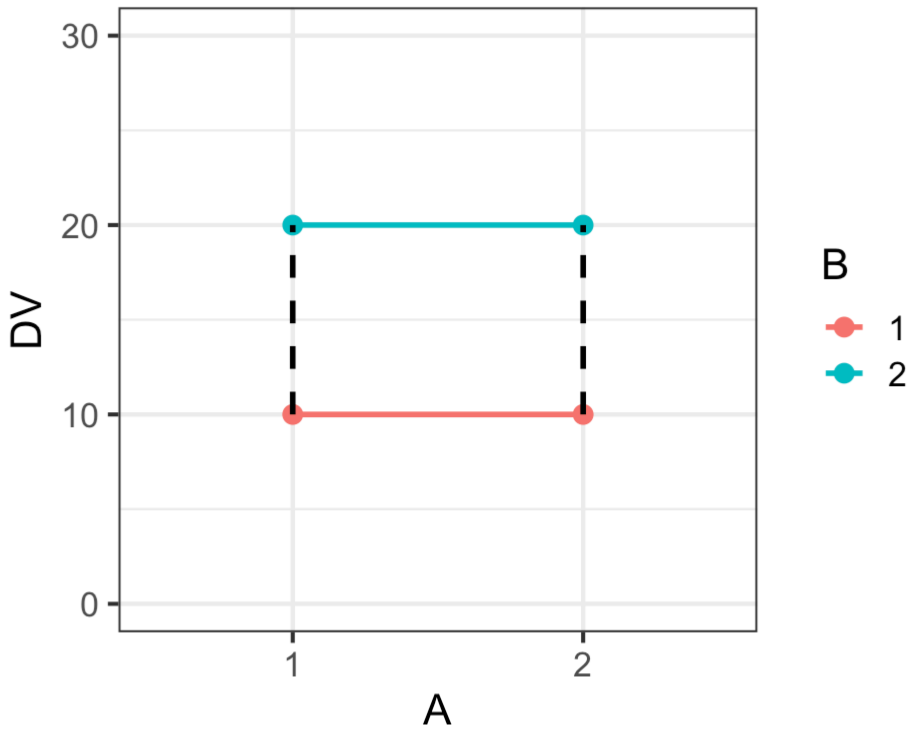

- A:B interaction: no
- effect of A (A1-A2) is:
  - (10-10) = 0 for B1
  - (20-20) = 0 for B2
  - interaction effect is (0 - 0) = 0
  - interaction effect is 0 so no interaction

Example 1

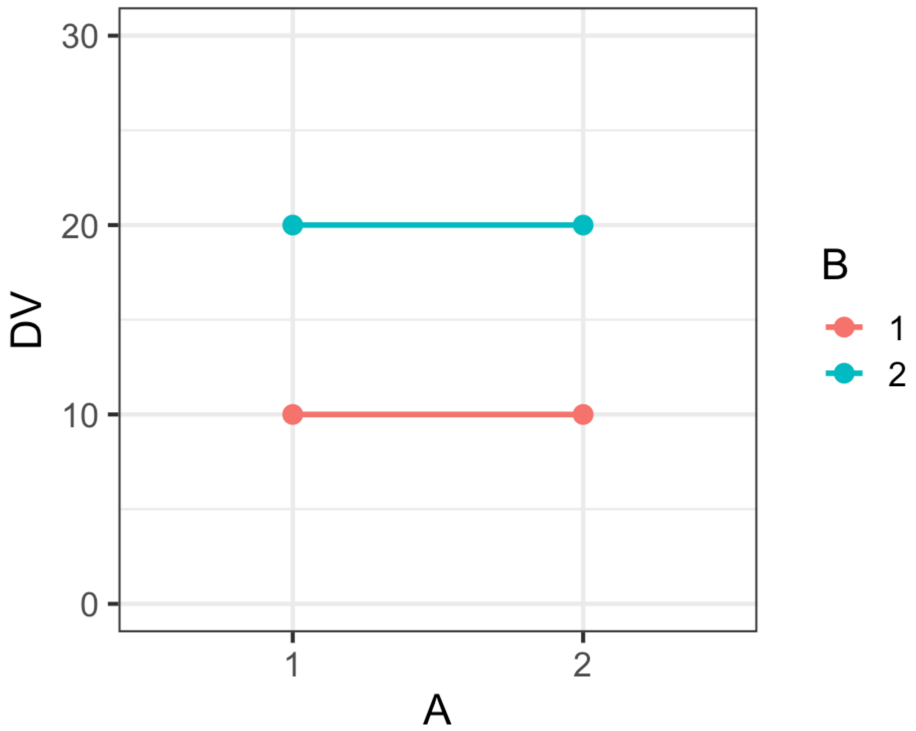

a general rule:

- **parallel lines** = no interaction

Example 2

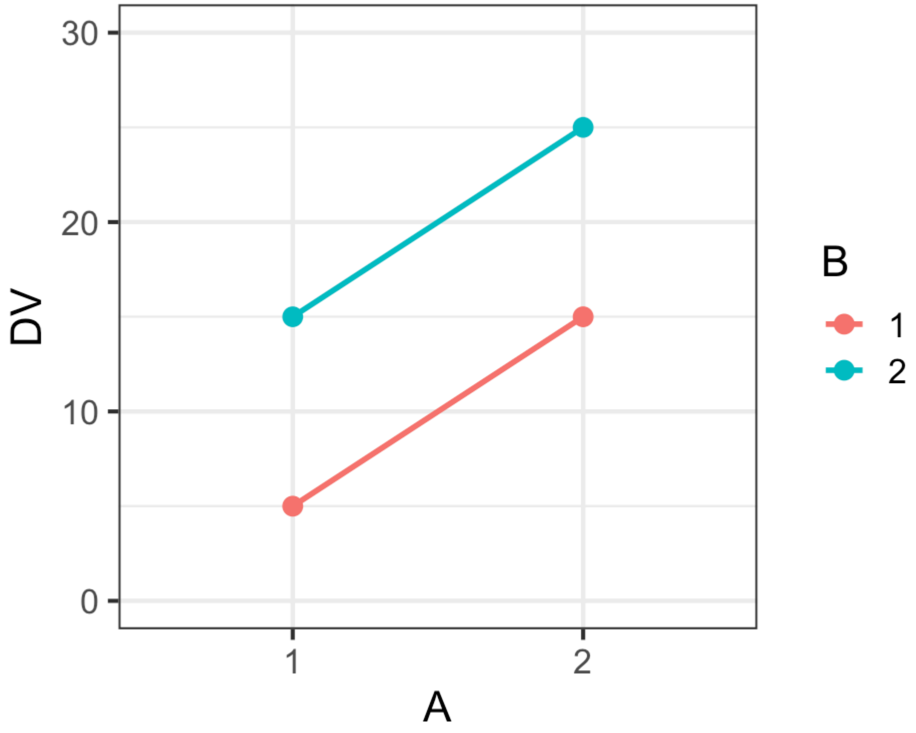

Example 2

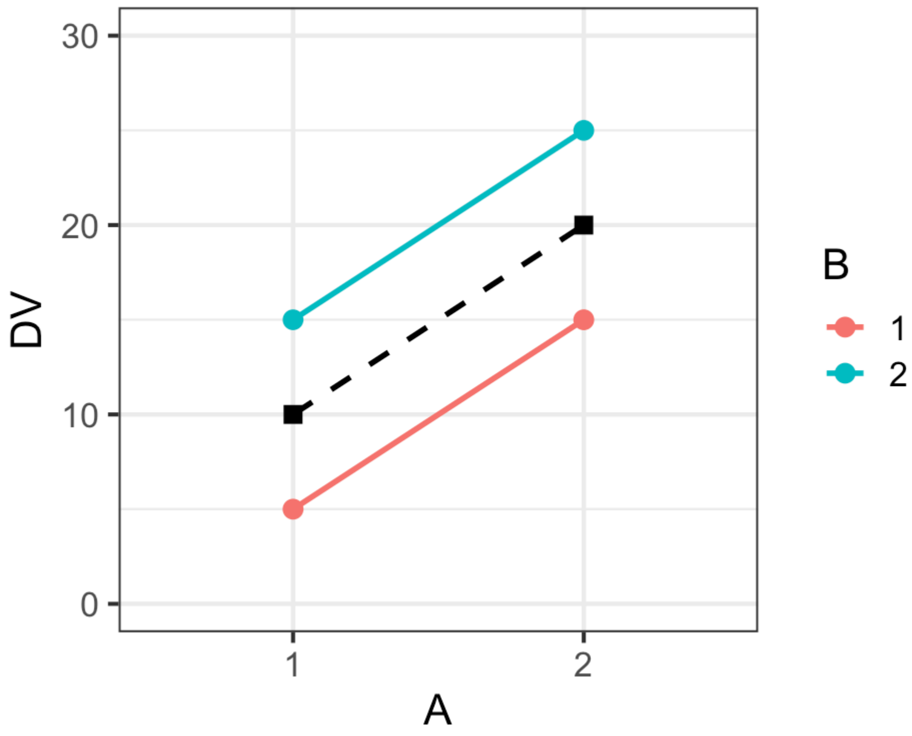

- main effect of A: **yes**
- A1 - A2 ≠ 0
- (15+5)/2 - (25+15)/2 ≠ 0
- 10 - 20 ≠ 0
- -10 ≠ 0

Example 2

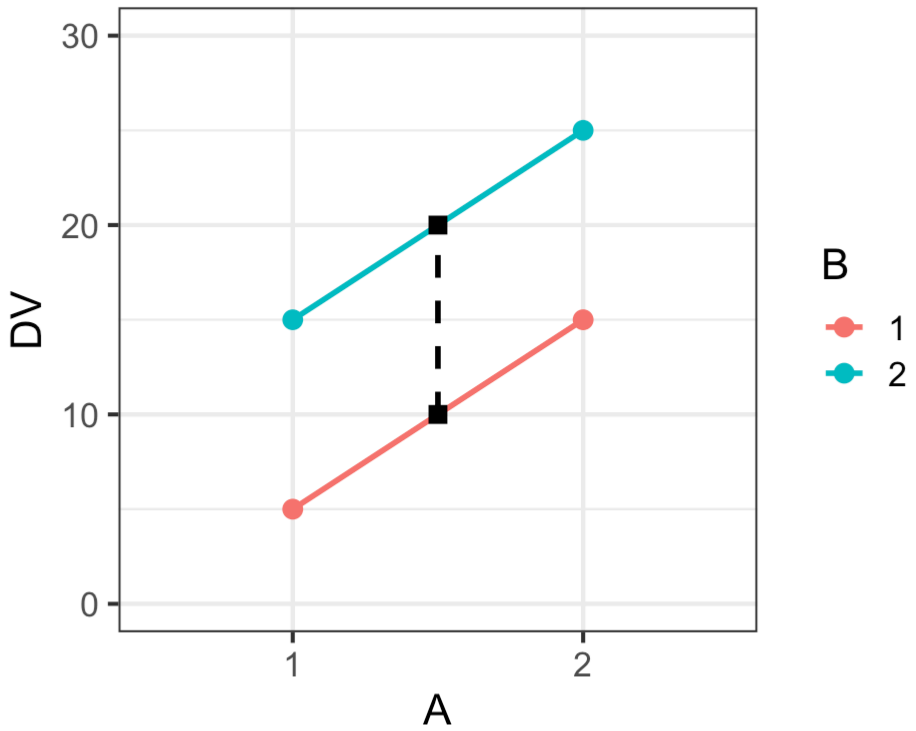

- main effect of B: **yes**
- B1 - B2 ≠ 0
- (15+25)/2 - (5+15)/2 ≠ 0
- 20 - 10 ≠ 0
- 10 ≠ 0

Example 2

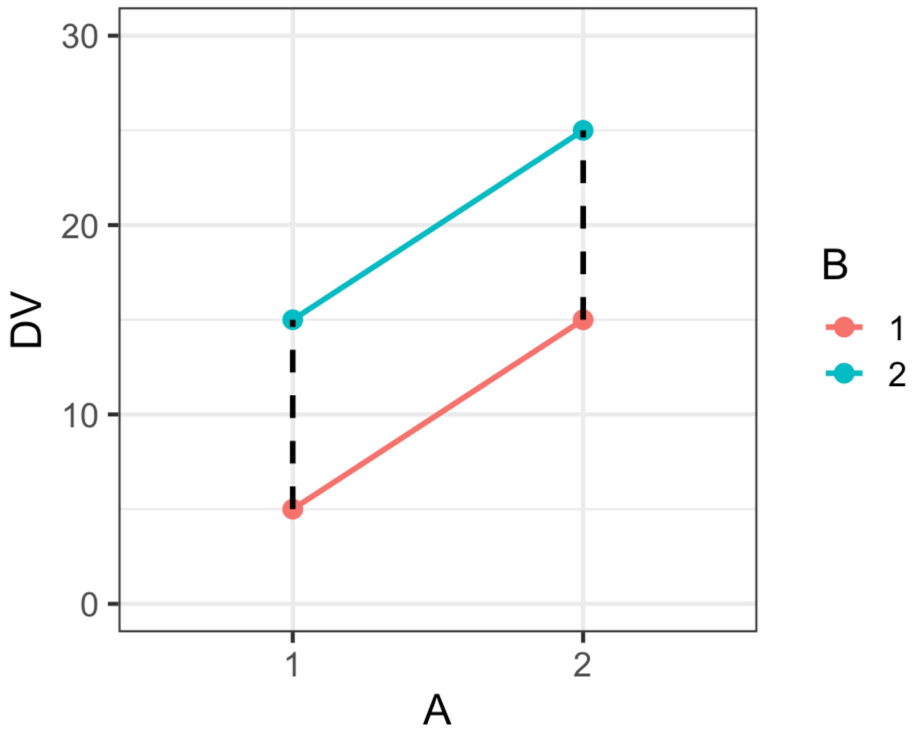

- A:B interaction: no
- B1-B2 within A1 is same as B1-B2 within A2

Example 3

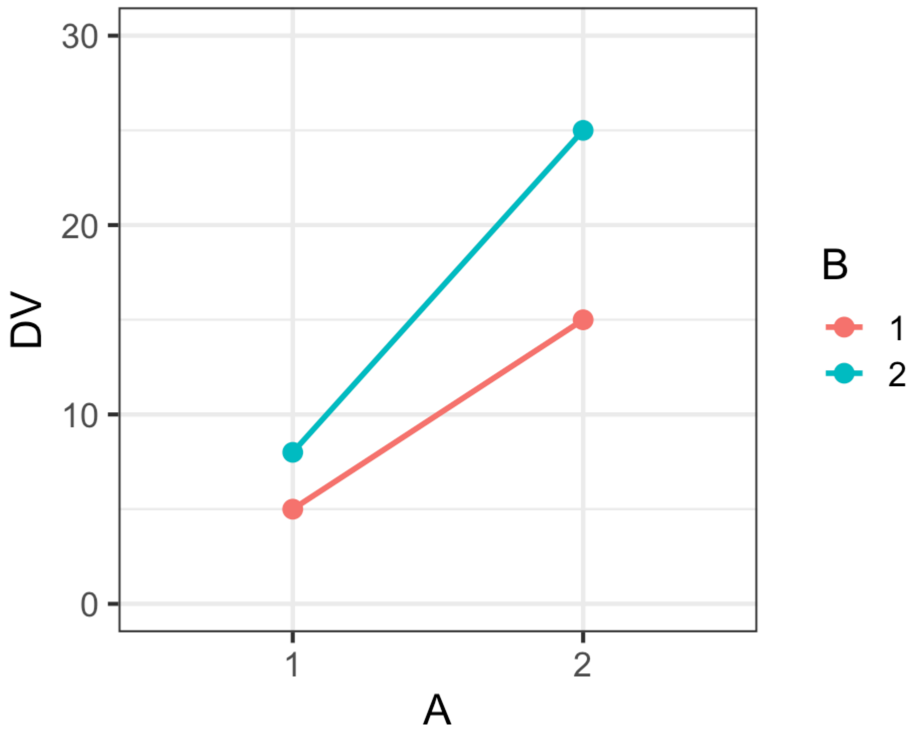

Example 3

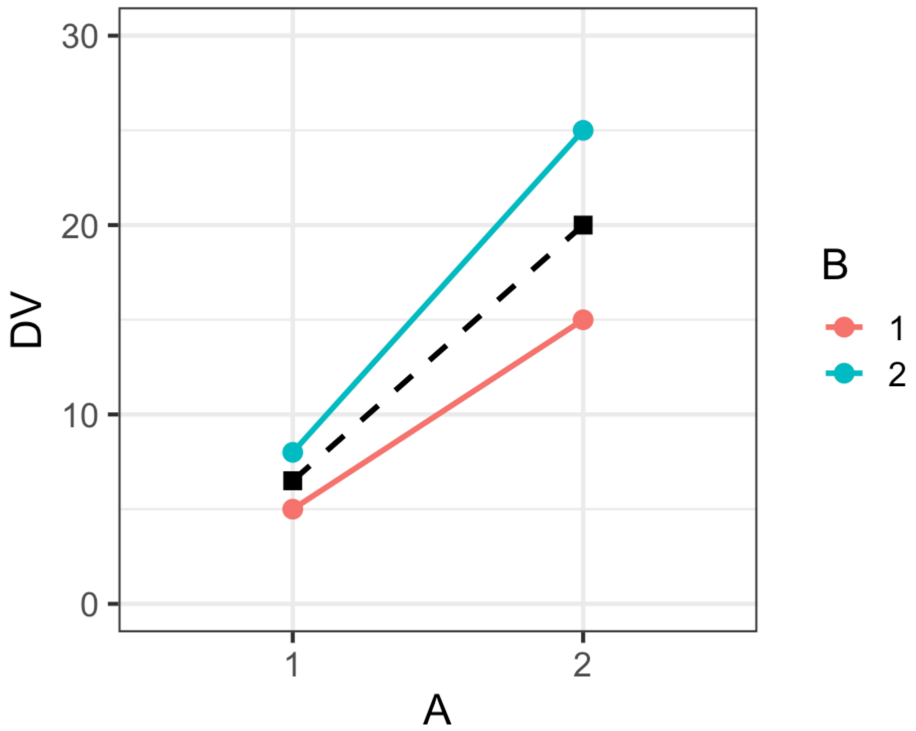

- main effect of A: **yes**

Example 3

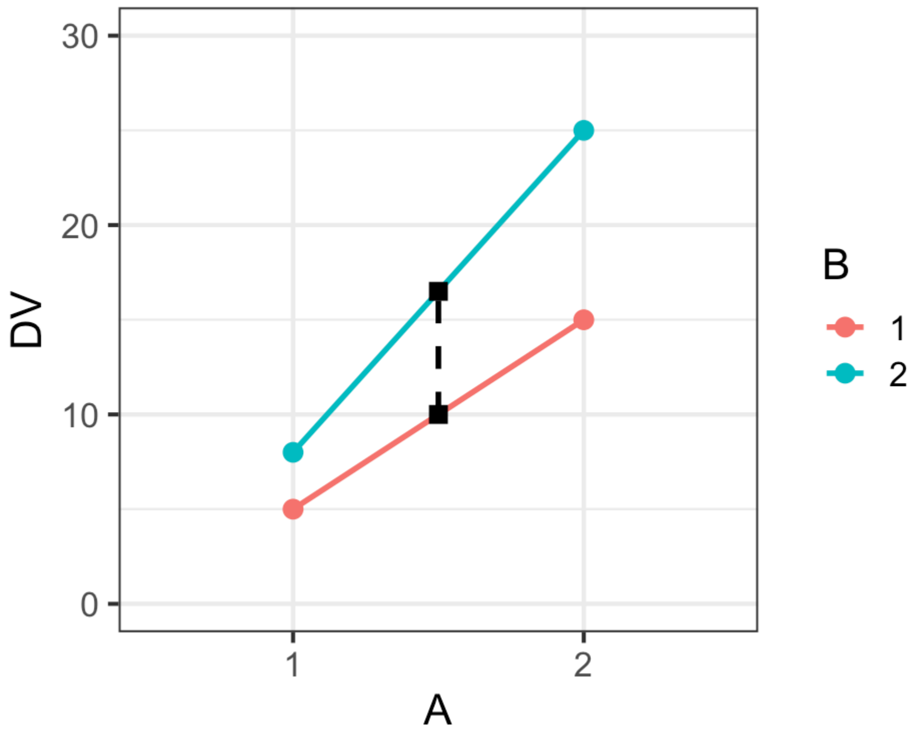

- main effect of B: **yes**

Example 3

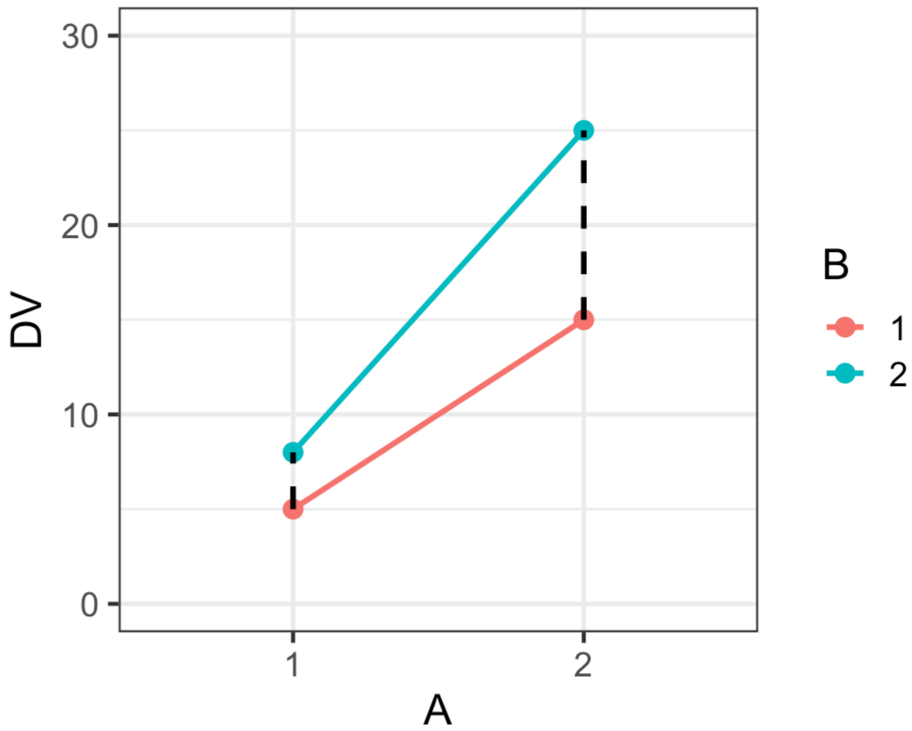

- A:B interaction: **yes**

Example 4

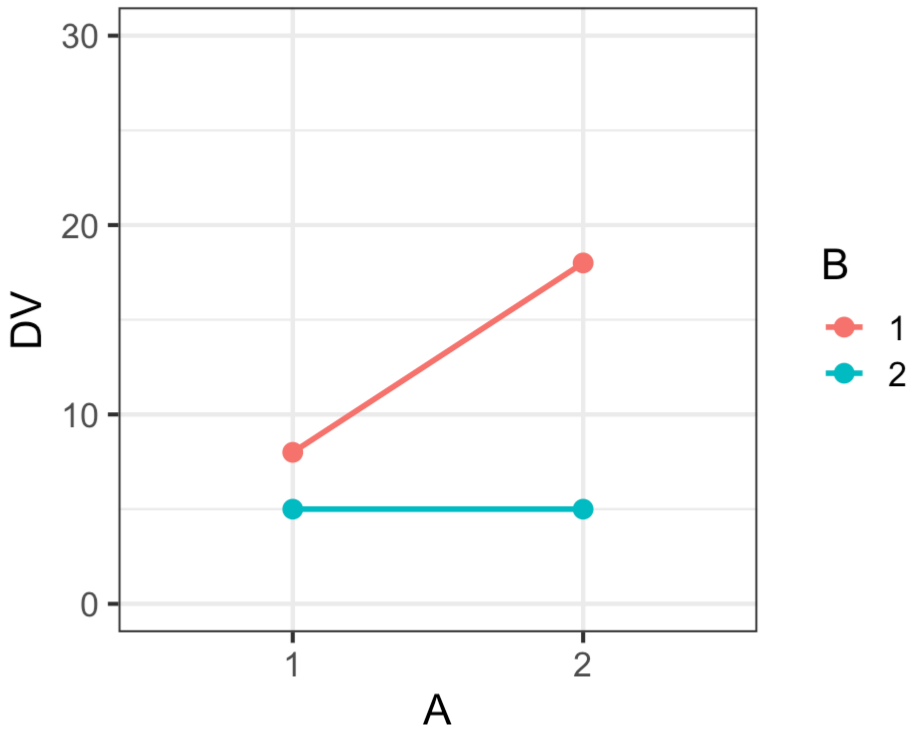

Example 4

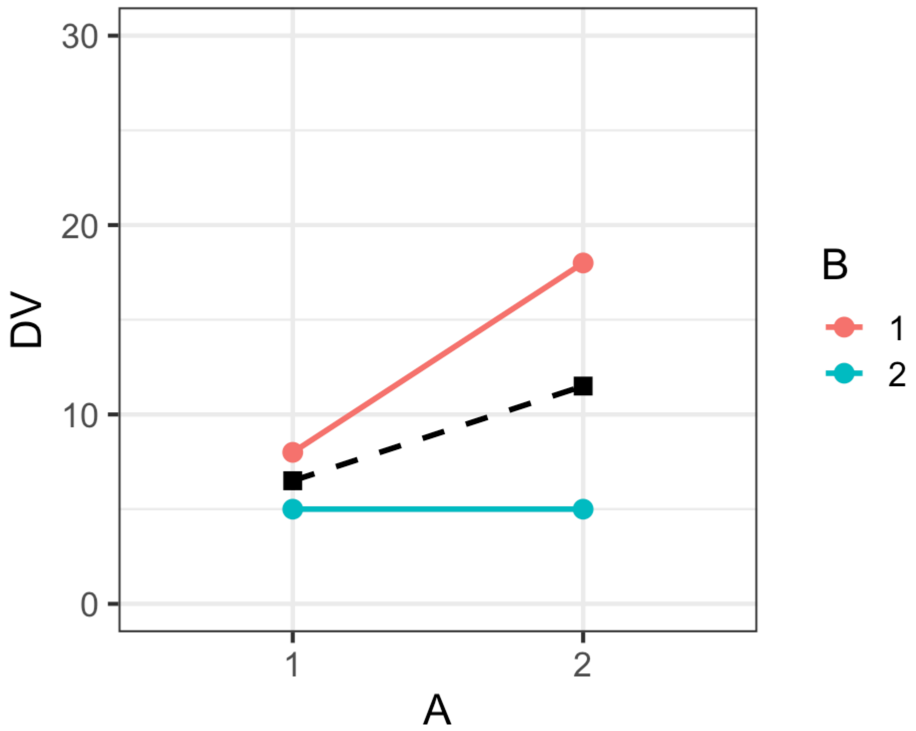

- main effect of A: **yes**

Example 4

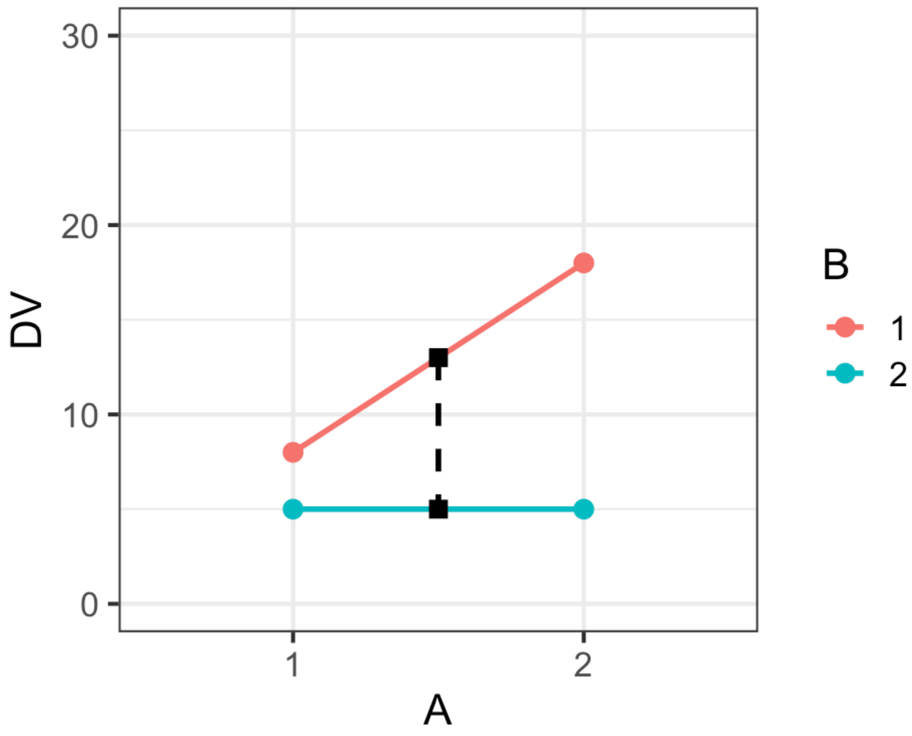

- main effect of B: **yes**

Example 4

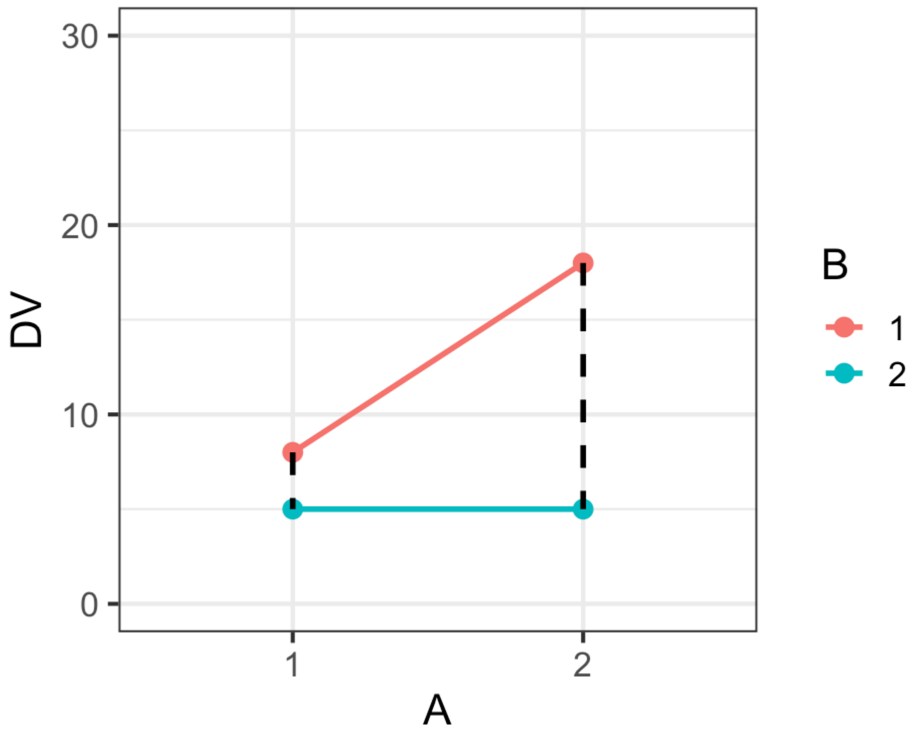

- A:B interaction: **yes**

Example 5

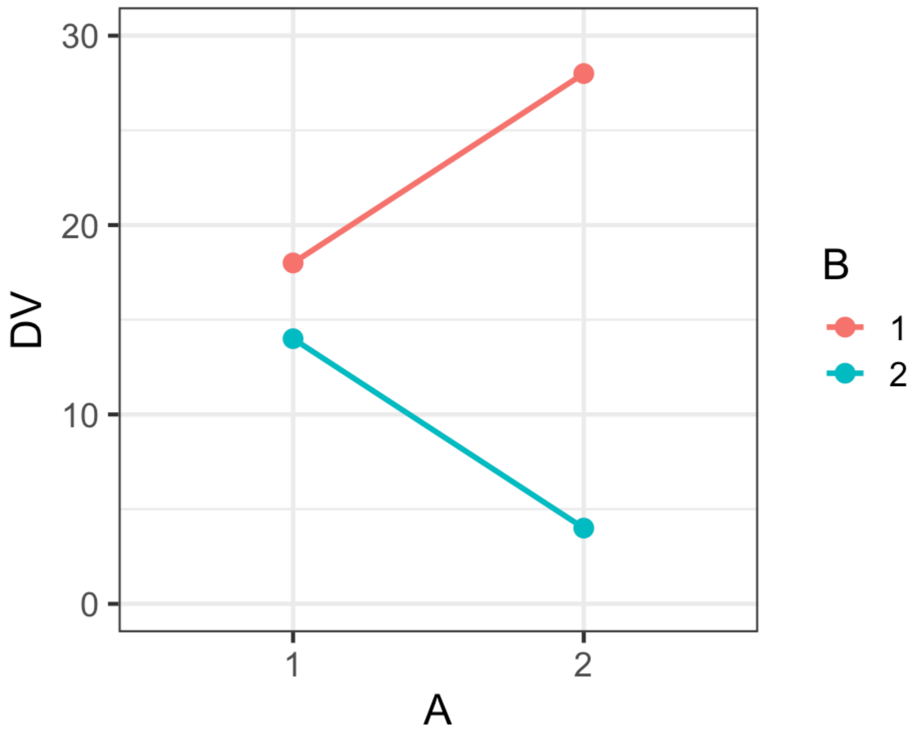

Example 5

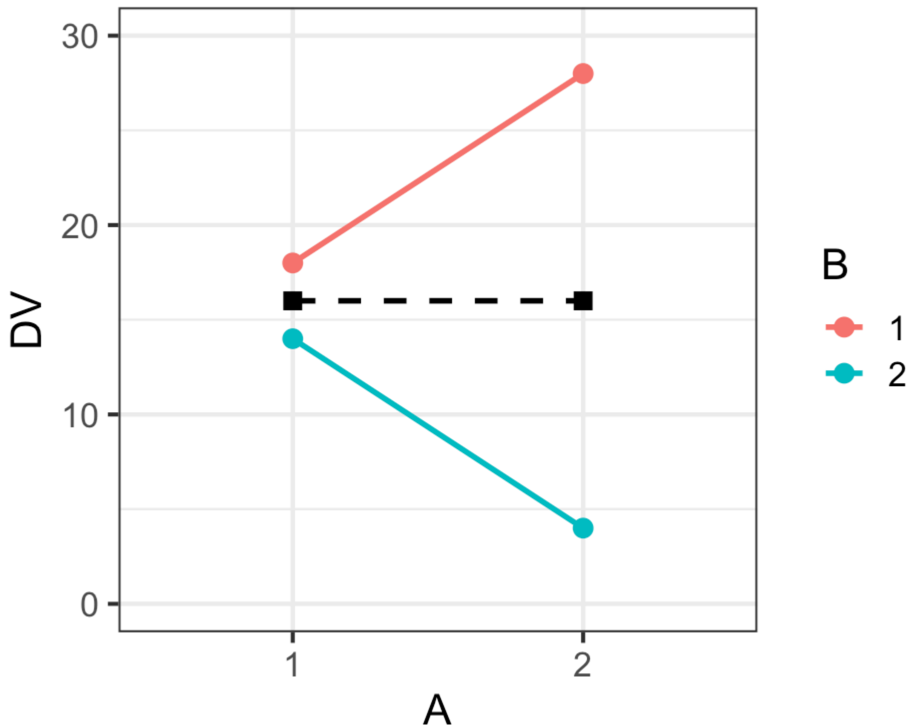

- main effect of A: no
  - !! is this an ok conclusion? !!
  - is it true A has no effect on the DV? (no!)

Example 5

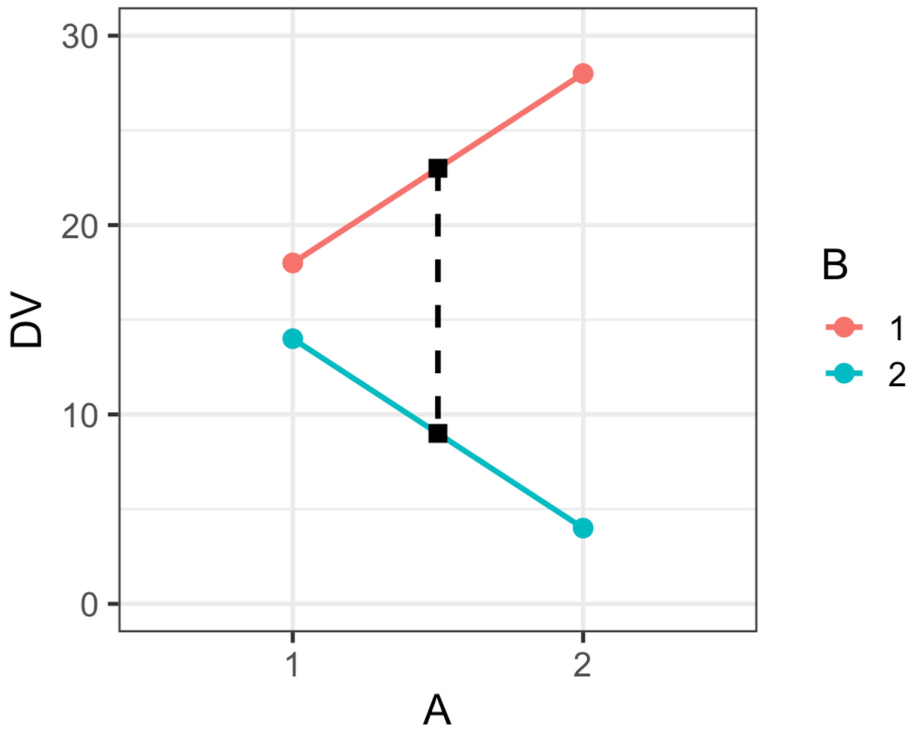

- main effect of B: **yes**

Example 5

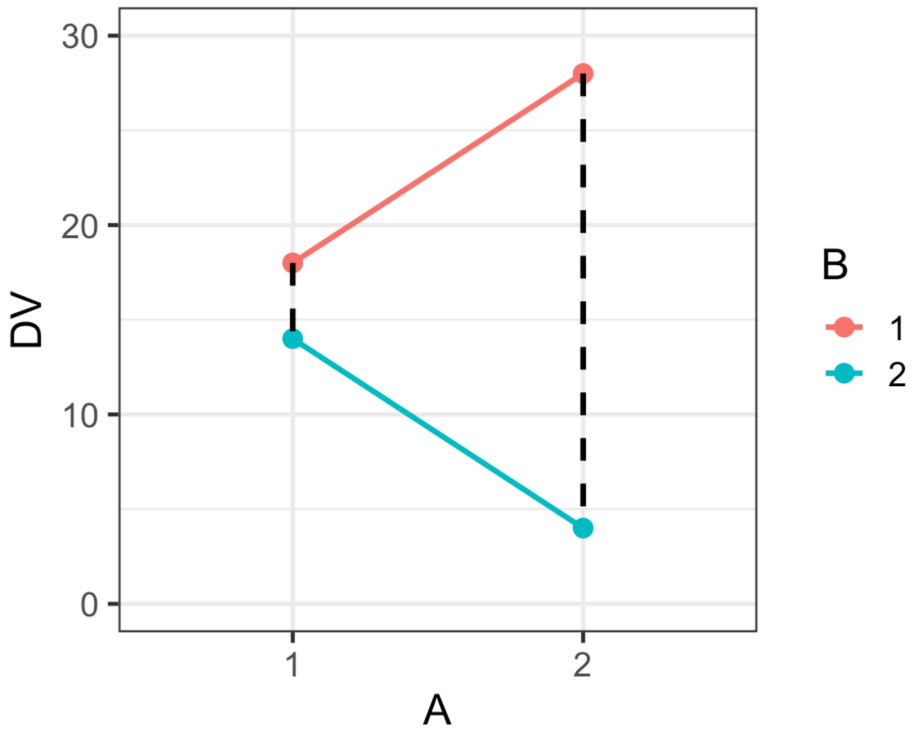

- A:B interaction: **yes**

Example 6

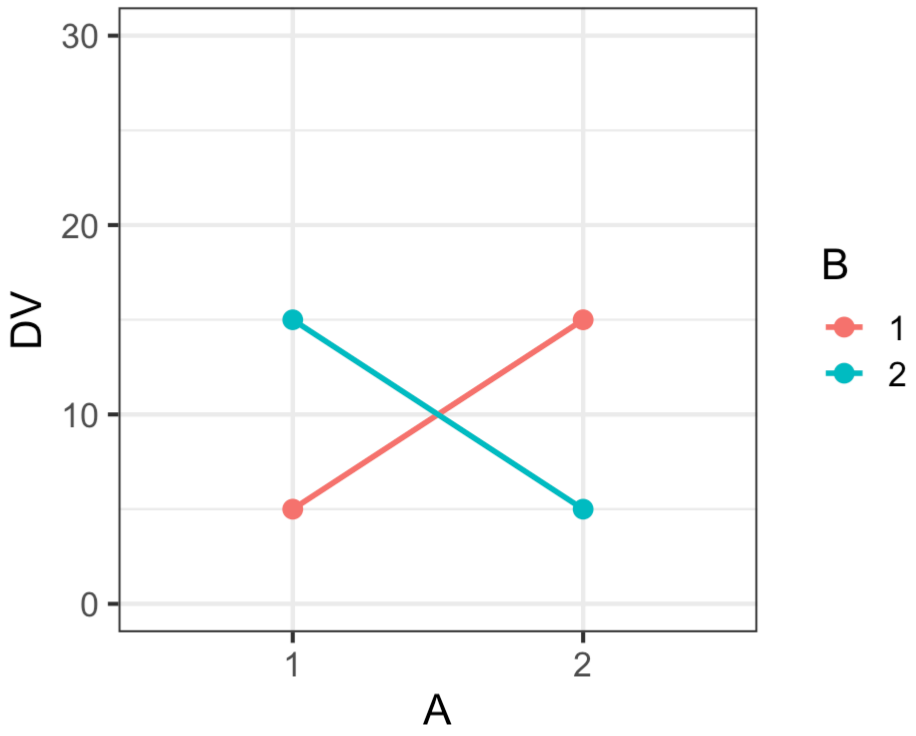

Example 6

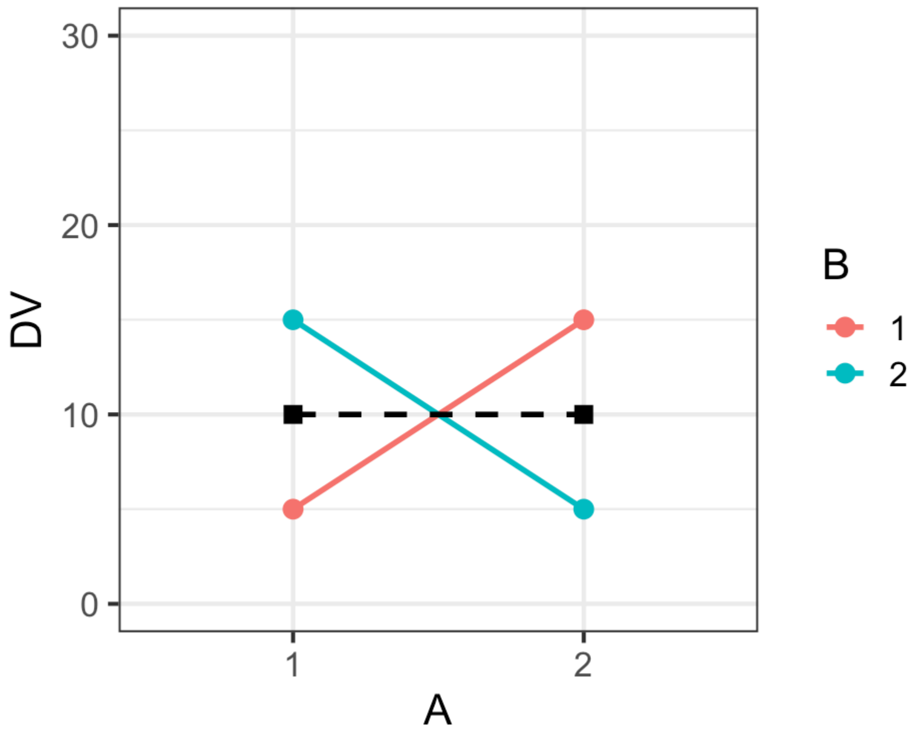

- main effect of A: no

Example 6

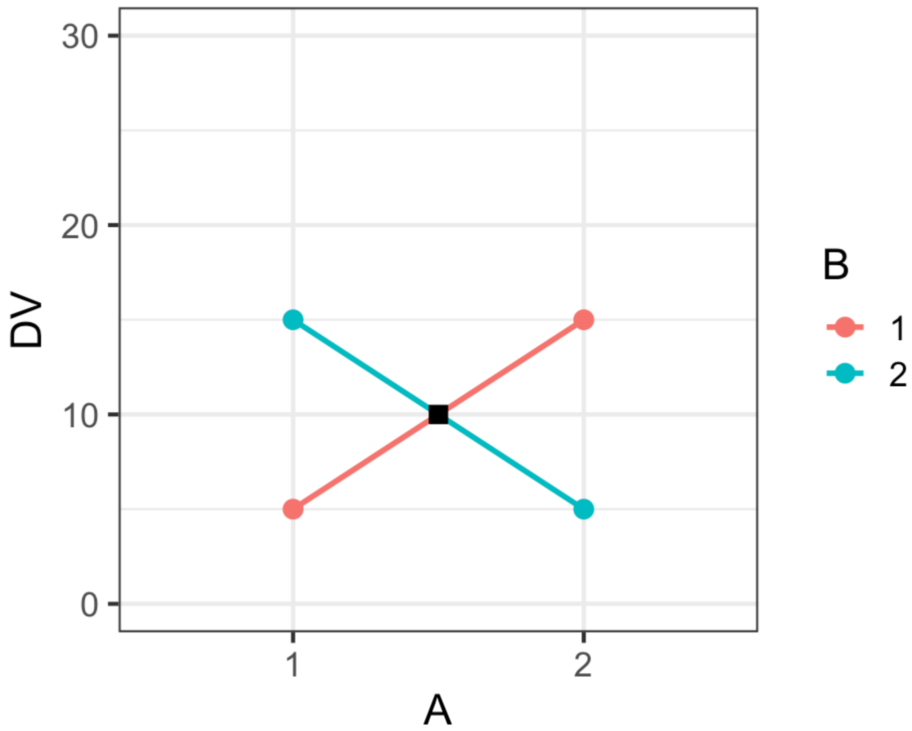

- main effect of B: no

Example 6

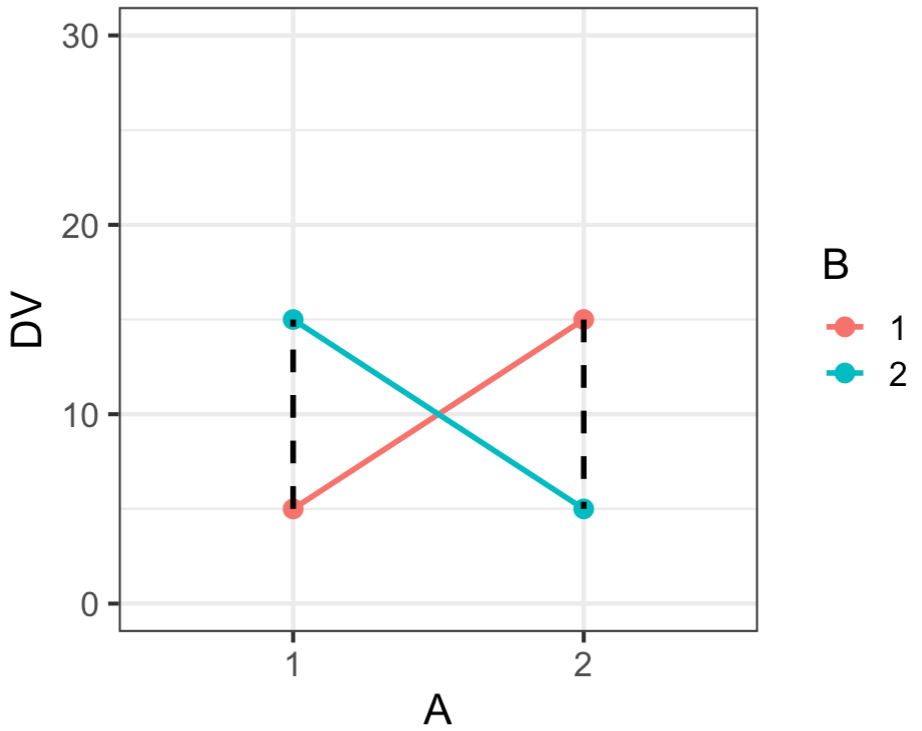

- A:B interaction: **yes**!
- effect of A (A1-A2) is:
  - (5-15) = -10 for B1
  - (15-5) = +10 for B2
- -10 - +10 = -20
  - -20 is not zero
  - interaction is present

Possible Outcomes

For a two-way factorial design:

1. IV1: no, IV2: no, IV1:IV2: no
2. IV1: yes, IV2: no, IV1:IV2: no
3. IV1: no, IV2: yes, IV1:IV2: no
4. IV1: yes, IV2: yes, IV1:IV2: no
5. IV1: no, IV2: no, IV1:IV2: yes
6. IV1: yes, IV2: no, IV1:IV2: yes
7. IV1: no, IV2: yes, IV1:IV2: yes
8. IV1: yes, IV2: yes, IV1:IV2: yes

Possible Outcomes: examples

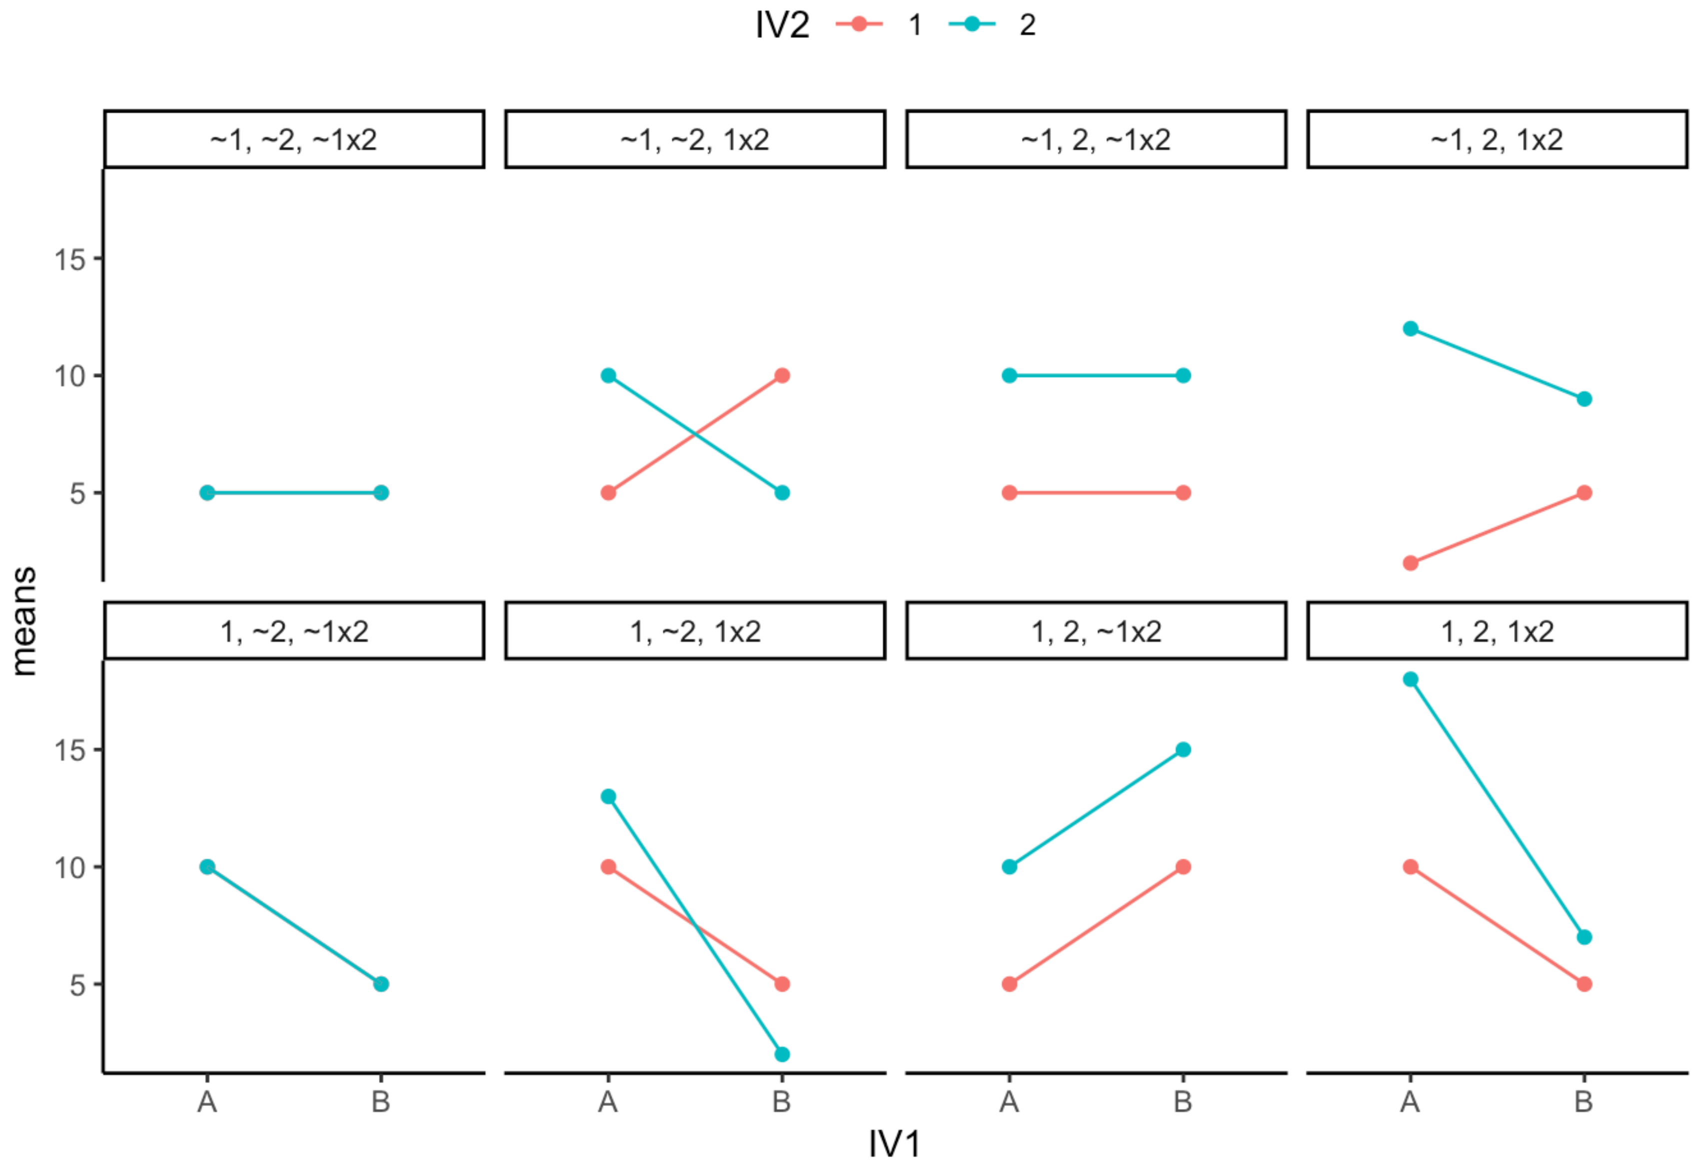

Interepretation of Effects

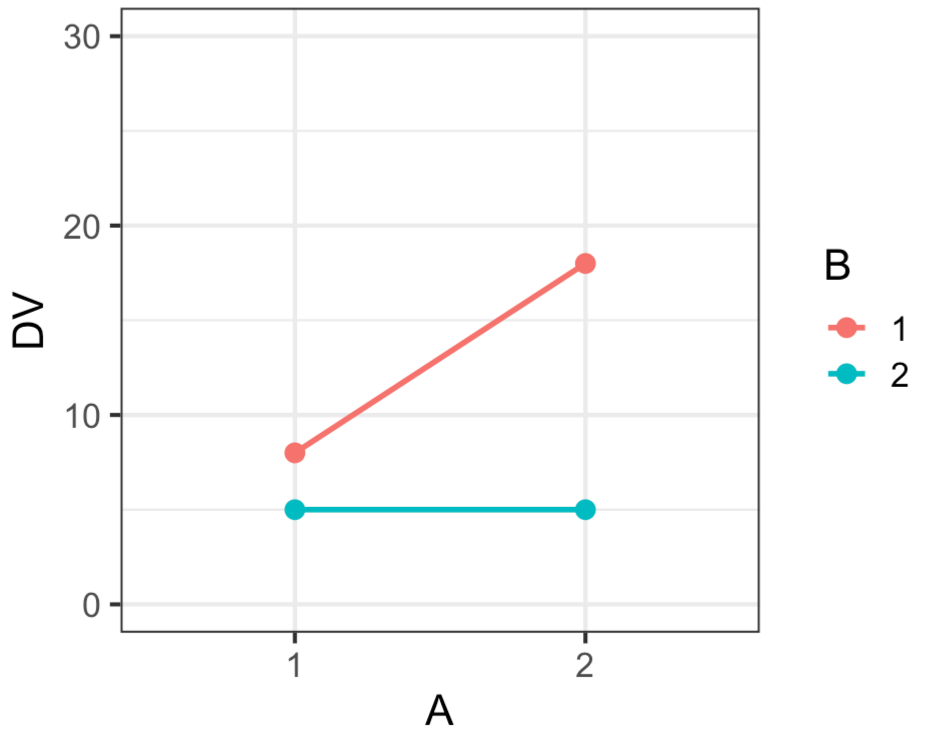

- main effect of A: **yes**
- main effect of B: **yes**
- A:B interaction: **yes**

Question:

- Is the **main effect of A** meaningful?

Interepretation of Effects

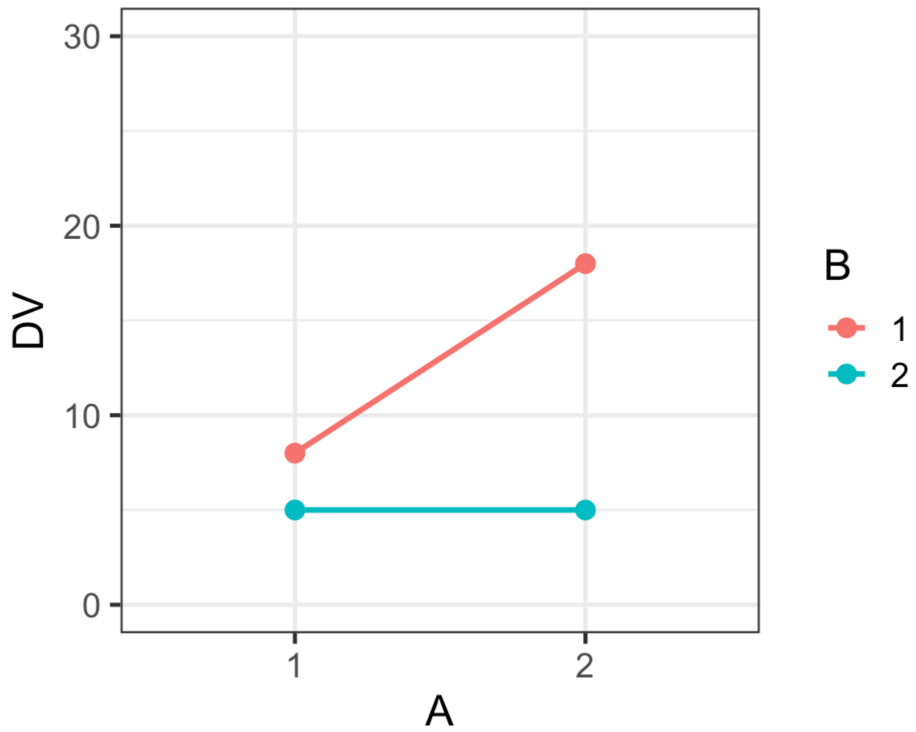

- main effect of A: **yes**
- main effect of B: **yes**
- A:B interaction: **yes**

Advice:

- if the **interaction effect** is significant, then the **main effects** are not meaningful
- the main effect ignores the levels of the other factor(s)
  - but if there is an interaction, that means by definition the effect of A is different for different levels of B
  - so **averaging over B does not make sense**

Interepretation of Effects

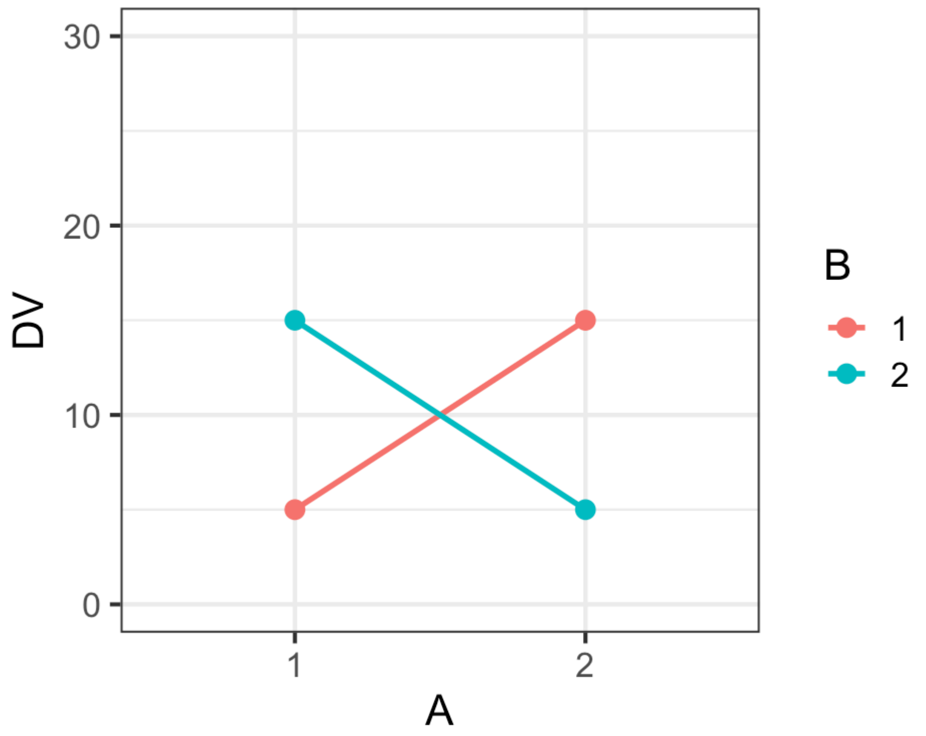

- main effect of A: no
- main effect of B: no
- A:B interaction: **yes**

Question:

- Is it true that A has no effect and B has no effect?
  - (main effect of A is not significant)
  - (main effect of B is not significant)

Interepretation of Effects

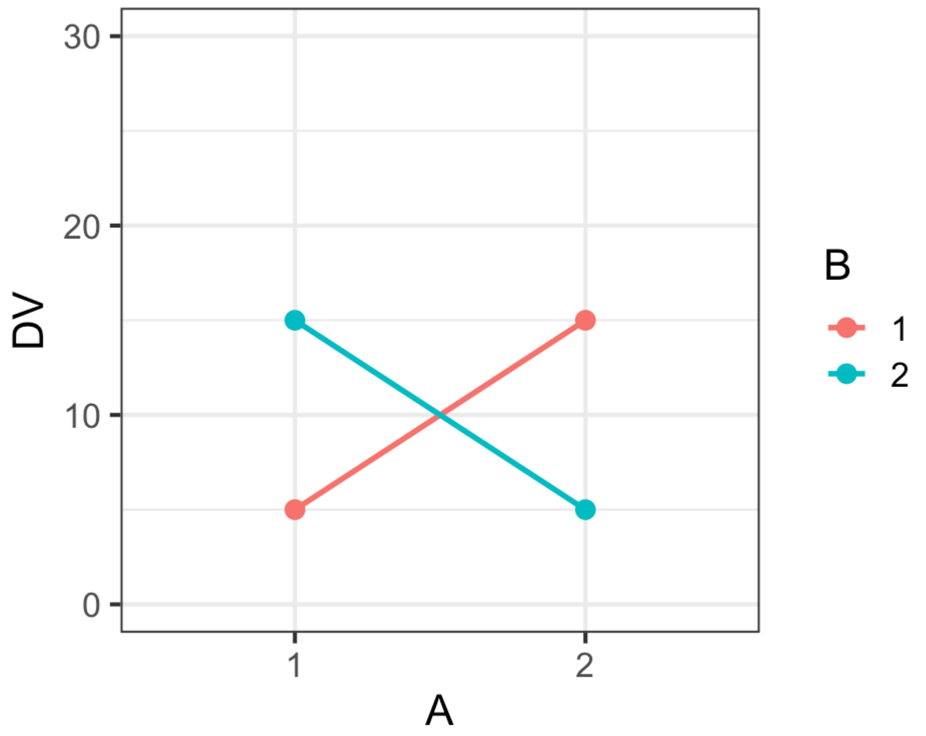

- main effect of A: no
- main effect of B: no
- A:B interaction: **yes**

Answer:

- of course not!
- the interaction effect is significant
- so the main effects are not meaningful
- A has an effect, but it is different for different levels of B
- *follow-up* tests are needed to test the effect of A at each level of B
  - we will cover this topic next week