



Factorial designs and the GLM

Professor Andy Field

WHOA
 @profandyfield

 www.youtube.com/user/ProfAndyField/

 www.discoveringstatistics.com

 www.milton-the-cat.rocks

 www.discovr.rocks



ANDY FIELD



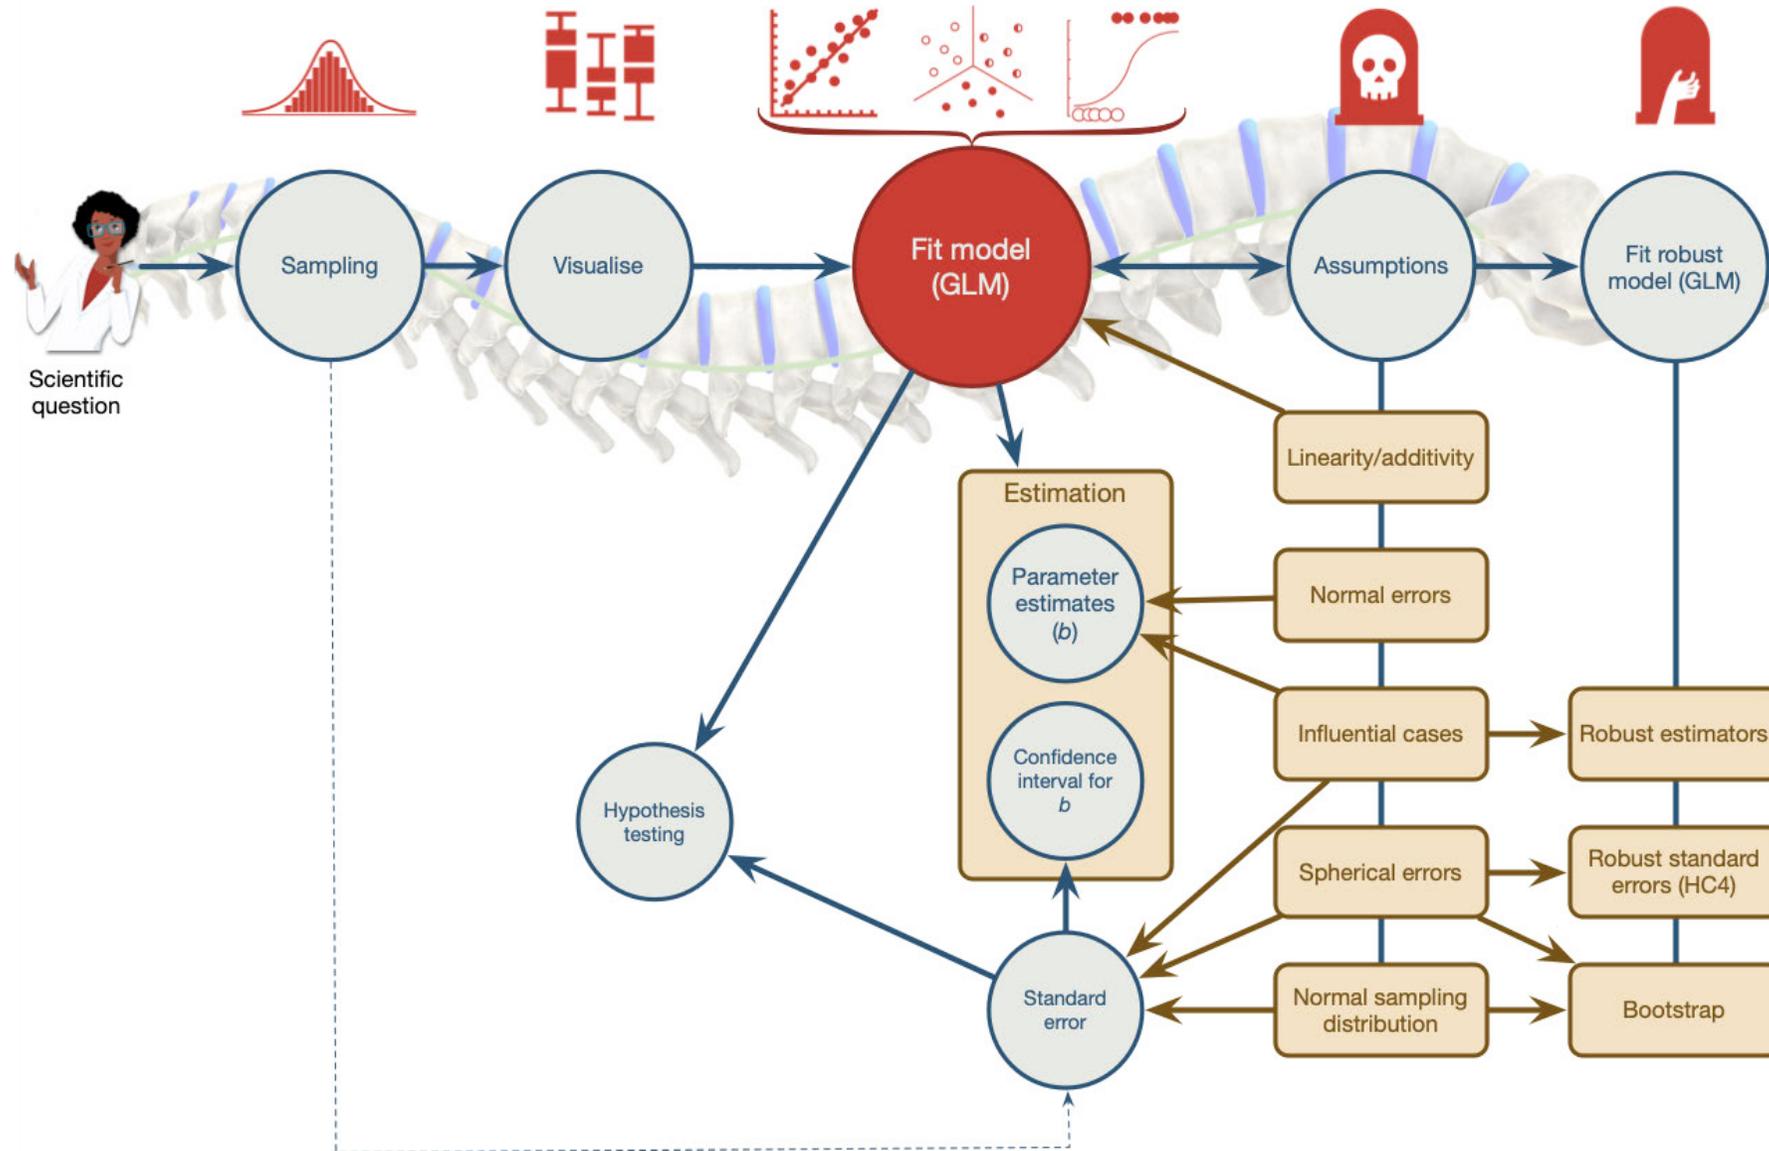
Learning outcomes

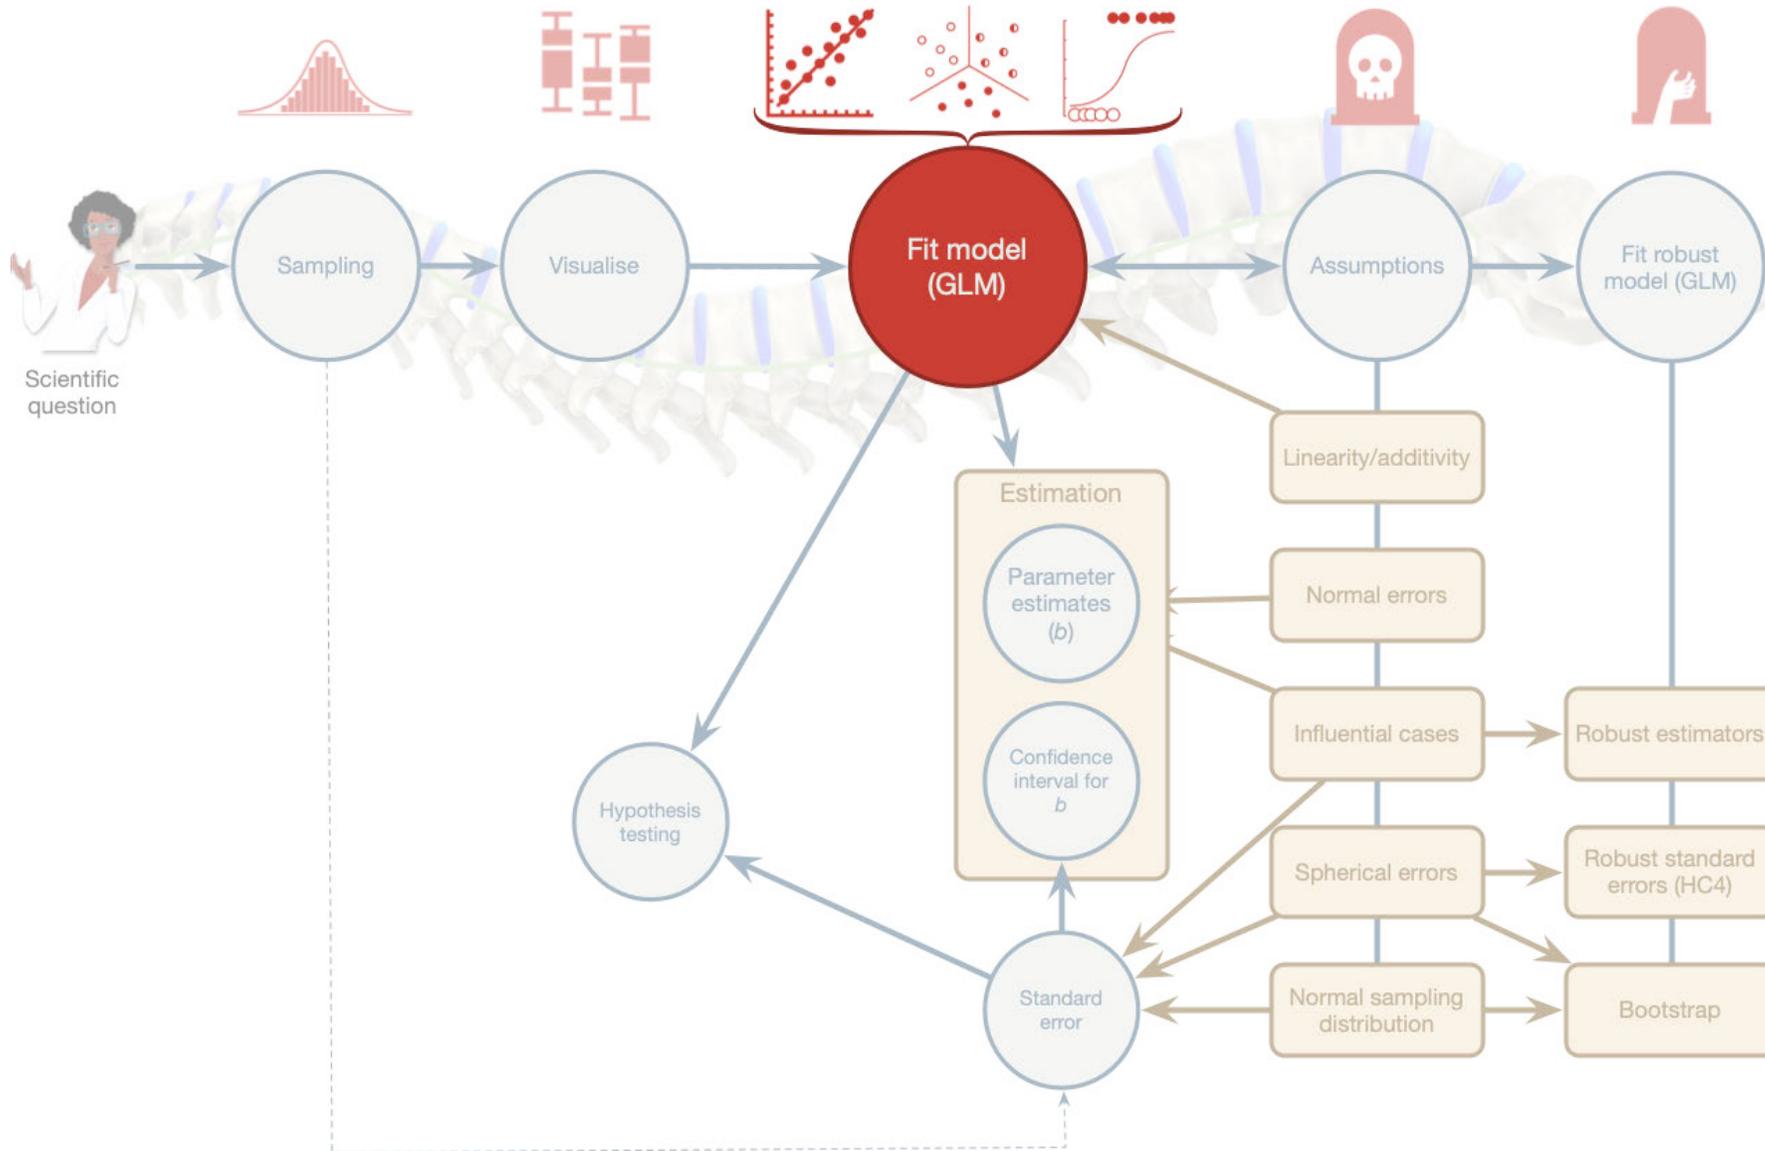
- What is a factorial design?
- Factorial designs and the GLM
 - Checking assumptions
 - Robust models
- Moderation and interaction effects
- Breaking down interactions
 - Plots
 - Simple effects analysis



ANDY FIELD







Terminology

- Variables in experimental designs
 - Predictor variables referred to as **independent variables (IV)** because their value is independent of (is not being predicted from) other variables.
 - Outcome variable referred to as the **dependent variables (DV)**, because their value 'depends' upon (is being predicted from) other variables.
- **Factorial design**: Two or more predictor variables/IV have been manipulated
- **n-way design**: The number of predictor variables/IVs manipulated
 - Two-way = 2 predictor variables/IVs manipulated
 - Three-way = 3 predictor variables/IVs manipulated
- The allocation of participants
 - **Independent design** = different entities in all conditions
 - **Repeated measures design** = the same entities in all conditions
 - **Mixed design** = different entities in all conditions of at least one predictor/IV, the same entities in all conditions of at least one other predictor variable/IV



Moderation: Do *video* games lead to aggression?



ANDY FIELD



Moderation

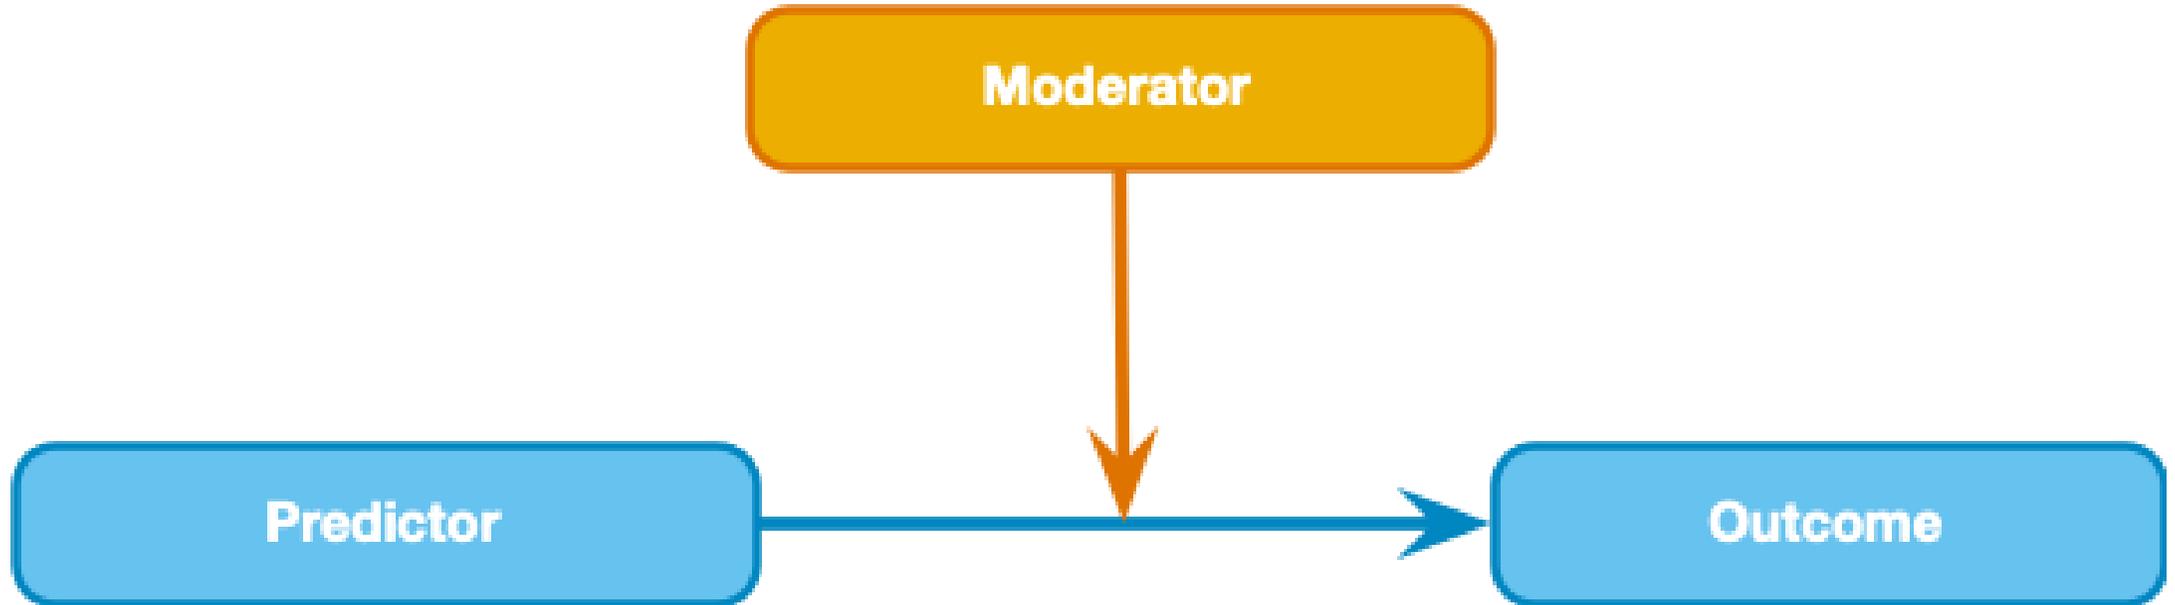
- Is there a link between video games and aggression?
 - Outcome = aggressive behaviour (**aggress**)
 - Callus Unemotional Traits (**caunts**)
 - Video game Use (**vid_game**)



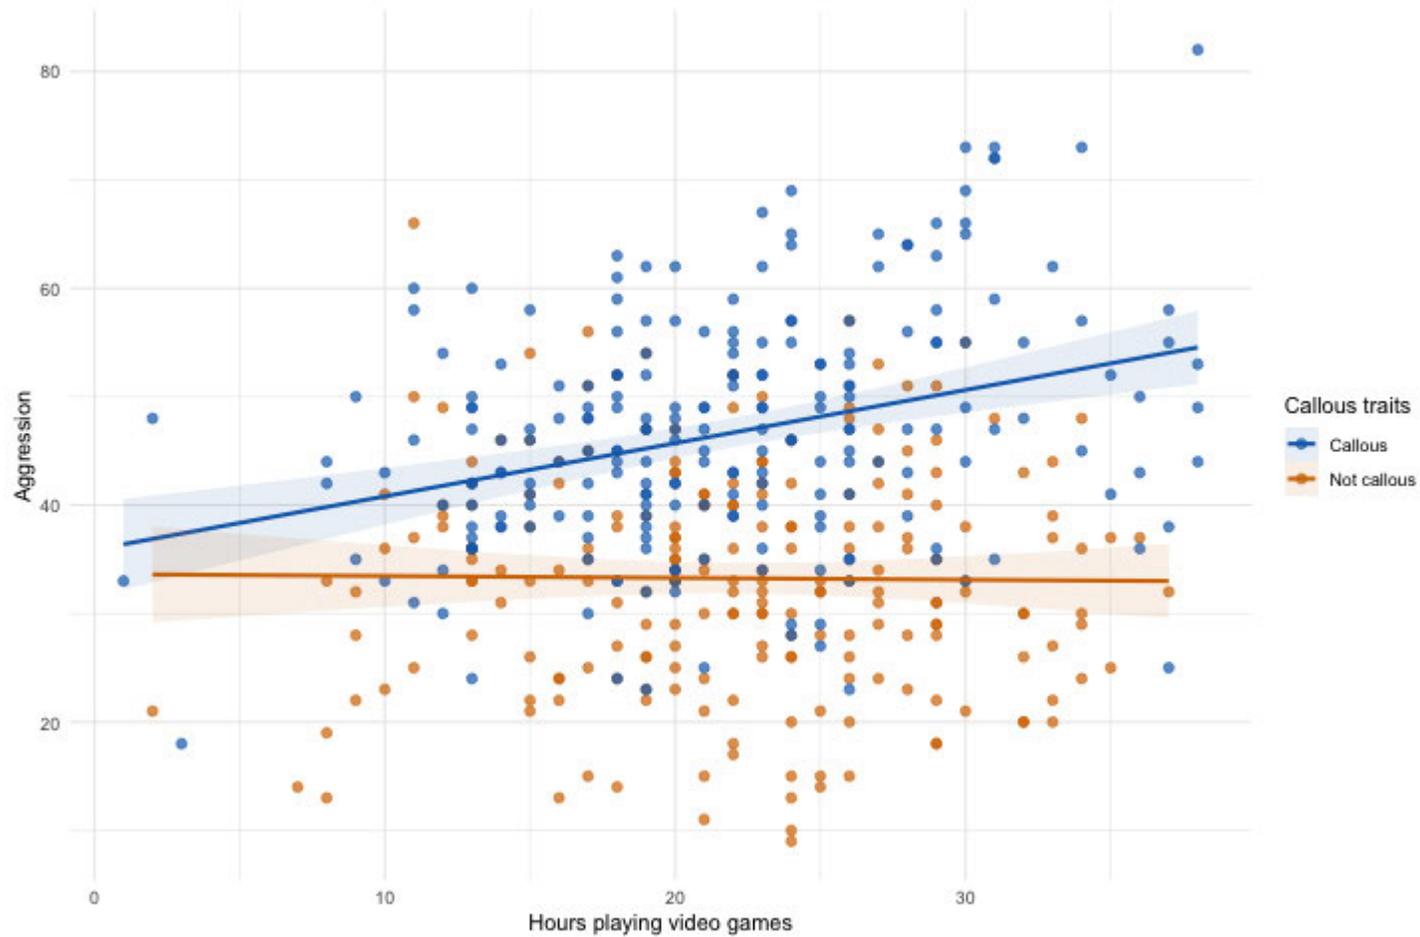
ANDY FIELD



Moderation: The theoretical model

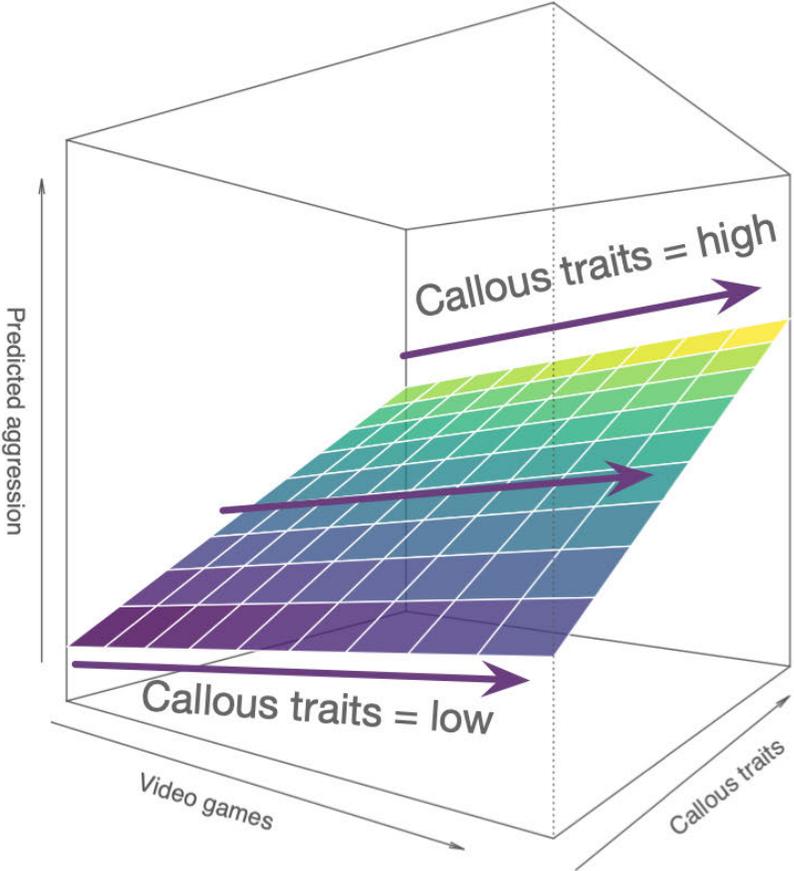


Categorical moderator variable

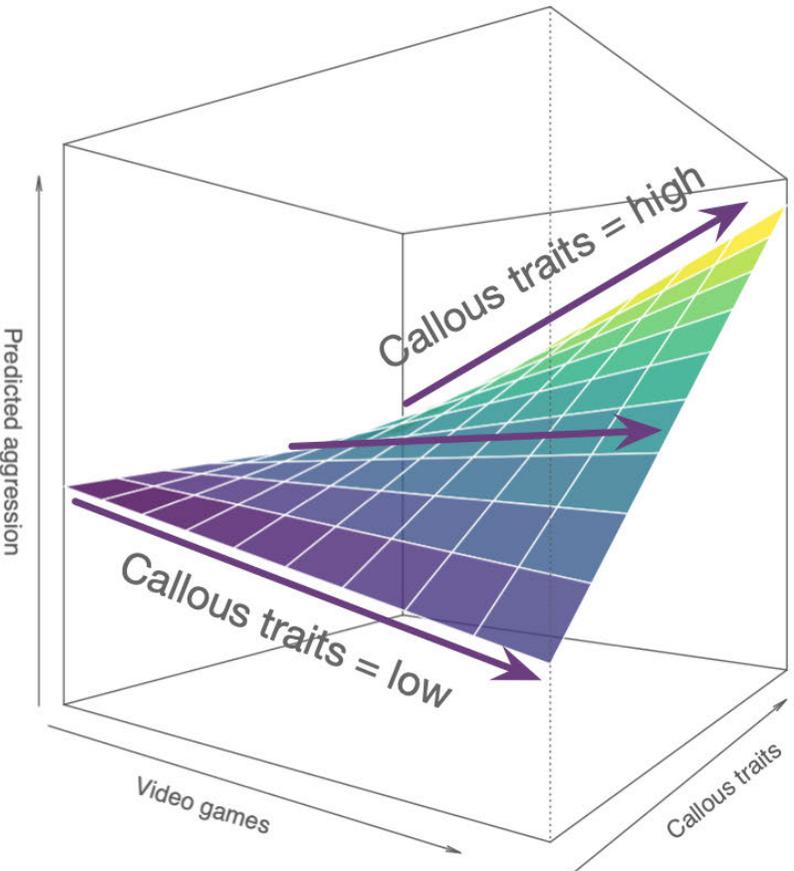


Continuous moderator variable

No moderation/interaction



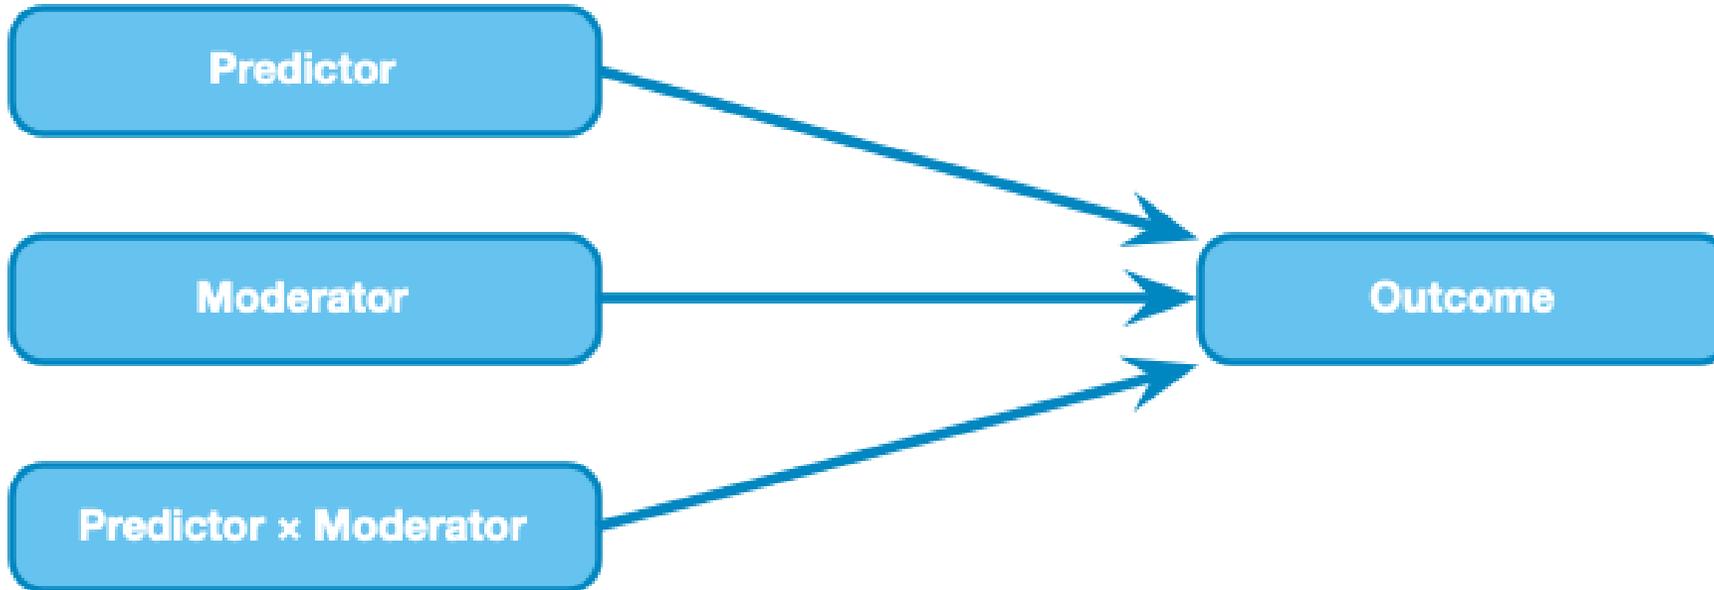
Moderation/interaction



ANDY FIELD



Moderation: The statistical model



$$\widehat{\text{aggression}}_i = \hat{b}_0 + \hat{b}_1 \text{video games}_i + \hat{b}_2 \text{caunts}_i + \hat{b}_3 \text{video games} \times \text{caunts}_i + e_i$$



Interactions in the linear model

- An interaction term is literally the product of the two variables:

$$\text{interaction} = \text{video_games} \times \text{caunts}$$

- Interpreting interactions
 - If the interaction term is significant, we have a significant moderation effect
 - The parameter estimate (b) for the interaction quantifies the size of the interaction effect
 - The effect of one predictor is stronger at some levels of another predictor
 - In factorial designs, the effect of one predictor is stronger in certain categories of the other predictor
 - Simple effects analysis



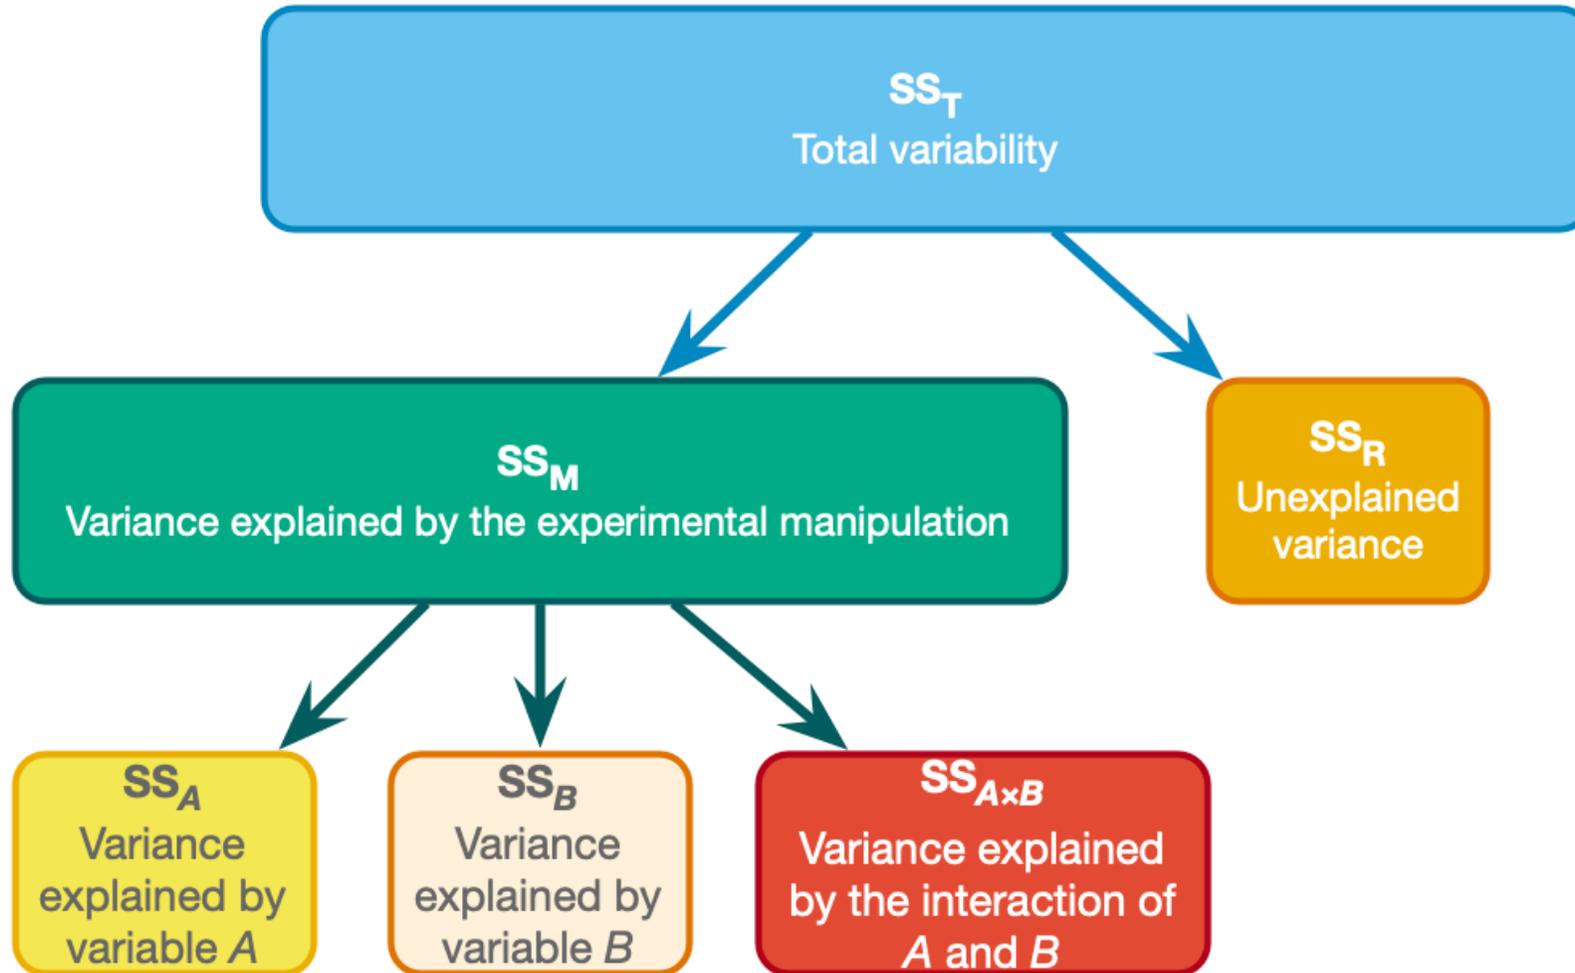
A dangerous example

- Injury severity (0-20) measured during video games
- 40 Participants
 - 10 = Xbox One kinect (Static game)
 - 10 = Xbox One kinect (Active game)
 - 10 = Nintendo Switch (Static game)
 - 10 = Nintendo Switch (Active game)
- Variables:
 - Predictor = Console (Xbox or Switch)
 - Predictor = Type of game (active or static)
 - Outcome = injury severity (0 = no injury, 20 = severe injury)
- The model

$$\widehat{\text{injuries}}_i = \hat{b}_0 + \hat{b}_1 \text{game}_i + \hat{b}_2 \text{console}_i + \hat{b}_3 \text{game} \times \text{console}_i + e_i$$



Partitioning variance



Fitting the model

- Remember that for F -statistics we typically want Type III sums of squares
- The `aov_4()` function from the `afex` package
 - An easier option
 - Automatically sets contrasts
 - Built in interaction plot with `afex_plot()`
 - But ... no parameter estimates, diagnostic plots, or robust methods
- Use `lm()`
 - A trickier option (but you *can* do it)
 - Manually set contrasts
 - Easy parameter estimates, diagnostic plots, and robust methods

```
xbox_afx <- afex::aov_4(injury ~ game*console  
+ (1|id), data = xbox_tib)
```

```
active_vs_static <- c(-1/2, 1/2)  
switch_vs_xbox <- c(-1/2, 1/2)  
contrasts(xbox_tib$game) <- active_vs_static  
contrasts(xbox_tib$console) <- switch_vs_xbox  
xbox_lm <- lm(injury ~ game*console, data =  
xbox_tib)  
car::Anova(xbox_lm, type = 3)
```

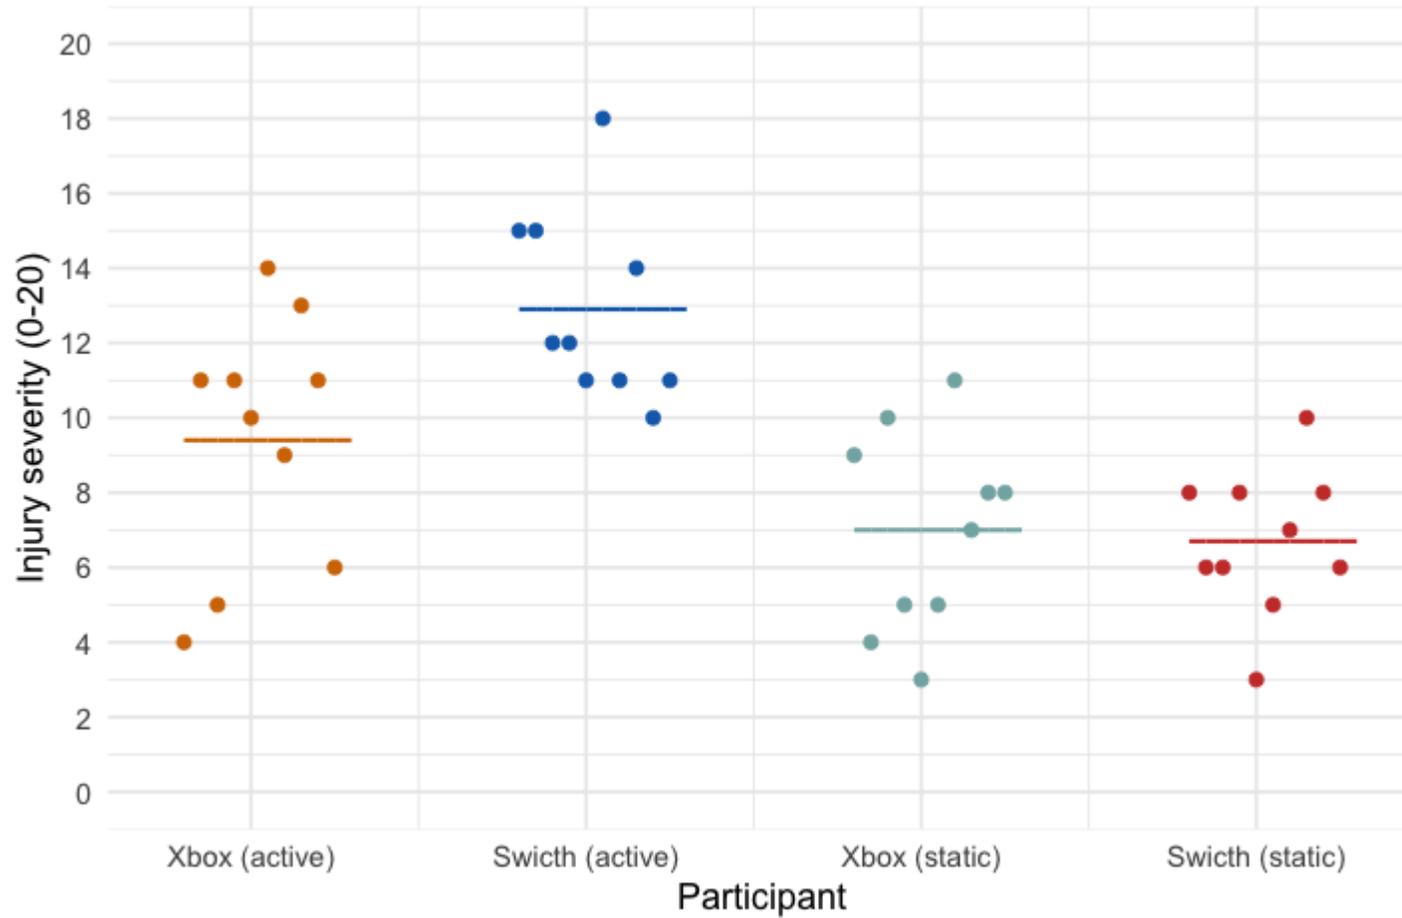
Overall model summary

```
xbox_afx <- afex::aov_4(injury ~ game*console + (1|id), data = xbox_tib)
```

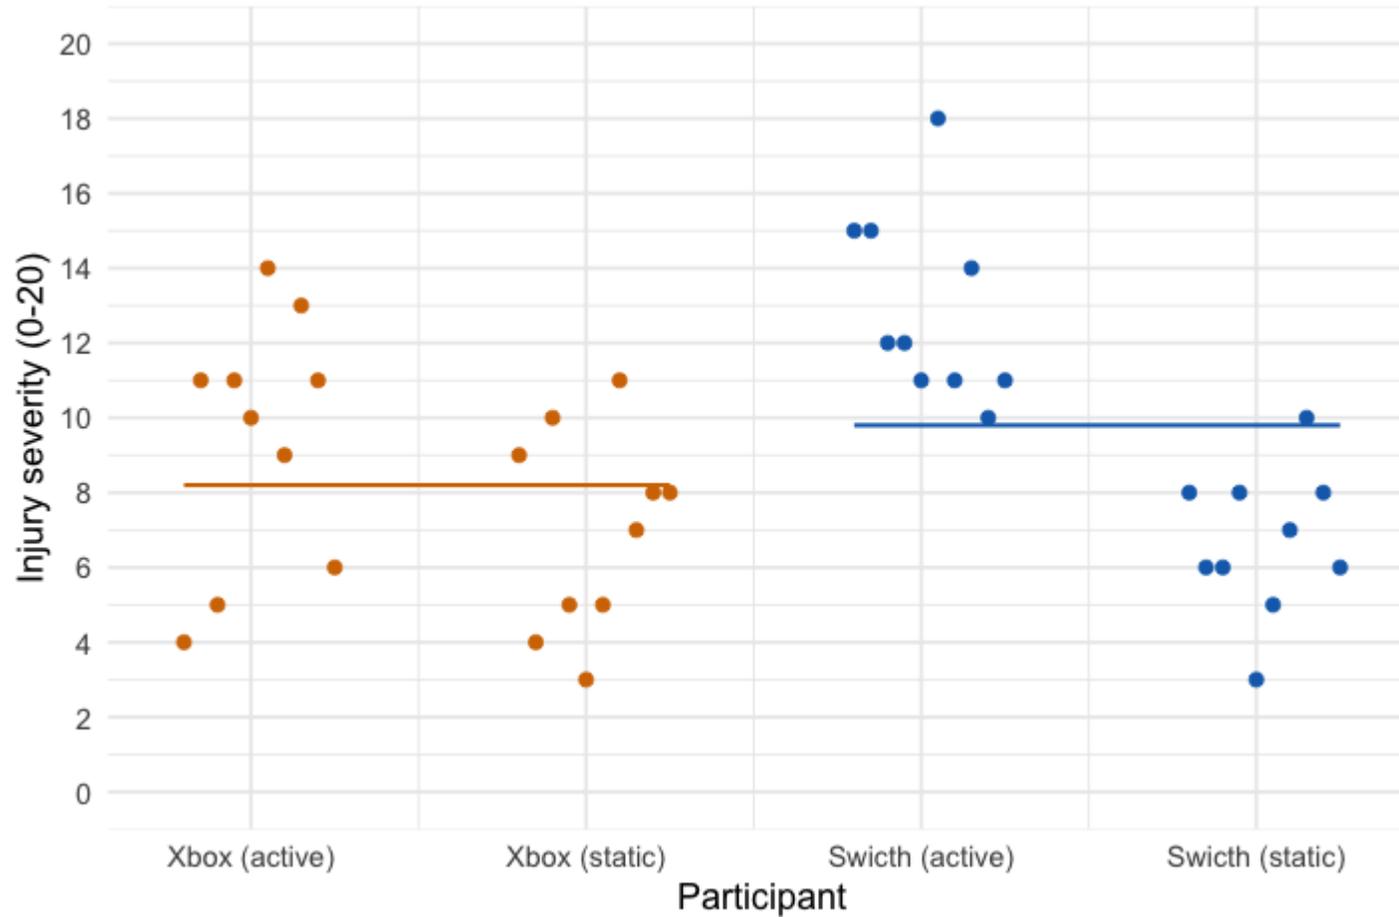
```
active_vs_static <- c(-1/2, 1/2)  
switch_vs_xbox <- c(-1/2, 1/2)  
contrasts(xbox_tib$game) <- active_vs_static  
contrasts(xbox_tib$console) <- switch_vs_xbox  
xbox_lm <- lm(injury ~ game*console, data = xbox_tib)  
car::Anova(xbox_lm, type = 3)
```

	Sum Sq	Df	F value	Pr(>F)
(Intercept)	3240.0	1	453.147	0.000
game	184.9	1	25.860	0.000
console	25.6	1	3.580	0.067
game:console	36.1	1	5.049	0.031
Residuals	257.4	36	NA	NA

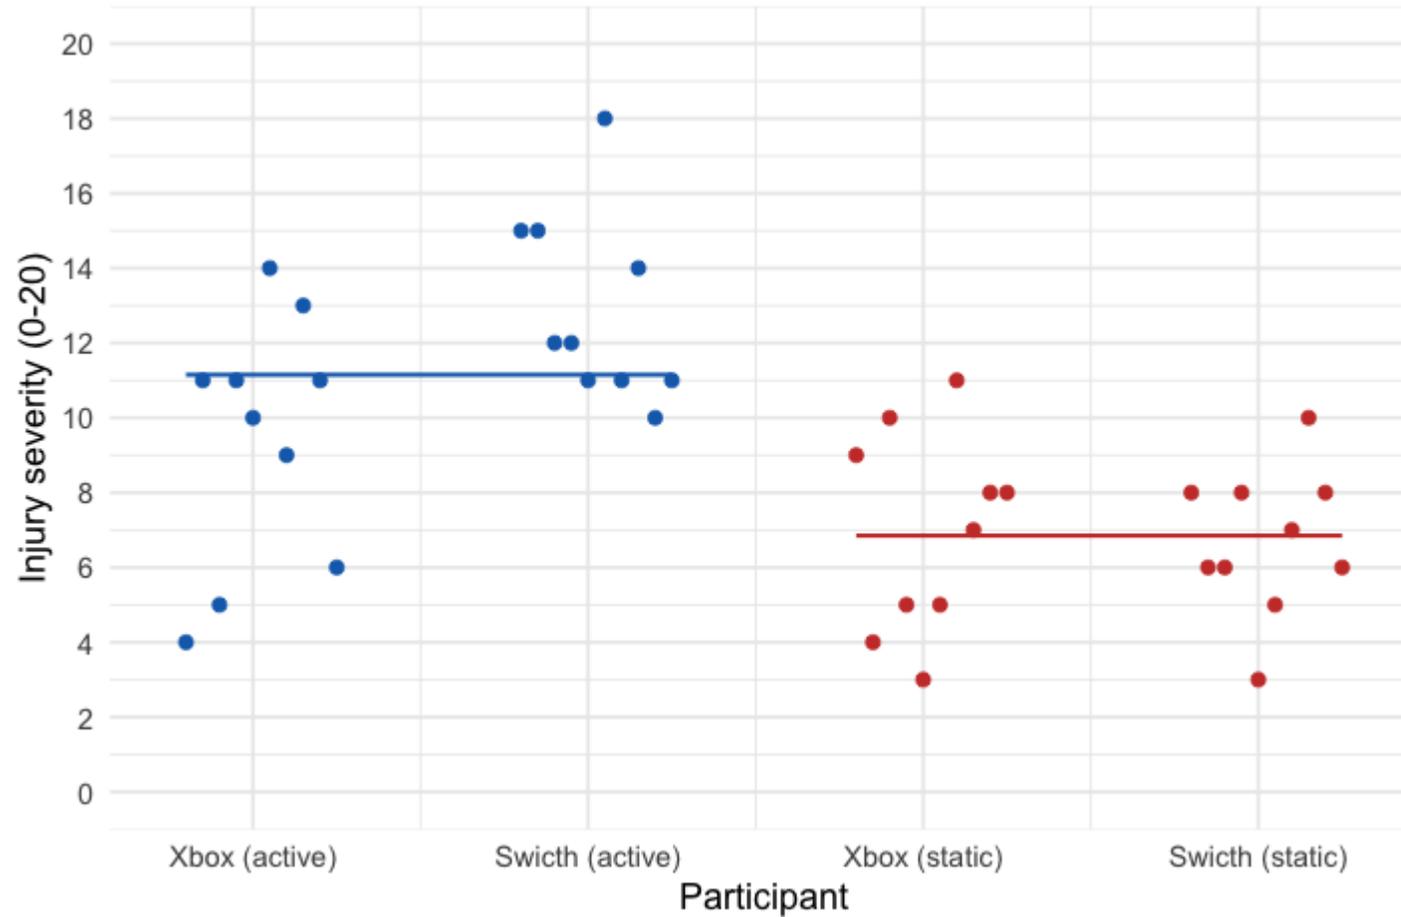
The data



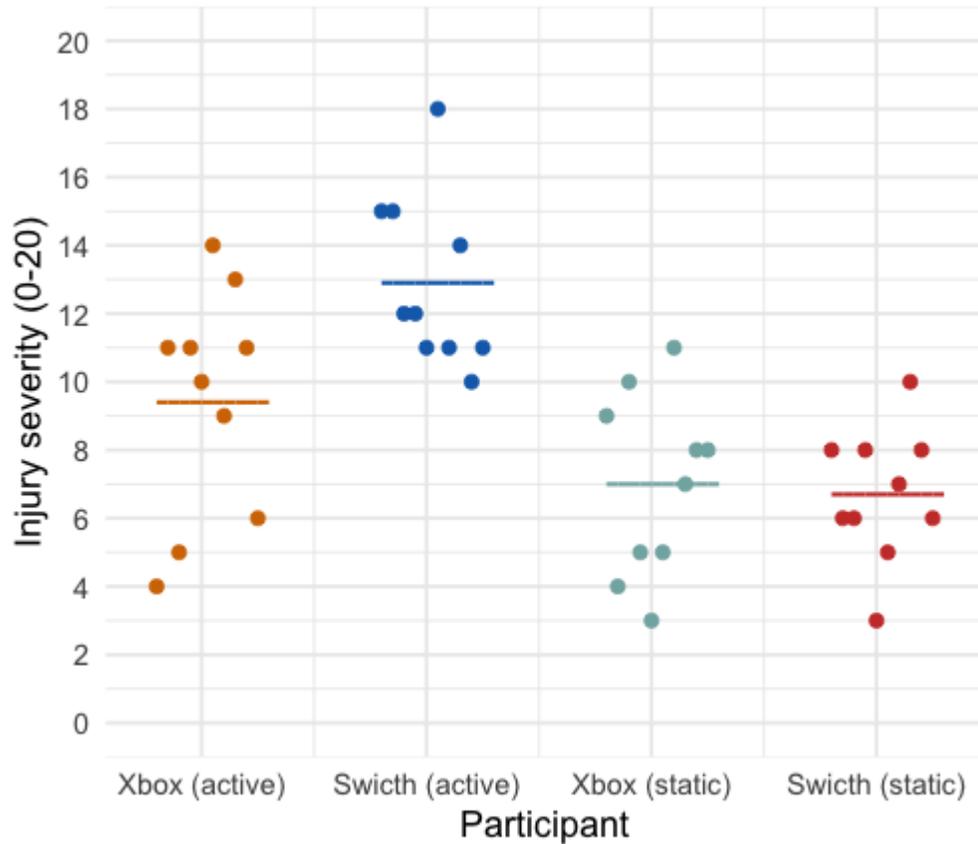
The main effect of console



The main effect of game



The interaction (moderation effect)

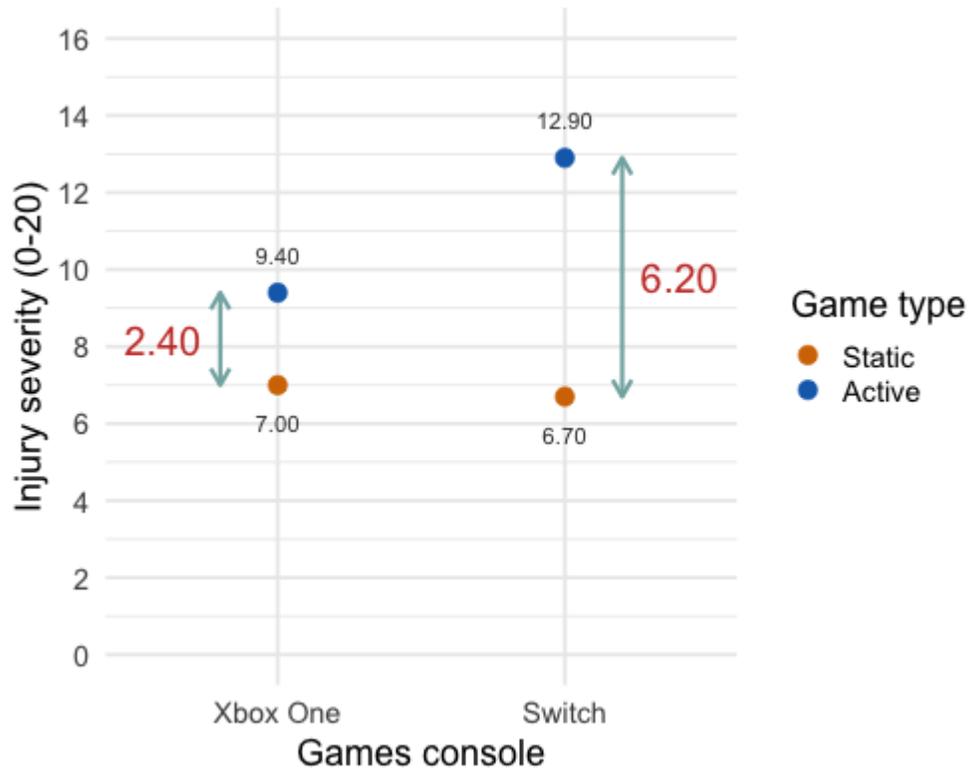


```
broom::tidy(xbox_lm)
```

term	estimate	std.error	statistic	p.value
(Intercept)	9.0	0.42	21.29	0.00
game1	4.3	0.85	5.09	0.00
console1	1.6	0.85	1.89	0.07
game1:console1	3.8	1.69	2.25	0.03

Interpreting interactions

Console moderating the effect of game



$$\begin{aligned} \text{game}_{\text{Switch}} &= \bar{X}_{\text{active, Switch}} - \bar{X}_{\text{static, Switch}} \\ &= 12.90 - 6.70 \\ &= 6.20 \end{aligned}$$

$$\begin{aligned} \text{game}_{\text{Xbox}} &= \bar{X}_{\text{active, Xbox}} - \bar{X}_{\text{static, Xbox}} \\ &= 9.40 - 7.00 \\ &= 2.40 \end{aligned}$$

$$\begin{aligned} \text{Interaction} &= \text{game}_{\text{Switch}} - \text{game}_{\text{Xbox}} \\ &= 6.2 - 2.40 \\ &= 3.8 \end{aligned}$$



Simple effects analysis

The effect of type of game within consoles

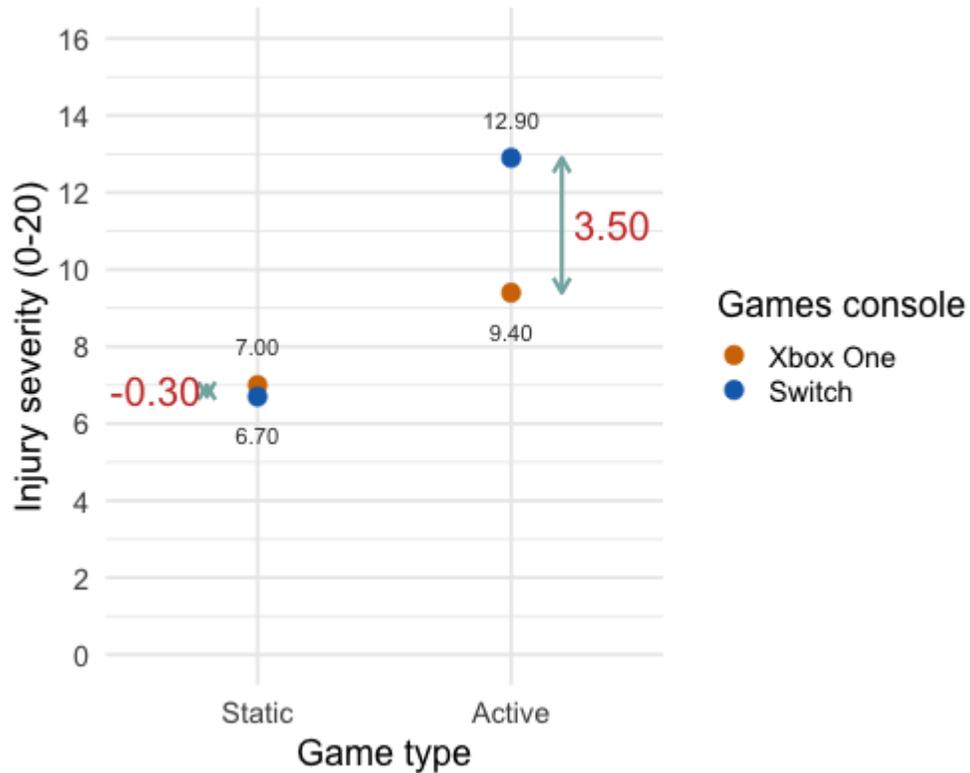
```
emmeans::joint_tests(xbox_lm, "console")
```

```
## console = Xbox One:  
## model term df1 df2 F.ratio p.value  
## game      1  36   4.028 0.0523  
##  
## console = Switch:  
## model term df1 df2 F.ratio p.value  
## game      1  36  26.881 <.0001
```

- The effect of playing active vs static games is less significant for the Xbox than the switch.
- The difference in mean injuries between active and static games for the **Xbox** ($M_{\text{difference}} = 2.40$) is not significantly different from 0 ($p = 0.052$) [but note the inherent daftness in this statement].
- The difference in mean injuries between active and static games for the the Nintendo **Switch** ($M_{\text{difference}} = 6.20$) is significantly different from 0 ($p < 0.0001$).

Interpreting interactions

Game moderating the effect of console



$$\begin{aligned}\text{console}_{\text{active}} &= \bar{X}_{\text{active, Switch}} - \bar{X}_{\text{active, Xbox}} \\ &= 12.90 - 9.4 \\ &= 3.5\end{aligned}$$

$$\begin{aligned}\text{console}_{\text{static}} &= \bar{X}_{\text{static, Switch}} - \bar{X}_{\text{static, Xbox}} \\ &= 6.7 - 7.00 \\ &= -0.3\end{aligned}$$

$$\begin{aligned}\text{Interaction} &= \text{console}_{\text{active}} - \text{console}_{\text{static}} \\ &= 3.5 - (-0.3) \\ &= 3.8\end{aligned}$$



Simple effects analysis

The effect of console within the type of game

```
emmeans::joint_tests(xbox_lm, "game")
```

```
## game = Static:
## model term df1 df2 F.ratio p.value
## console      1  36   0.063 0.8033
##
## game = Active:
## model term df1 df2 F.ratio p.value
## console      1  36   8.566 0.0059
```

- The effect of playing Xbox vs Switch games is less significant for static games than active ones.
- The difference in mean injuries between Xbox and Switch games that are **static** ($M_{\text{difference}} = -0.30$) is not significantly different from 0 ($p = 0.803$).
- The difference in mean injuries between Xbox and Switch games that are **active** ($M_{\text{difference}} = 3.50$) is significantly different from 0 ($p = 0.006$).

 Repeat the following mantra:

"It is never sensible to interpret main effects in the presence of a significant interaction effect."



Effect sizes

Raw effect size

- We've seen that the difference in the mean number of injuries when playing active games compared to static ones on the Switch (6.20) is 3.8 injuries more than for the Xbox (2.40).
- Alternatively, when playing Switch compared to Xbox there were 3.8 more injuries for active games compared to static ones.

Partial Omega squared ω_p^2

```
xbox_afx %>%  
  effectsize::omega_squared(., ci = 0.95)
```

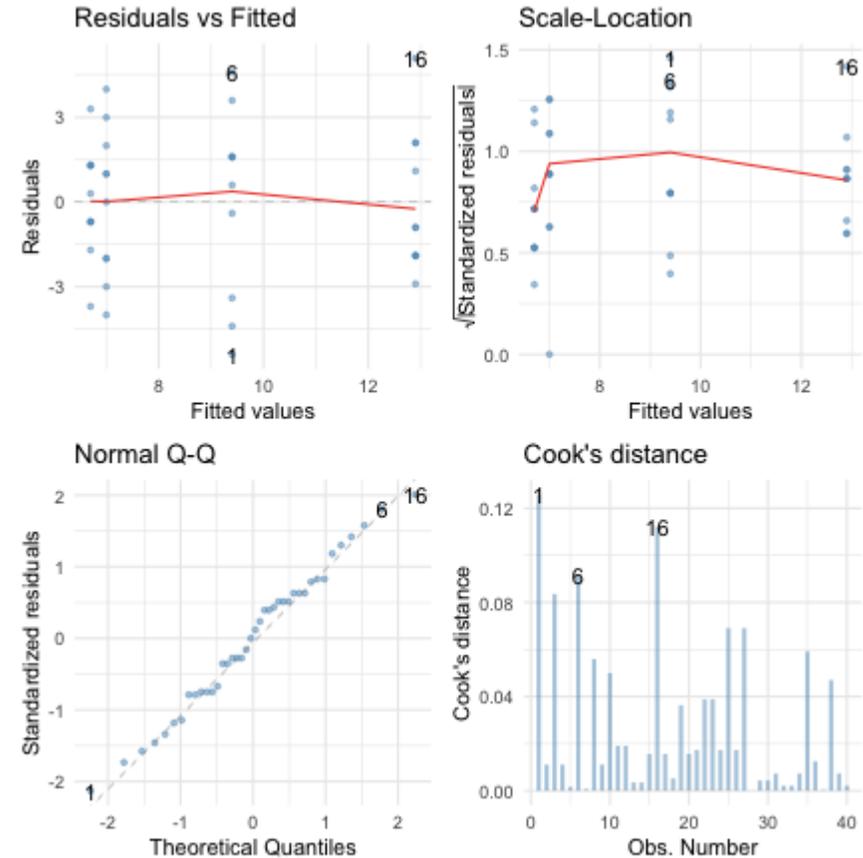
```
car::Anova(xbox_lm, type = 3) %>%  
  effectsize::omega_squared(., ci = 0.95)
```

- The interaction explains 9% of the variance in injuries not attributable to other predictors

	Parameter	Omega_Sq_partial	CI	CI_low	CI_high
2	game	0.38	0.95	0.15	0.58
3	console	0.06	0.95	-0.03	0.27
4	game:console	0.09	0.95	-0.03	0.31

Diagnostic plots

```
library(ggfortify)
ggplot2::autoplot(xbox_lm,
                  which = c(1, 3, 2, 4),
                  colour = "#5c97bf",
                  smooth.colour = "#ef4836",
                  alpha = 0.5,
                  size = 1) +
  theme_minimal()
```



Robust models

- If you use `lm()` you can fit any of the robust models we have discussed.

Example: HC4 standard errors

```
parameters::model_parameters(xbox_lm, robust = TRUE, vcov.type = "HC4", digits = 3)
```

Parameter	Coefficient	SE	CI_low	CI_high	t	df_error	p
(Intercept)	9.0	0.45	8.10	9.90	20.19	36	0.00
game1	4.3	0.89	2.49	6.11	4.82	36	0.00
console1	1.6	0.89	-0.21	3.41	1.80	36	0.08
game1:console1	3.8	1.78	0.18	7.42	2.13	36	0.04

The beer goggles effect



ANDY FIELD



A science reflecting art example

- The beer goggles effect: subjective perceptions of physical attractiveness become inaccurate after drinking alcohol
- Chen et al., (2014)¹
 - Alcohol consumption reduces accuracy in symmetry judgements
 - Symmetric faces have been shown to be rated as more attractive.
 - If the beer-goggles effect is driven by alcohol impairing symmetry judgements then you'd expect a stronger effect for unattractive (asymmetric) faces than attractive (symmetric) ones
- Fictional data but matches Chen et al. findings



Chen, et al., (2014). The moderating effect of stimulus attractiveness on the effect of alcohol consumption on attractiveness ratings. *Alcohol and Alcoholism*, 49, 515–519. <https://doi.org/10.1093/alcalc/agu026>



ANDY FIELD

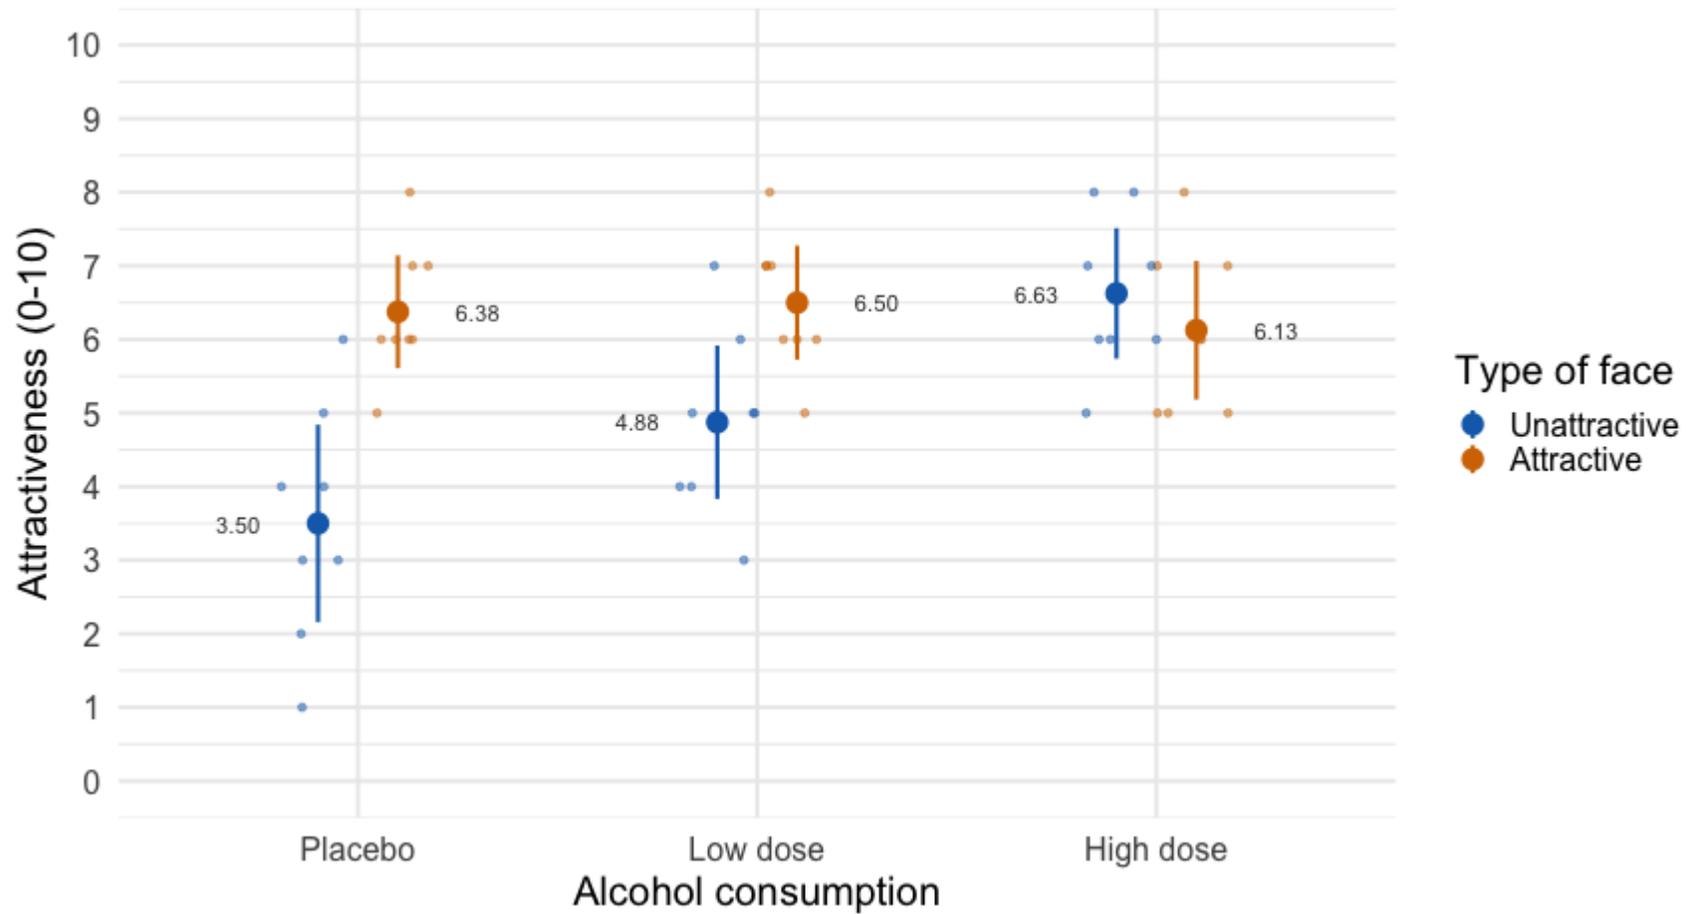


Study design

- 48 Participants (8 per group)
- Predictor: **Alcohol**
 - **Placebo** group: 500 ml of alcohol-free beer
 - **Low-dose** group: 500 ml of beer (4% ABV);
 - **High-dose** group: 500 ml of beer (7% ABV).
- Predictor: **Face type**
 - Rated 50 **unattractive** (asymmetric) faces
 - Rated 50 **attractive** (symmetric) faces
- Outcome:
 - Median rating of the 50 photos on a scale from 0 (pass me a paper bag) to 10 (pass me their phone number)



The data



Overall model summary

```
goggles_afx <- afex::aov_4(attractiveness ~ facetype*alcohol + (1|id), data = goggles_tib)
```

```
alcohol_vs_none <- c(-2/3, 1/3, 1/3)
low_vs_high <- c(0, -1/2, 1/2)
contrasts(goggles_tib$alcohol) <- cbind(alcohol_vs_none, low_vs_high)
contrasts(goggles_tib$facetype) <- c(-1/2, 1/2)
goggles_lm <- lm(attractiveness ~ facetype*alcohol, data = goggles_tib)
car::Anova(goggles_lm, type = 3)
```

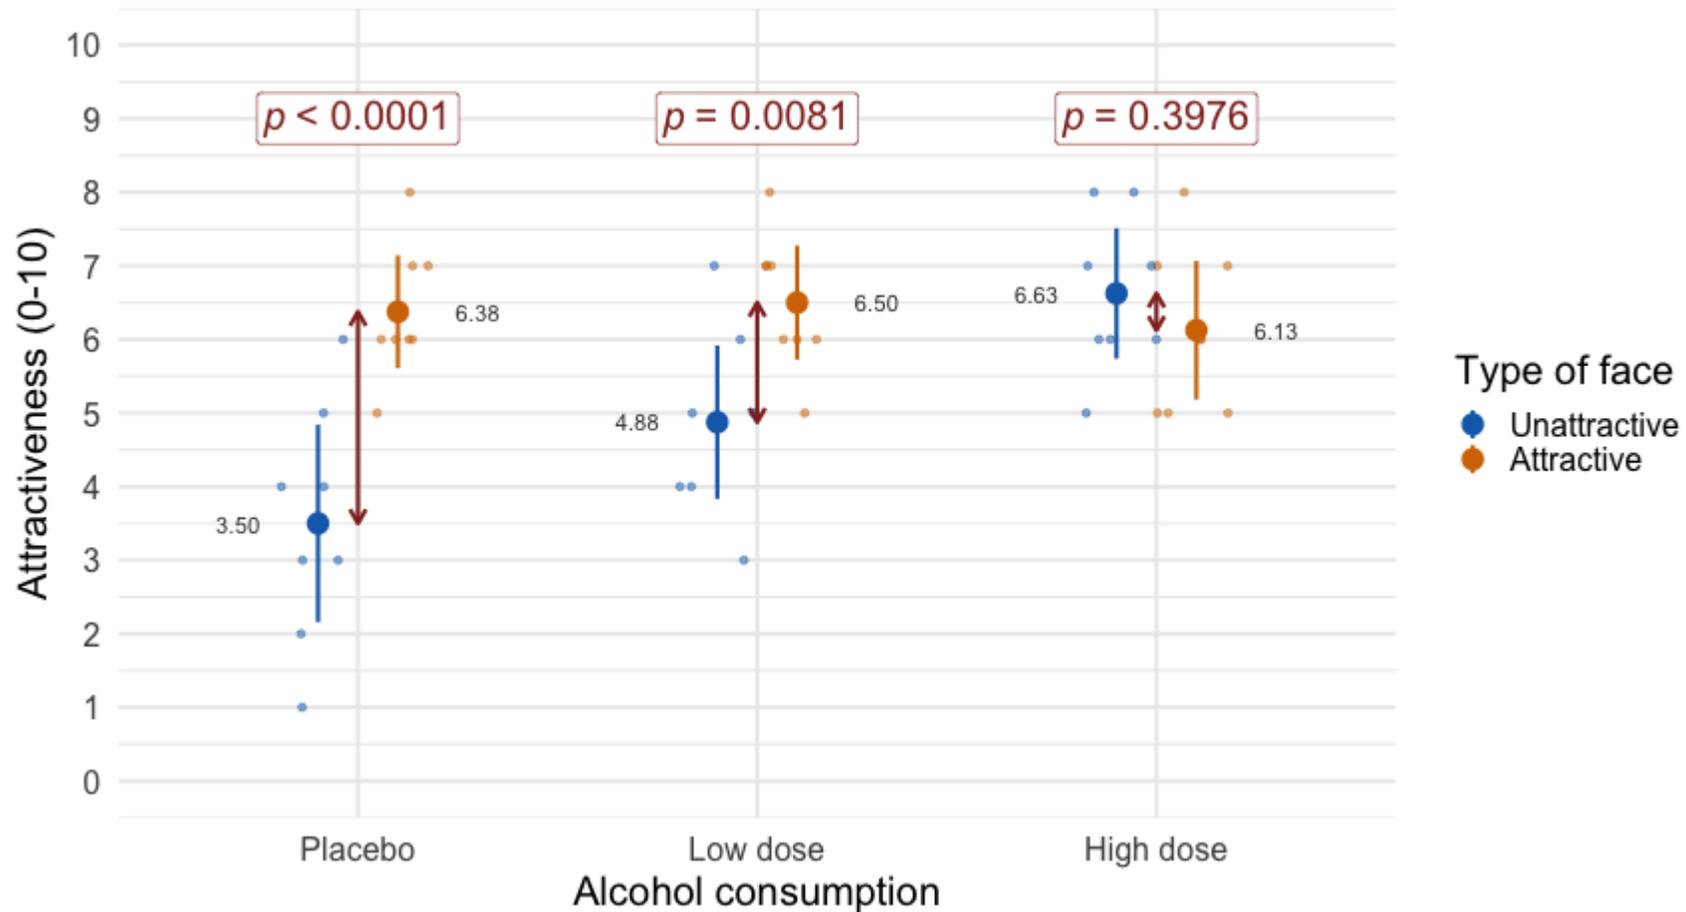
	Sum Sq	Df	F value	Pr(>F)
(Intercept)	1541.333	1	1125.843	0.000
facetype	21.333	1	15.583	0.000
alcohol	16.542	2	6.041	0.005
facetype:alcohol	23.292	2	8.507	0.001
Residuals	57.500	42	NA	NA

Simple effects: Facetype within alcohol group

```
## alcohol = Placebo:  
## model term df1 df2 F.ratio p.value  
## facetype    1  42  24.150 <.0001  
##  
## alcohol = Low dose:  
## model term df1 df2 F.ratio p.value  
## facetype    1  42   7.715 0.0081  
##  
## alcohol = High dose:  
## model term df1 df2 F.ratio p.value  
## facetype    1  42   0.730 0.3976
```



Simple effects: Facetype within alcohol group

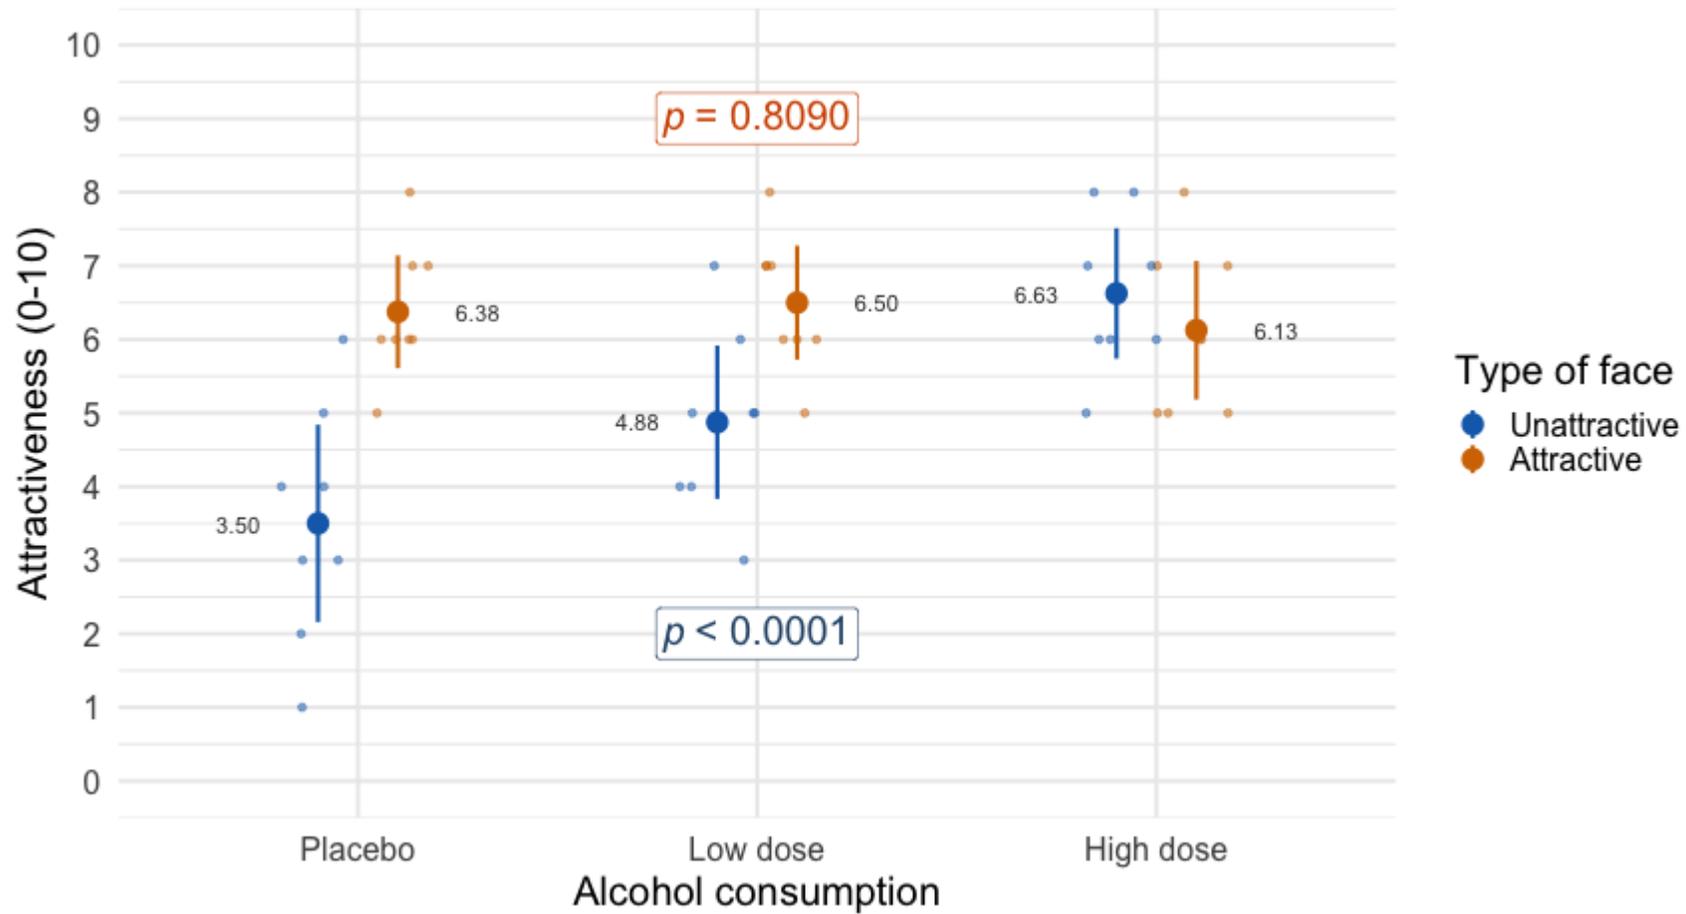


Simple effects: alcohol within facetype

```
## facetype = Unattractive:  
## model term df1 df2 F.ratio p.value  
## alcohol      2  42  14.335 <.0001  
##  
## facetype = Attractive:  
## model term df1 df2 F.ratio p.value  
## alcohol      2  42   0.213 0.8090
```

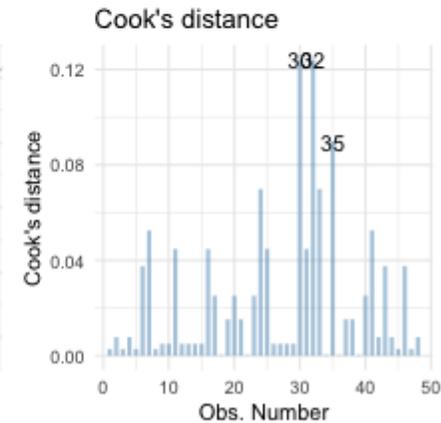
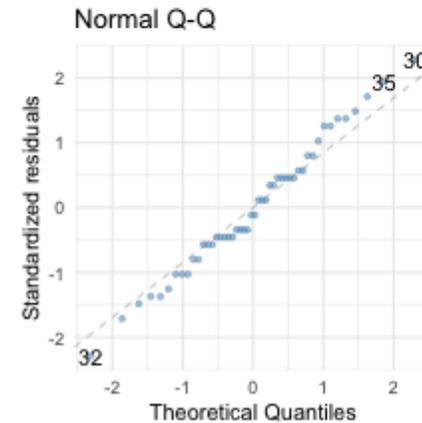
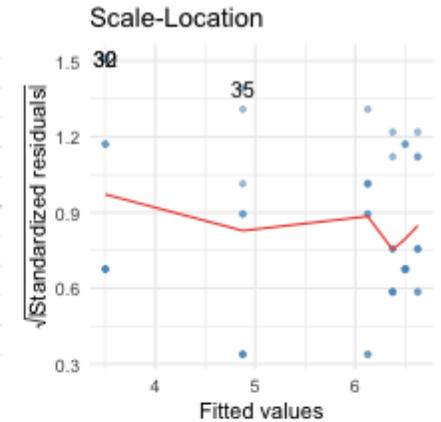
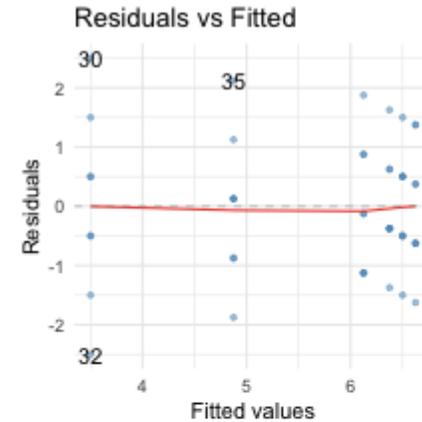


Simple effects: alcohol within facetype



Diagnostic plots

```
library(ggfortify)
ggplot2::autoplot(goggles_lm,
                  which = c(1, 3, 2, 4),
                  colour = "#5c97bf",
                  smooth.colour = "#ef4836",
                  alpha = 0.5,
                  size = 1) +
  theme_minimal()
```



Summary

- Factorial designs: two or more predictor variables are manipulated
- Moderation: Where the effect of one predictor differs at levels of another
- Testing moderation (interaction effects)
 - The model: a linear model in which predictors are entered as well as their interaction
 - Interpret interactions using plots or simple effects analysis, which quantifies the effect of one predictor at each level of another
 - Test model assumptions with diagnostic plots
 - Fit robust variants
- Main effects
 - Show the effect of one predictor variable collapsing across levels of other predictors
 - In the presence of a significant interaction main effects are uninteresting

