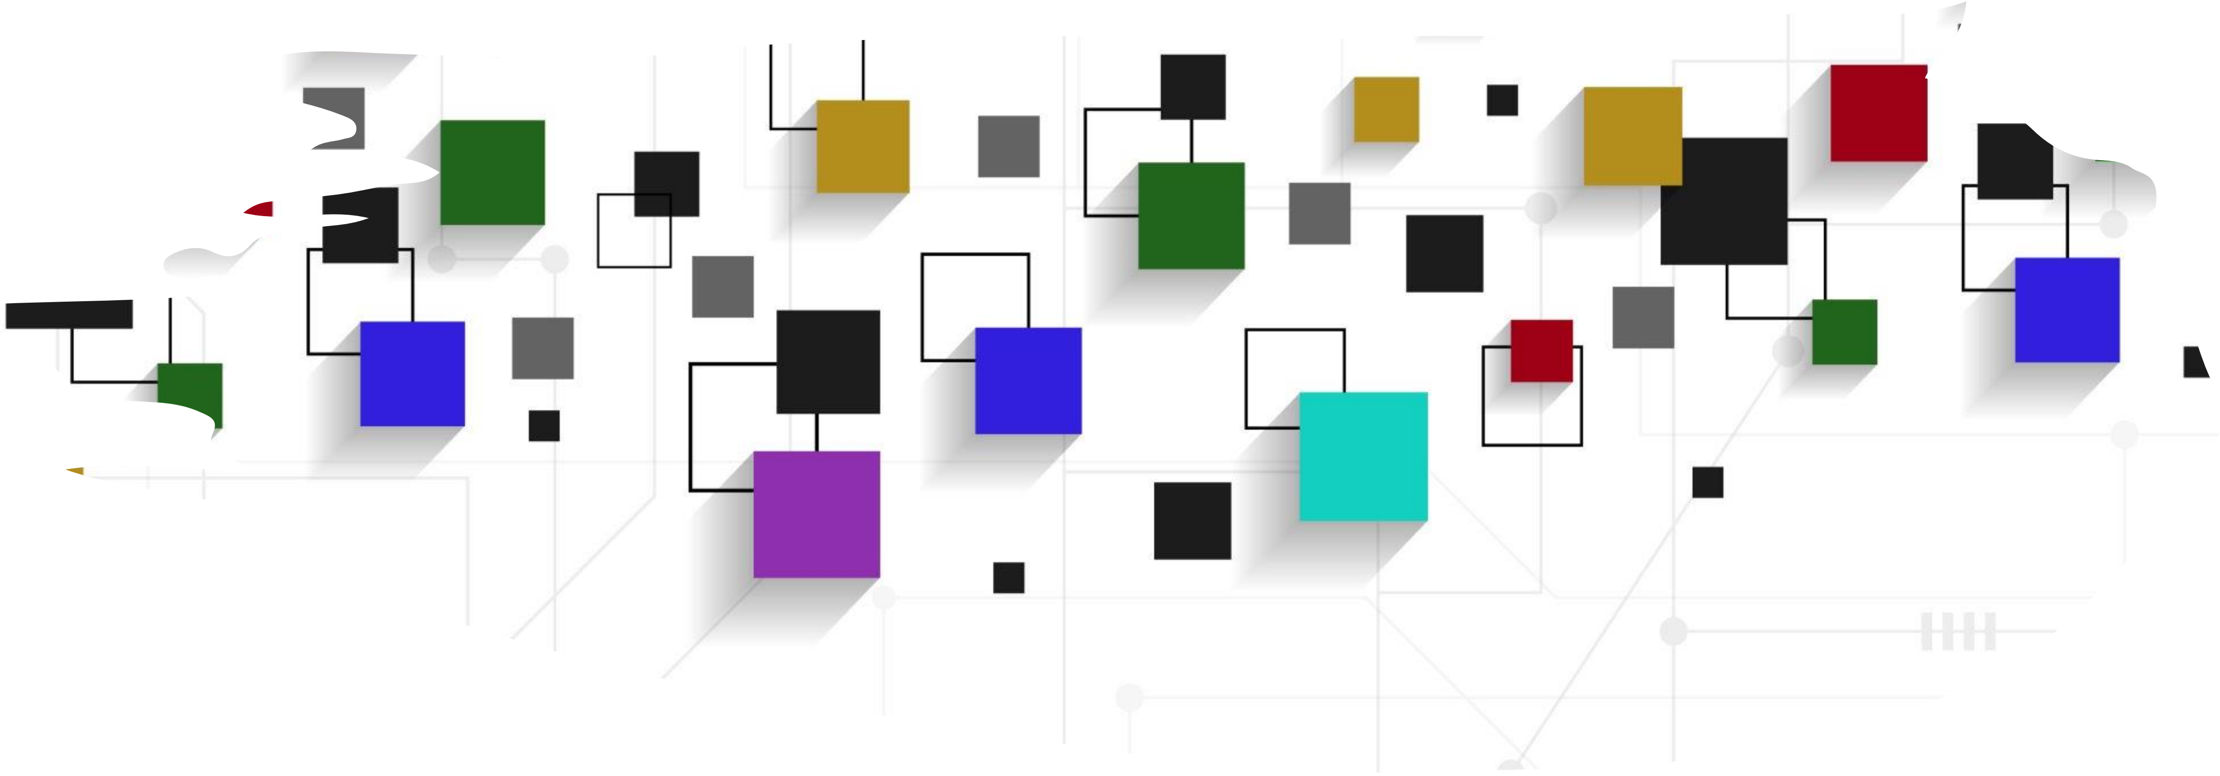

CogLab: Complex Statistical Models

WEEK 10

project checklist

- project checklist

	Pilot 1	Pilot 2	Pilot 3
Which browser were you using?			
Which operating system (Mac / Windows / iPad, etc.)			
Date of piloting			
Were instructions clear? Please note down which instructions had typos / were unclear			
How long did the task take you?			
Was there a consent form?			
Was the demographic survey displayed correctly?			
Did you see the data being displayed at the end of the study?			
What do you think the experiment was about?			
Any other comments?			

CogLab Project Checklist

Task	Check if done
Sanity Check ☐ Is the attention check response being recorded? ☐ Is the free association response being recorded? ☐ Can you differentiate between training / attention / association / prime / target? ☐ Can you differentiate between prime and target trials? ☐ Can you differentiate practice and test trials? ☐ Is subject ID being recorded? ☐ Is RT being recorded?	☐
For the **demographic survey**, how are you showing these questions? Are there multiple answers people can pick or is it a binary choice? Are people able to select multiple answers when they should not be?	☐
For the **demographic survey,** what questions are being shown on the same screen? What questions should be on different screens?	☐
For the **demographic survey**, how are the data being recorded and is it being recorded? Also, do you have all the questions you need?	☐

Before Pre-Registration:
☐ Are you providing accuracy feedback on priming practice trials?
☐ Have you addressed ALL the feedback from Milestone 4?
    ☐ Feedback 1
    ☐ Feedback 2
    ☐ Feedback 3
☐ Are you recording IP addresses?
☐ Are you commenting the condition definition inside cognition.run
☐ Have you piloted your experiment with Uma + other group + 5 friends)
☐ Have they completed the pilot feedback sheet?
☐ Have you sent the cognition.run link by Nov 10?
☐ Have you finalized the analysis plan + sample size?
☐ Have you created and submitted a pre-registration draft?
	☐

Analysis ☐ Did you confirm/correct all datatypes? ☐ Did you figure out how to "filter" certain types of trials? ☐ Did you fix all typos in attention responses? ☐ Have you computed mean attention accuracy? ☐ Have you applied exclusions based on accuracy AND RTs? ☐ Have you created an RT bar graph? ☐ Have you fit a statistical model?	☐

pre-registration + project checklist

- milestone #5: pre-registration + project checklist + piloting (Nov 10)

1. **Data Collection**: Have any data been collected for this study already?
2. **Main Question**: What is the main question being asked or hypothesis being tested in this study?
3. **Dependent Variable(s)**: Describe the key dependent variable(s) specifying how they will be measured.
4. **Condition(s)**: How many and which conditions will participants be assigned to? Please include an example trial of each type of condition you have in your experiment. Please also specify which independent variable will be within-participants or between-participants.
5. **Analyses**: Specify exactly which analyses you will conduct to examine the main question/hypothesis.
6. **Outliers & Exclusions**: Describe exactly how outliers will be defined and handled, and your precise rule(s) for excluding observations.
7. **Predicted Plot**: Please submit a predicted plot for your study based on what you expect the pattern to look like for your main hypothesis.
8. **Sample Size**: How many observations will be collected or what will determine sample size? No need to justify the decision, but be precise about exactly how the number will be determined.
9. **Exploratory details**: Anything else you would like to pre-register? (e.g., secondary analyses, variables collected for exploratory purposes, unusual analyses planned?)

recap

- what we covered:
  - linear regression and one-way ANOVAs
- your to-do's were:
  - *submit:* formative assignment #2
  - *work on:* pre-registration + project checklist

today’s agenda

- two-way ANOVA
- linear mixed effects models

complex models

- what if the variation in our data (Y) could be explained further?
- data = model + error
  - *one IV*: Y = a + bX + error
  - *multiple IVs*: $Y = a + b_1X_1 + b_2X_2 + \ldots + error$
- central idea remains the same, but more complex relationships are possible

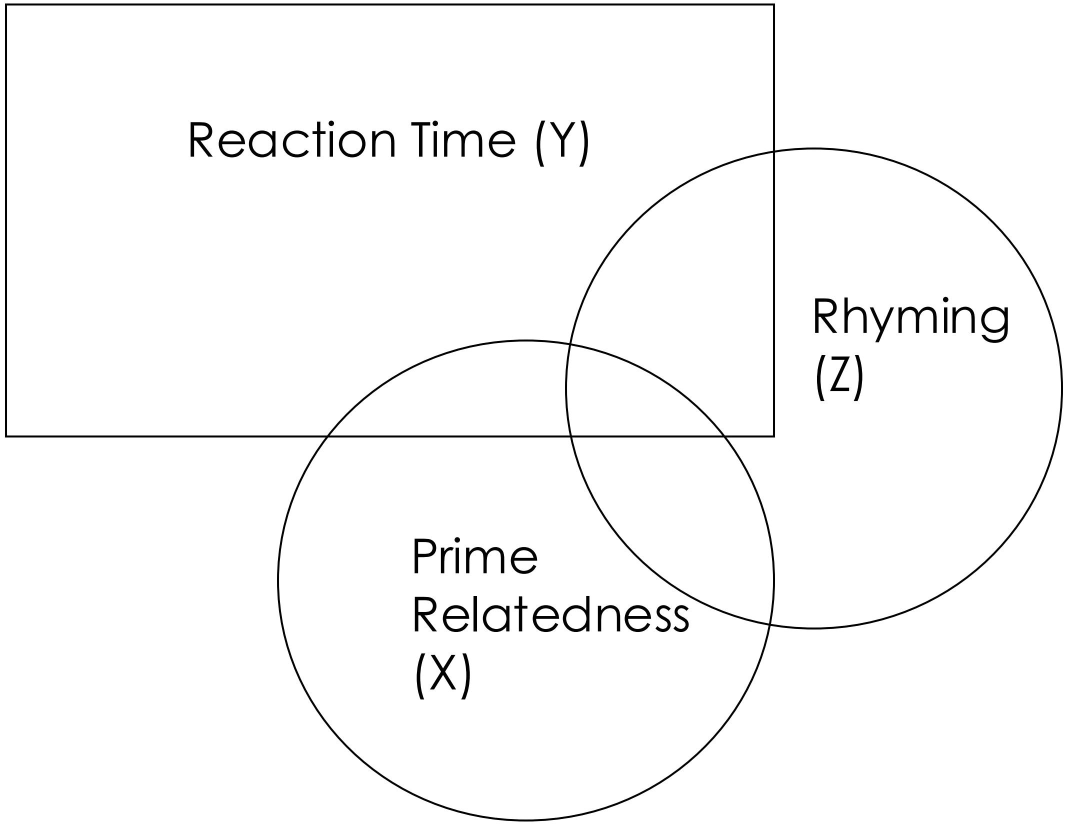

multiple IVs: data

- we will use the jobsatisfaction dataset from the datarium package
- install the package datarium
- new heading (# multiple IVs) & code chunk
- load and view the jobsatisfaction dataset

```
data("jobsatisfaction", package = "datarium")
View(jobsatisfaction)
```

id	gender	education_level	score
1	male	school	5.51
2	male	school	5.65
3	male	school	5.07
4	male	school	5.51
5	male	school	5.94
6	male	school	5.80
7	male	school	5.22
8	male	school	5.36
9	male	school	4.78
10	male	college	6.01
11	male	college	6.01
12	male	college	6.45

multiple IVs: exploration

- let's explore the data:
  - visualize the pattern via a boxplot

multiple IVs: exploration

- let's explore the data:
  - visualize the pattern via a boxplot
  - do you see differences in job satisfaction?

```
jobsatisfaction %>%
  ggplot()+
  geom_boxplot(aes(x = gender, y = score, color = education_level))
```

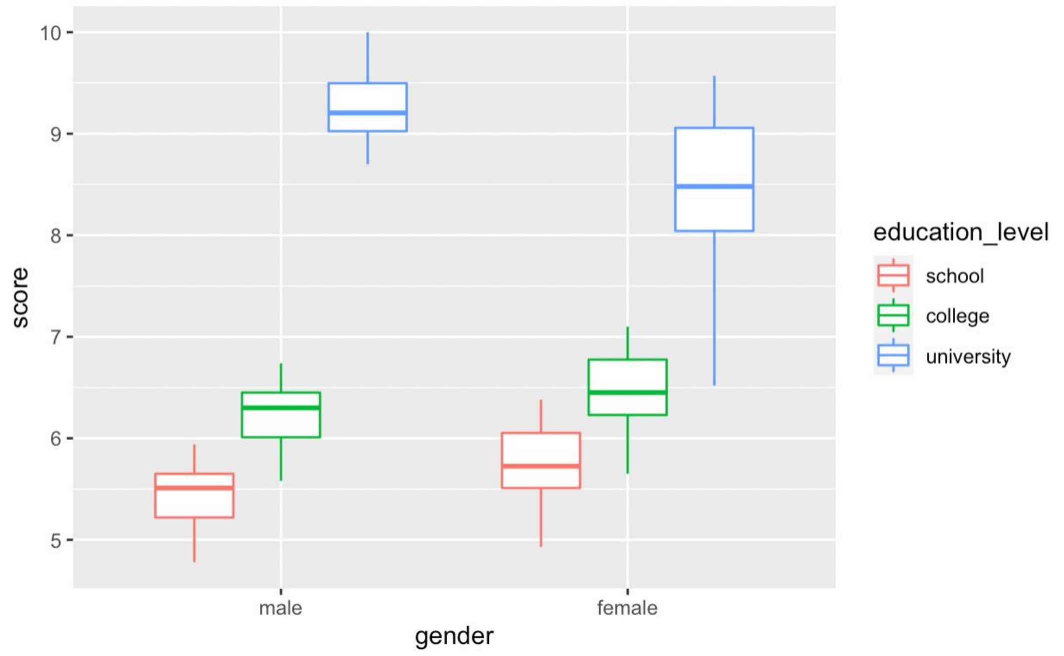

multiple IVs: research question

- does job satisfaction vary as a function of gender and education level?
- dependent variable?
- independent variable?

```
jobsatisfaction %>%
  ggplot()+
  geom_boxplot(aes(x = gender, y = score, color = education_level))
```

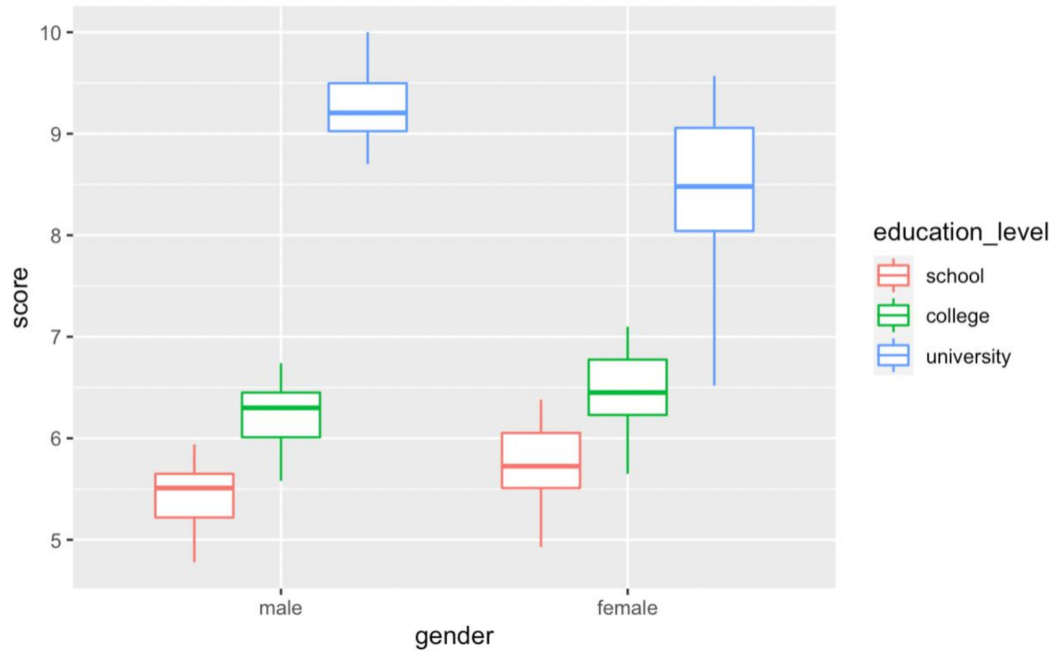

main effects

- when you have multiple variables in your experiment design, there are few different possibilities for how the pattern of data might look
- you could have the dependent variable vary as a function of IV1 and/or IV2 (main effects), and these effects might interact with each other
- main effects refer to differences in means of levels of an independent variable
- what is an example of a main effect for the jobsatisfaction dataset?
- what would the plot of this main effect look like?

id	gender	education_level	score
1	male	school	5.51
2	male	school	5.65
3	male	school	5.07
4	male	school	5.51
5	male	school	5.94
6	male	school	5.80
7	male	school	5.22
8	male	school	5.36
9	male	school	4.78
10	male	college	6.01
11	male	college	6.01
12	male	college	6.45

interactions

- interactions refer to situations when the difference in means between IV1's levels differs based on the levels of IV2, i.e., you cannot simply infer a difference in means
- what is an example of an interaction for the jobsatisfaction dataset?
- what would the plot of this interaction look like?

id	gender	education_level	score
1	male	school	5.51
2	male	school	5.65
3	male	school	5.07
4	male	school	5.51
5	male	school	5.94
6	male	school	5.80
7	male	school	5.22
8	male	school	5.36
9	male	school	4.78
10	male	college	6.01
11	male	college	6.01
12	male	college	6.45

visually…

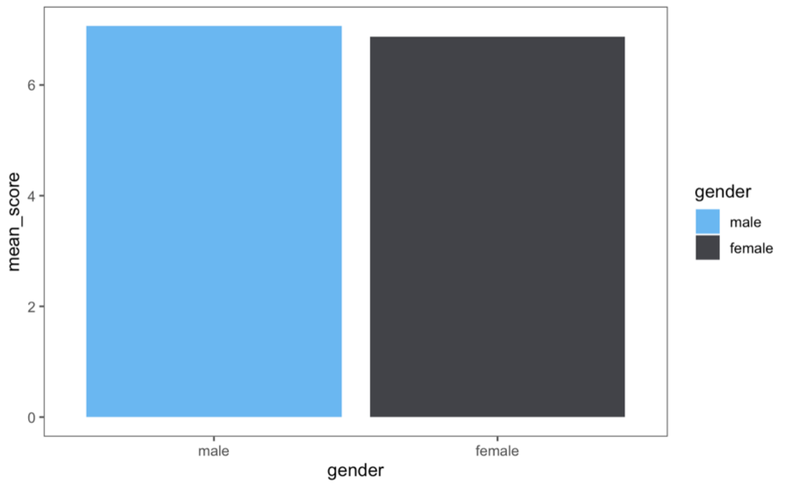

0
2
4
6
mean_score
male
female
gender
gender
male
female

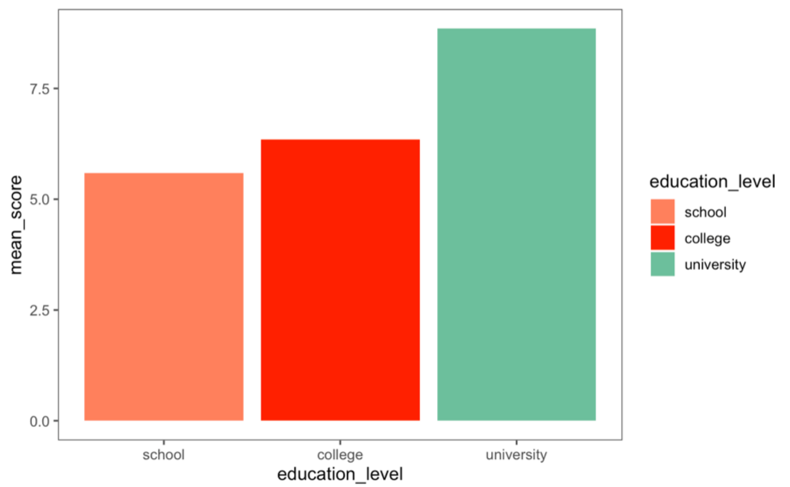

0.0
2.5
5.0
7.5
mean_score
school
college
university
education_level
education_level
school
college
university

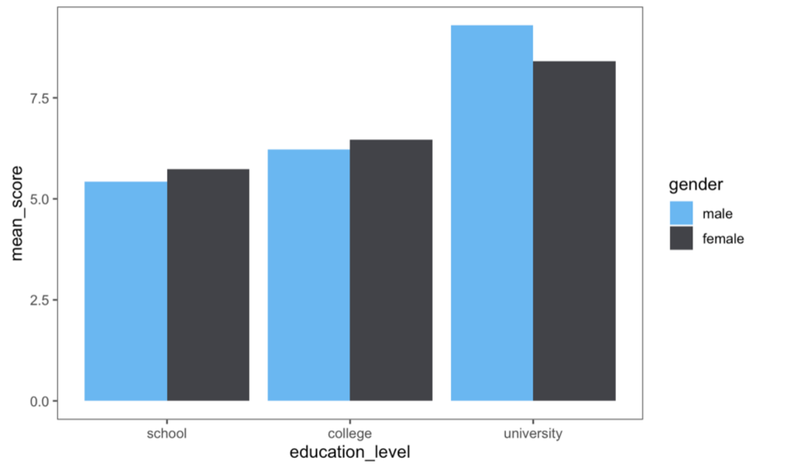

0.0
2.5
5.0
7.5
mean_score
school
college
university
education_level
gender
male
female

6.90
6.95
7.00
7.05
mean_score
male
female
gender
gender
male
female

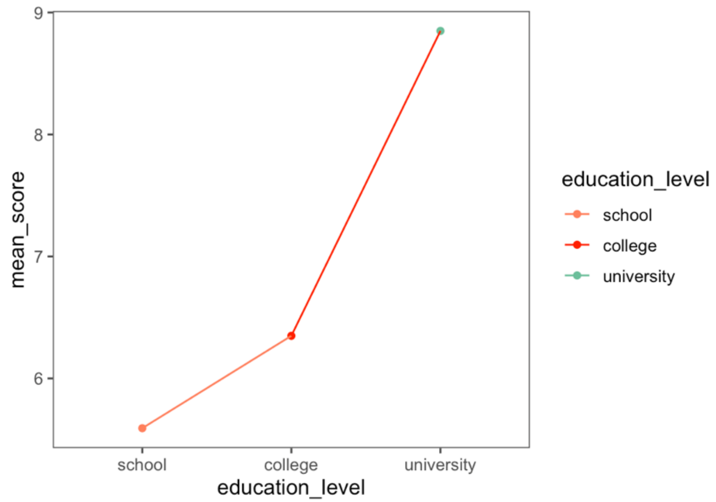

6
7
8
9
mean_score
school
college
university
education_level
education_level
school
college
university

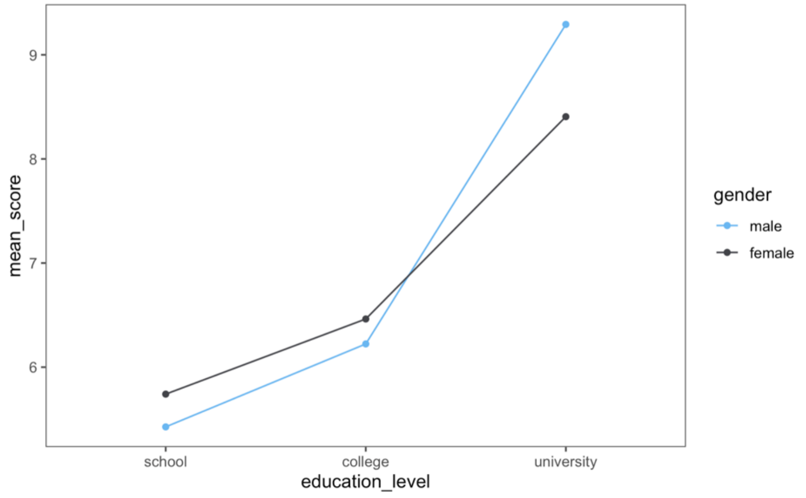

6
7
8
9
mean_score
school
college
university
education_level
gender
male
female

mathematically…

- main effect of gender:
  - mean (male) – mean (female)
- main effect of education level
  - mean(school) – mean (college)
  - mean(college) – mean (university)
  - mean(university) – mean(school)
- interaction (difference of differences)
  - $\text{diff(male-female)}_{school} - \text{diff(male-female)}_{college}$
  - $\text{diff(male-female)}_{university} - \text{diff(male-female)}_{college}$
  - $\text{diff(male-female)}_{school} - \text{diff(male-female)}_{university}$

gender <fctr>	education_level <fctr>	mean <dbl>	sd <dbl>
male	school	5.426667	0.3638681
male	college	6.223333	0.3396322
male	university	9.292000	0.4445422
female	school	5.741000	0.4744225
female	college	6.463000	0.4746941
female	university	8.406000	0.9379078

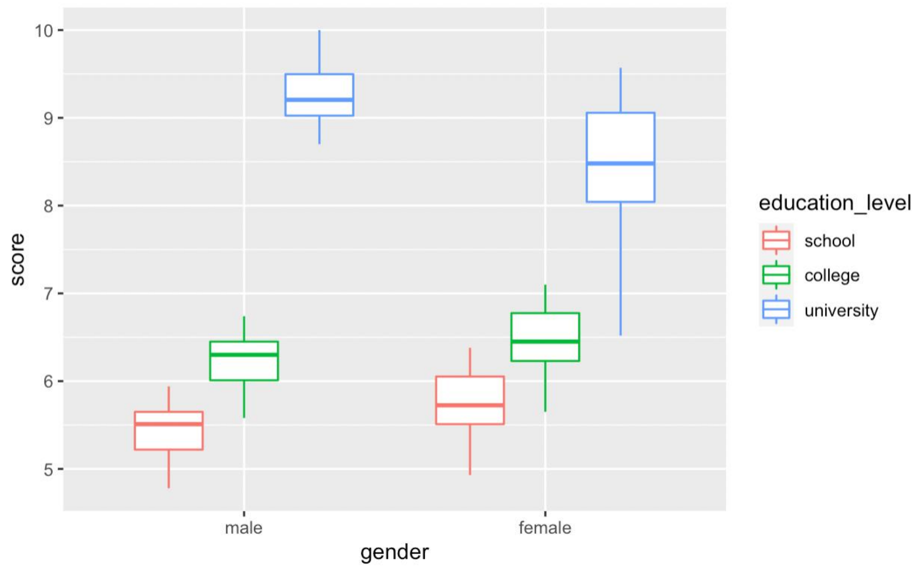

complex models in R

- data = model + error
- we define a job_model that uses a linear model as before, with separate terms for main effects and interactions
- how do we view the results of this model?

```
job_model = lm(data = jobsatisfaction,
               score ~ gender + education_level + gender:education_level)
```

```
car::Anova(job_model)
```

understanding main effects & interactions

- viewing the car::Anova() result is useful to understand broad patterns
- we see a main effect of education level, but it is qualified by the interaction with gender
- what does this mean?

```
car::Anova(job_model)
```

```
> car::Anova(job_model)
Anova Table (Type II tests)

Response: score
                         Sum Sq Df  F value    Pr(>F)
gender                    0.225  1   0.7447  0.392115
education_level         113.684  2 187.8921 < 2.2e-16 ***
gender:education_level    4.440  2   7.3379  0.001559 **
Residuals                15.731 52
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
```

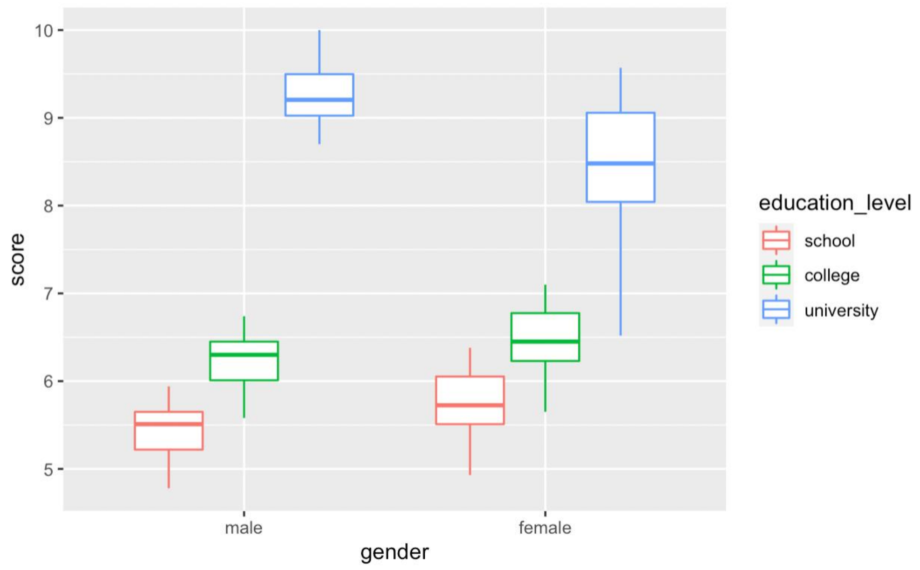

decomposing an interaction

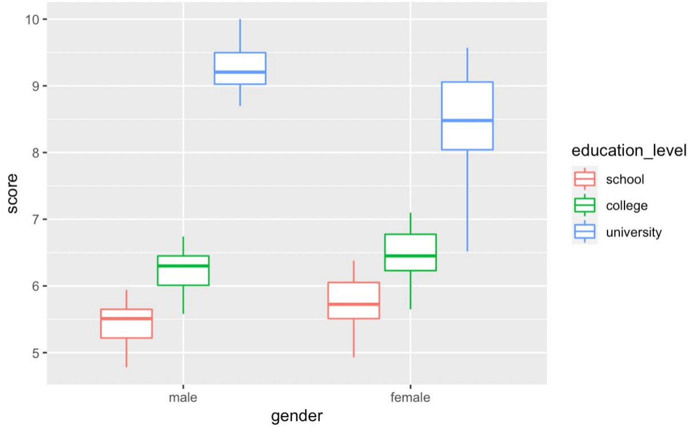

- when there is a significant interaction, we want to go in and understand the nature of this interaction
- interaction (difference of differences)
  - diff(male-female)$_{school}$– diff(male-female)$_{college}$
  - diff(male-female)$_{university}$– diff(male-female)$_{college}$
  - diff(male-female)$_{school}$– diff(male-female)$_{university}$
- where do we think the difference may be?

using emmeans

- we use emmeans as before, except now we specify a conditional effect
- what do the contrasts tell us?

```r
emmeans::emmeans(job_model,
                 pairwise ~ gender | education_level,
                 adjust="tukey")
```

```
$emmeans
education_level = school:
 gender emmean    SE df lower.CL upper.CL
 male     5.43 0.183 52     5.06     5.79
 female   5.74 0.174 52     5.39     6.09

education_level = college:
 gender emmean    SE df lower.CL upper.CL
 male     6.22 0.183 52     5.86     6.59
 female   6.46 0.174 52     6.11     6.81

education_level = university:
 gender emmean    SE df lower.CL upper.CL
 male     9.29 0.174 52     8.94     9.64
 female   8.41 0.174 52     8.06     8.76

Confidence level used: 0.95

$contrasts
education_level = school:
 contrast      estimate    SE df t.ratio p.value
 male - female   -0.314 0.253 52  -1.244  0.2191

education_level = college:
 contrast      estimate    SE df t.ratio p.value
 male - female   -0.240 0.253 52  -0.948  0.3473

education_level = university:
 contrast      estimate    SE df t.ratio p.value
 male - female    0.886 0.246 52   3.602  0.0007
```

ANOVA: assumptions

- "all models are wrong, but some are useful" (Box, 1976)
- the model does not know where the data come from or whether they are appropriate for the model; that is your responsibility as a researcher
  - linearity
  - normality of residuals (i.e., leftover error term)
  - homogeneity of variance
  - independence of observations

inspecting the model

```
install.packages("performance", dependencies = TRUE)
install.packages("see", dependencies = TRUE)
install.packages("patchwork", dependencies = TRUE)
```

```
library(performance)
```

```
check_model(job_model)
```

- first we install the performance, see, and patchwork packages
- load performance
- check the model
- minor variations are ok, major variations are warnings!

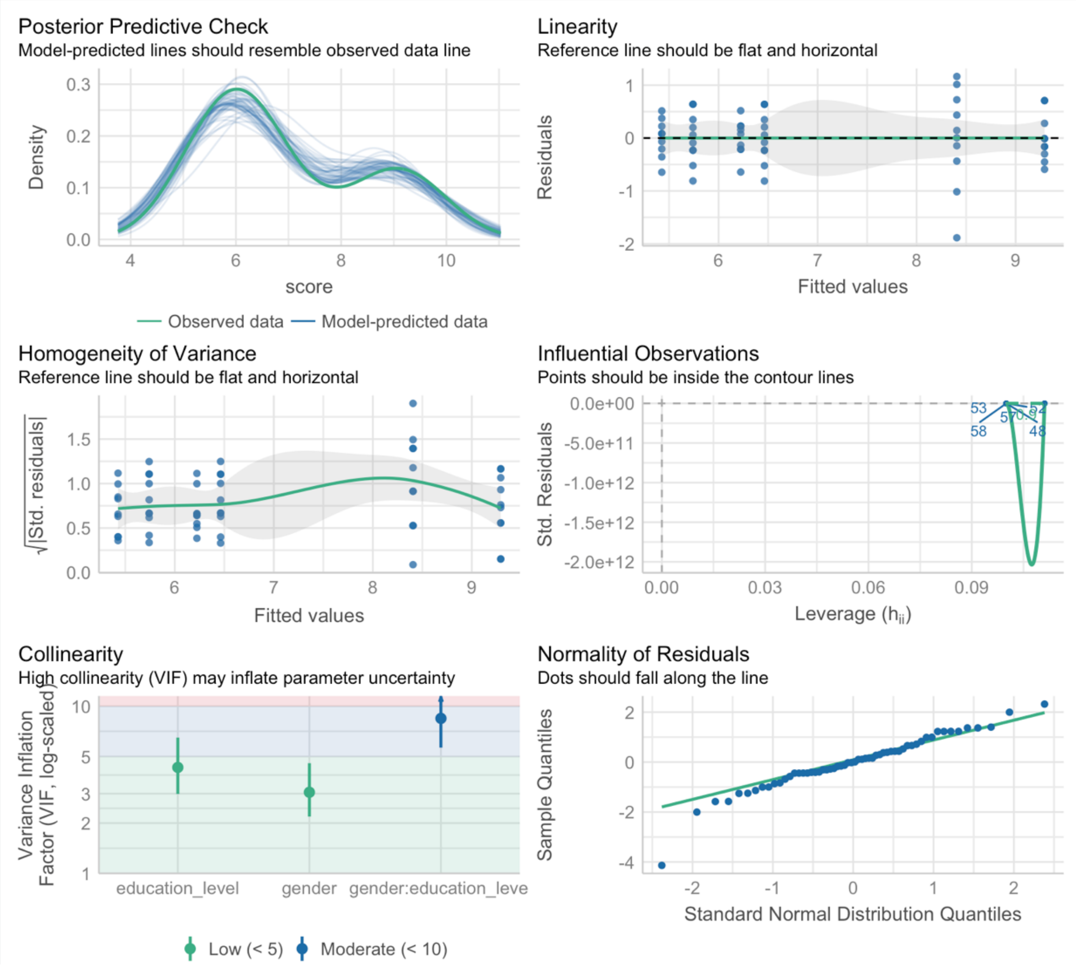

revisiting class data

- run all chunks
- view the priming data
  - 492 rows
- what are the IVs?
- what is the DV?
- double check data types for IV/DV

```r
priming_data = savic %>% filter(typeoftrial == "target") %>%
  select(ID, rt, relatedness, prime, response, type, correct,
         block_number, target, correct_key)%>%
  filter(!is.na(rt), rt > 200 , rt < 1500, correct == "TRUE", block_number == 1) %>%
  filter(relatedness %in% c("related", "unrelated")  & type %in% c("direct", "shared")) %>%
  filter(!ID %in% low_acc_IDs)
```

first-R-notebook.Rmd × | priming_data × | iris ×

Filter

	ID	rt	relatedness	prime	response	type	correct	block_number	target	correct_key
1	958013040	647	related	dodish	a	direct	TRUE	1	horse	A
mipp	958013040	990	unrelated	foobly	l	direct	TRUE	1	horse	L
3	958013040	673	unrelated	geck	a	shared	TRUE	1	apple	A
4	958013040	506	unrelated	foobly	l	direct	TRUE	1	horse	L
5	958013040	549	related	mipp	a	shared	TRUE	1	apple	A
6	958013040	406	unrelated	mipp	a	shared	TRUE	1	horse	A
7	958013040	505	unrelated	mipp	l	shared	TRUE	1	horse	L

```
priming_data          492 obs. of 10 variables
  $ ID          : num [1:492] 9.58e+08 9.58e+08 9.58e+08 9.58e+0
  $ rt          : num [1:492] 647 990 673 506 549 406 505 739 62
  $ relatedness : Factor w/ 3 levels "novel","related",..: 2 3 3
  $ prime       : chr [1:492] "dodish" "foobly" "geck" "foobly"
  $ response    : chr [1:492] "a" "l" "a" "l" ...
  $ type        : Factor w/ 3 levels "direct","novel",..: 1 1 3
  $ correct     : chr [1:492] "TRUE" "TRUE" "TRUE" "TRUE" ...
  $ block_number: chr [1:492] "1" "1" "1" "1" ...
  $ target      : chr [1:492] "horse" "horse" "apple" "horse" ..
  $ correct_key : chr [1:492] "A" "L" "A" "L" ...
```

priming model

- fit a model
- check the assumptions
  - linearity
  - normality of residuals
  - homogeneity of variance
  - independence of observations

```r
rt_lm_model = lm(data = priming_data, rt ~ relatedness + type + relatedness:type)
```

```r
car::Anova(rt_lm_model)
```

```
> car::Anova(rt_lm_model)
Anova Table (Type II tests)

Response: rt
                  Sum Sq  Df F value Pr(>F)
relatedness        35181   1  2.1385 0.1443
type                2098   1  0.1275 0.7212
relatedness:type   15773   1  0.9588 0.3280
Residuals        8028021 488
```

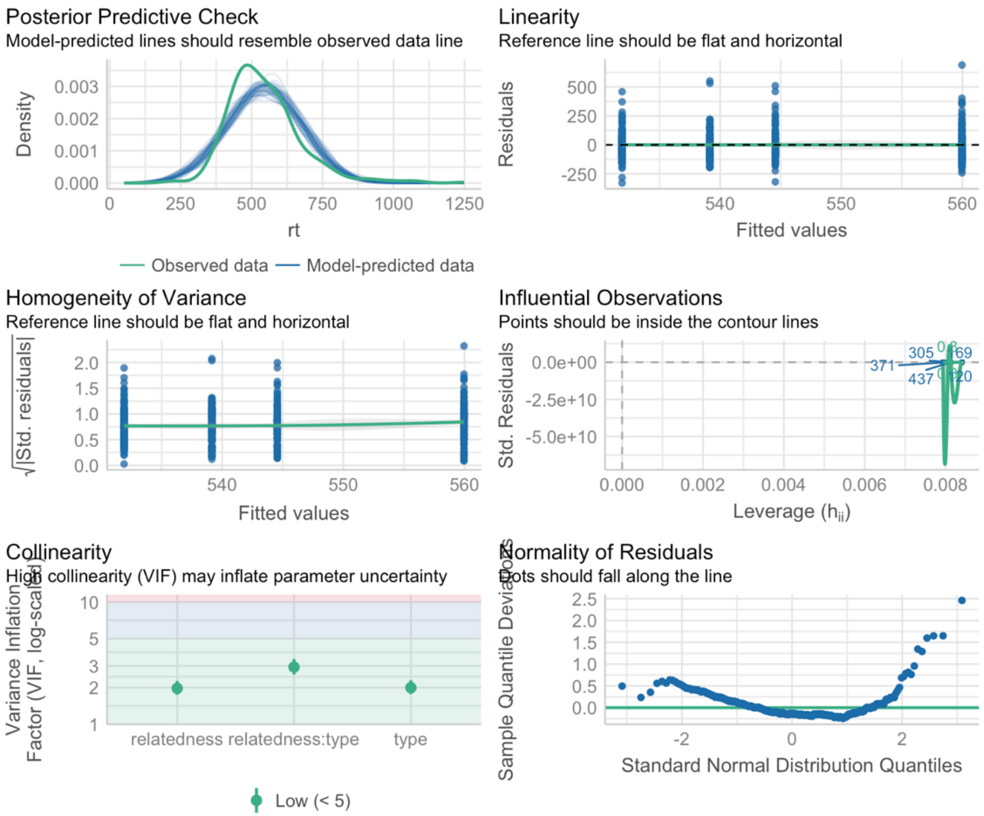

non-independent designs

- whenever multiple observations are collected from participants, especially in within-subject designs, we cannot use a typical ANOVA
- usual solution: repeated measures ANOVA on the means per condition per ID

revisiting class priming data

- RT ~ relatedness x type
- how many trials before exclusion per ID per level of each IV?
- what about after exclusions?
- this is an unbalanced design
- typically we would compute means per level and then run a repeated measures anova
- problem: losing data in aggregation

```
savic %>%
  filter(typeoftrial == "target") %>%
  filter(relatedness %in% c("related", "unrelated")  & type %in% c("direct", "shared")) %>%
  group_by(ID, relatedness, type) %>%
  count()
```

```
# A tibble: 160 × 4
# Groups:   ID, relatedness, type [160]
        ID relatedness type       n
     <dbl> <fct>       <fct>  <int>
1  5418680 related     direct    16
2  5418680 related     shared    16
3  5418680 unrelated   direct    16
4  5418680 unrelated   shared    16
5 46356924 related     direct    16
6 46356924 related     shared    16
```

```
# A tibble: 36 × 4
# Groups:   ID, relatedness, type [36]
        ID relatedness type       n
     <dbl> <fct>       <fct>  <int>
1 46356924 related     direct    16
2 46356924 related     shared    16
3 46356924 unrelated   direct    16
4 46356924 unrelated   shared    16
5 52271504 related     direct    14
6 52271504 related     shared    14
```

```
priming_data %>%
  group_by(ID, relatedness, type) %>%
  count()
```

ANOVA: limitations

- limited to continuous DVs
  - examples of non-continuous/categorical DVs?
  - changes the distribution of your variables, violates the normality assumption
  - common distributions: binomal (yes/no, correct/incorrect), multinomial (know,don't know, TOT, other), poisson (counting number of website visitors)
- limited to categorical IVs
  - examples of continuous IVs?
- cannot deal with missing data
- cannot handle nested/clustered design
  - male/females in sectors in cities
  - trials in subjects in conditions
- cannot handle unbalanced design
  - different number of trials (after exclusion) for each subject?

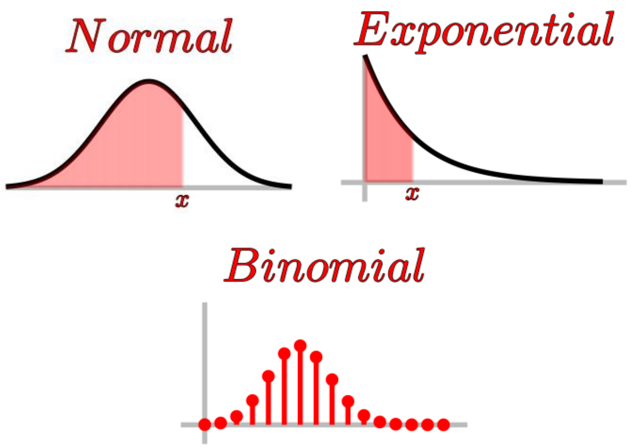

a flexible model

- linear/generalized mixed effects models!
- these models consider the variability due to:
  - missing data
  - categorical/continuous IVs and DVs
  - unbalanced designs
  - clustered designs (no collapsing into means)
- think of them as the parent models from which special cases such as t-tests and ANOVAs are derived
- different 'lines/curves' are fit for each individual and for each item, with their own slope and intercept, instead of "averaging" across everyone

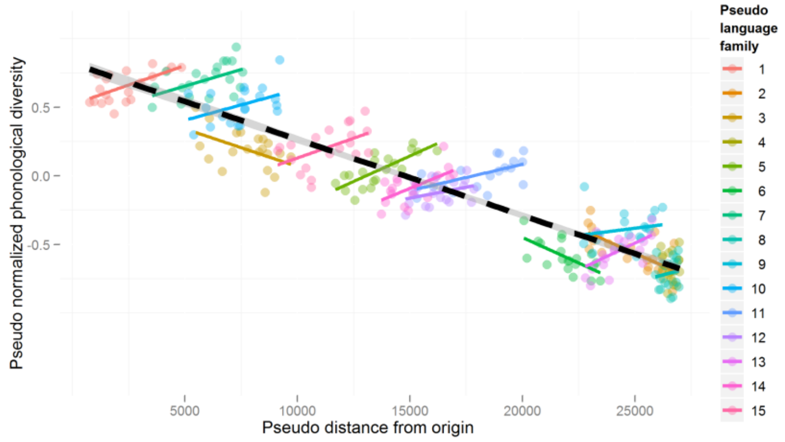

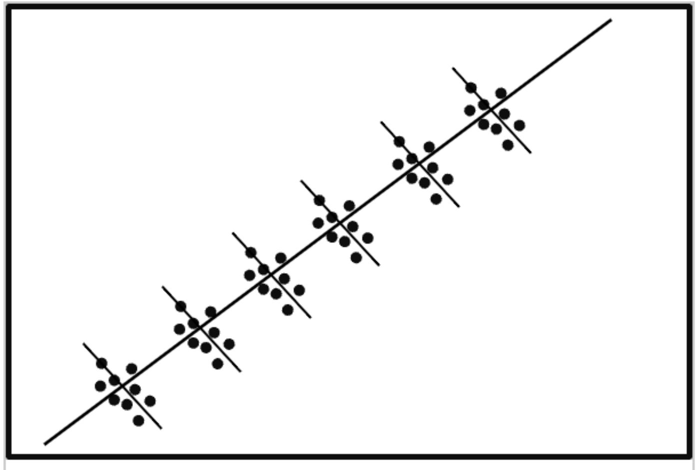

mixed linear model in R

- install lmerTest
- lmer() function
  - fixed part: your IVs
  - random part: your within-subject variables

```
library(lmerTest)
rt_model = lmer(data = priming_data,
                rt ~ relatedness*type + (1|ID))

summary(rt_model)
```

main effects and interactions

- car::Anova()

```
car::Anova(rt_model)
```

```
> car::Anova(rt_model)
Analysis of Deviance Table (Type II Wald chisquare tests)

Response: rt
                   Chisq Df Pr(>Chisq)
relatedness       1.8072  1     0.1788
type              1.8490  1     0.1739
relatedness:type  1.2565  1     0.2623
```

assumptions check

- same as before

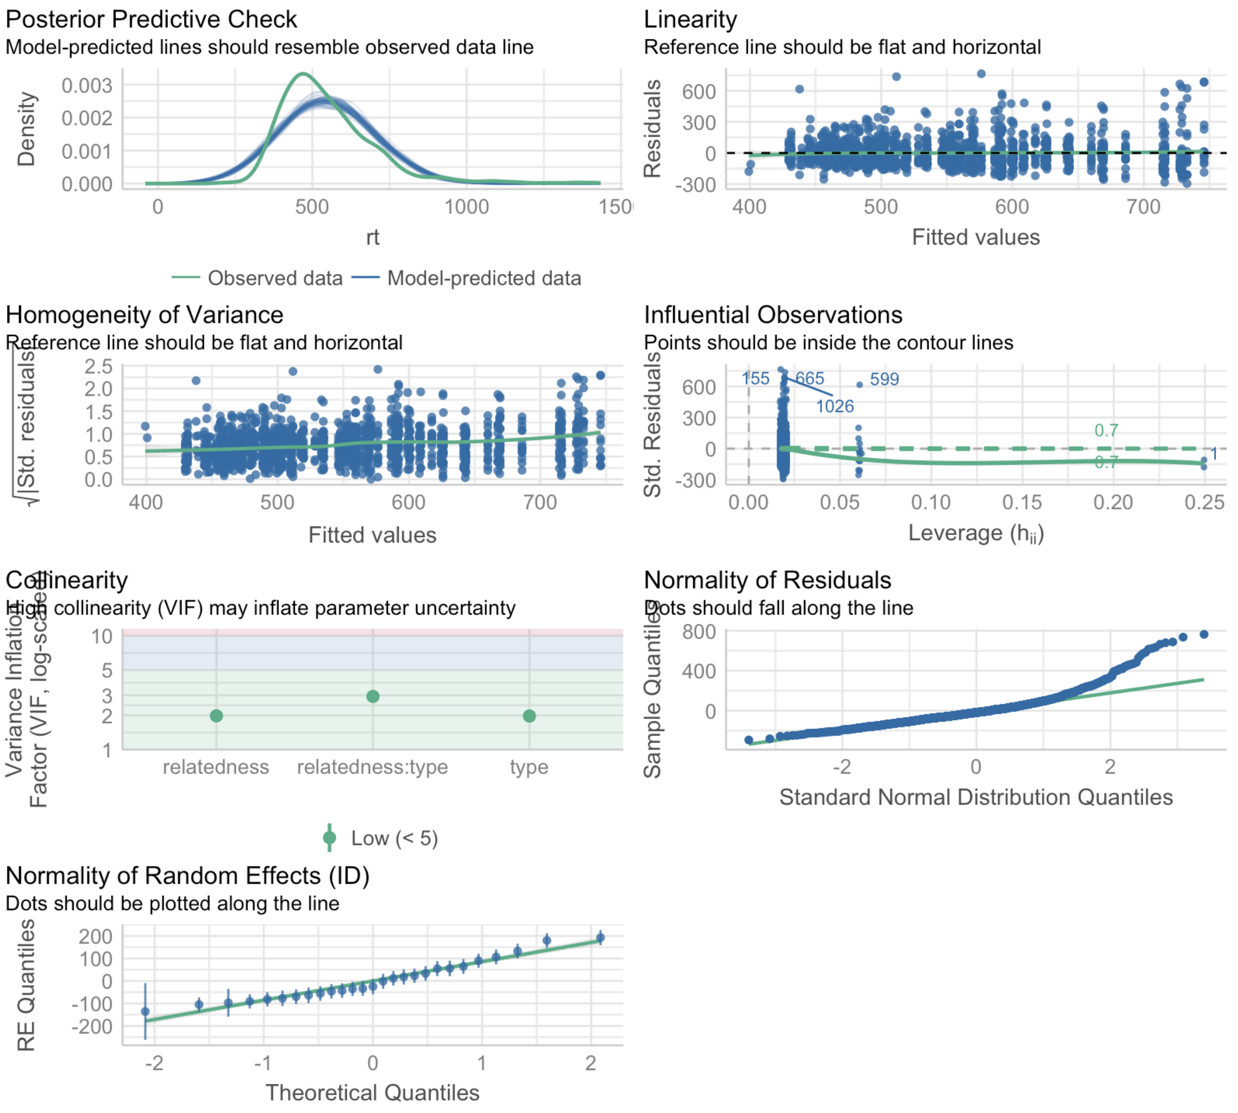

pre-registration + project checklist

- milestone #5: pre-registration + project checklist + piloting (Nov 10)

1. **Data Collection**: Have any data been collected for this study already?
2. **Main Question**: What is the main question being asked or hypothesis being tested in this study?
3. **Dependent Variable(s)**: Describe the key dependent variable(s) specifying how they will be measured.
4. **Condition(s)**: How many and which conditions will participants be assigned to? Please include an example trial of each type of condition you have in your experiment. Please also specify which independent variable will be within-participants or between-participants.
5. **Analyses**: Specify exactly which analyses you will conduct to examine the main question/hypothesis.
6. **Outliers & Exclusions**: Describe exactly how outliers will be defined and handled, and your precise rule(s) for excluding observations.
7. **Predicted Plot**: Please submit a predicted plot for your study based on what you expect the pattern to look like for your main hypothesis.
8. **Sample Size**: How many observations will be collected or what will determine sample size? No need to justify the decision, but be precise about exactly how the number will be determined.
9. **Exploratory details**: Anything else you would like to pre-register? (e.g., secondary analyses, variables collected for exploratory purposes, unusual analyses planned?)

next time

- **before** class
  - *work on:* project checklist and pre-registration
- **during** class
  - analysis review
  - sona + prolific data collection