Introduction to mixed-effects models

(for longitudinal data)

Nora Wickelmaier

Leibniz-Institut für Wissensmedien

2026-04-28

Plan for today

- We will walk through an example for a longitudinal data set (6 time points per subject)

Plan for today

- We will walk through an example for a longitudinal data set (6 time points per subject)
- I will explain the general concepts with the slides

Plan for today

- We will walk through an example for a longitudinal data set (6 time points per subject)
- I will explain the general concepts with the slides
- We will switch to R and use the lme4 package to fit the models

Plan for today

- We will walk through an example for a longitudinal data set (6 time points per subject)
- I will explain the general concepts with the slides
- We will switch to R and use the lme4 package to fit the models
- You will use R to fit the model and try out different specifications

Plan for today

- We will walk through an example for a longitudinal data set (6 time points per subject)
- I will explain the general concepts with the slides
- We will switch to R and use the lme4 package to fit the models
- You will use R to fit the model and try out different specifications
- We will discuss the results

Plan for today

- We will walk through an example for a longitudinal data set (6 time points per subject)
- I will explain the general concepts with the slides
- We will switch to R and use the lme4 package to fit the models
- You will use R to fit the model and try out different specifications
- We will discuss the results
- All the materials are here: https://gitea.iwm-tuebingen.de/nwickelmaier/lead_longitudinal

Plan for today

- We will walk through an example for a longitudinal data set (6 time points per subject)
- I will explain the general concepts with the slides
- We will switch to R and use the lme4 package to fit the models
- You will use R to fit the model and try out different specifications
- We will discuss the results
- All the materials are here: https://gitea.iwm-tuebingen.de/nwickelmaier/lead_longitudinal

→ Try to code along in R! Ask as many questions as possible

Outline

1. Introduction to longitudinal data
2. Example: Depression and Imipramin

1 Introduction to longitudinal data

Longitudinal data

- Consist of repeated measurements on the same subject taken over time
- Are a frequent use case for mixed-effects models
- Contain time as a predictor: time trends within and between subjects are of interest

```
library("lme4")
data("sleepstudy")
?sleepstudy
str(sleepstudy)
summary(sleepstudy)
head(sleepstudy)
```

Sleep study

- Average reaction time per day for subjects in a sleep deprivation study
- On day 0, the subjects had their normal amount of sleep
- Starting that night they were restricted to 3 hours of sleep per night
- Observations represent the average reaction time on a series of tests given each day to each subject

A data frame with 180 observations on the following 3 variables

`Reaction`	Average reaction time (ms)
`Days`	Number of days of sleep deprivation
`Subject`	Subject number on which the observation was made

Visualization of data

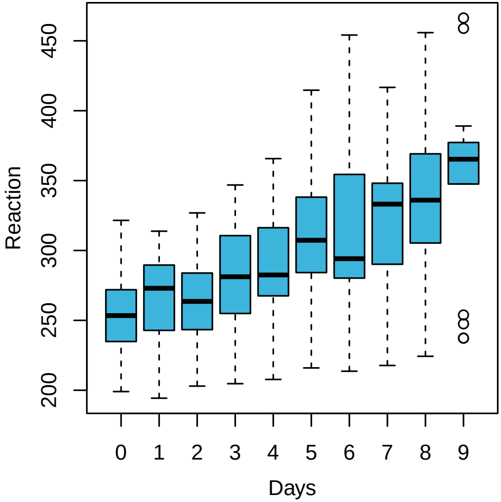


```
boxplot(Reaction ~ Days, sleepstudy)
```

Visualization of individual data

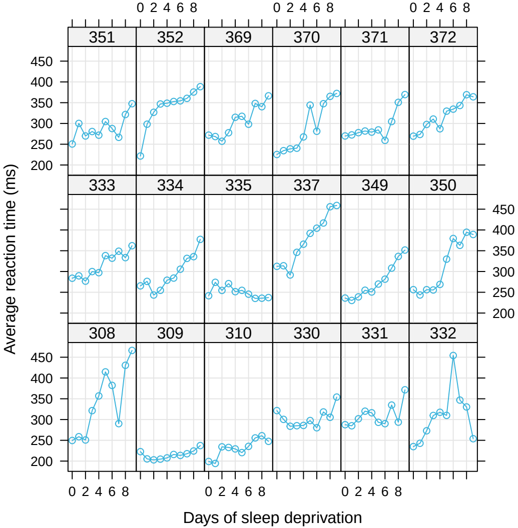


```
library("lattice")

xyplot(Reaction ~ Days | Subject,
  data = sleepstudy,
  type = c("g", "b"),
  xlab = "Days of sleep deprivation",
  ylab = "Average reaction time (ms)",
  aspect = "xy")
```

Random intercept model

- The random intercept model adds a random intercept for each subject

$$y_{ij} = \beta_0 + \beta_1\, Days_{ij} + \upsilon_{0i} + \varepsilon_{ij}$$

with $\upsilon_{0i} \overset{iid}{\sim} N(0, \sigma_\upsilon^2)$, $\varepsilon_{ij} \overset{iid}{\sim} N(0, \sigma_\varepsilon^2)$

- The slope is identical for each subject (and the population)

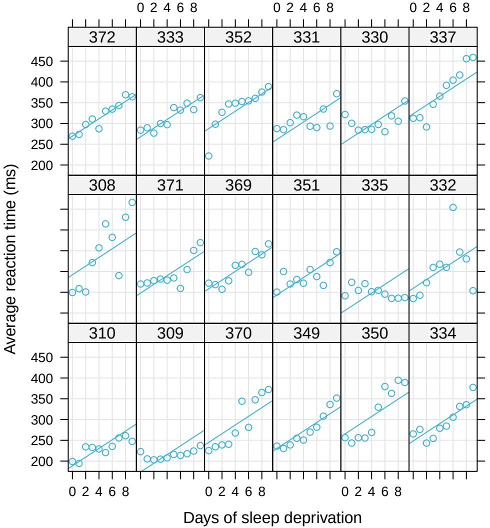

Fixed and random effects for random intercept model

Subject	Days	Reaction	Fixed		Random	
			Intercept	Slope	SubInt	Resid
308	0	250	251.4	10.5	40.8	−42.6
308	1	259	251.4	10.5	40.8	−44.0
308	2	251	251.4	10.5	40.8	−62.3
⋮	⋮	⋮	⋮	⋮	⋮	⋮
308	8	431	251.4	10.5	40.8	54.7
308	9	466	251.4	10.5	40.8	80.0
309	0	223	251.4	10.5	−77.8	49.2
309	1	205	251.4	10.5	−77.8	21.2
309	2	203	251.4	10.5	−77.8	8.5
⋮	⋮	⋮	⋮	⋮	⋮	⋮
310	0	199	251.4	10.5	−63.1	10.8
310	1	194	251.4	10.5	−63.1	−4.4
310	2	234	251.4	10.5	−63.1	25.1
⋮	⋮	⋮	⋮	⋮	⋮	⋮
					σ_υ^2	σ_ε^2

Random slope model

- The random slope model adds a random intercept and a random slope for each subject

$$y_{ij} = \beta_0 + \beta_1\, Days_{ij} + \upsilon_{0i} + \upsilon_{1i}\, Days_{ij} + \varepsilon_{ij}$$

with

$$\begin{pmatrix} \upsilon_{0i} \\ \upsilon_{1i} \end{pmatrix} \overset{iid}{\sim} N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \boldsymbol{\Sigma}_\upsilon = \begin{pmatrix} \sigma^2_{\upsilon_0} & \sigma_{\upsilon_0 \upsilon_1} \\ \sigma_{\upsilon_0 \upsilon_1} & \sigma^2_{\upsilon_1} \end{pmatrix} \right)$$

$$\varepsilon_{ij} \overset{iid}{\sim} N(0, \sigma^2_\varepsilon)$$

- Individual slopes for each subject

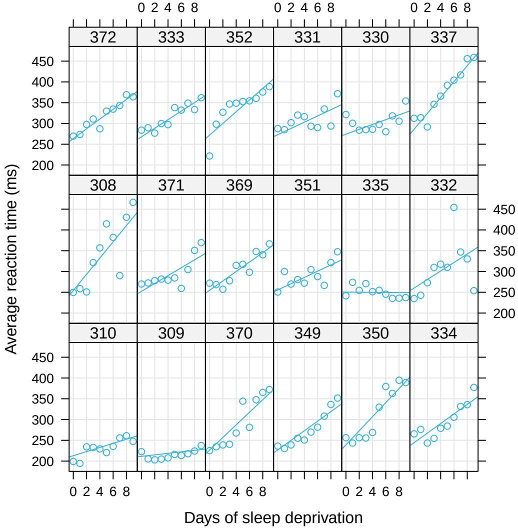

Fixed and random effects for random slope model

Subject	Days	Reaction	Fixed		Random		
			Intercept	Slope	SubInt	SubSlp	Resid
308	0	250	251.4	10.5	2.3	9.2	−4.1
308	1	259	251.4	10.5	2.3	9.2	−14.6
308	2	251	251.4	10.5	2.3	9.2	−42.2
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
308	8	431	251.4	10.5	2.3	9.2	19.6
308	9	466	251.4	10.5	2.3	9.2	35.7
309	0	223	251.4	10.5	−40.4	−8.6	11.7
309	1	205	251.4	10.5	−40.4	−8.6	−7.6
309	2	203	251.4	10.5	−40.4	−8.6	−11.7
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
310	0	199	251.4	10.5	−39.0	−5.4	−13.4
310	1	194	251.4	10.5	−39.0	−5.4	−23.1
310	2	234	251.4	10.5	−39.0	−5.4	11.8
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
					$\sigma^2_{\upsilon_0}$ $\sigma_{\upsilon_0\upsilon_1}$	$\sigma^2_{\upsilon_1}$	σ^2_{ε}

Partial pooling

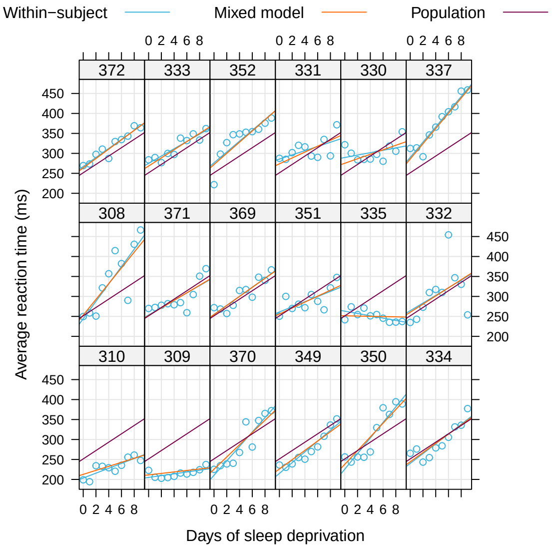

- Within-subject regression line shows regression line fitted to data for each individual
- Population regression line shows fixed effects for mixed-effects model
- Mixed model regression line shows individual regression lines as predicted by mixed-effects models

2 Example: Depression and Imipramin

Quadratic time trends

- A lot of times the assumption of a linear time trend is too simple
- Change is not happening unbraked linearly but flattens out

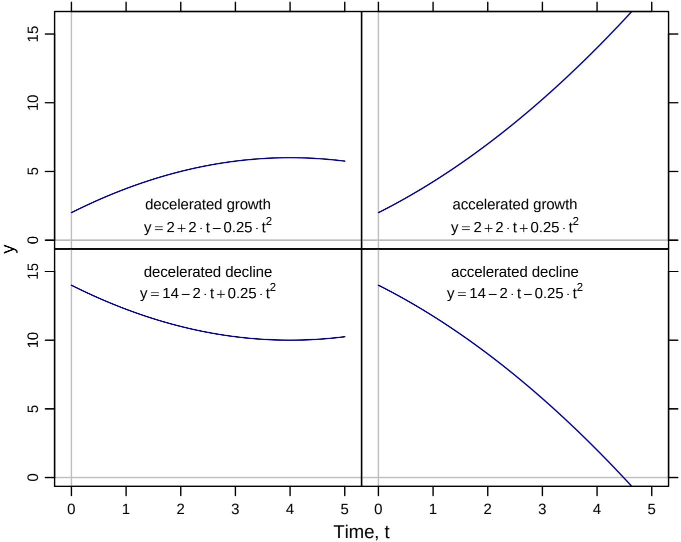

Quadratic time trends

- Quadratic regression model

$$
\begin{aligned}
y_{ij} &= b_{0i} + b_{1i}\,t_{ij} + b_{2i}\,t_{ij}^2 + \varepsilon_{ij} \\
&= b_{0i} + (b_{1i} + b_{2i}\,t_{ij})t_{ij} + \varepsilon_{ij}
\end{aligned}
$$

- The linear change depends on time t

$$
\frac{\partial y}{\partial t} = b_{1i} + 2b_{2i}\,t
$$

- The intercept $t = -b_{1i}/(2b_{2i})$ is the point in time when a positive (negative) trend becomes negative (postive)

Depression and Imipramin (Reisby et al., 1977)

- Reisby et al. (1977) studied the effect of Imipramin on 66 inpatients treated for depression
- Depression was measured with the Hamilton depression rating scale (HDRS)
- Additionally, the concentration of Imipramin and its metabolite Desipramin was measured in their blood plasma
- Patients were classified into endogenous and non-endogenous depressed
- Depression was measured weekly for 6 time points; the effect of the antidepressant was observed starting at week 2 for four weeks

Descriptive statistics

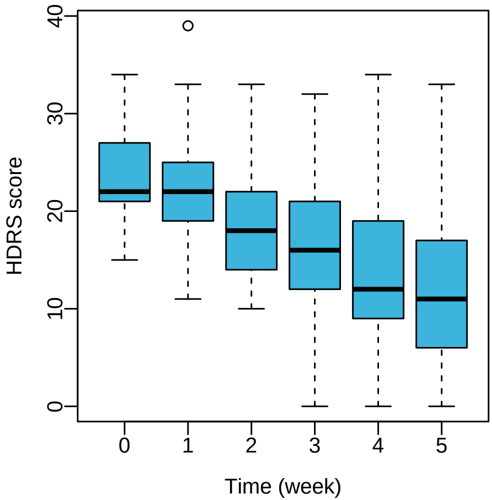

HDRS score

t	W0	W1	W2	W3	W4	W5
M	23.44	21.84	18.31	16.42	13.62	11.95
SD	4.53	4.70	5.49	6.42	6.97	7.22
n	61	63	65	65	63	58

Empirical correlation matrix of HDRS score

	W0	W1	W2	W3	W4	W5
Week 0	1	.49	.41	.33	.23	.18
Week 1	.49	1	.49	.41	.31	.22
Week 2	.41	.49	1	.74	.67	.46
Week 3	.33	.41	.74	1	.82	.57
Week 4	.23	.31	.67	.82	1	.65
Week 5	.18	.22	.46	.57	.65	1

Predictions random slope model

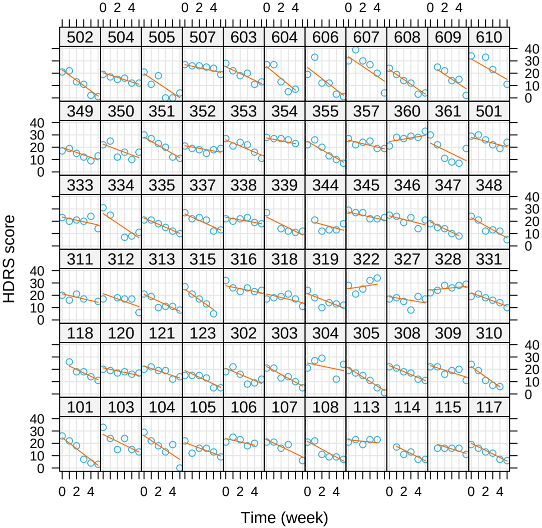

$$y_{ij} = \beta_0 + \beta_1 t_{ij} + \upsilon_{0i} + \upsilon_{1i} t_{ij} + \varepsilon_{ij}$$

with

$$\begin{pmatrix} \upsilon_{0i} \\ \upsilon_{1i} \end{pmatrix} \overset{iid}{\sim} N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \mathbf{\Sigma}_\upsilon = \begin{pmatrix} \sigma^2_{\upsilon_0} & \sigma_{\upsilon_0 \upsilon_1} \\ \sigma_{\upsilon_0 \upsilon_1} & \sigma^2_{\upsilon_1} \end{pmatrix} \right)$$

$$\varepsilon_{ij} \overset{iid}{\sim} N(0, \sigma^2_\varepsilon)\mathbf{I}_{n_i})$$

Exercise

- What would be a possible extension of this model?

Exercise

- What would be a possible extension of this model?
- Write down the model equations
    - What changes for the fixed effects?
    - How do the variance components for the random effects change?

Exercise

- What would be a possible extension of this model?
- Write down the model equations
  - What changes for the fixed effects?
  - How do the variance components for the random effects change?
- How can we interpret the random quadratic effects for this model?

Exercise

- What would be a possible extension of this model?
- Write down the model equations
    - What changes for the fixed effects?
    - How do the variance components for the random effects change?
- How can we interpret the random quadratic effects for this model?
- How do we add (random) quadratic effects to a random slope model using `lme4::lmer()`?

Exercise

- What would be a possible extension of this model?
- Write down the model equations
  - What changes for the fixed effects?
  - How do the variance components for the random effects change?
- How can we interpret the random quadratic effects for this model?
- How do we add (random) quadratic effects to a random slope model using `lme4::lmer()`?
- Fit this growth curve model in R

Exercise

- What would be a possible extension of this model?
- Write down the model equations
  - What changes for the fixed effects?
  - How do the variance components for the random effects change?
- How can we interpret the random quadratic effects for this model?
- How do we add (random) quadratic effects to a random slope model using `lme4::lmer()`?
- Fit this growth curve model in R

```
dat      <- read.table("data/reisby.txt", header = TRUE)
dat$id   <- factor(dat$id)
dat$diag <- factor(dat$diag, levels = c("nonen", "endog"))
dat      <- na.omit(dat)     # drop missing values
```

Model with quadratic trend

- Model with quadratic individual and quadratic group trend

$$y_{ij} = \beta_0 + \beta_1\, t_{ij} + \beta_2\, t_{ij}^2 + \upsilon_{0i} + \upsilon_1\, t_{ij} + \upsilon_2\, t_{ij}^2 + \varepsilon_{ij}$$

with

$$\begin{pmatrix} \upsilon_{0i} \\ \upsilon_{1i} \\ \upsilon_{2i} \end{pmatrix} \overset{iid}{\sim} N\left(\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma^2_{\upsilon_0} & \sigma_{\upsilon_0\upsilon_1} & \sigma_{\upsilon_0\upsilon_2} \\ \sigma_{\upsilon_0\upsilon_1} & \sigma^2_{\upsilon_1} & \sigma_{\upsilon_1\upsilon_2} \\ \sigma_{\upsilon_0\upsilon_2} & \sigma_{\upsilon_1\upsilon_2} & \sigma^2_{\upsilon_2} \end{pmatrix} \right)$$

$$\varepsilon_{ij} \overset{iid}{\sim} N(0, \sigma^2_\varepsilon)$$

Model predictions

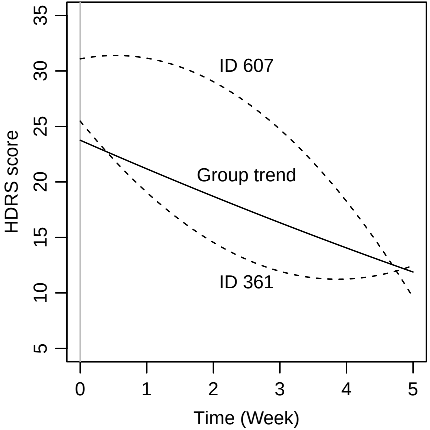

- Averaged over persons an approximately linear trend is obtained, $\hat{\beta}_1 = -2.63$, $\hat{\beta}_2 = 0.05$
- Some of the predicted individual trends are strongly nonlinear

- Test against a model without individual quadratic trends

$$H_0\colon\ \sigma^2_{\upsilon_2} = \sigma_{\upsilon_0\upsilon_2} = \sigma_{\upsilon_1\upsilon_2} = 0 \qquad G^2(3) = 10.98,\ p = .012$$

Model predictions

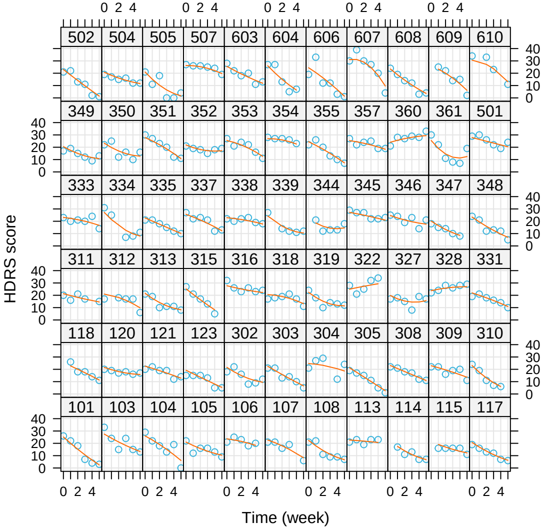


```r
xyplot(
  hamd + predict(lme3)
      ~ week | id,
  data = dat,
  type = c("p", "l", "g"),
  distribute.type = TRUE,
  ylab = "HDRS score",
  xlab = "Time (Week)")
```

Centering variables

- If multiples of the time variables (t, t^2, t^3, etc.) are entered into the regression equation, multicollinearity can become a problem
- For example, $t = 0, 1, 2, 3$ and $t^2 = 0, 1, 4, 9$ correlate almost perfectly
- By centering the variables, this problem can be diminished: $(t - \bar{t}) = -1.5, -0.5, 0.5, 1.5$ and $(t - \bar{t})^2 = 2.25, 0.25, 0.25, 2.25$ are uncorrelated
- By centering variables the interpretation of the intercept in a linear model changes:
    - Uncentered intercepts represent the difference to the first time point ($t = 0$)
    - Centered intercepts represent the difference after half of the time

Exercise

- Center time

Exercise

- Center time
- Compare the correlation between time and time^2 with and without centering

Exercise

- Center time
- Compare the correlation between time and time^2 with and without centering
- Refit the model with quadratic effects when time is centered

Exercise

- Center time
- Compare the correlation between time and time^2 with and without centering
- Refit the model with quadratic effects when time is centered
- Which parameters change?
  - What changes for the fixed effects?
  - How do the variance components for the random effects change?

Exercise

- Center time
- Compare the correlation between time and time^2 with and without centering
- Refit the model with quadratic effects when time is centered
- Which parameters change?
    - What changes for the fixed effects?
    - How do the variance components for the random effects change?
- *Why* do the variance components for the random effects change?

Investigating random effects structure

- In order to get a better understanding of the necessary random effects it might be a good idea to take a closer look at them
- Two plots often used are the so-called caterpillar and shrinkage plots
- Play around with different models and compare how, e. g., the caterpillar plots change with and without covariances in the model!

```
# model with quadratic time trend
m <- lmer(hamd ~ week + I(week^2) + (week + I(week^2) | id),
          data = dat, REML = FALSE)

library("lattice")
dotplot(ranef(m), scales = list( x = list(relation = "free")))$id
```

Caterpillar plot

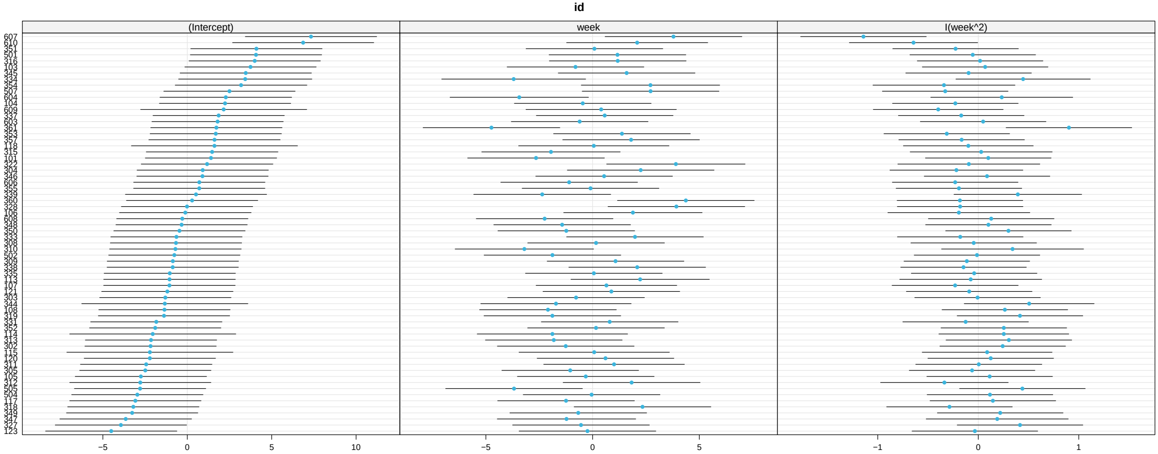

Shrinkage plots

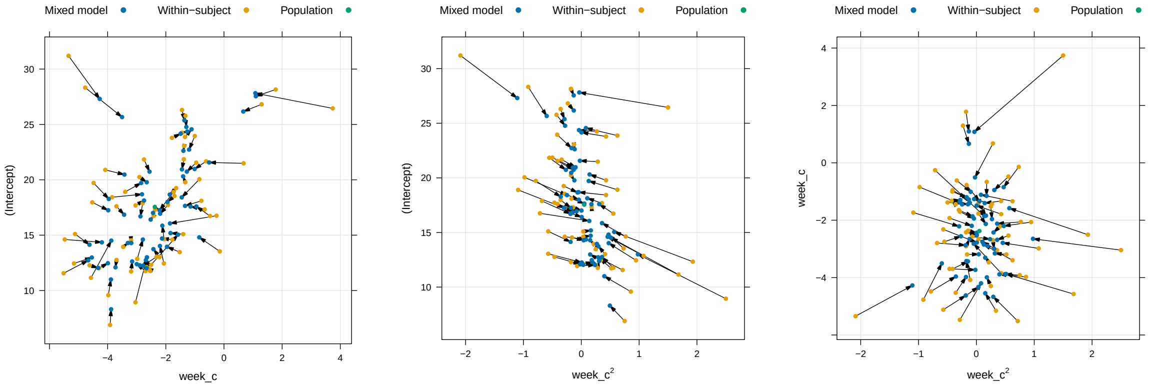

What we learned today...

...and how to go on

What we learned today...

...and how to go on

1. We learned

What we learned today...

...and how to go on

1. We learned
   - The basic concept of random effects and why to include them in a model

What we learned today...

...and how to go on

1. We learned
    - The basic concept of random effects and why to include them in a model
    - How to model data collected over several time points

What we learned today...

...and how to go on

1. We learned
    - The basic concept of random effects and why to include them in a model
    - How to model data collected over several time points
    - How to compute mixed-effects model with quadratic time trends in R using `lmer()` from the lme4 package

What we learned today...

...and how to go on

1. We learned
   - The basic concept of random effects and why to include them in a model
   - How to model data collected over several time points
   - How to compute mixed-effects model with quadratic time trends in R using `lmer()` from the lme4 package
   - How to interpret parameters in mixed-effects model with quadratic effects

What we learned today...

...and how to go on

1. We learned
   - The basic concept of random effects and why to include them in a model
   - How to model data collected over several time points
   - How to compute mixed-effects model with quadratic time trends in R using `lmer()` from the lme4 package
   - How to interpret parameters in mixed-effects model with quadratic effects
2. Next steps

What we learned today...

...and how to go on

1. We learned
   - The basic concept of random effects and why to include them in a model
   - How to model data collected over several time points
   - How to compute mixed-effects model with quadratic time trends in R using `lmer()` from the lme4 package
   - How to interpret parameters in mixed-effects model with quadratic effects
2. Next steps
   - Do this exercise https://gitea.iwm-tuebingen.de/nwickelmaier/lead_longitudinal/src/branch/master/exercises/schizo.md

What we learned today...

...and how to go on

1. We learned
   - The basic concept of random effects and why to include them in a model
   - How to model data collected over several time points
   - How to compute mixed-effects model with quadratic time trends in R using `lmer()` from the lme4 package
   - How to interpret parameters in mixed-effects model with quadratic effects
2. Next steps
   - Do this exercise https://gitea.iwm-tuebingen.de/nwickelmaier/lead_longitudinal/src/branch/master/exercises/schizo.md
   - It has a very similar structure than the depression dataset and this will help you to generalize the concepts we learned today

What we learned today. . .

. . . and how to go on

1. We learned
   - The basic concept of random effects and why to include them in a model
   - How to model data collected over several time points
   - How to compute mixed-effects model with quadratic time trends in R using `lmer()` from the lme4 package
   - How to interpret parameters in mixed-effects model with quadratic effects
2. Next steps
   - Do this exercise https://gitea.iwm-tuebingen.de/nwickelmaier/lead_longitudinal/src/branch/master/exercises/schizo.md
   - It has a very similar structure than the depression dataset and this will help you to generalize the concepts we learned today
   - You can send questions to me and even make an appointment with me to go over your solution

References

Bates, D. (2010). lme4: Mixed-effects modeling with R (book draft). https://lme4.r-forge.r-project.org/

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, *67*(1), 1–48. https://doi.org/10.18637/jss.v067.i01

Reisby, N., Gram, L. F., Bech, P., Nagy, A., Petersen, G. O., Ortmann, J., Ibsen, I., Dencker, S. J., Jacobsen, O., Krautwald, O., Sondergaard, I., & Christiansen, J. (1977). Imipramine: Clinical effects and pharmacokinetic variability. *Psychopharmacology*, *54*, 263–272. https://doi.org/10.1007/BF00426574