Linear mixed effects modeling in neuroimaging

Turku PET Centre neuroimaging course

Tuulia Malen,
PhD student, clinical psychologist

Image

PET image with tracer [11C]raclopride that binds to dopamine receptors

Image

PET image with tracer [11C]raclopride that binds to dopamine receptors

Image processing: from raw data to values

Binding potential: estimate of dopamine receptor availability in a brain region (e.g. putamen)

Image

PET image with tracer [11C]raclopride that binds to dopamine receptors

Image processing: from raw data to values

Binding potential: Estimate of dopamine receptor availability in a brain region (e.g. putamen)

Research question

What is the relationship between age and dopamine receptors?

Positive: increase with age

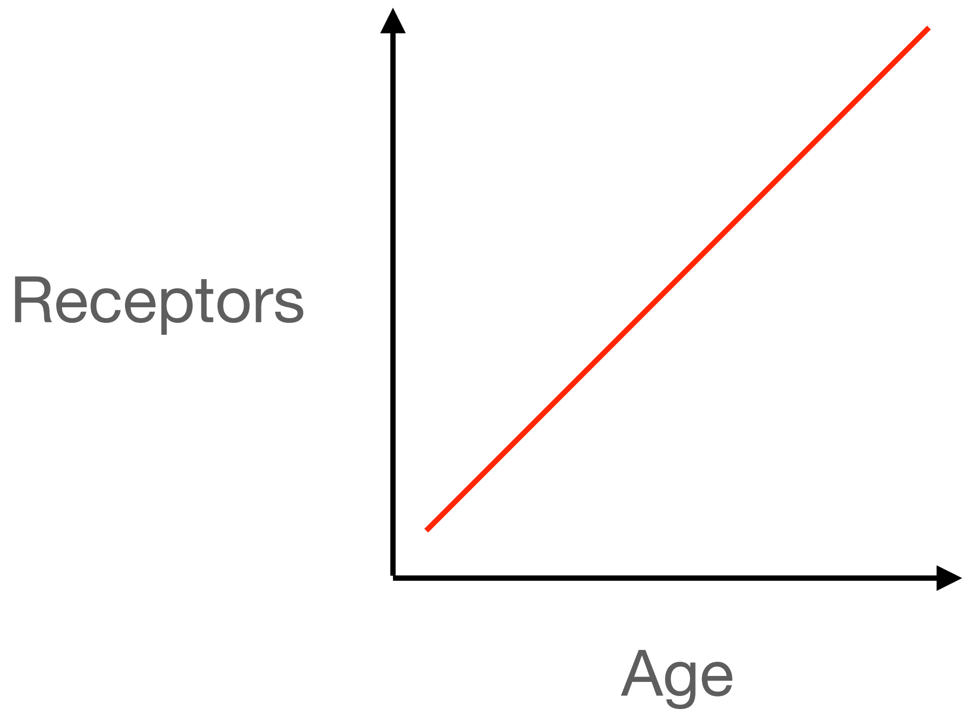

Negative: decrease with age?

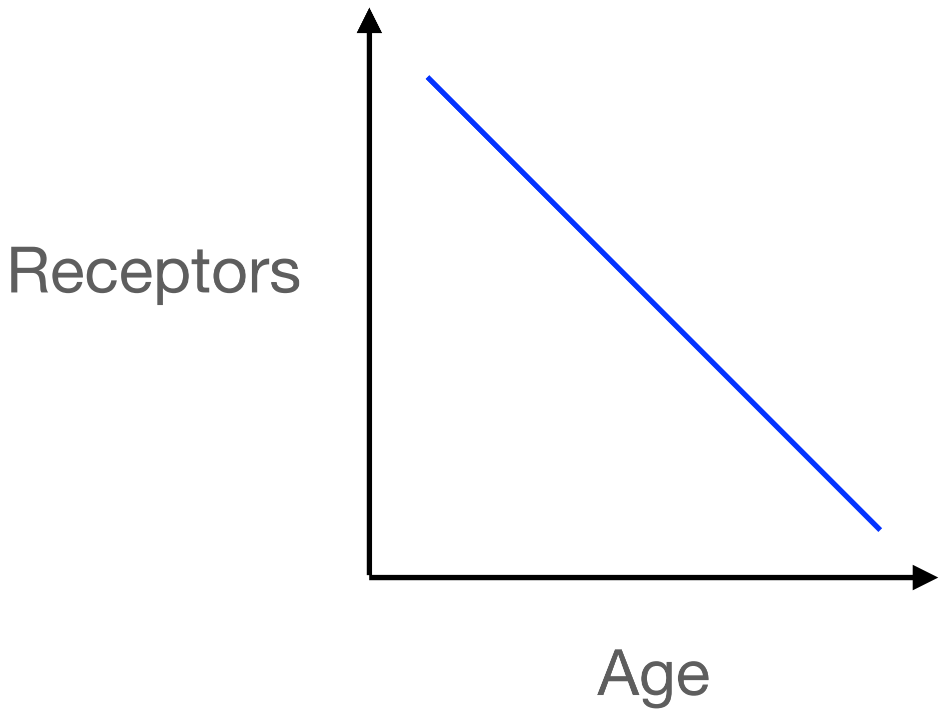

Neither: does not change with age

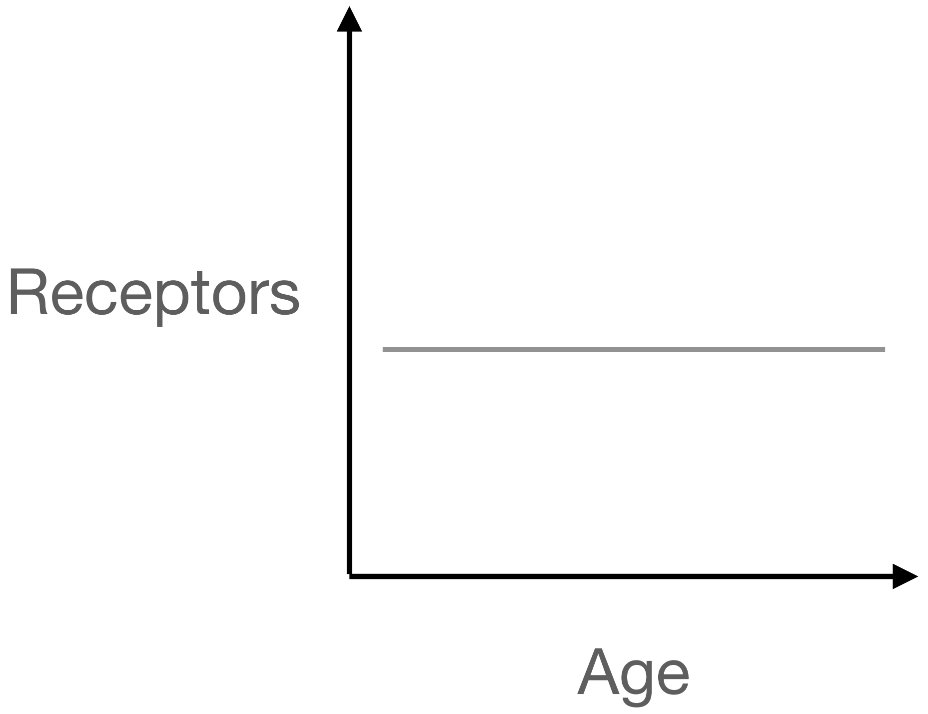

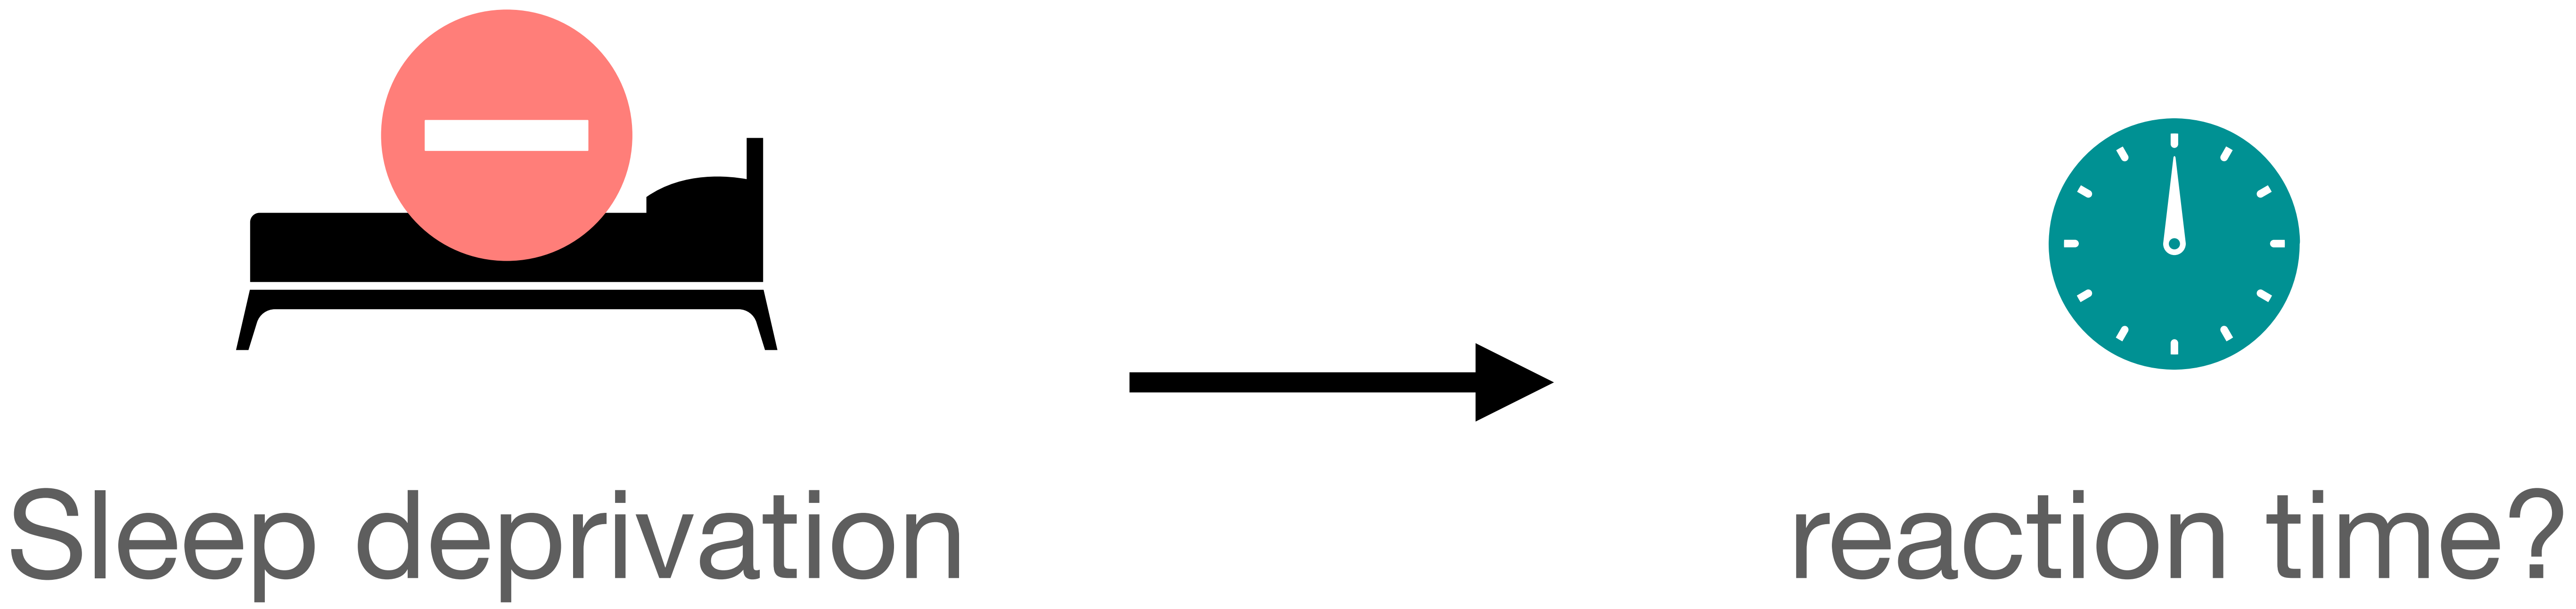

Sleepstudy dataset from lme4 package in R
Days= Number of days of sleep deprivation (3 hours/ night)
Reaction= Average reaction time (ms)

What is the effect of sleep deprivation on reaction time?

$$y = \alpha + \beta \cdot day$$

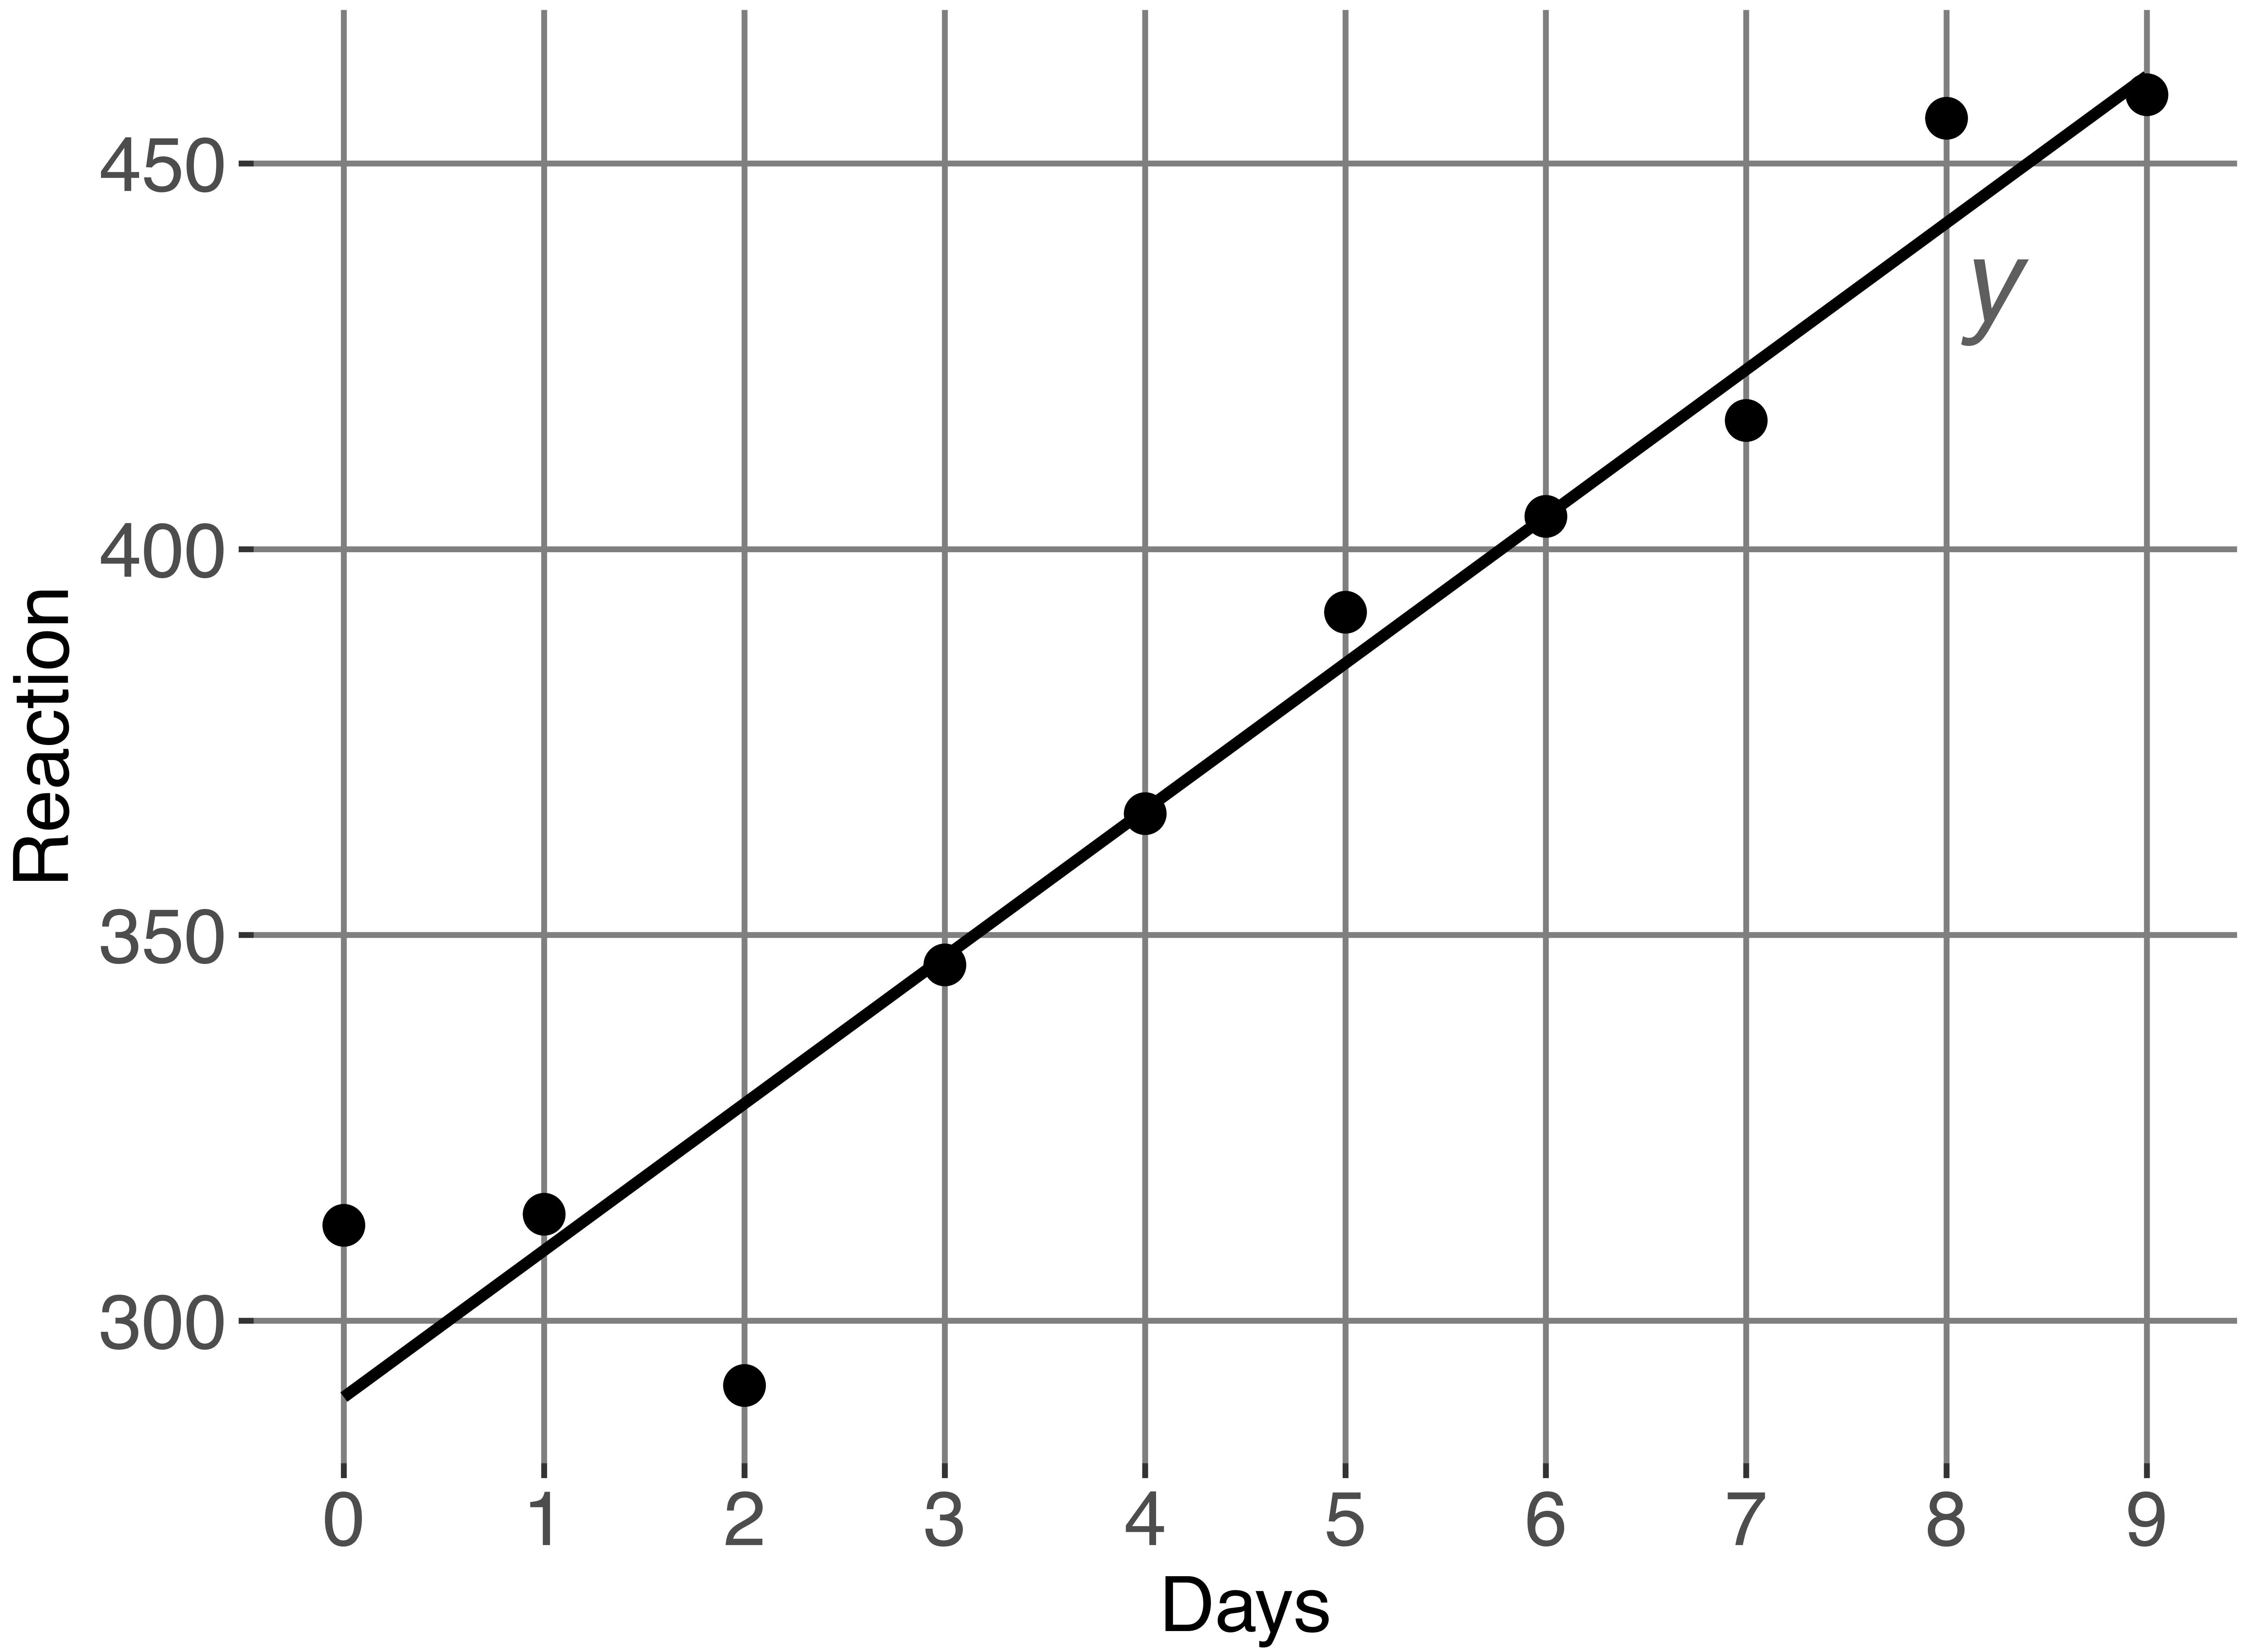

What is the effect of sleep deprivation on reaction time?

$$y = \alpha + \beta \cdot day$$

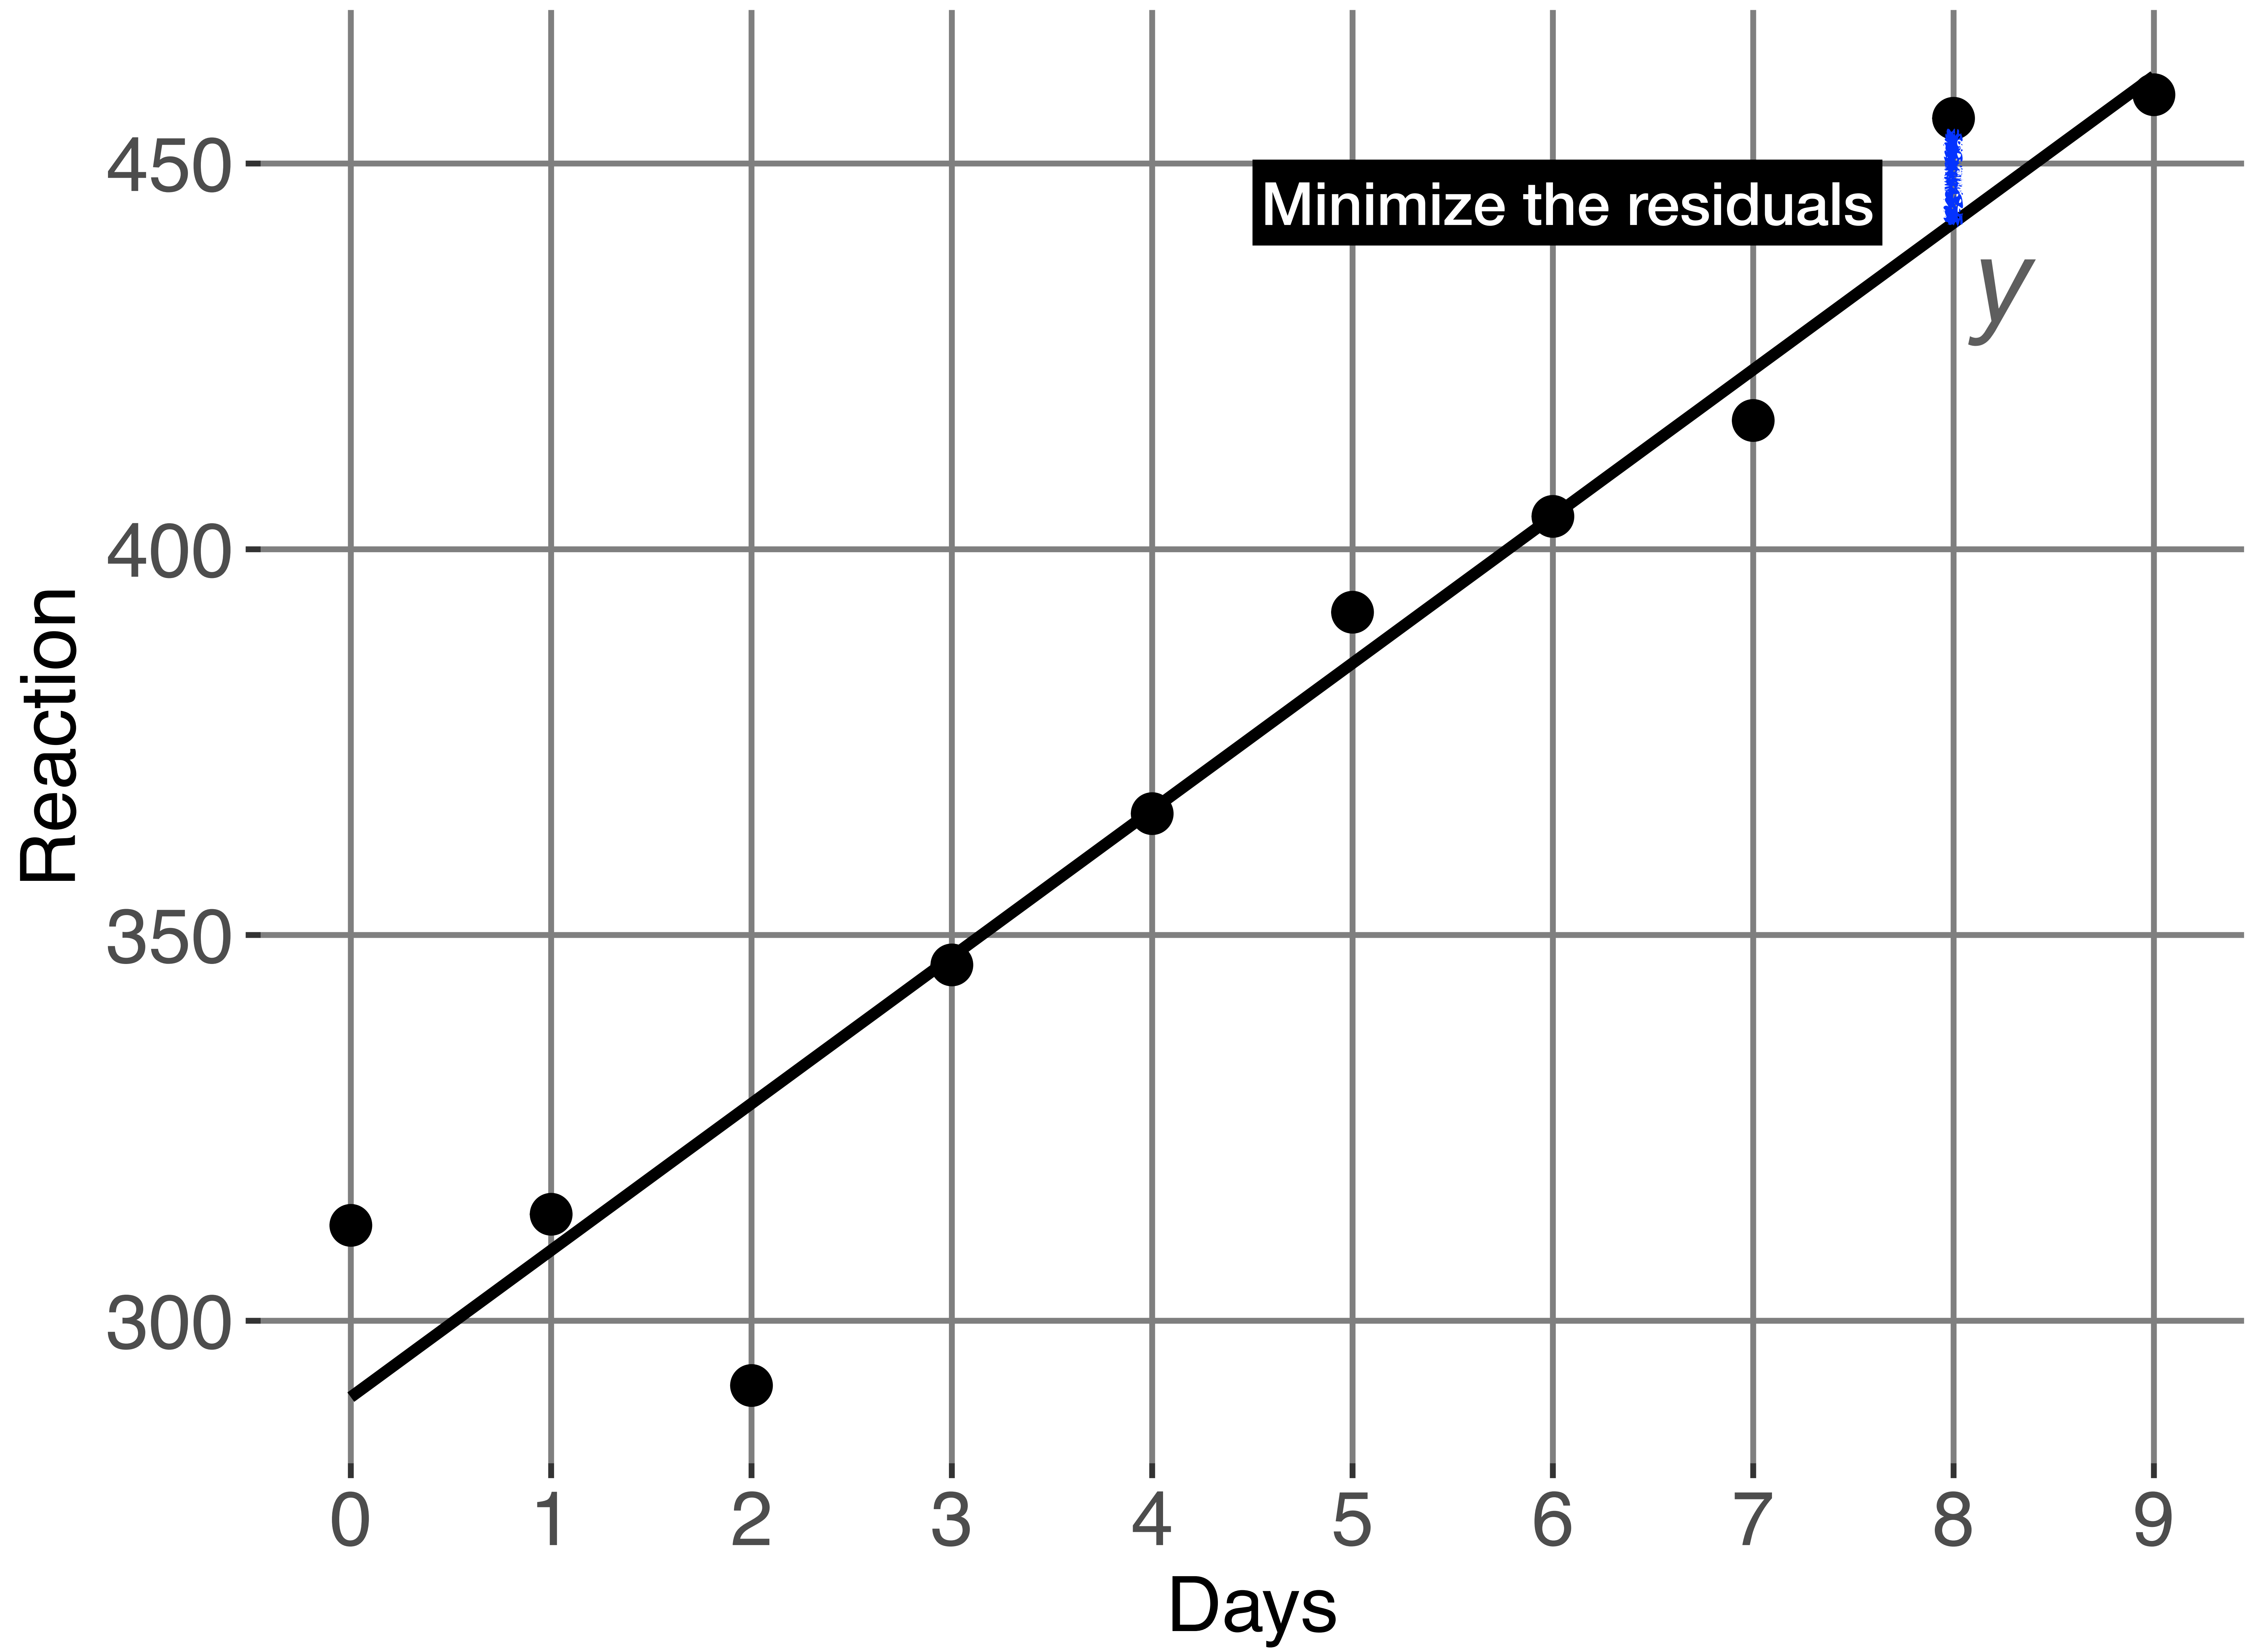

What is the effect of sleep deprivation on reaction time?

$$y = \alpha + \beta \cdot day$$

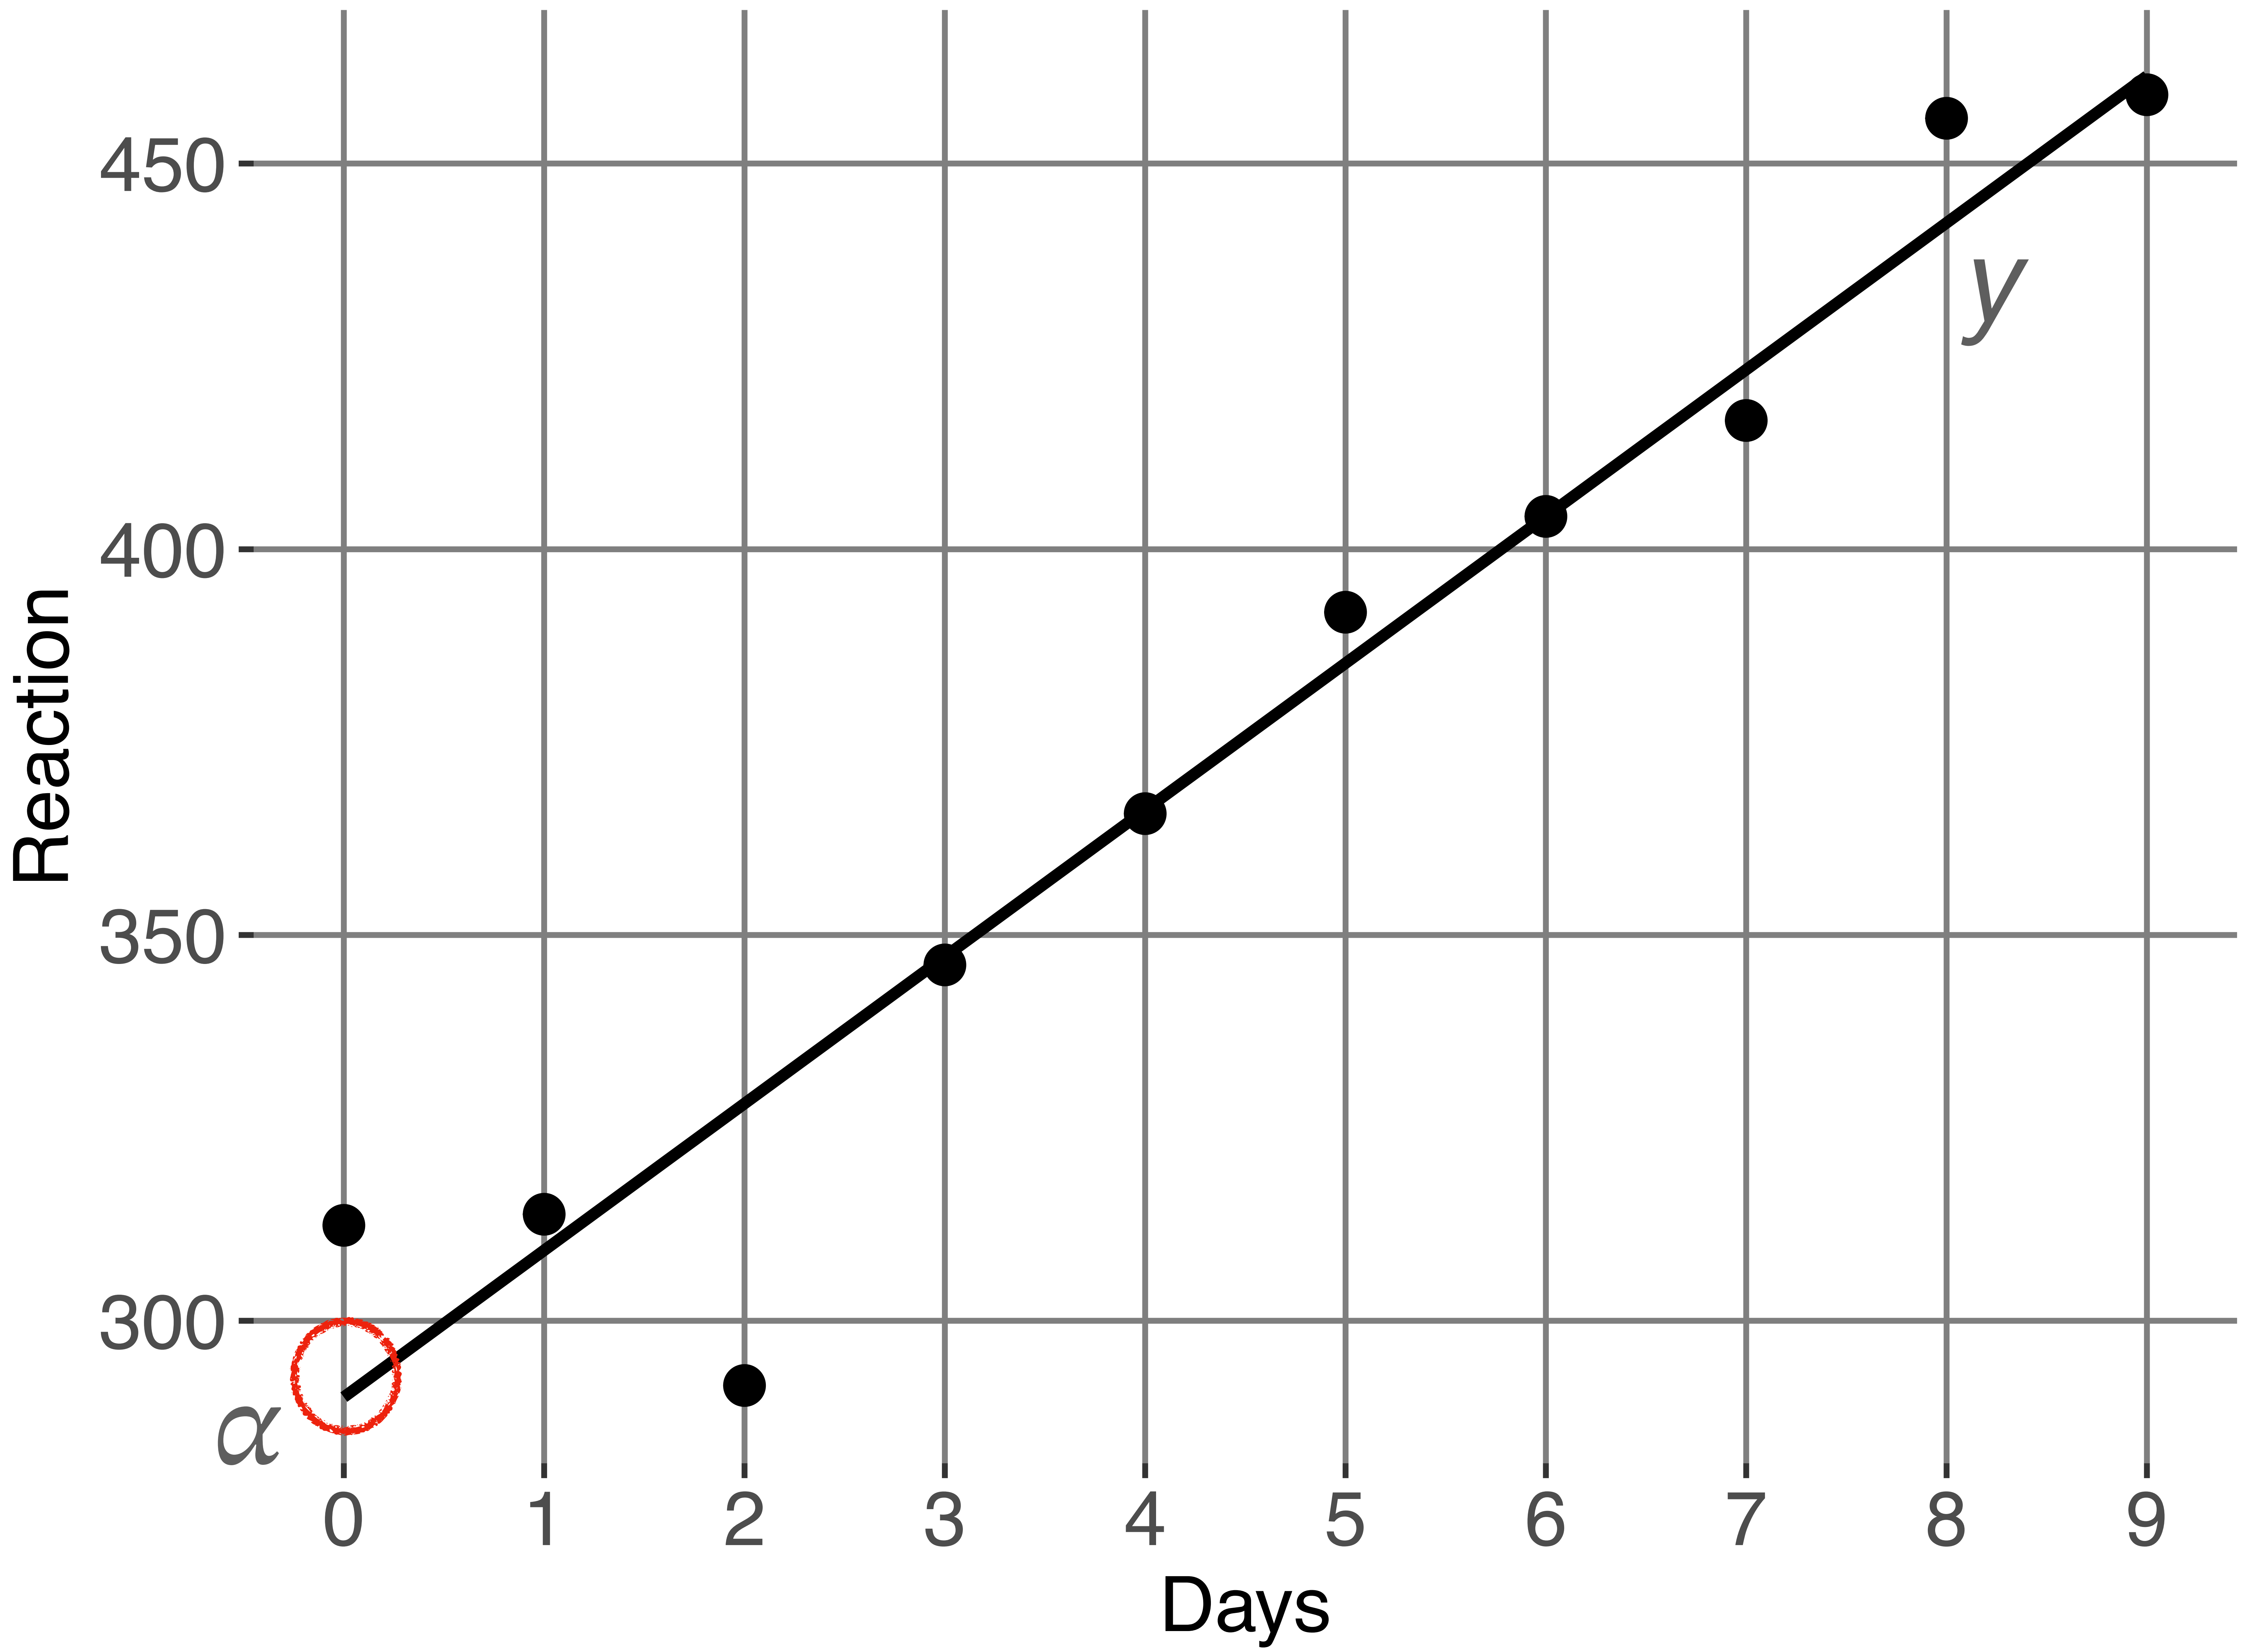

What is the effect of sleep deprivation on reaction time?

$$y = \alpha + \beta \cdot day$$

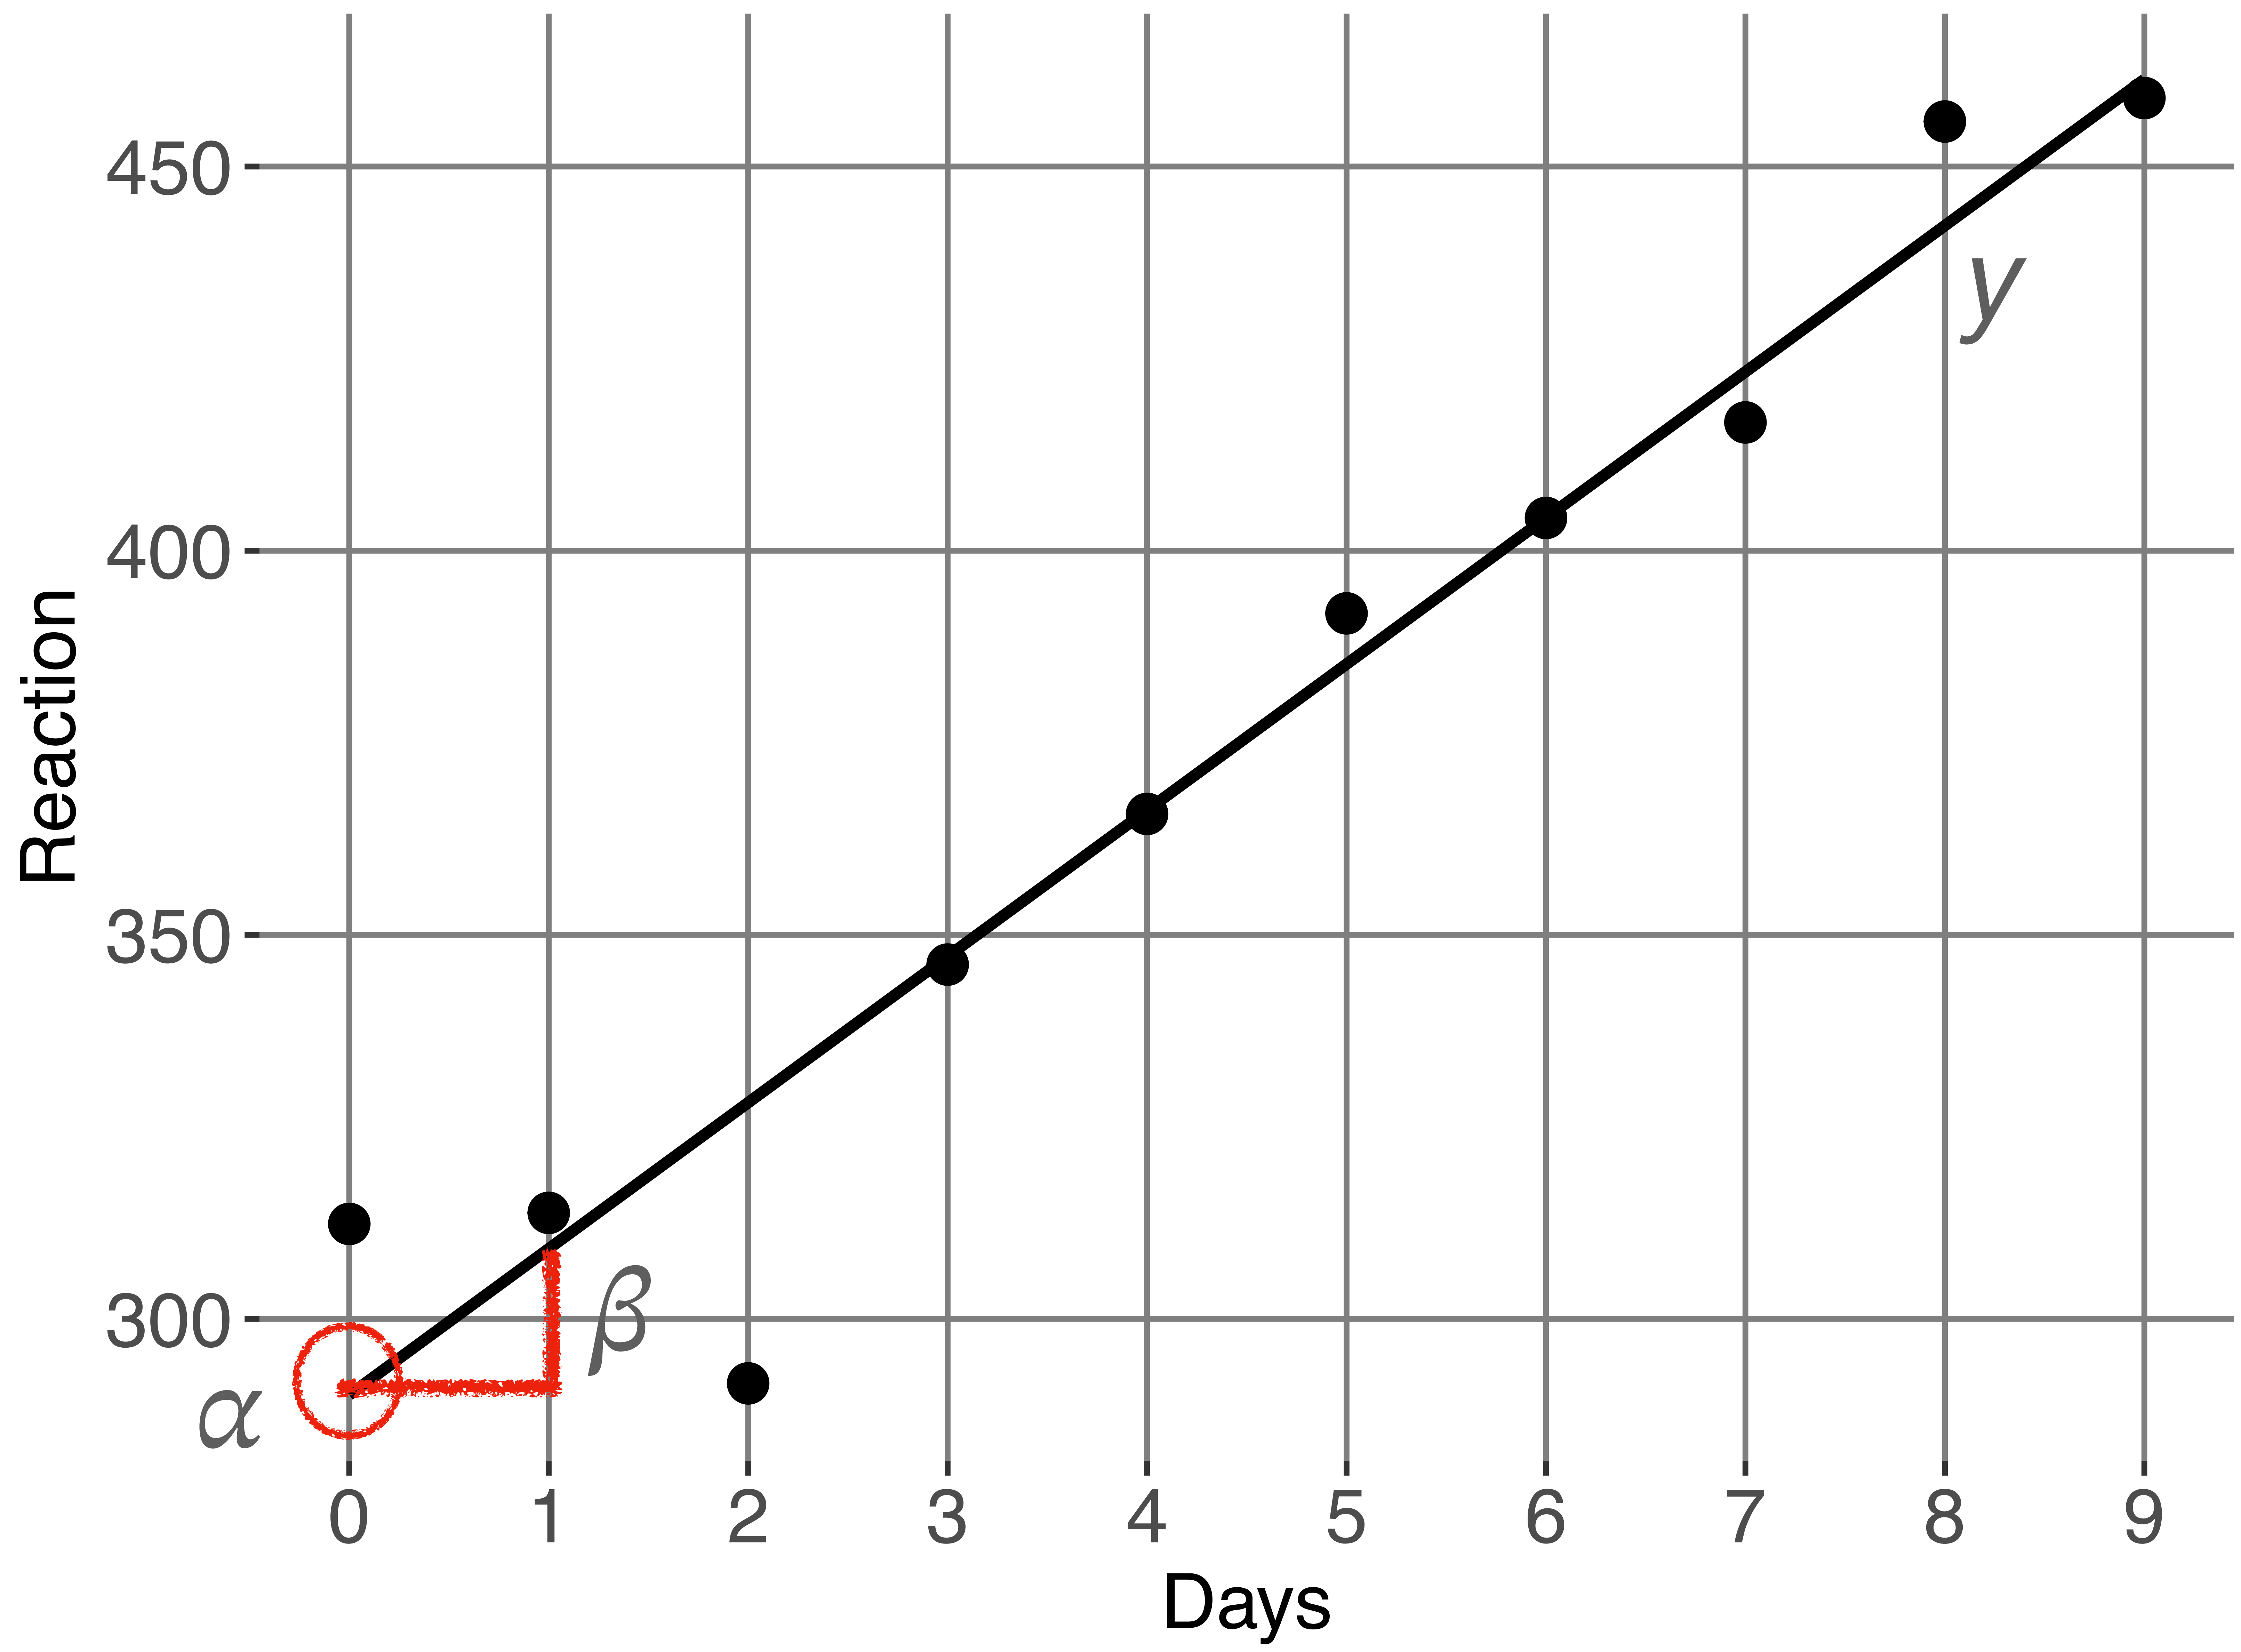

What is the effect of sleep deprivation on reaction time?

$$y = \alpha + \beta \cdot day$$

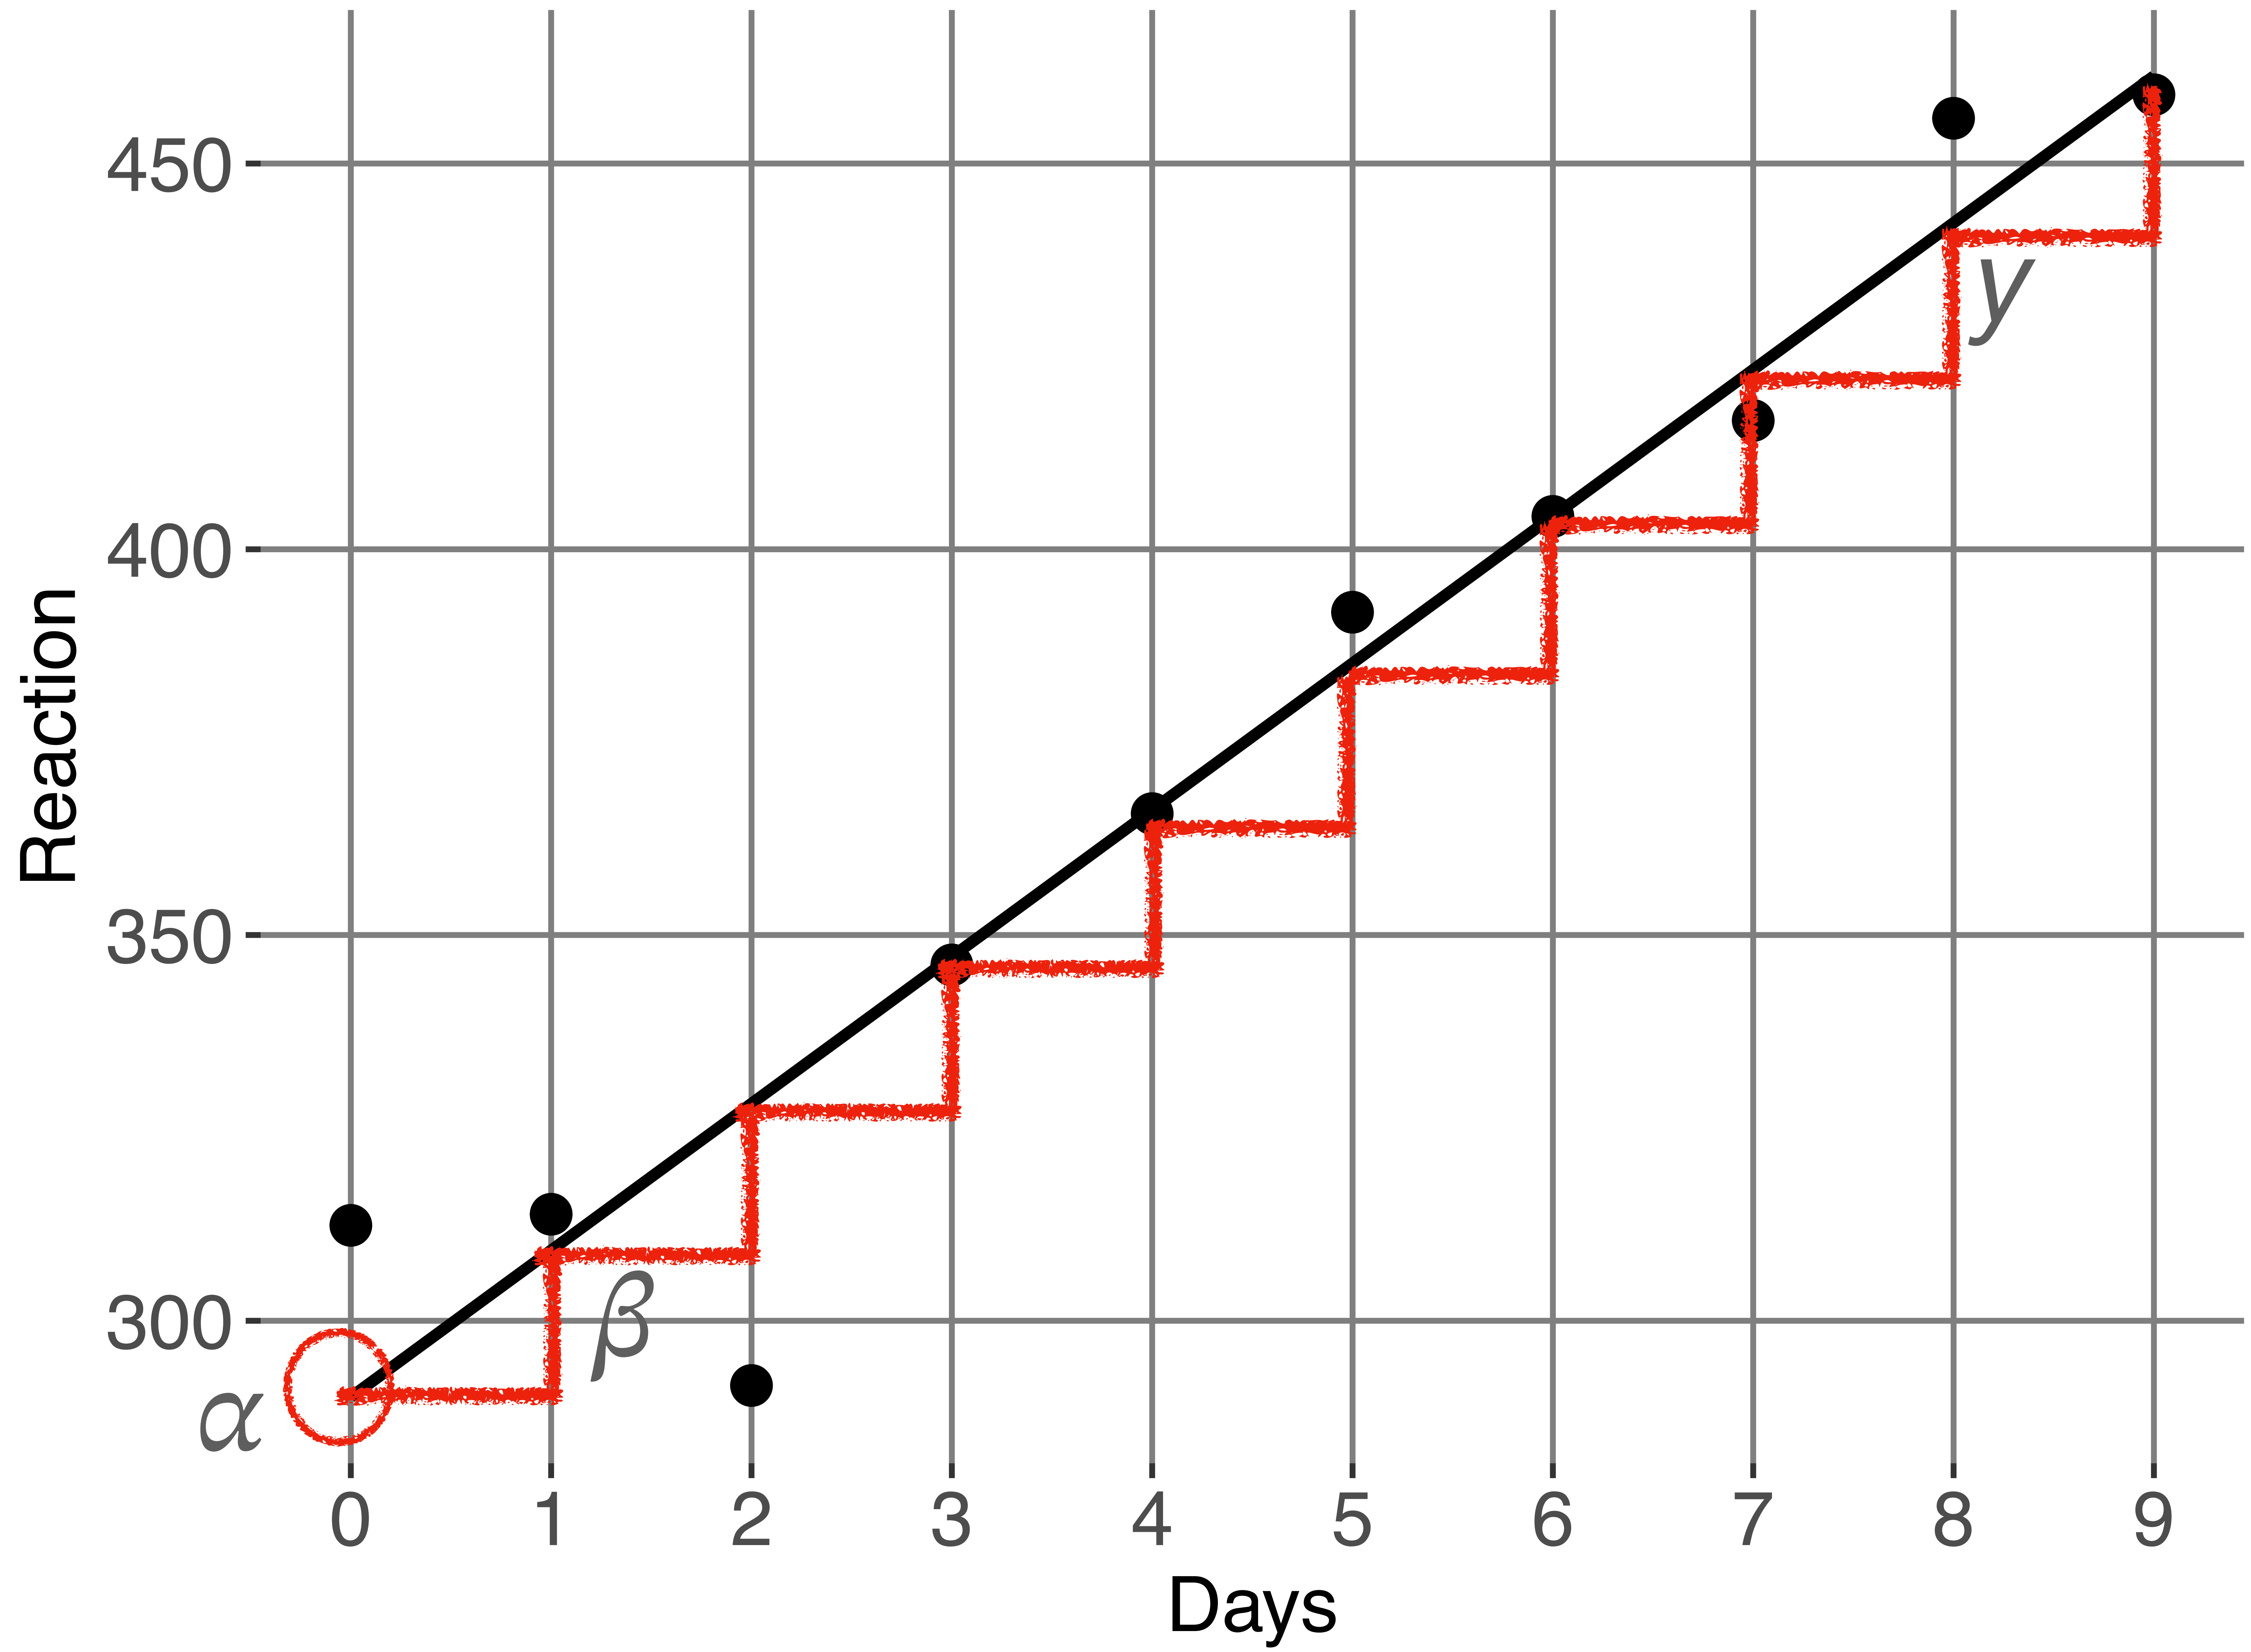

What is the effect of sleep deprivation on reaction time?

$$y = \alpha + \beta \cdot day$$

$$y_8 = \alpha + \beta \cdot day + \varepsilon_8$$

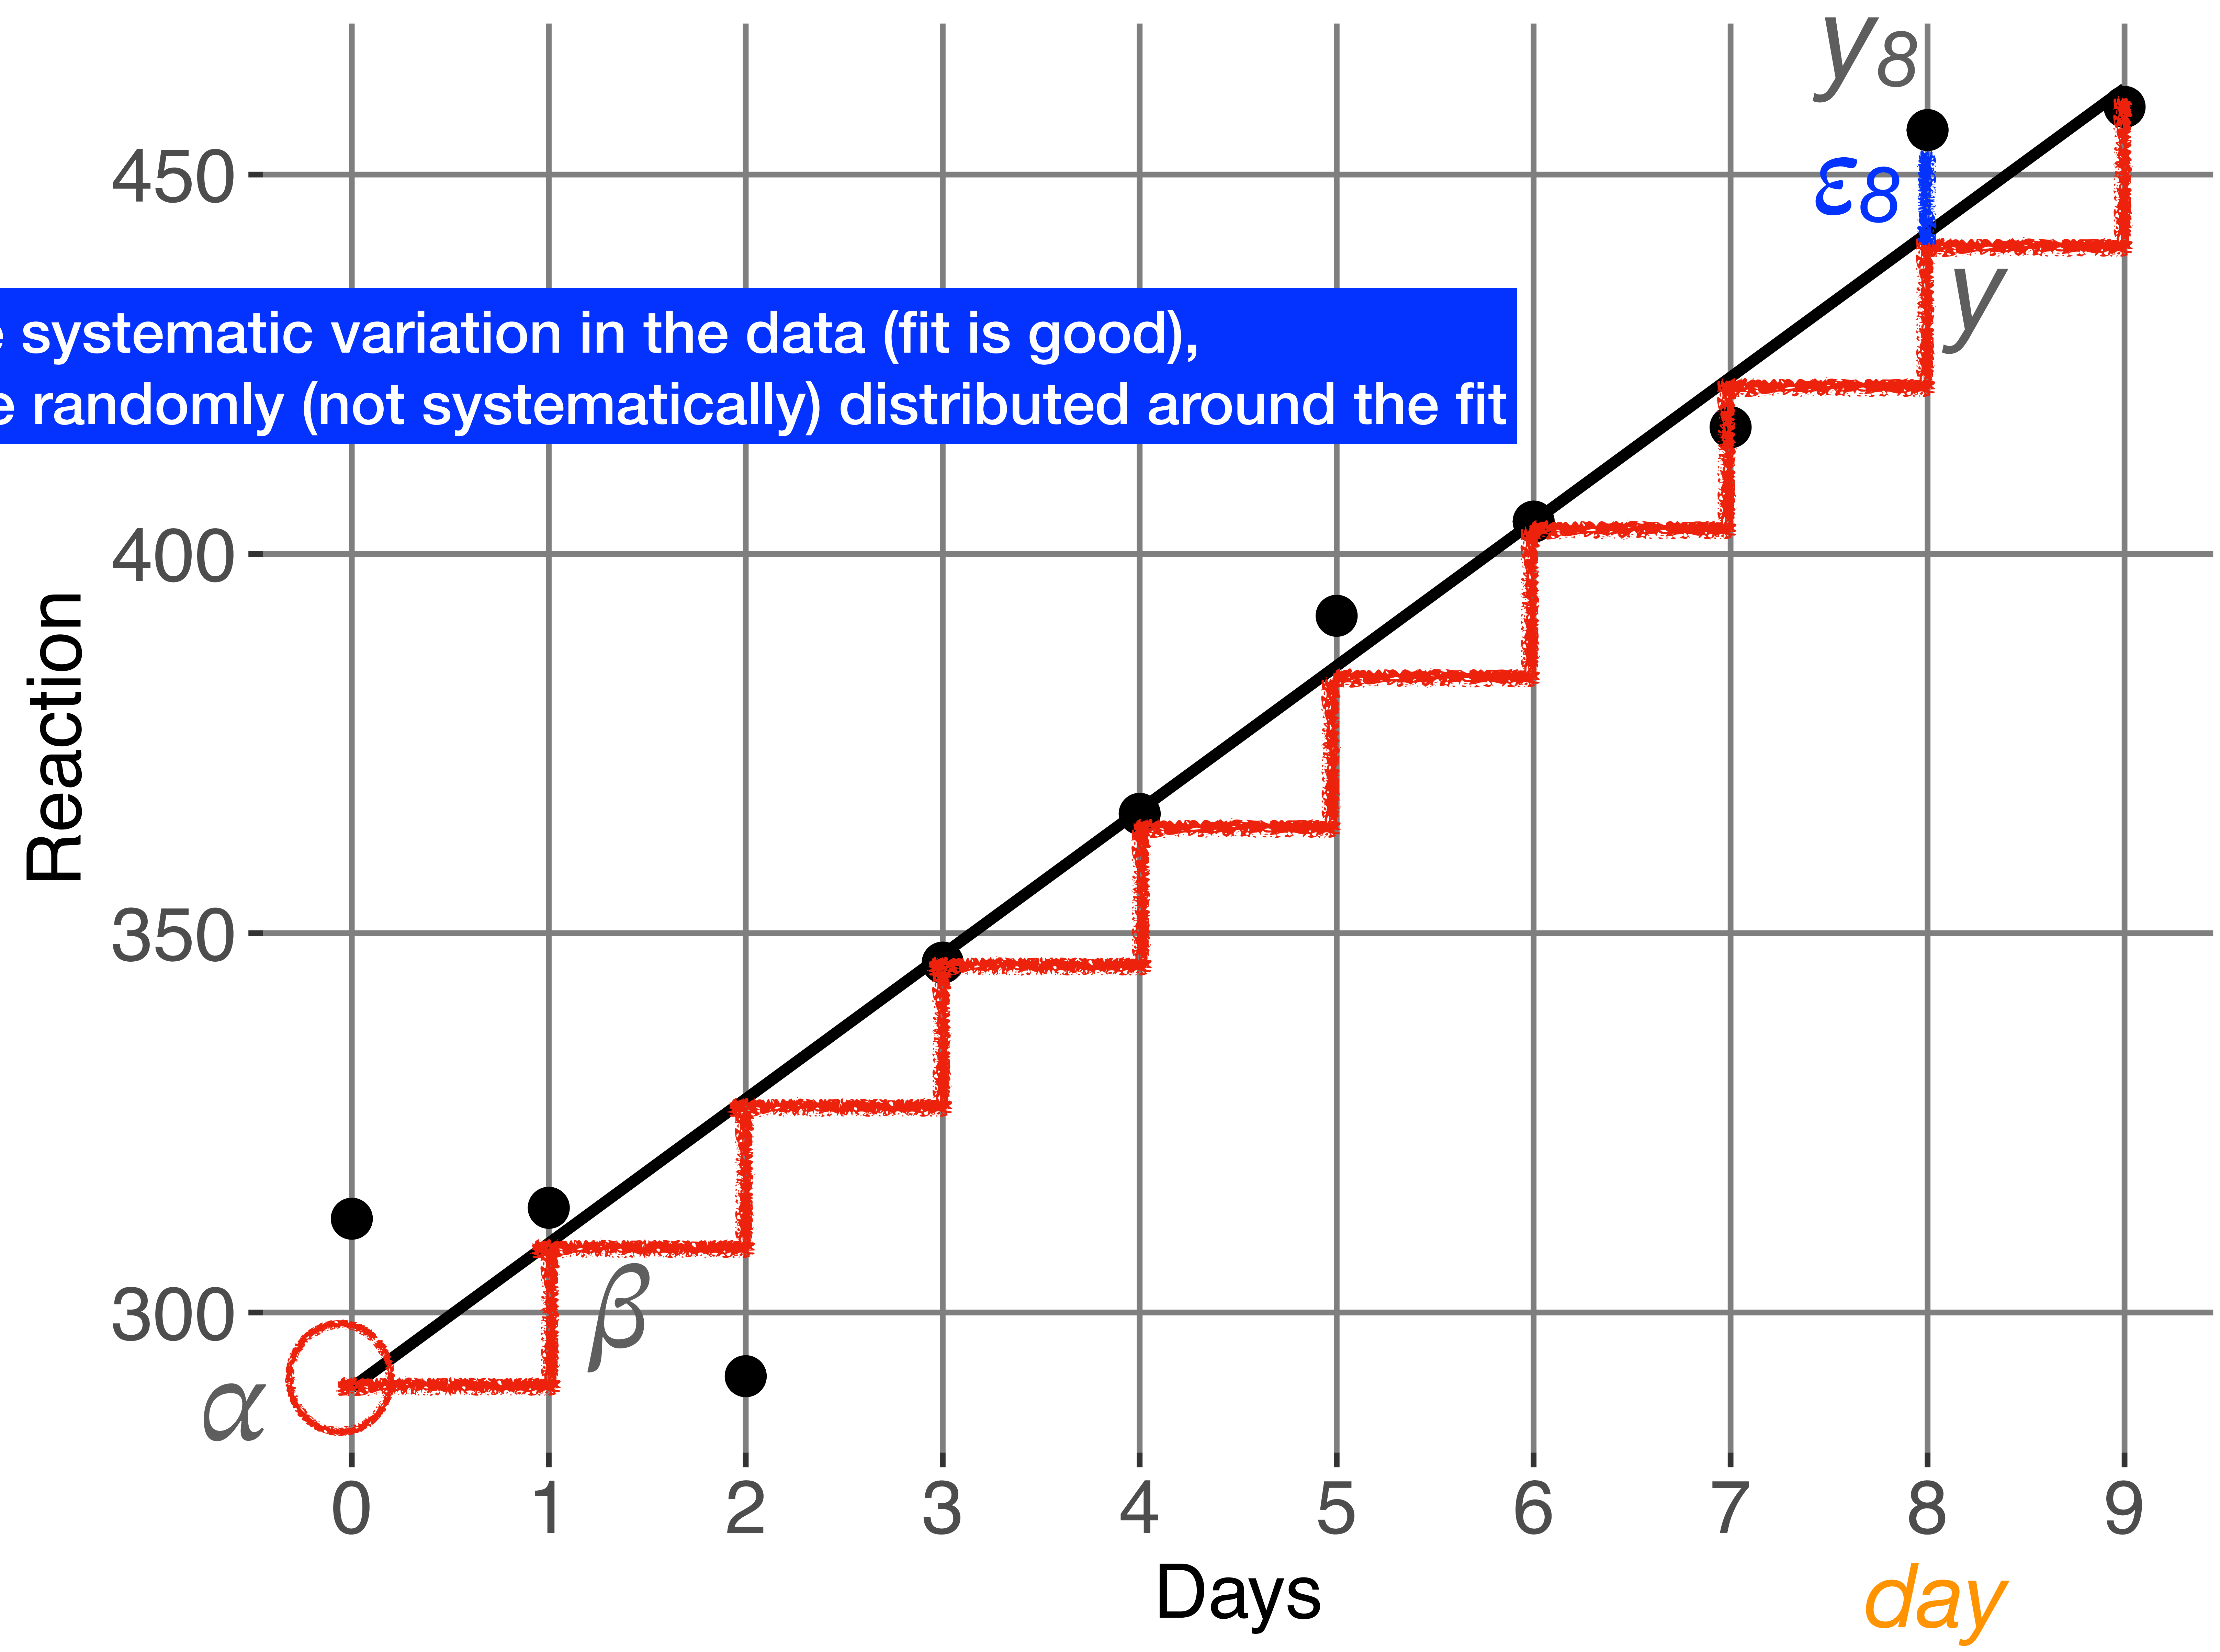

If the fit is catching the systematic variation in the data (fit is good), the residuals should be randomly (not systematically) distributed around the fit

β

Slope (regression coefficient)

What happens in the dependent variable when we move along 1 unit in the predictor?
How much reaction time increases when we have one bad night of sleep?
'The effect of sleep deprivation on reaction time'

β

Slope (regression coefficient)

What happens in the dependent variable when we move along 1 unit in the predictor?
How much reaction time increases when we have one bad night of sleep?
'The effect of sleep deprivation on reaction time'

α

Intercept

What is the overall level of the dependent variable
What is the reaction time with no sleep deprivation?

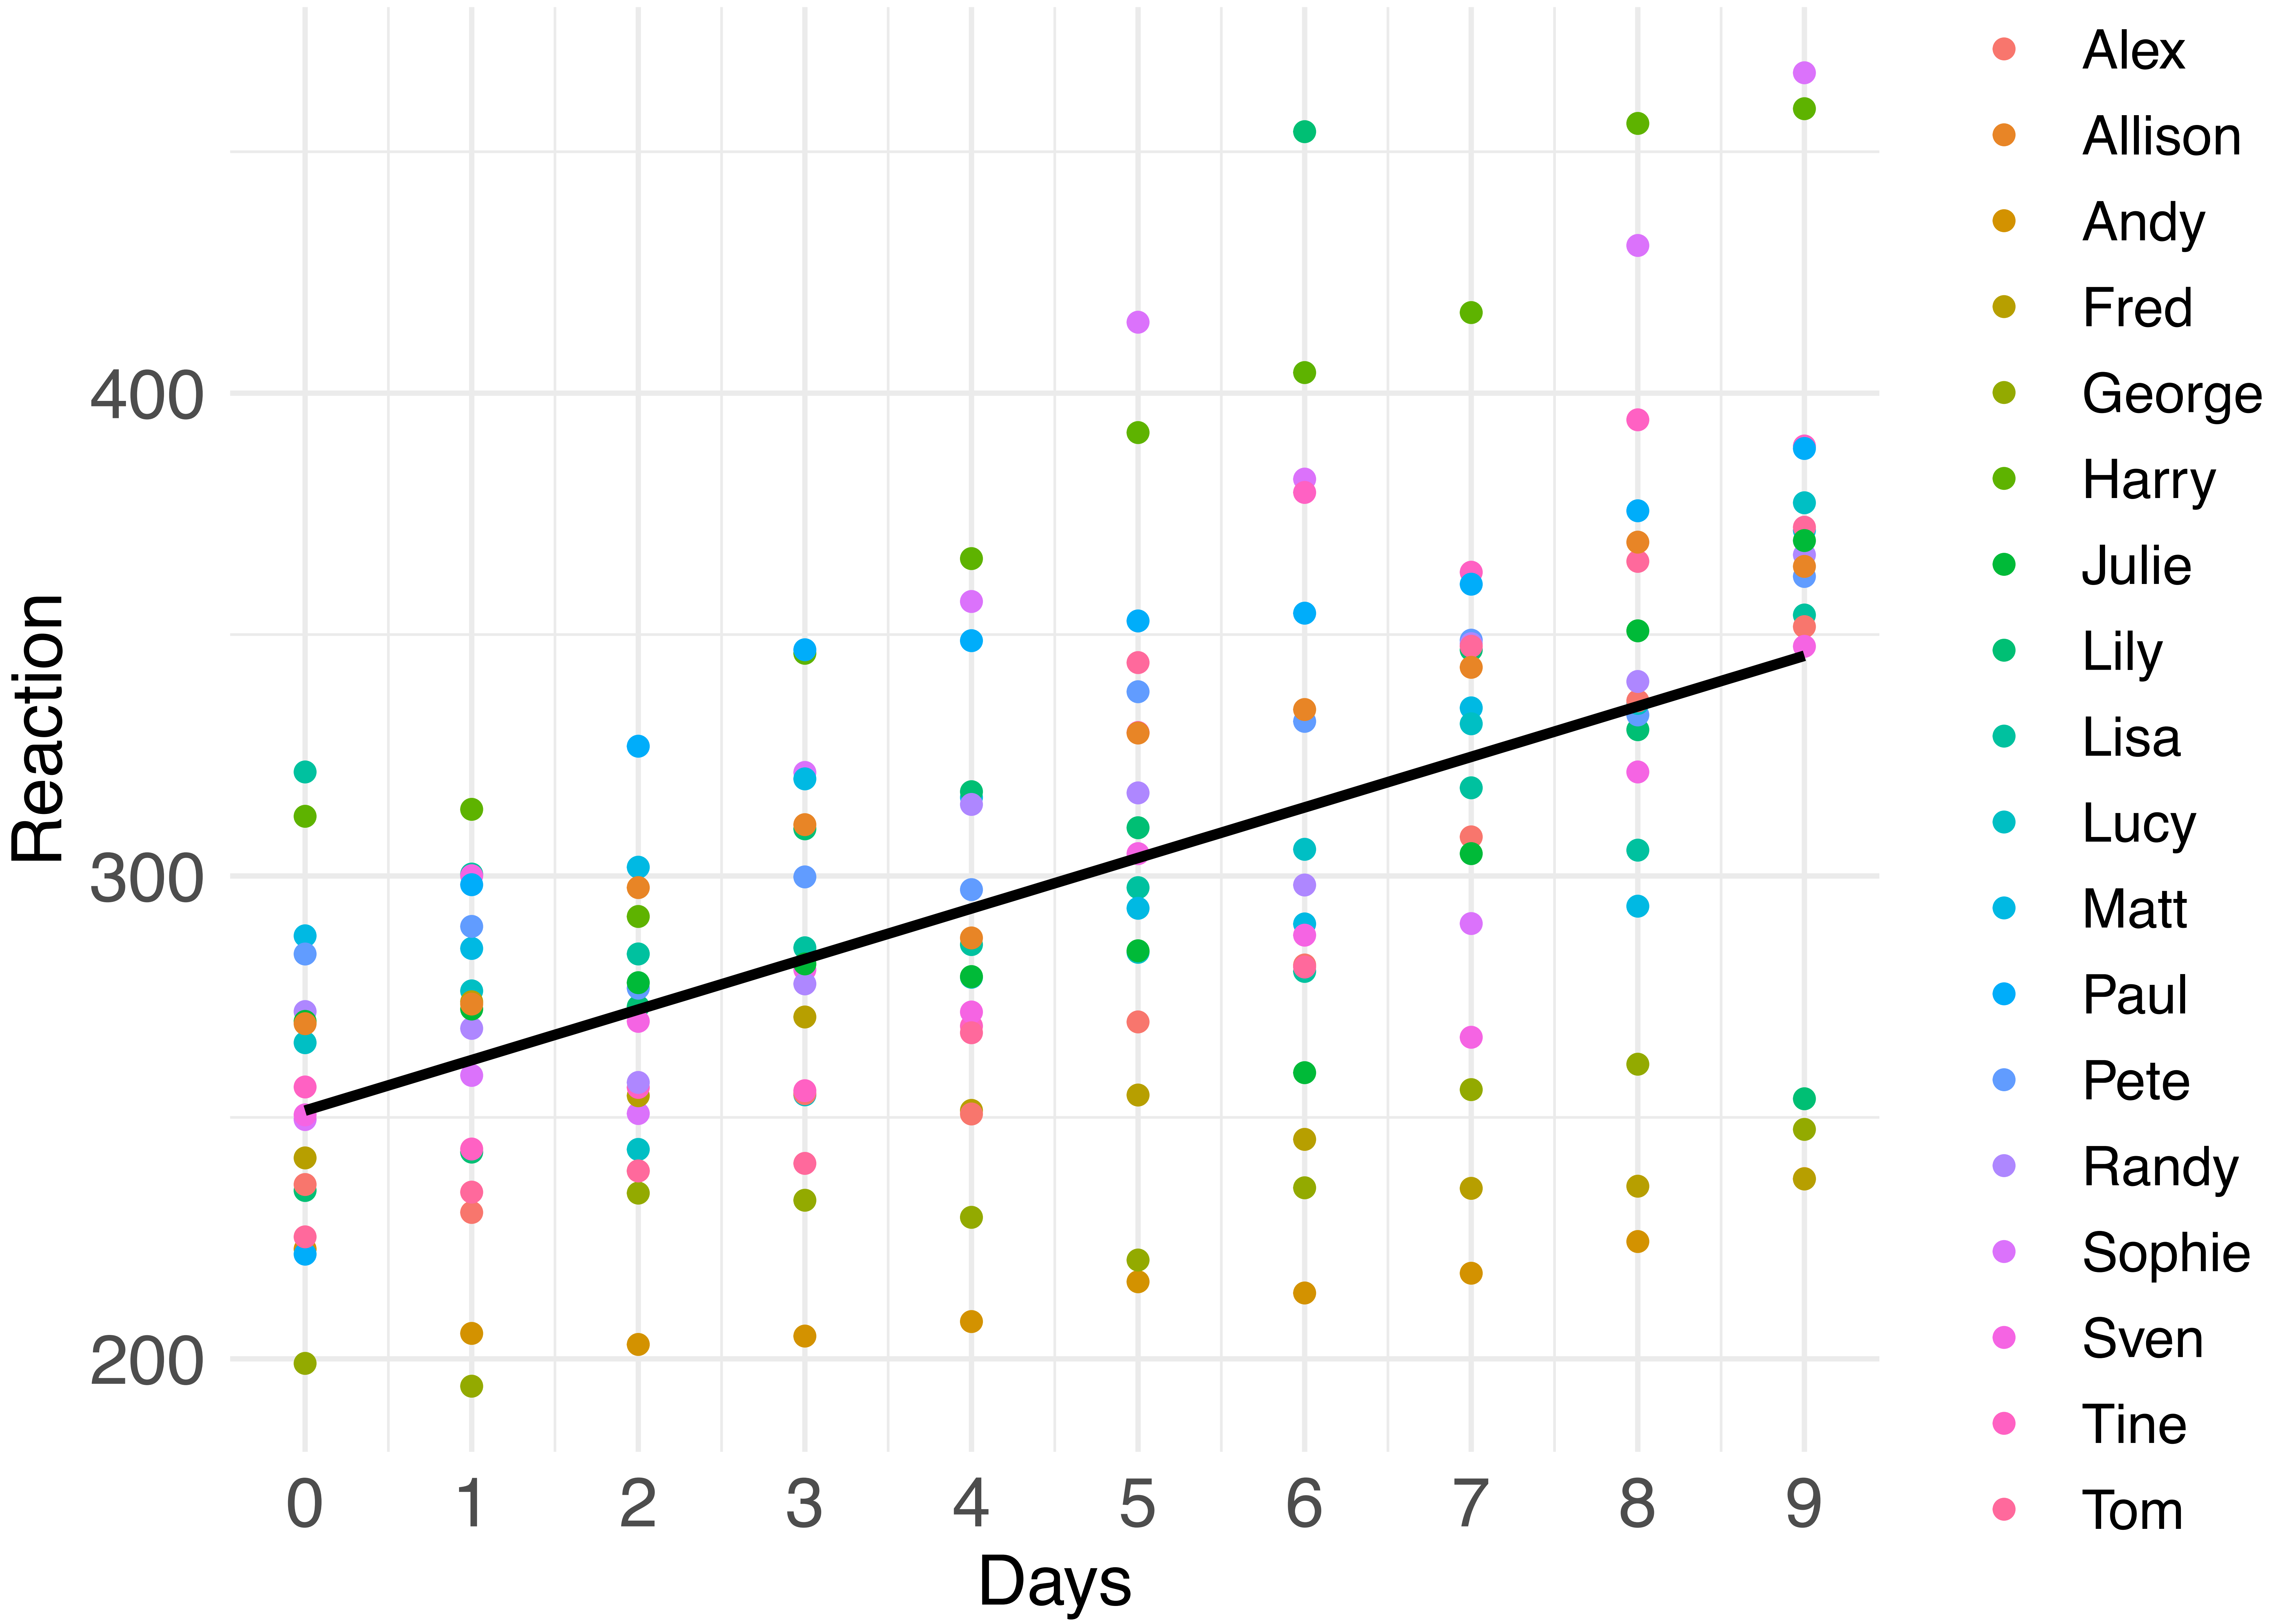

Sleepstudy dataset from lme4 package in R
Days= Number of days of sleep deprivation (3 hours/ night)
Reaction= Average reaction time (ms)

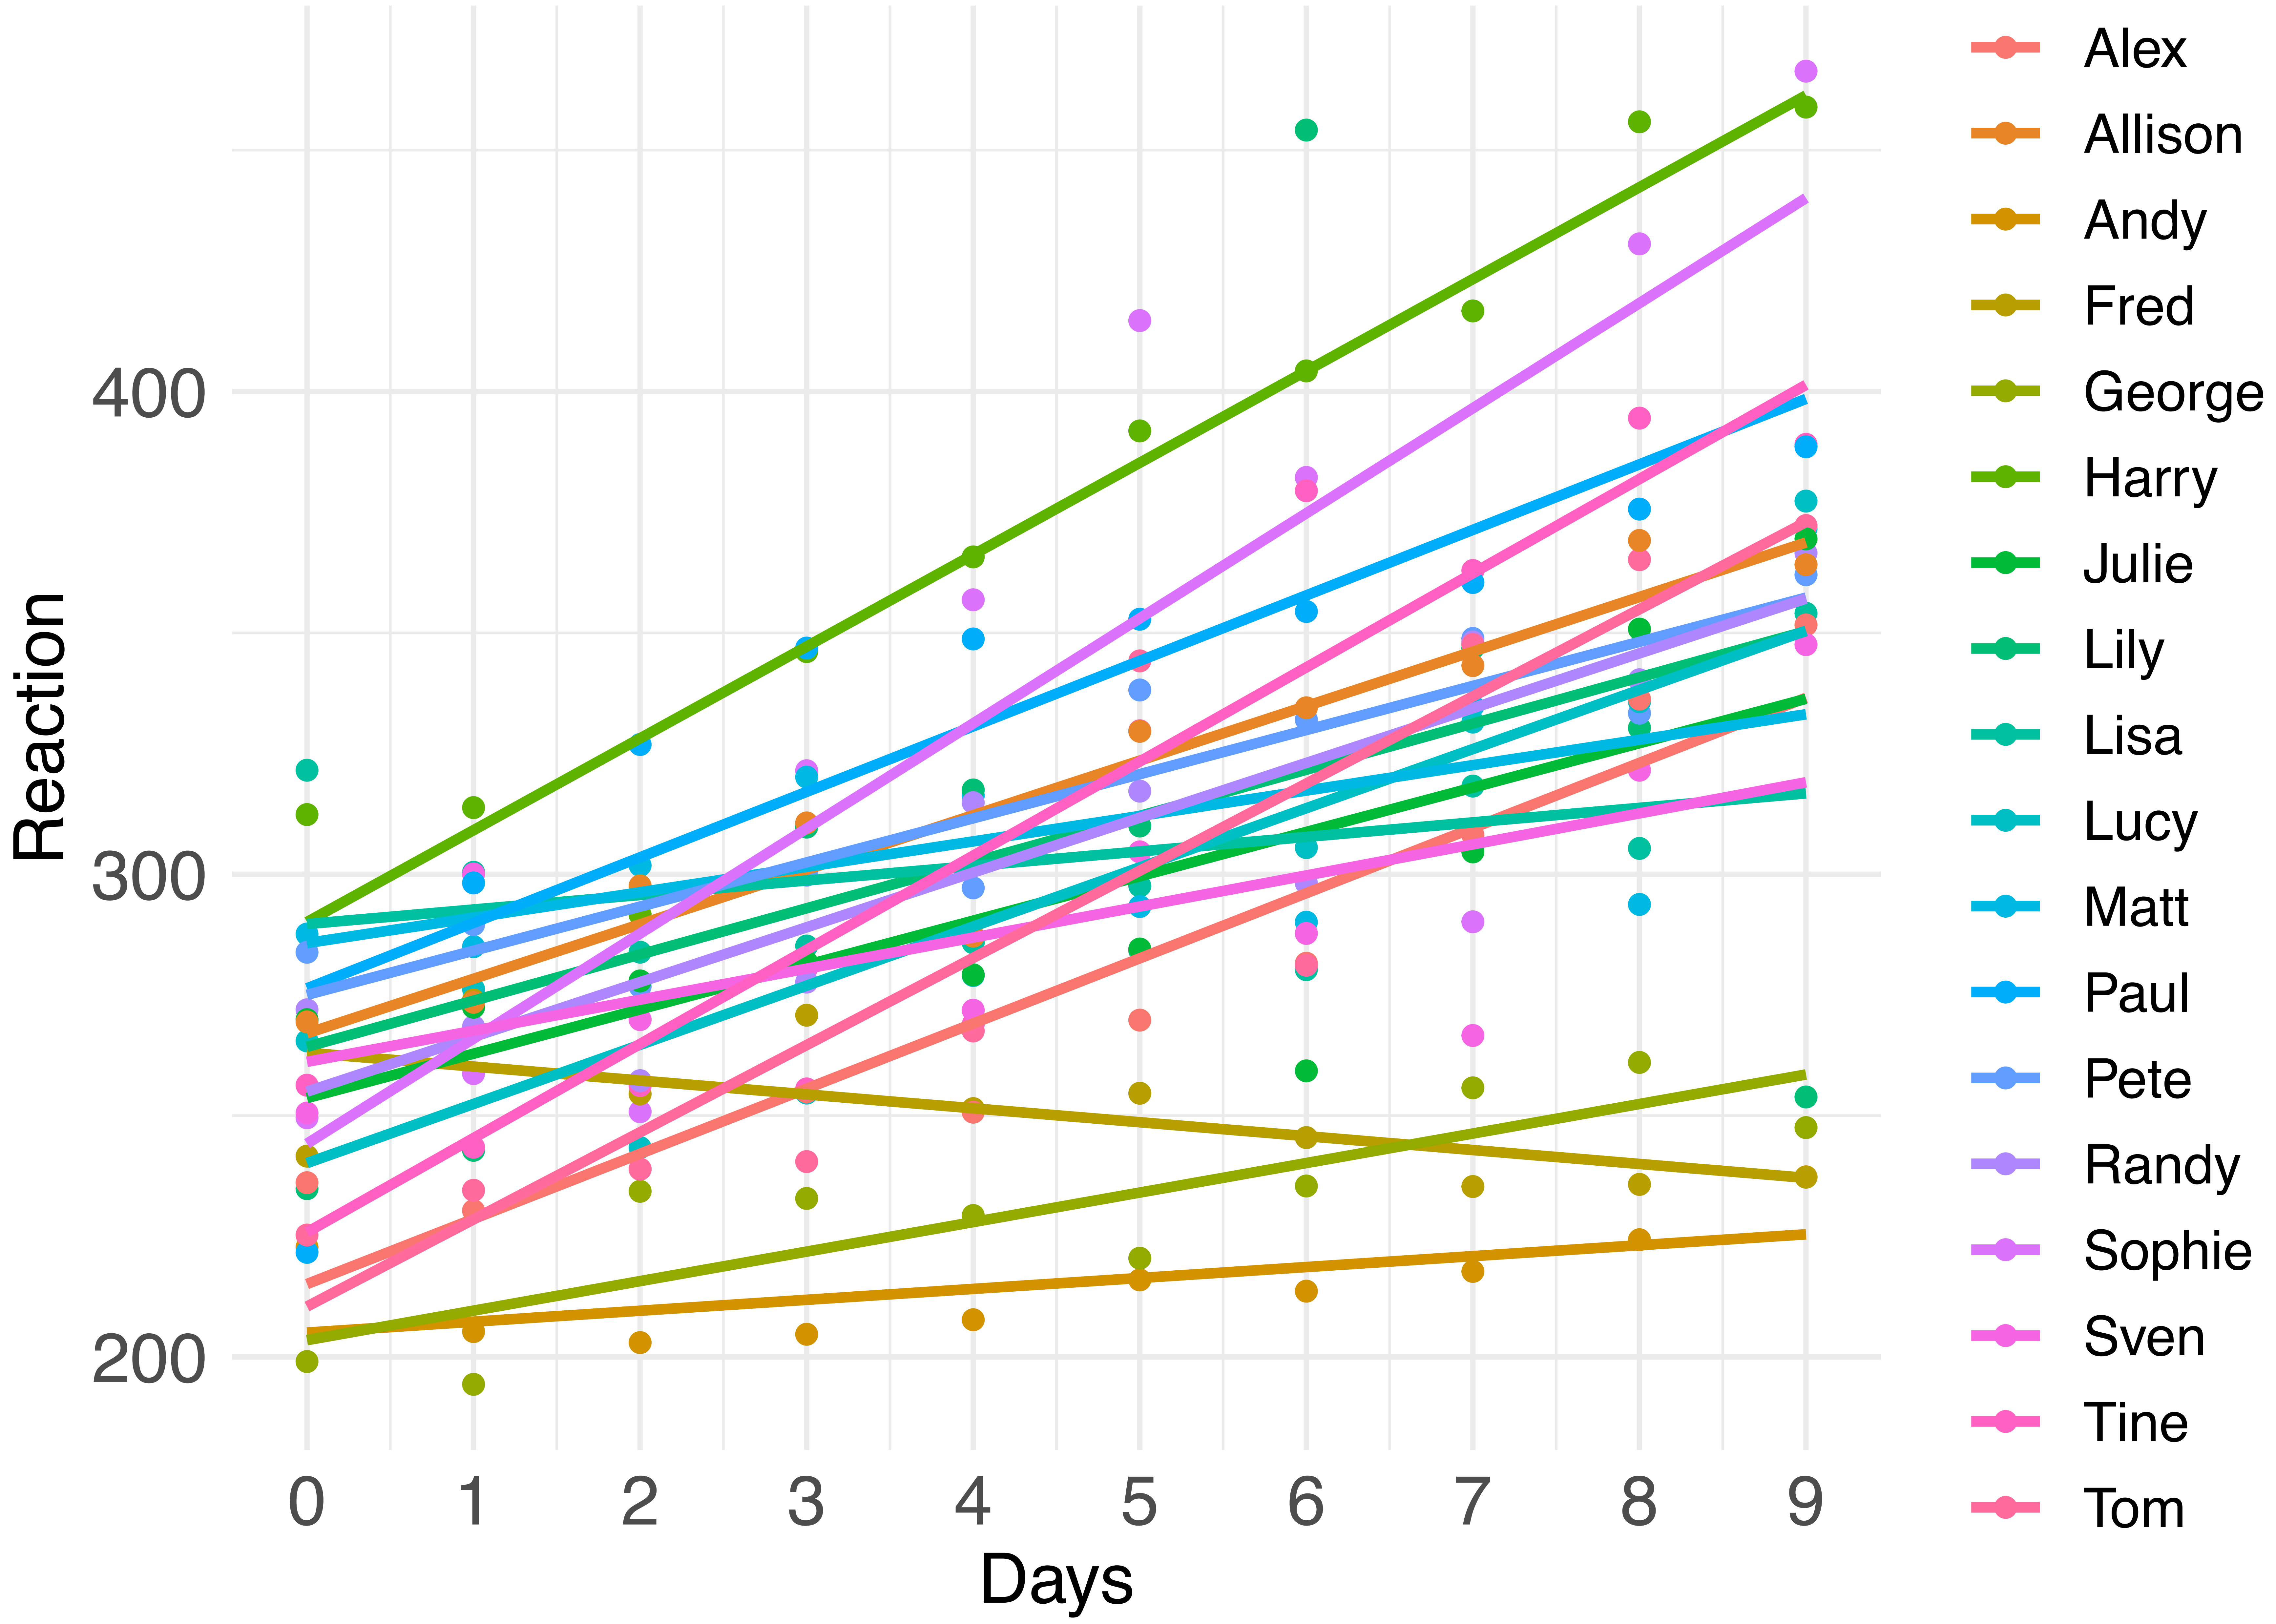
Reaction
400
300
200
0
1
2
3
4
5
6
7
8
9
Days
Alex
Allison
Andy
Fred
George
Harry
Julie
Lily
Lisa
Lucy
Matt
Paul
Pete
Randy
Sophie
Sven
Tine
Tom

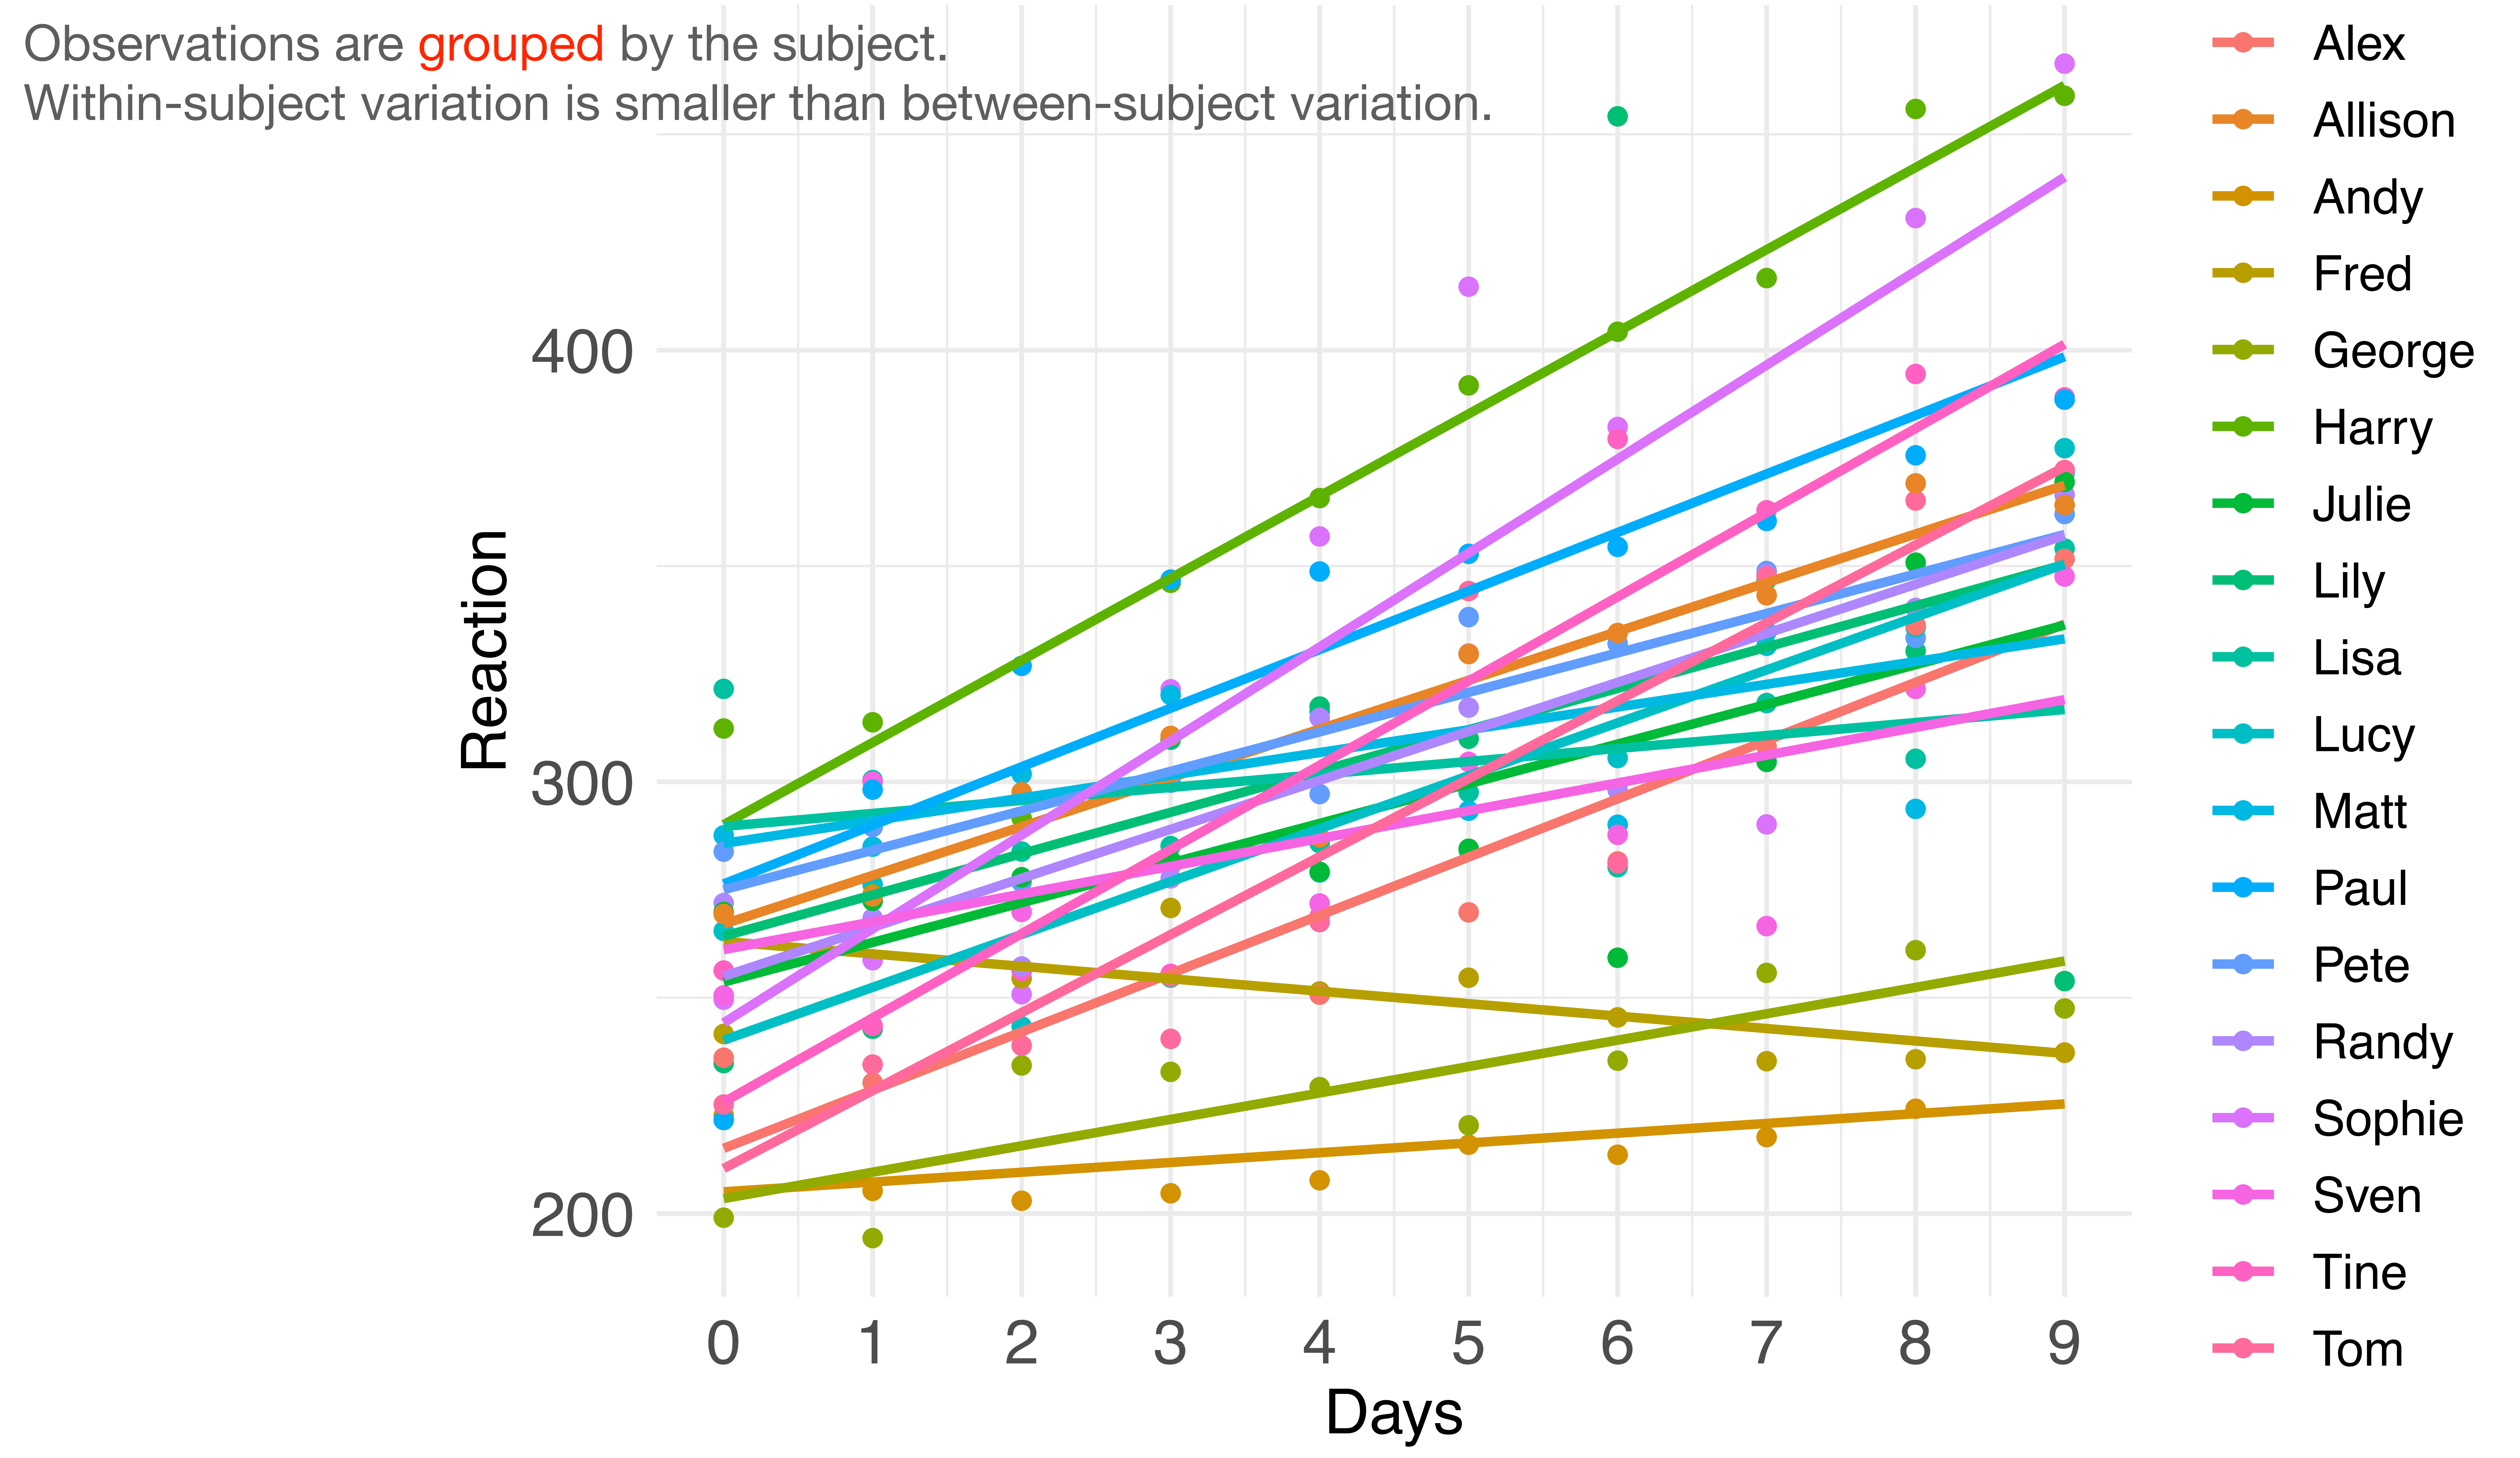
Observations are grouped by the subject.
Within-subject variation is smaller than between-subject variation.
Reaction
400
300
200
0
1
2
3
4
5
6
7
8
9
Days
Alex
Allison
Andy
Fred
George
Harry
Julie
Lily
Lisa
Lucy
Matt
Paul
Pete
Randy
Sophie
Sven
Tine
Tom

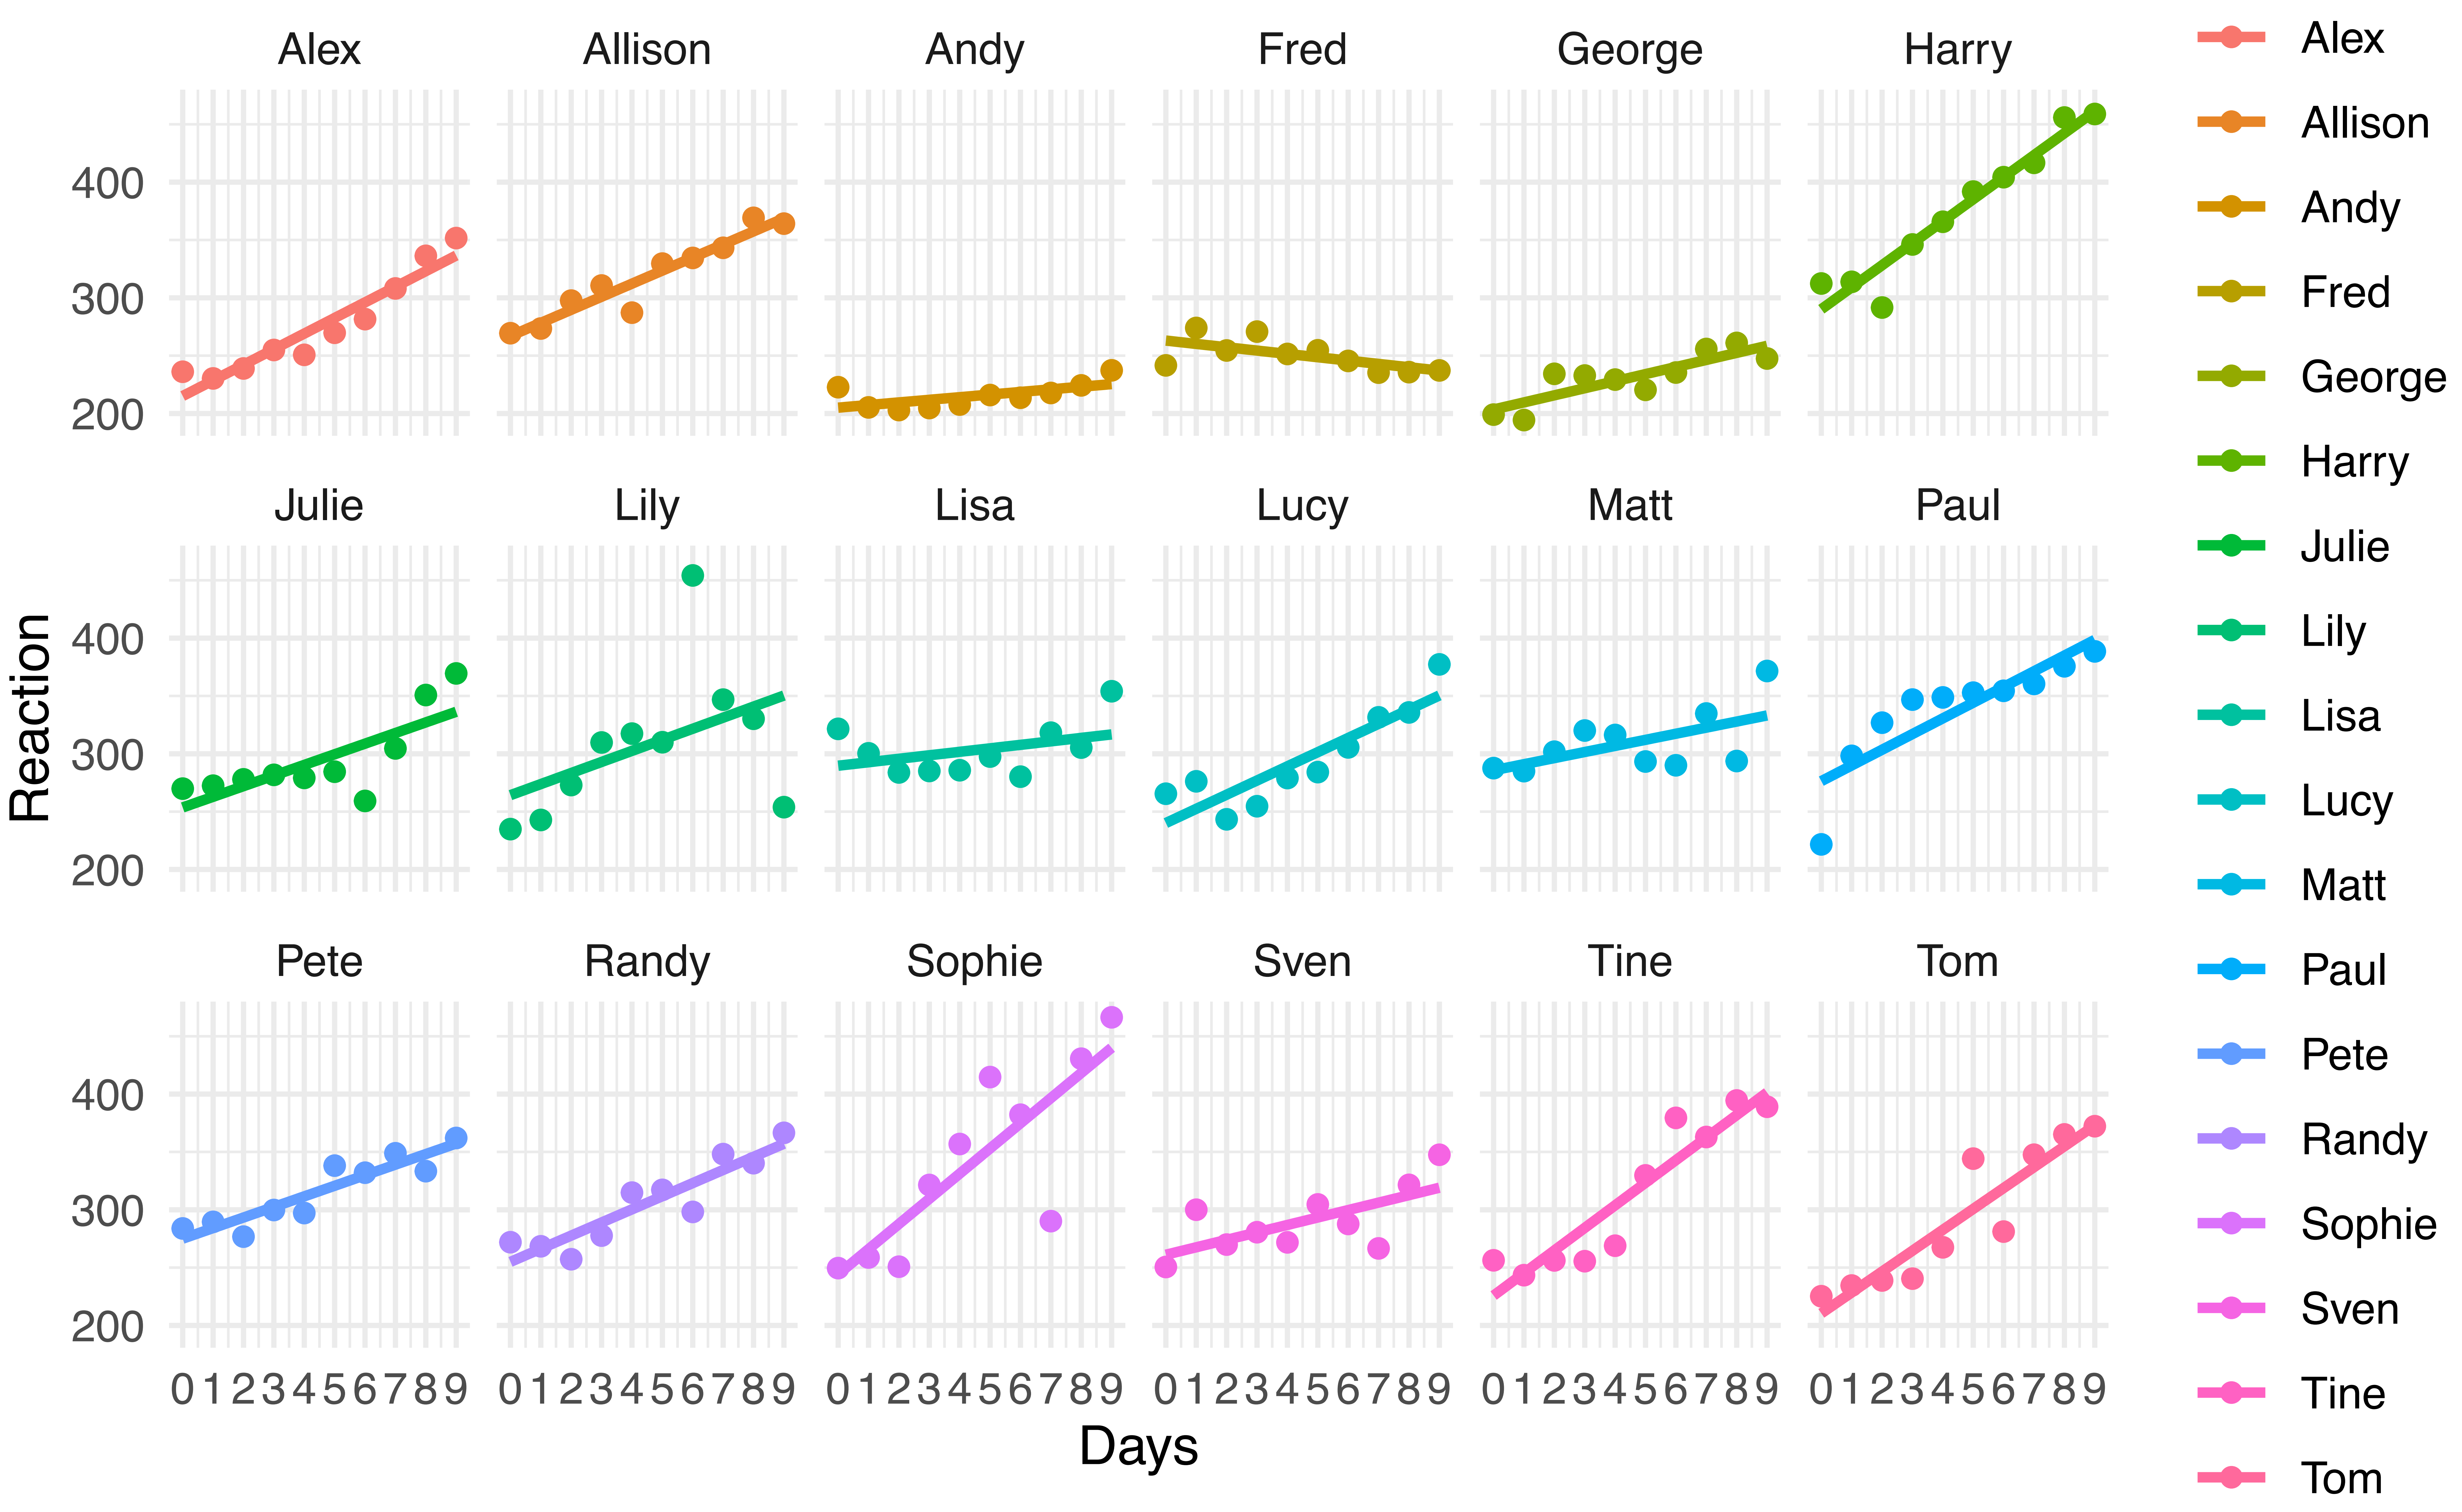

Alex
Allison
Andy
Fred
George
Harry
Julie
Lily
Lisa
Lucy
Matt
Paul
Pete
Randy
Sophie
Sven
Tine
Tom
Reaction
200
300
400
0123456789
Days

What is **generally** the effect of sleep deprivation on reaction time?

R: language and environment for statistical computing and graphics, freely available

R Core Team. (2021). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL https:/ / www.R-project.org/ .

$y = \alpha + \beta \cdot day$

```
fit_1 <- lm(Reaction ~ 1 + Days, data = data)
```

$y = \alpha + \beta \cdot day$

```
fit_1 <- lm(Reaction ~ 1 + Days, data = data)
```

Intercept = Reaction time when predictor (Days) is 0

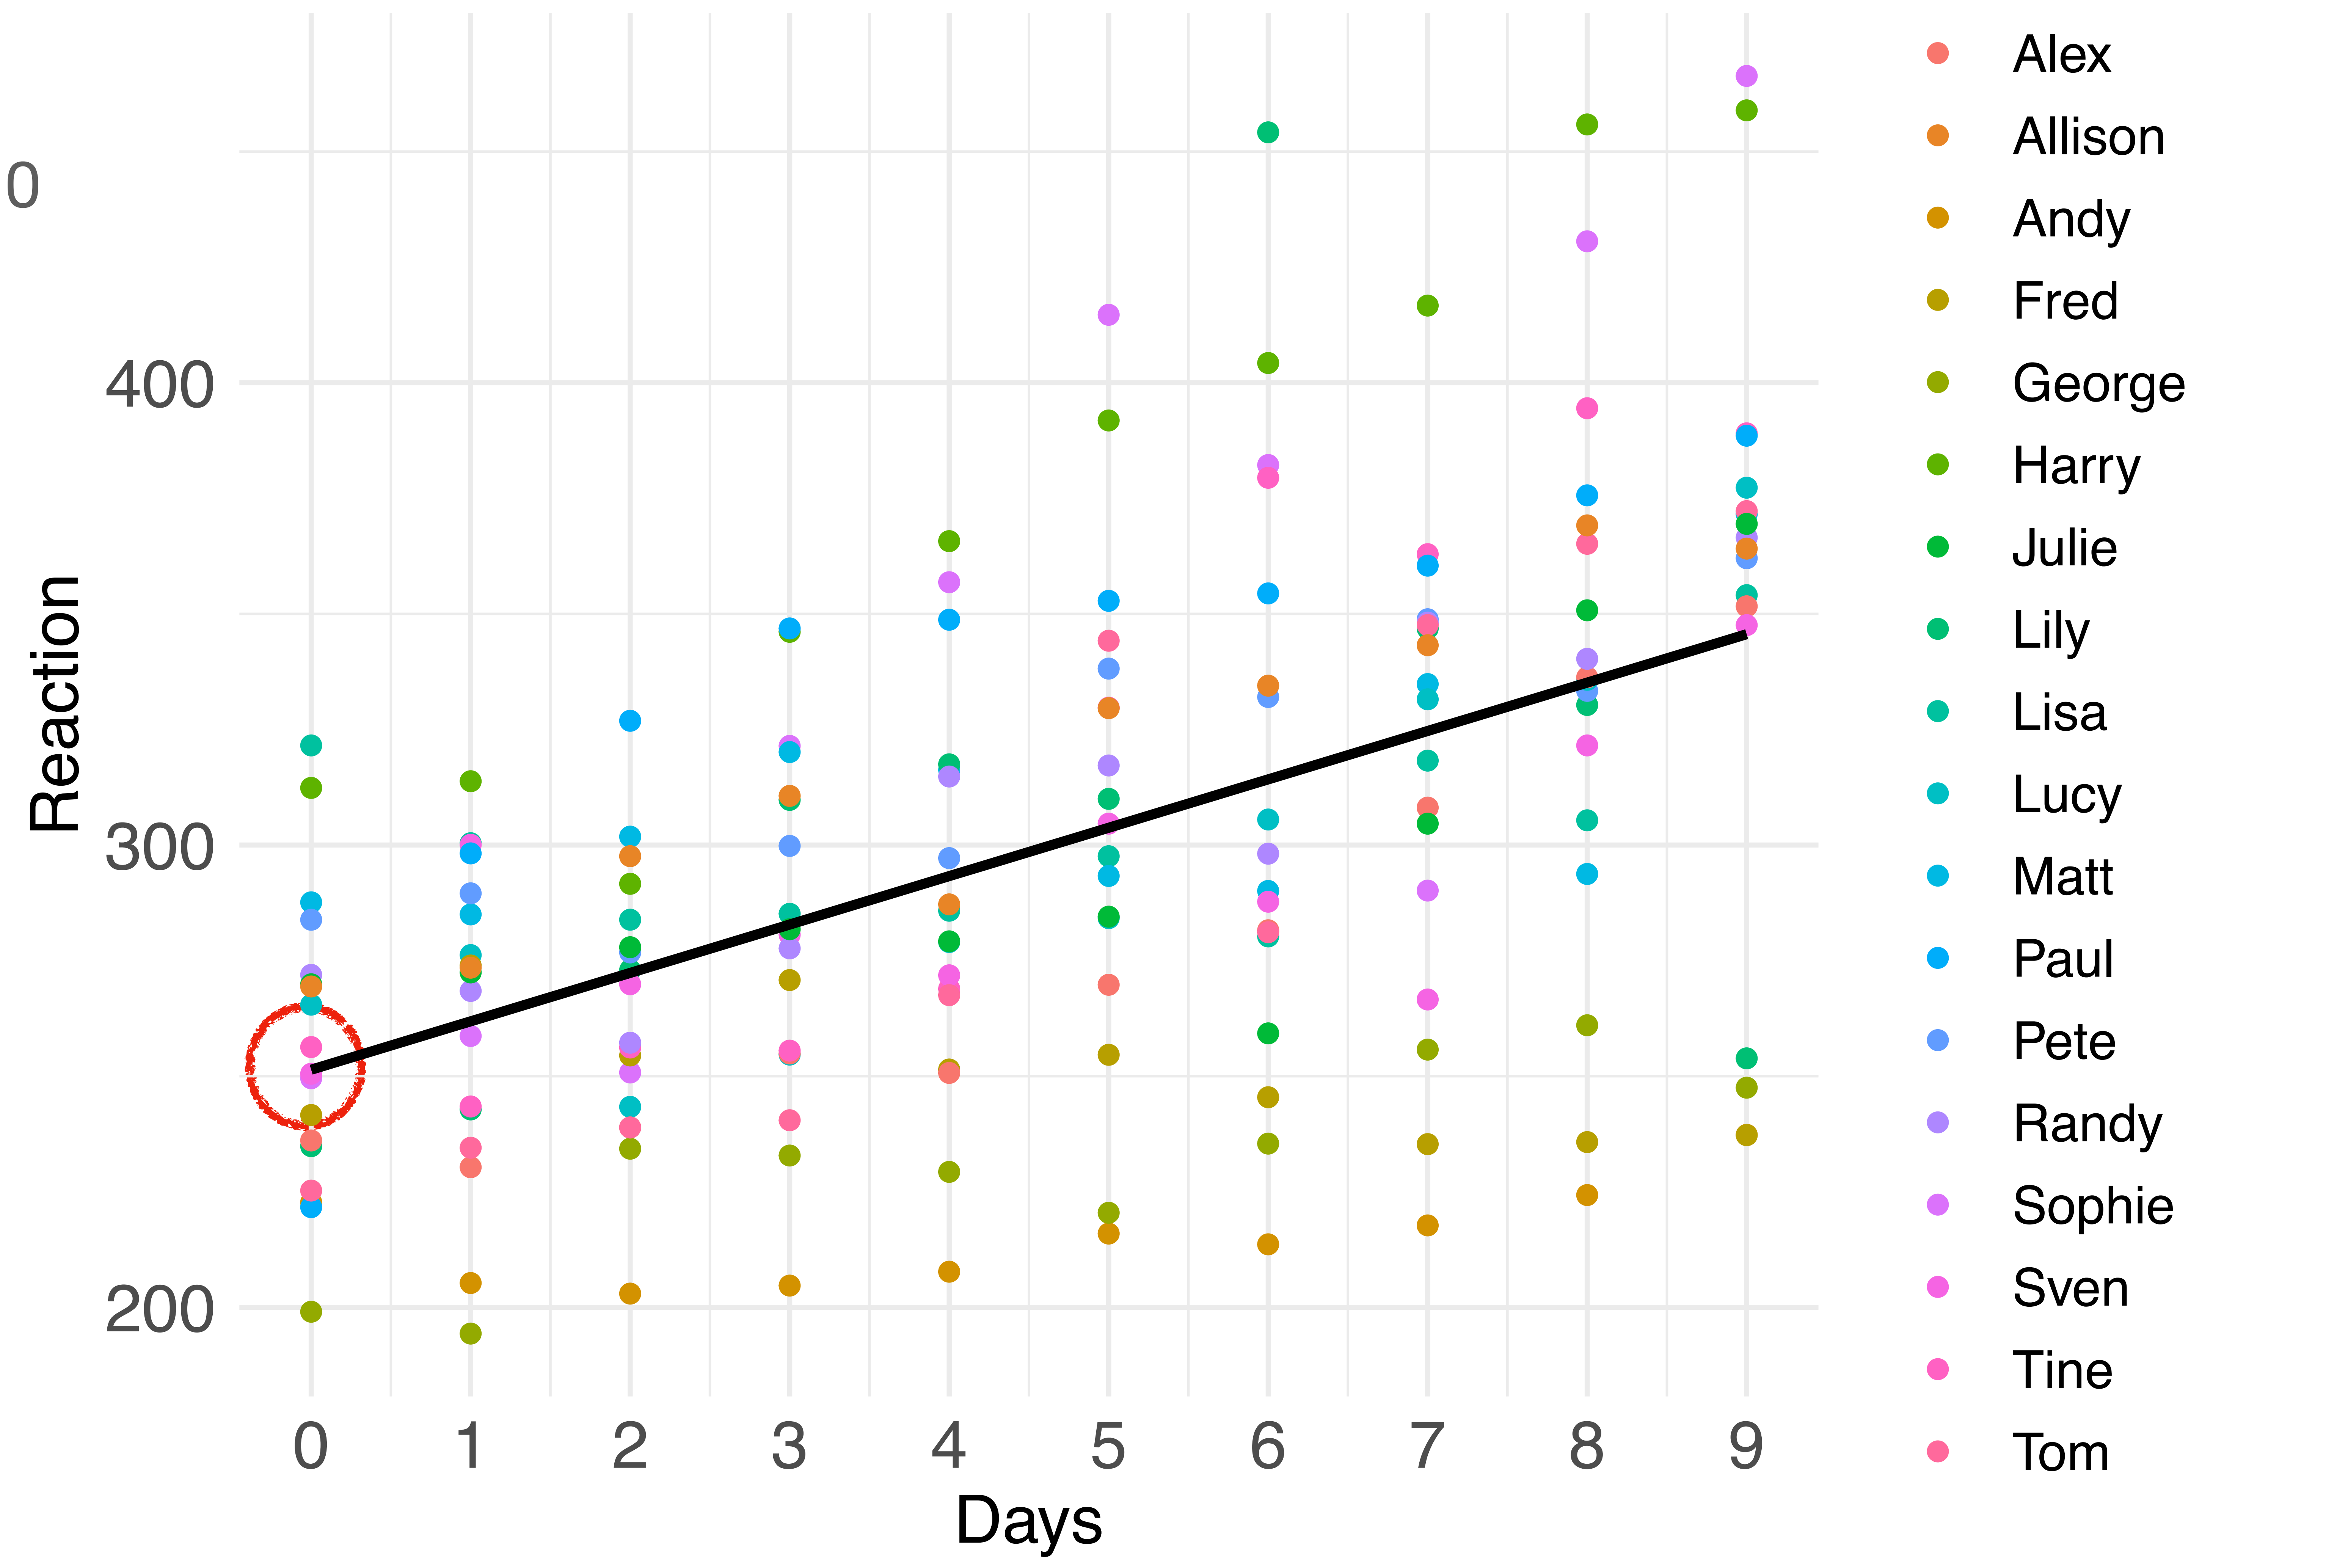

$y = \alpha + \beta \cdot day$

```
fit_1 <- lm(Reaction ~ 1 + Days, data = data)
```

Fixed effect of days (of deprivation) on reaction time = the slope
Fixed = Population level (applied for all together)

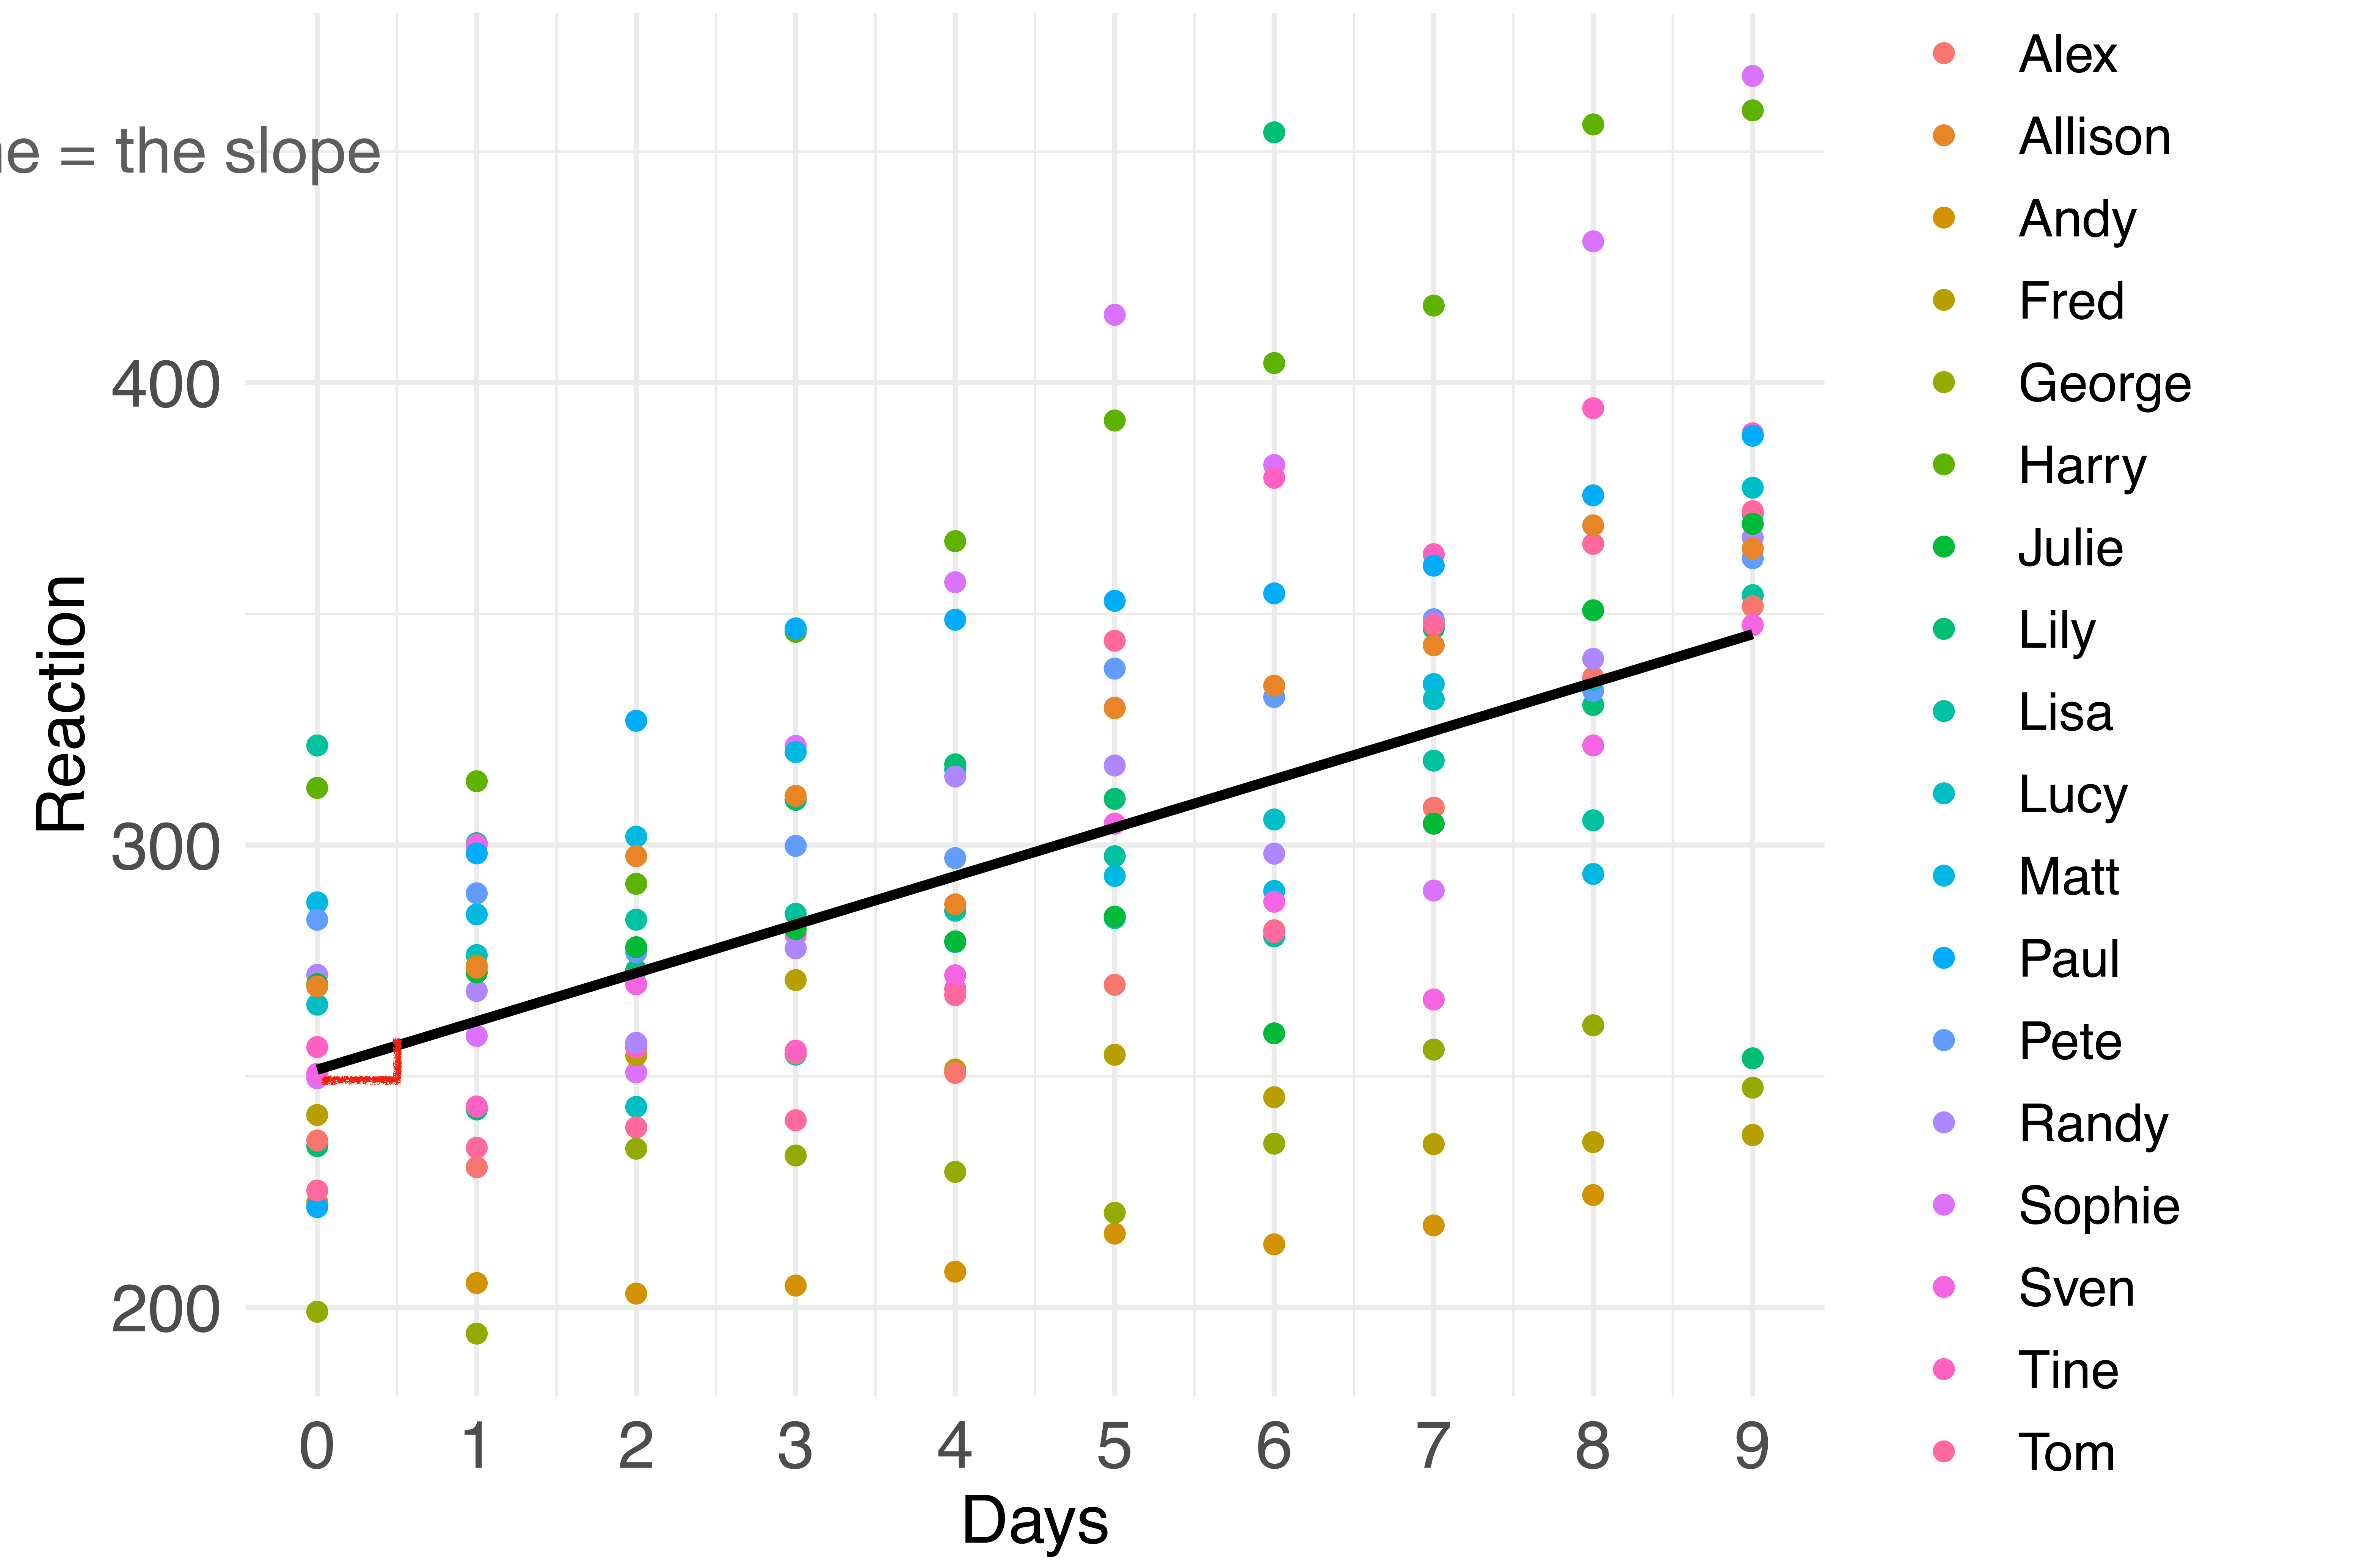

$y = \alpha + \beta \cdot day$

```
fit_1 <- lm(Reaction ~ 1 + Days, data = data)
```

The reaction times of the subjects each day

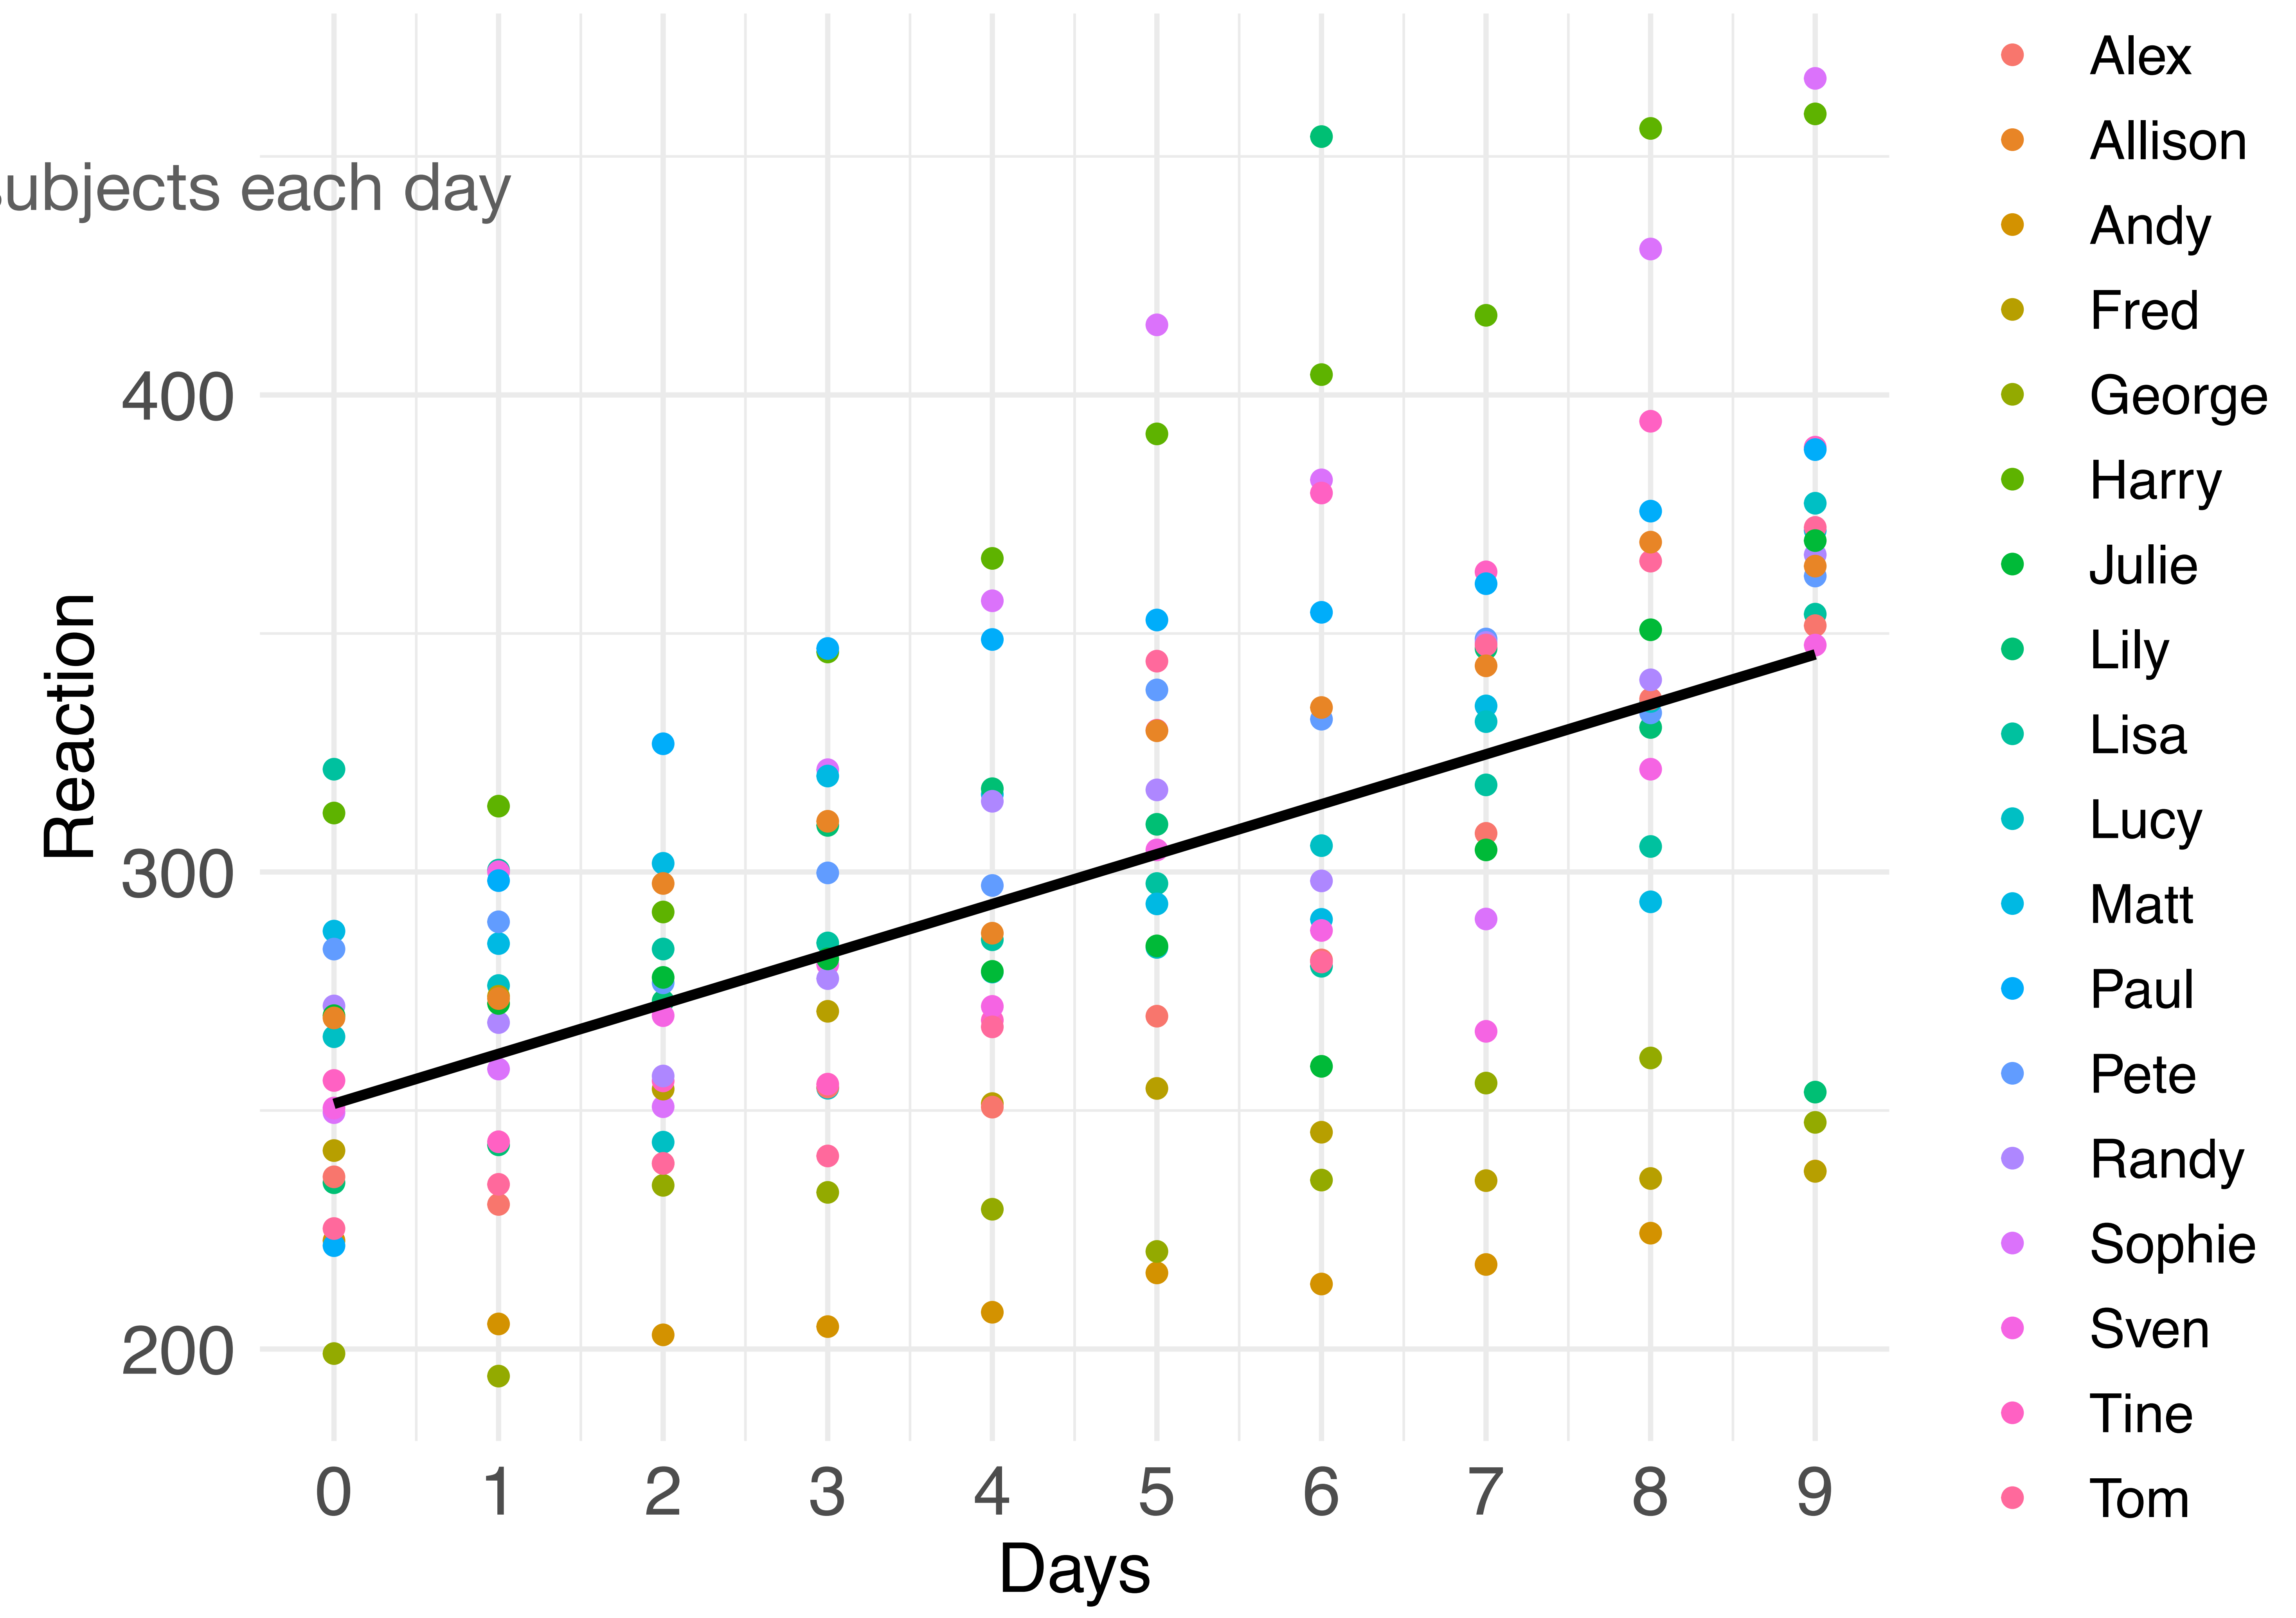

$y = \alpha + \beta \cdot day$

```
fit_1 <- lm(Reaction ~ 1 + Days, data = data)

summary(fit_1)

Coefficients:
             Estimate Std. Error t value Pr(>|t|)
(Intercept)   251.405      6.610  38.033  < 2e-16 ***
Days           10.467      1.238   8.454 9.89e-15 ***
```

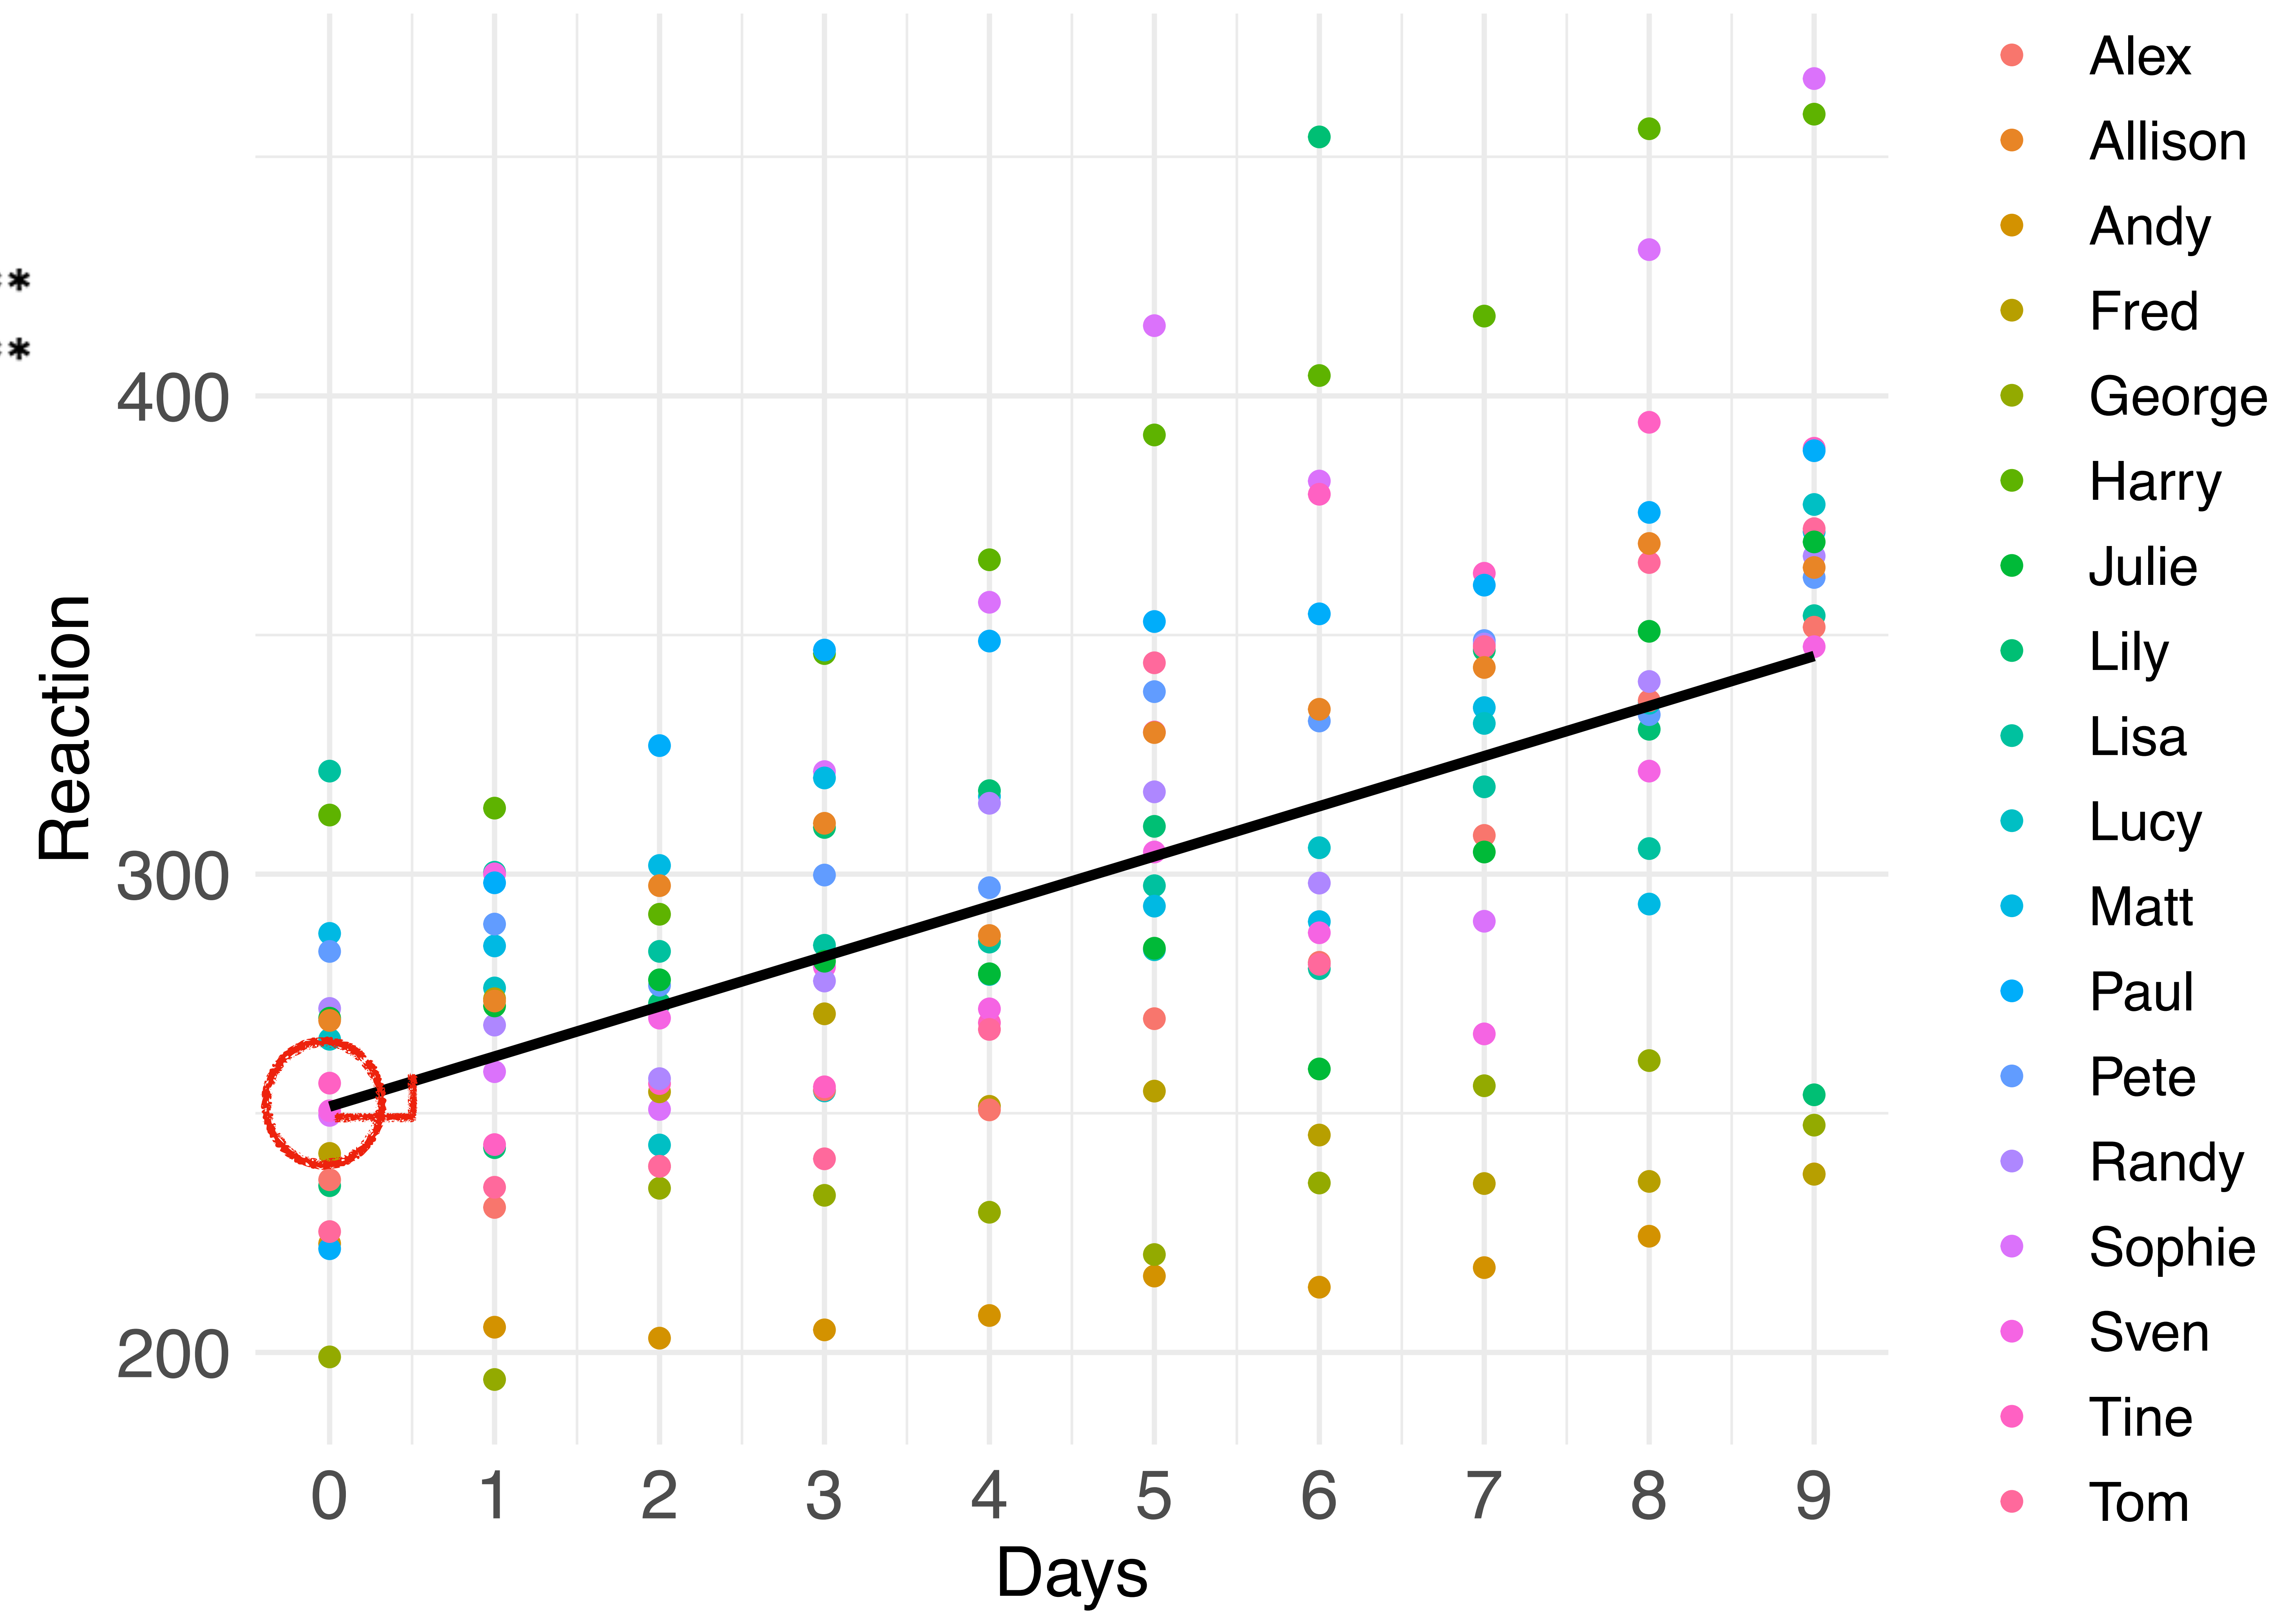

$y = \alpha + \beta \cdot day$

fit_1 <- lm(Reaction ~ 1 + Days, data = data)

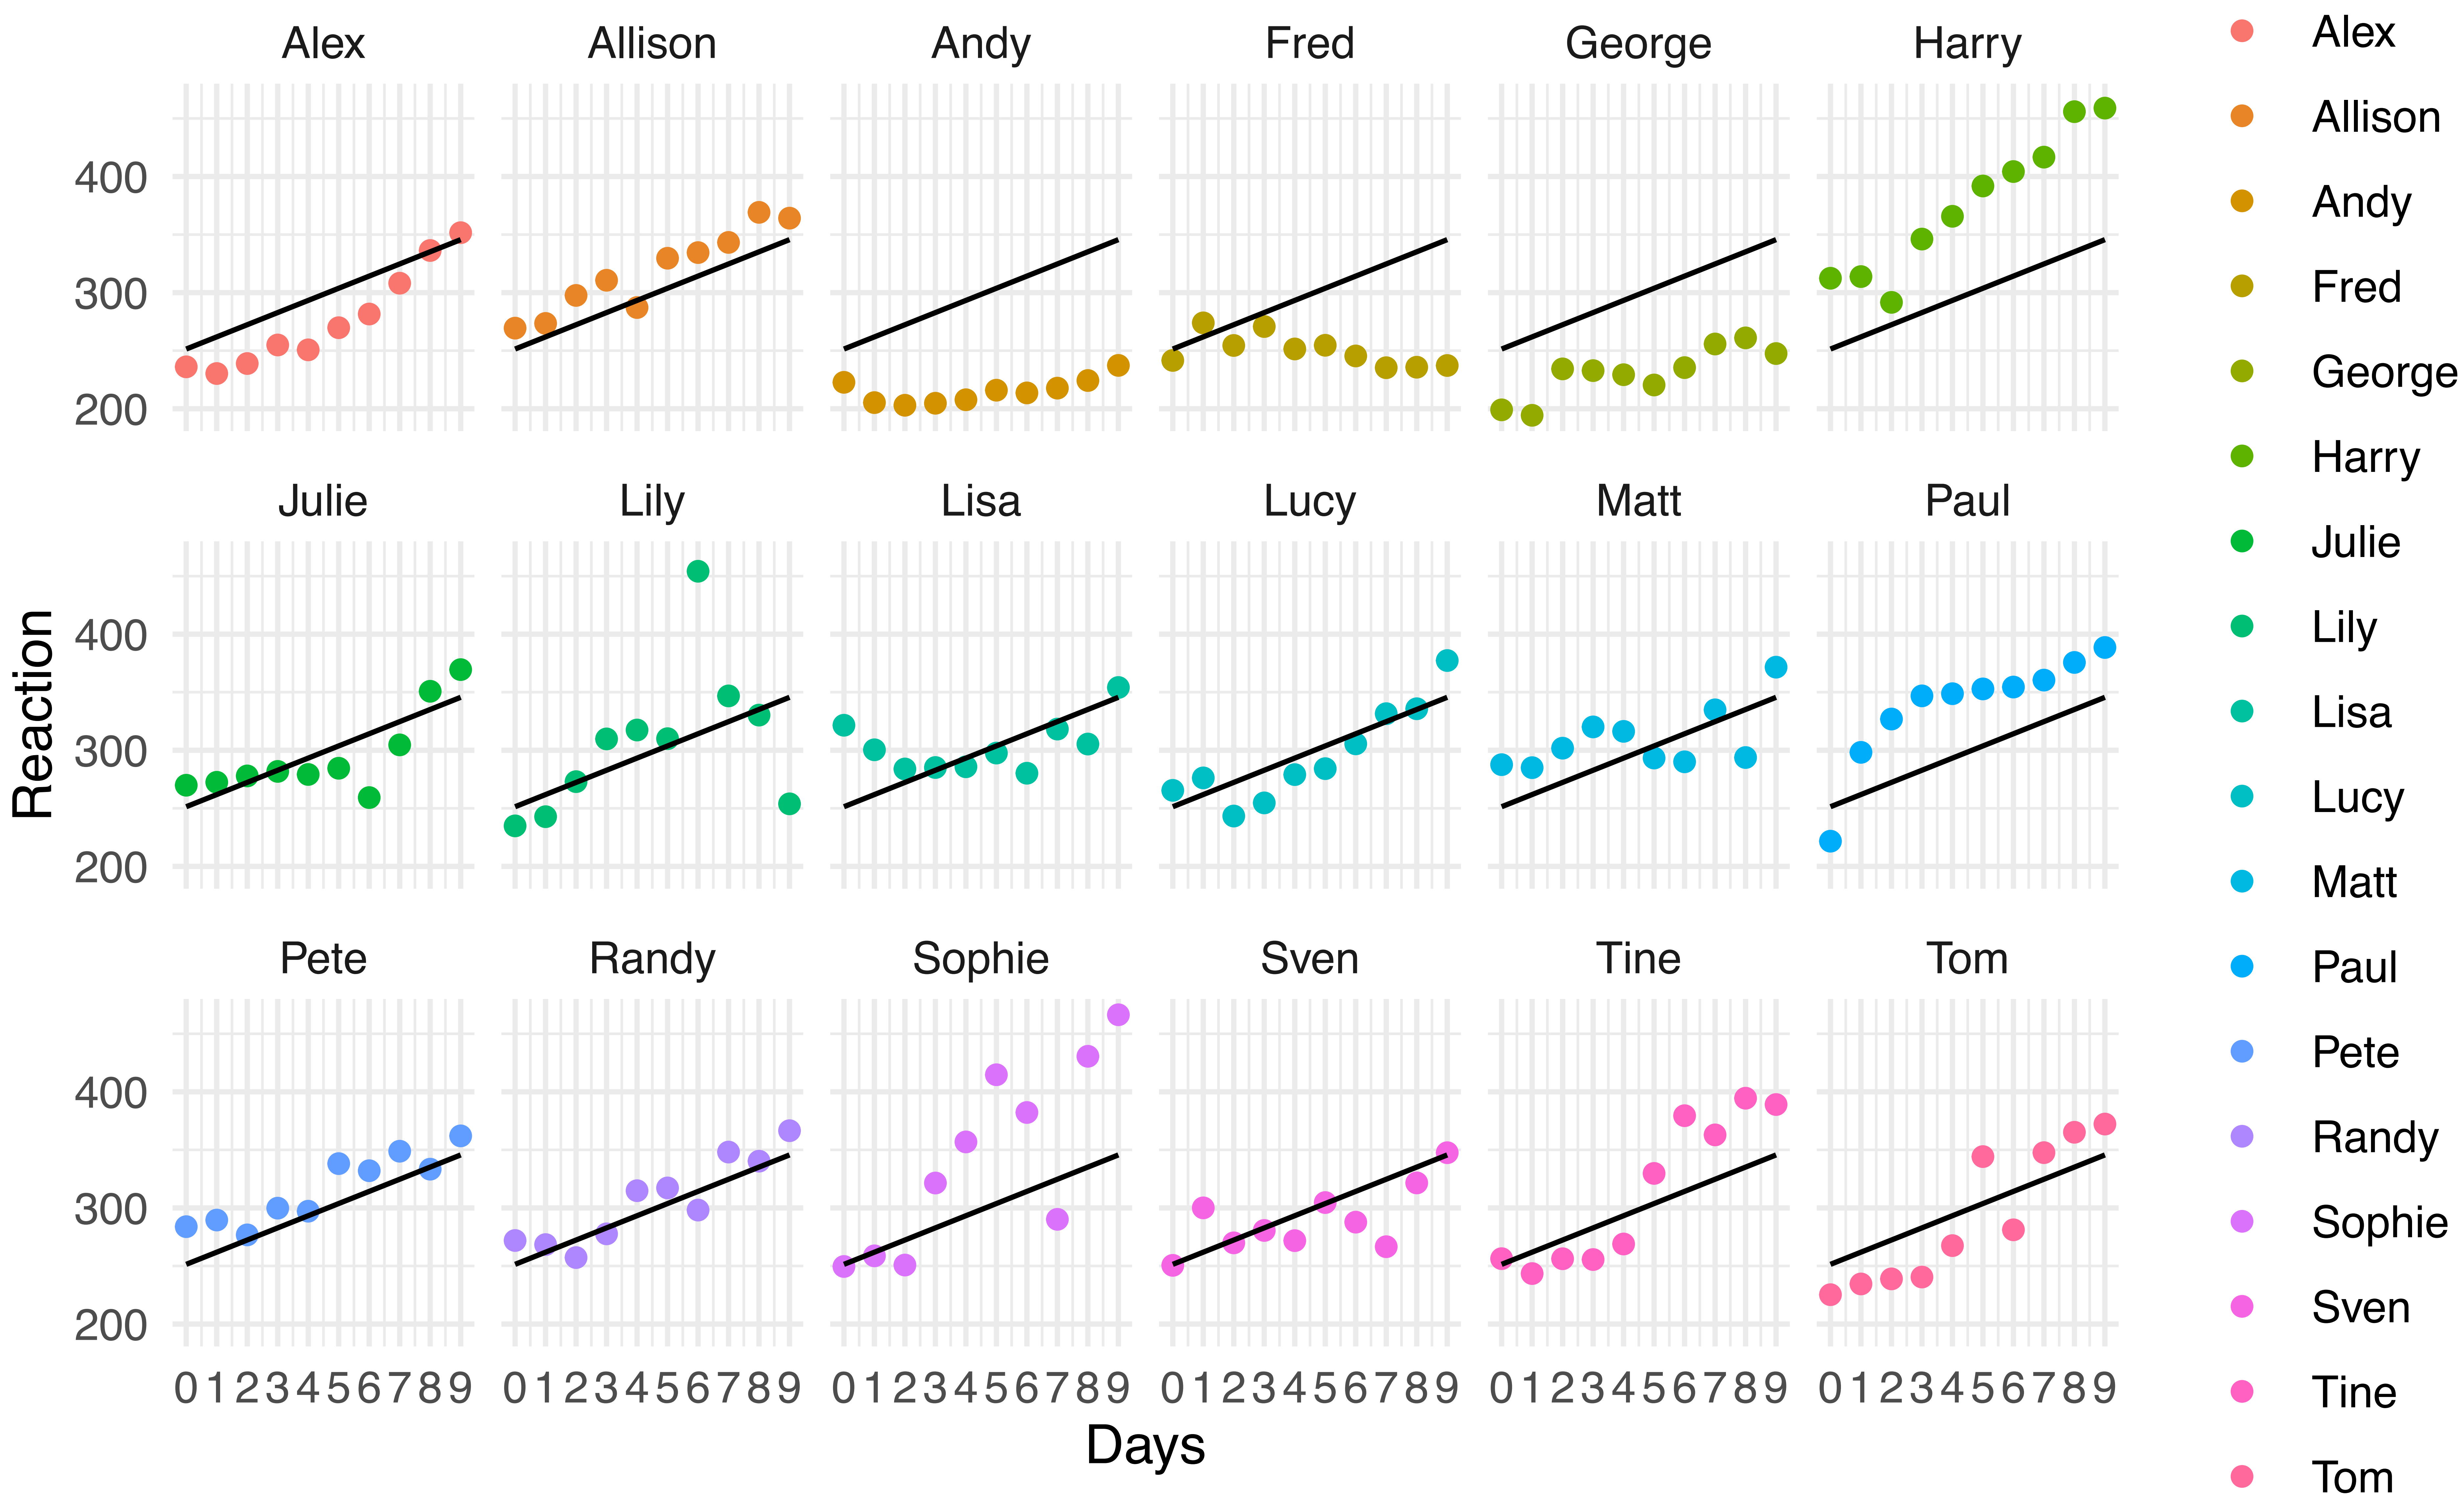

$y = \alpha + \beta \cdot day$

```
fit_1 <- lm(Reaction ~ 1 + Days, data = data)
```

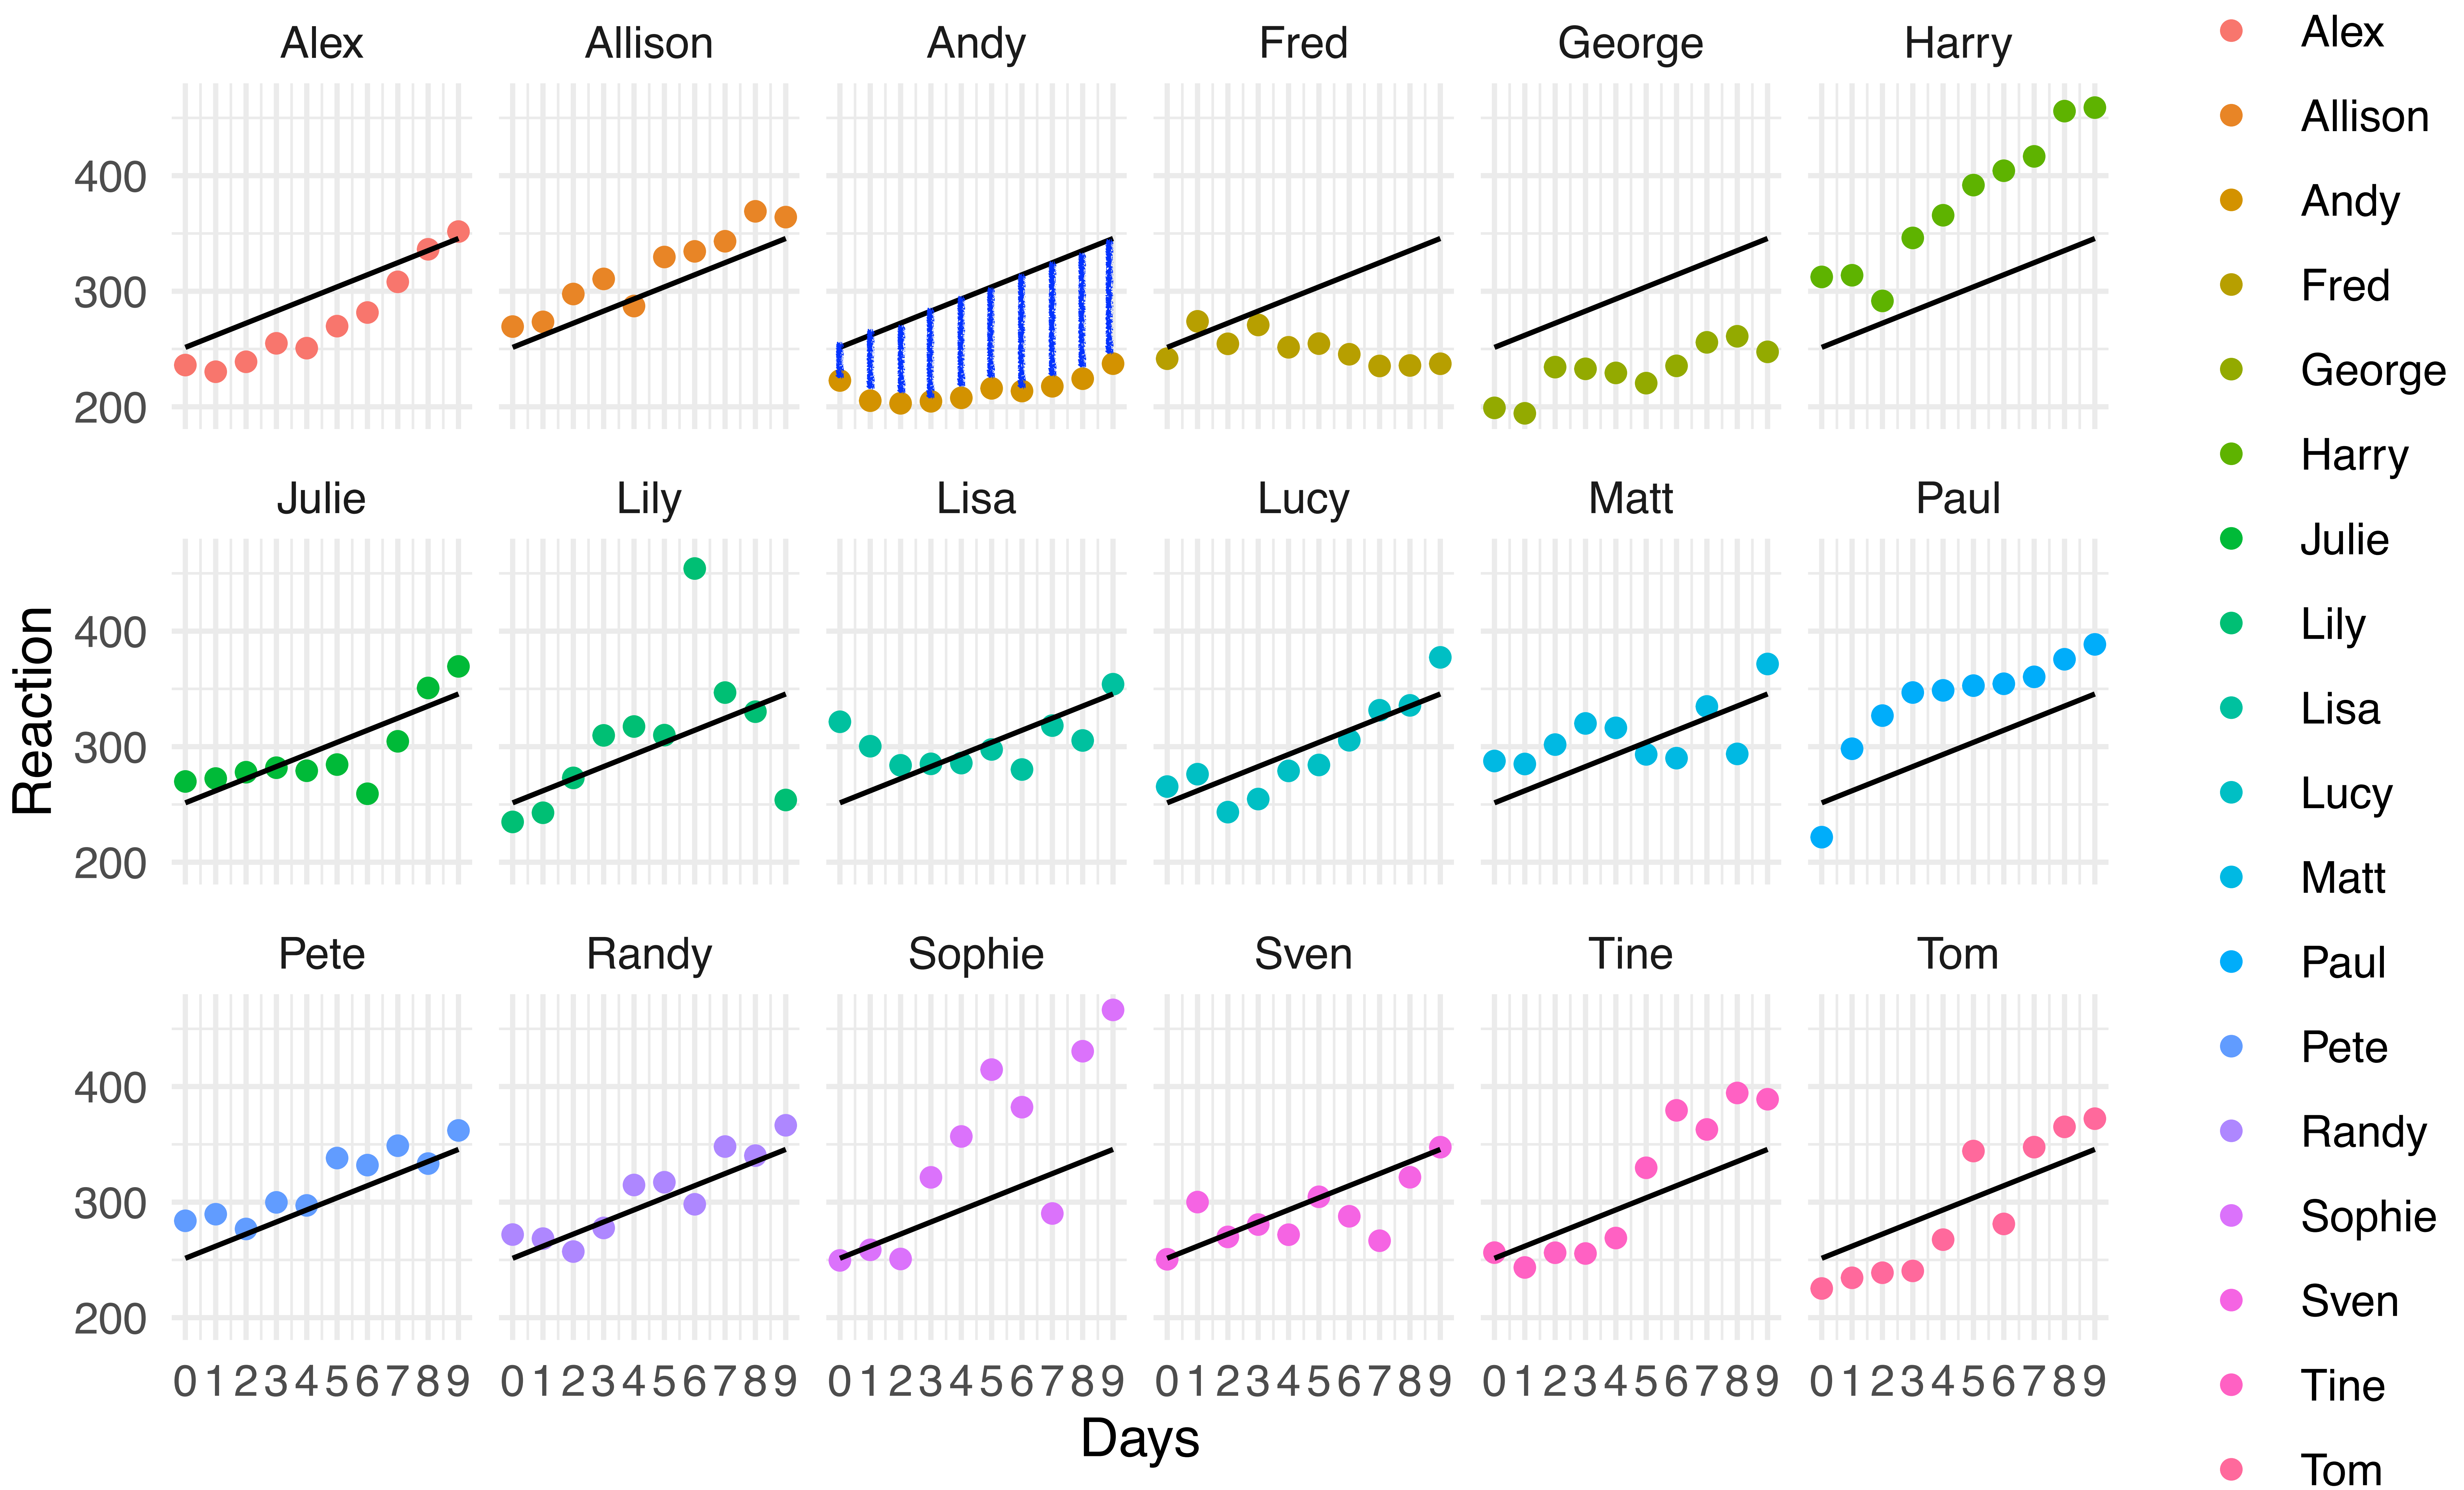

$$y = (\alpha + b_{\alpha,sub}) + \beta \cdot day$$

```
fit_1 <- lm(Reaction ~ 1 + Days, data = data)
fit_2 <- lmer(Reaction ~ 1 + Days + (1 | Subject), data = data)
```

R package 'lme4'

$$y = (\alpha + b_{\alpha,sub}) + \beta \cdot day$$

```
fit_1 <- lm(Reaction ~ 1 + Days, data = data)
fit_2 <- lmer(Reaction ~ 1 + Days + (1 | Subject), data = data)
```

Random intercept= subject specific intercept

$y = (\alpha + b_{\alpha,sub}) + \beta \cdot day$

```
fit_1 <- lm(Reaction ~ 1 + Days, data = data)
fit_2 <- lmer(Reaction ~ 1 + Days + (1 | Subject), data = data)
```

Random intercept= subject specific intercept

Subjects are a random sample of all the possible subjects we could have chosen
The subject-specific intercept is a random variable

$y = (\alpha + b_{\alpha,sub}) + \beta \cdot day$

fit_1 <- lm(Reaction ~ 1 + Days, data = data)

fit_2 <- lmer(Reaction ~ 1 + Days + (1 | Subject), data = data)

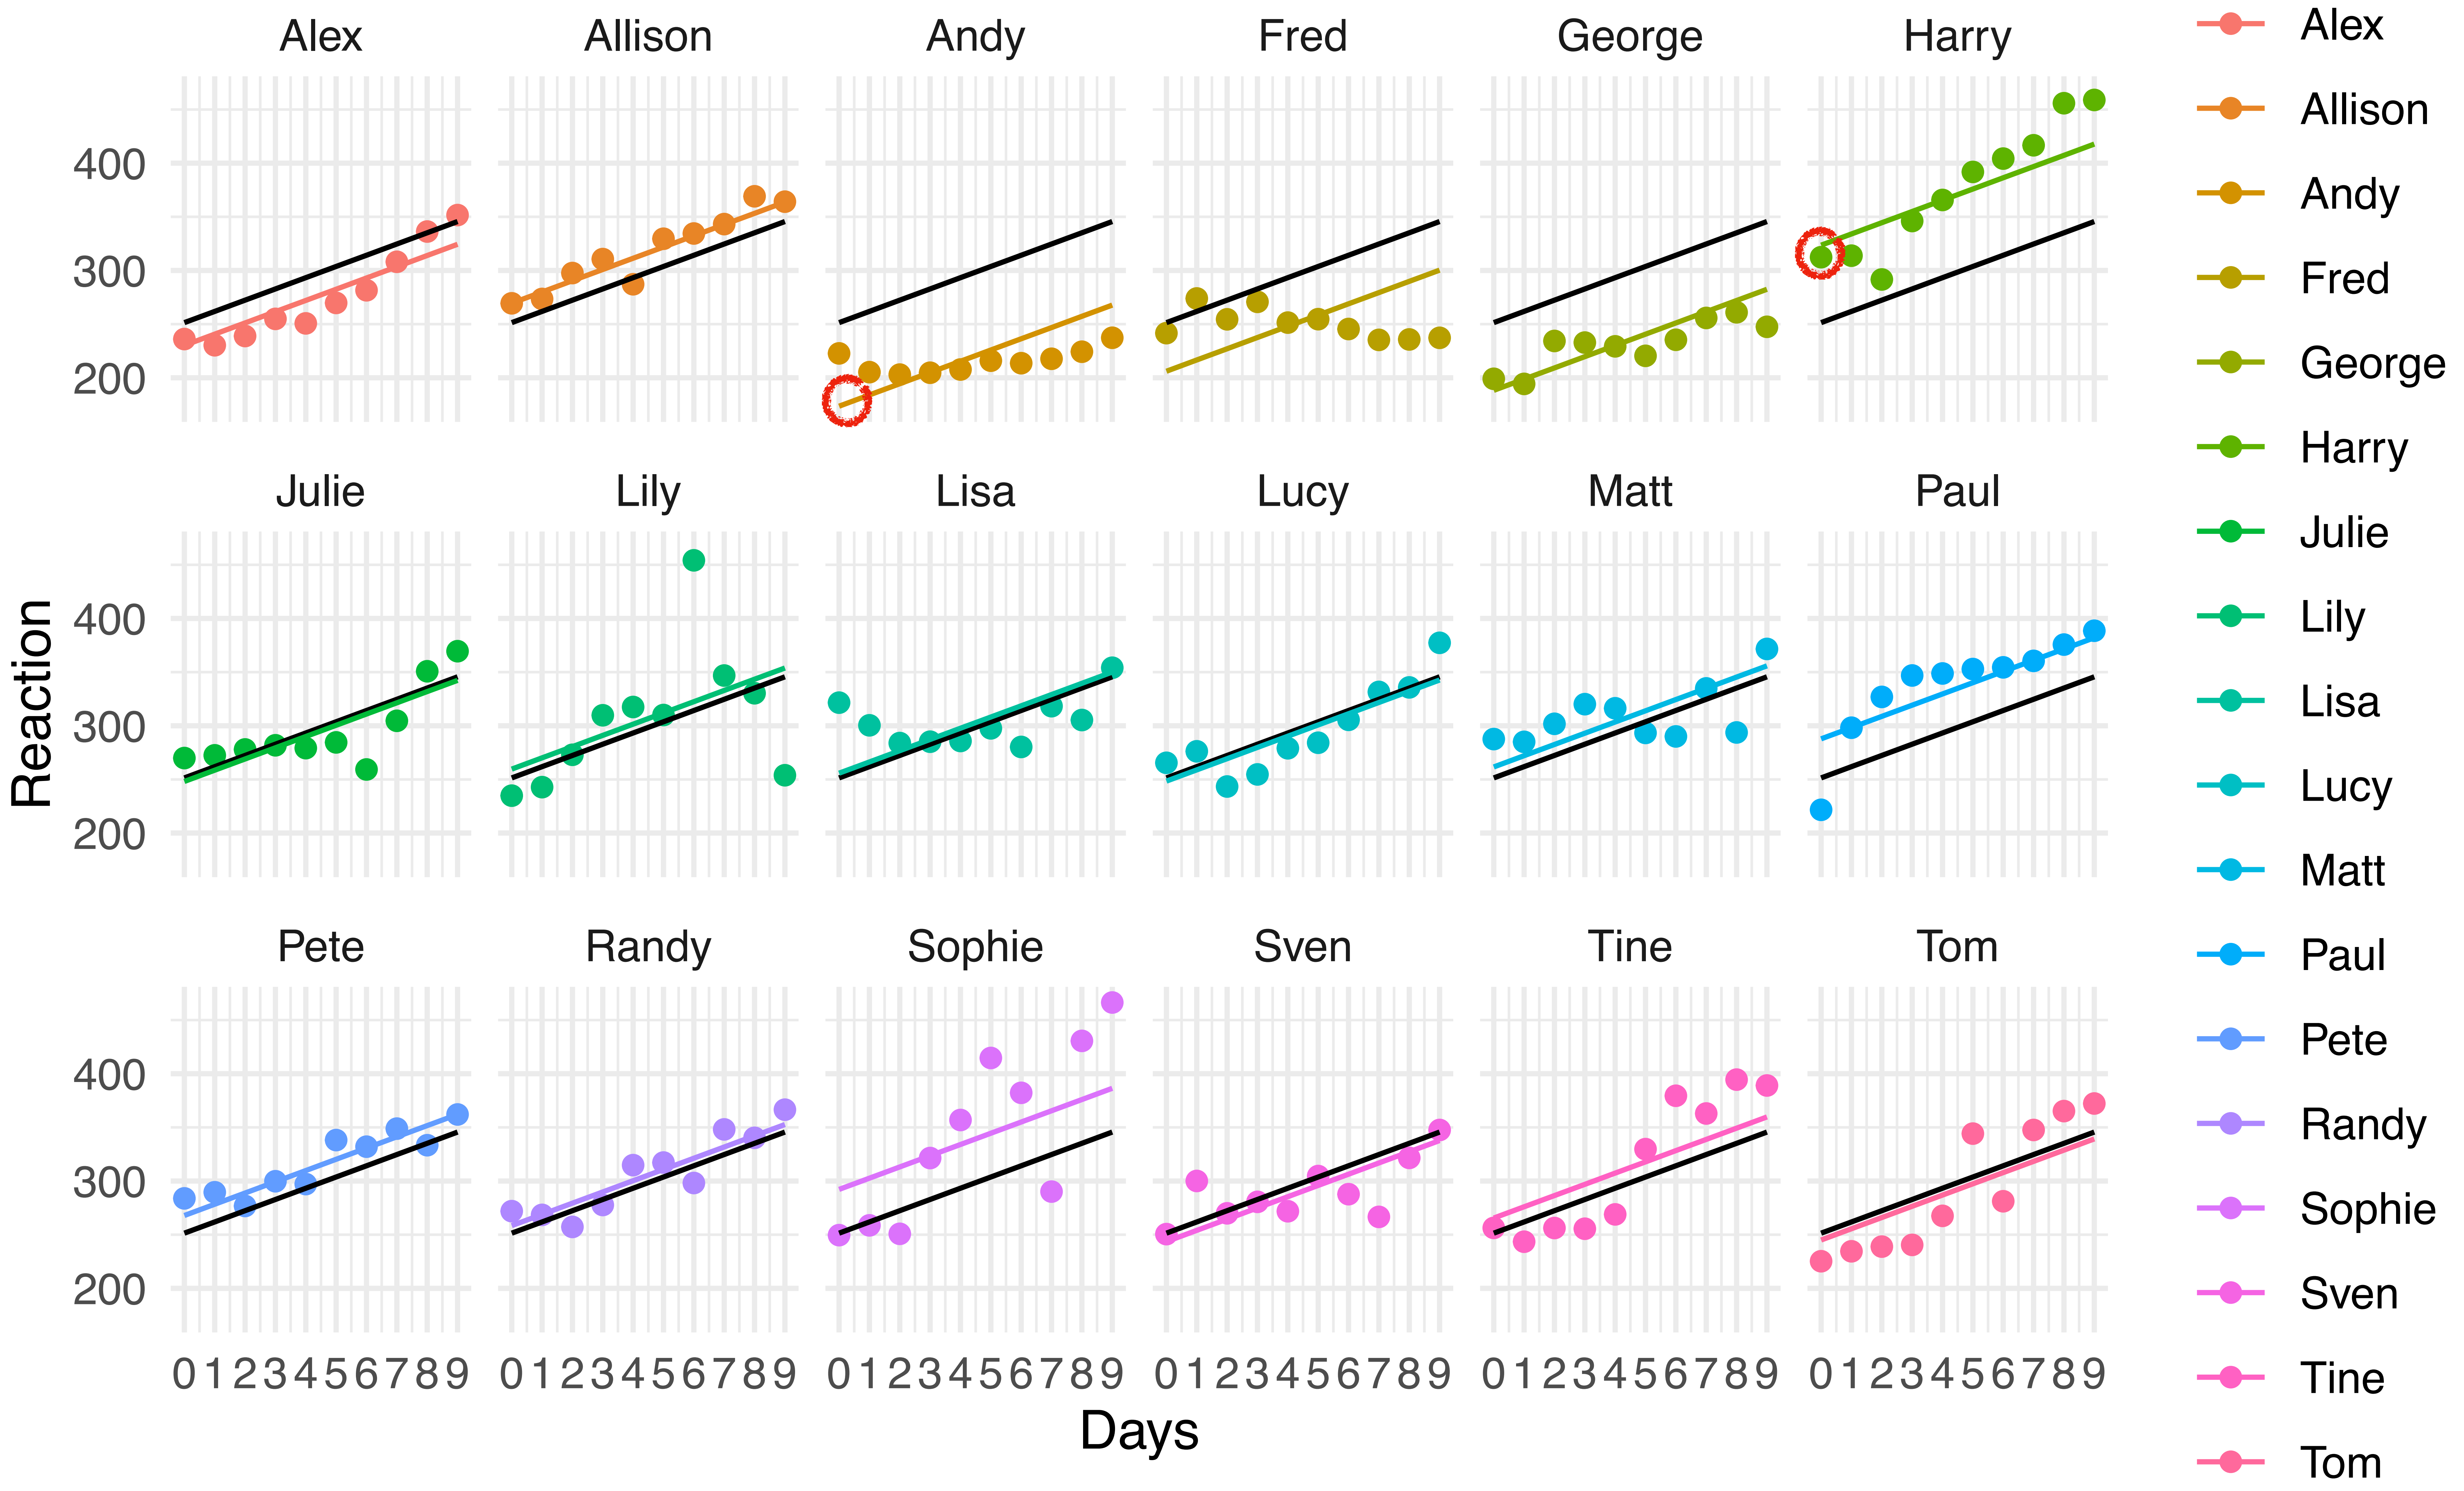

$$y = \alpha + (\beta + b_{\beta,sub}) \cdot day$$

```
fit_1 <- lm(Reaction ~ 1 + Days, data = data)
fit_2 <- lmer(Reaction ~ 1 + Days + (1 | Subject), data = data)
fit_3 <- lmer(Reaction ~ 1 + Days + (0 + Days | Subject) , data = data)
```

R package 'lme4'

$$y = \alpha + (\beta + b_{\beta,sub}) \cdot day$$

```
fit_1 <- lm(Reaction ~ 1 + Days, data = data)
fit_2 <- lmer(Reaction ~ 1 + Days + (1 | Subject), data = data)
fit_3 <- lmer(Reaction ~ 1 + Days + (0 + Days | Subject) , data = data)
```

Random slope = subject specific effect of days (of deprivation)

$$y = \alpha + (\beta + b_{\beta,sub}) \cdot day$$

```
fit_1 <- lm(Reaction ~ 1 + Days, data = data)
fit_2 <- lmer(Reaction ~ 1 + Days + (1 | Subject), data = data)
fit_3 <- lmer(Reaction ~ 1 + Days + (0 + Days | Subject) , data = data)
```

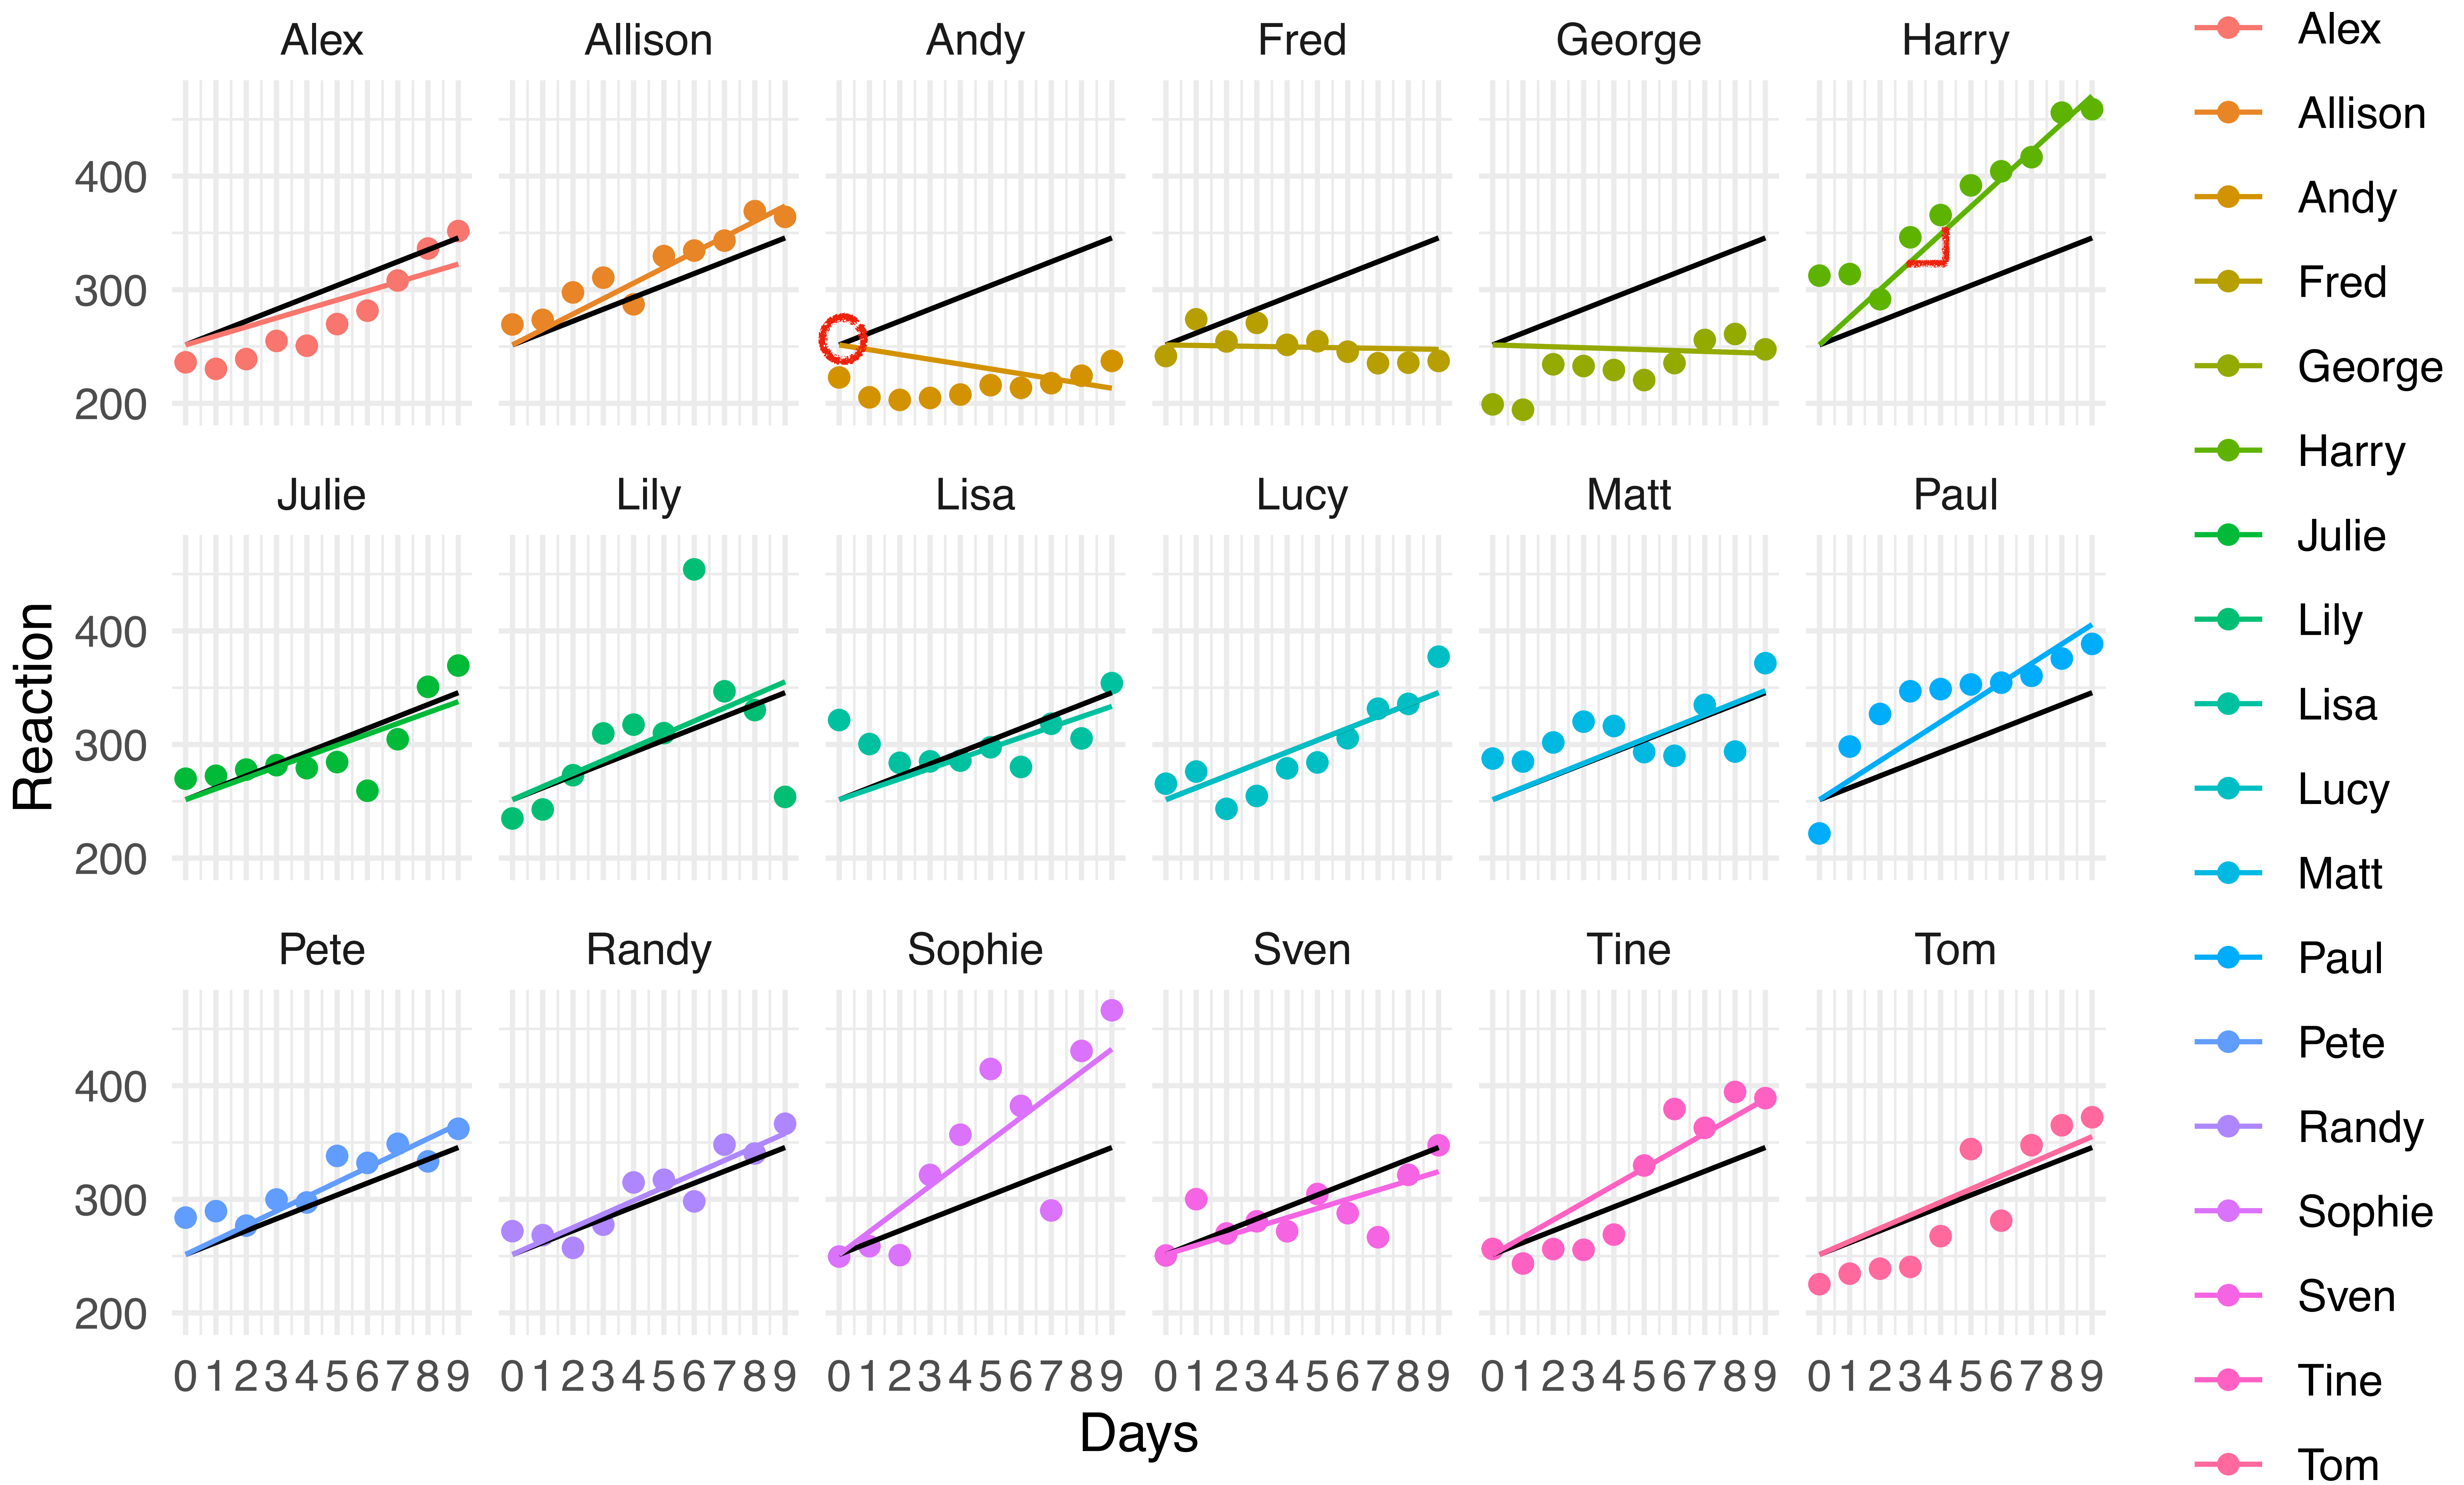

$$y = (\alpha + b_{\alpha,sub}) + (\beta + b_{\beta,sub}) \cdot day$$

```
fit_1 <- lm(Reaction ~ 1 + Days, data = data)
fit_2 <- lmer(Reaction ~ 1 + Days + (1 | Subject), data = data)
fit_3 <- lmer(Reaction ~ 1 + Days + (0 + Days | Subject) , data = data)
fit_4 <- lmer(Reaction ~ 1 + Days + (1 + Days | Subject) , data = data)
```

R package 'lme4'

$$y = (\alpha + b_{\alpha,sub}) + (\beta + b_{\beta,sub}) \cdot day$$

```
fit_1 <- lm(Reaction ~ 1 + Days, data = data)
fit_2 <- lmer(Reaction ~ 1 + Days + (1 | Subject), data = data)
fit_3 <- lmer(Reaction ~ 1 + Days + (0 + Days | Subject) , data = data)
fit_4 <- lmer(Reaction ~ 1 + Days + (1 + Days | Subject) , data = data)
```

Random intercept and slope = subject specific intercept and effect of days (of deprivation)

$$y = (\alpha + b_{\alpha,sub}) + (\beta + b_{\beta,sub}) \cdot day$$

```
fit_1 <- lm(Reaction ~ 1 + Days, data = data)
fit_2 <- lmer(Reaction ~ 1 + Days + (1 | Subject), data = data)
fit_3 <- lmer(Reaction ~ 1 + Days + (0 + Days | Subject) , data = data)
fit_4 <- lmer(Reaction ~ 1 + Days + (1 + Days | Subject) , data = data)
```

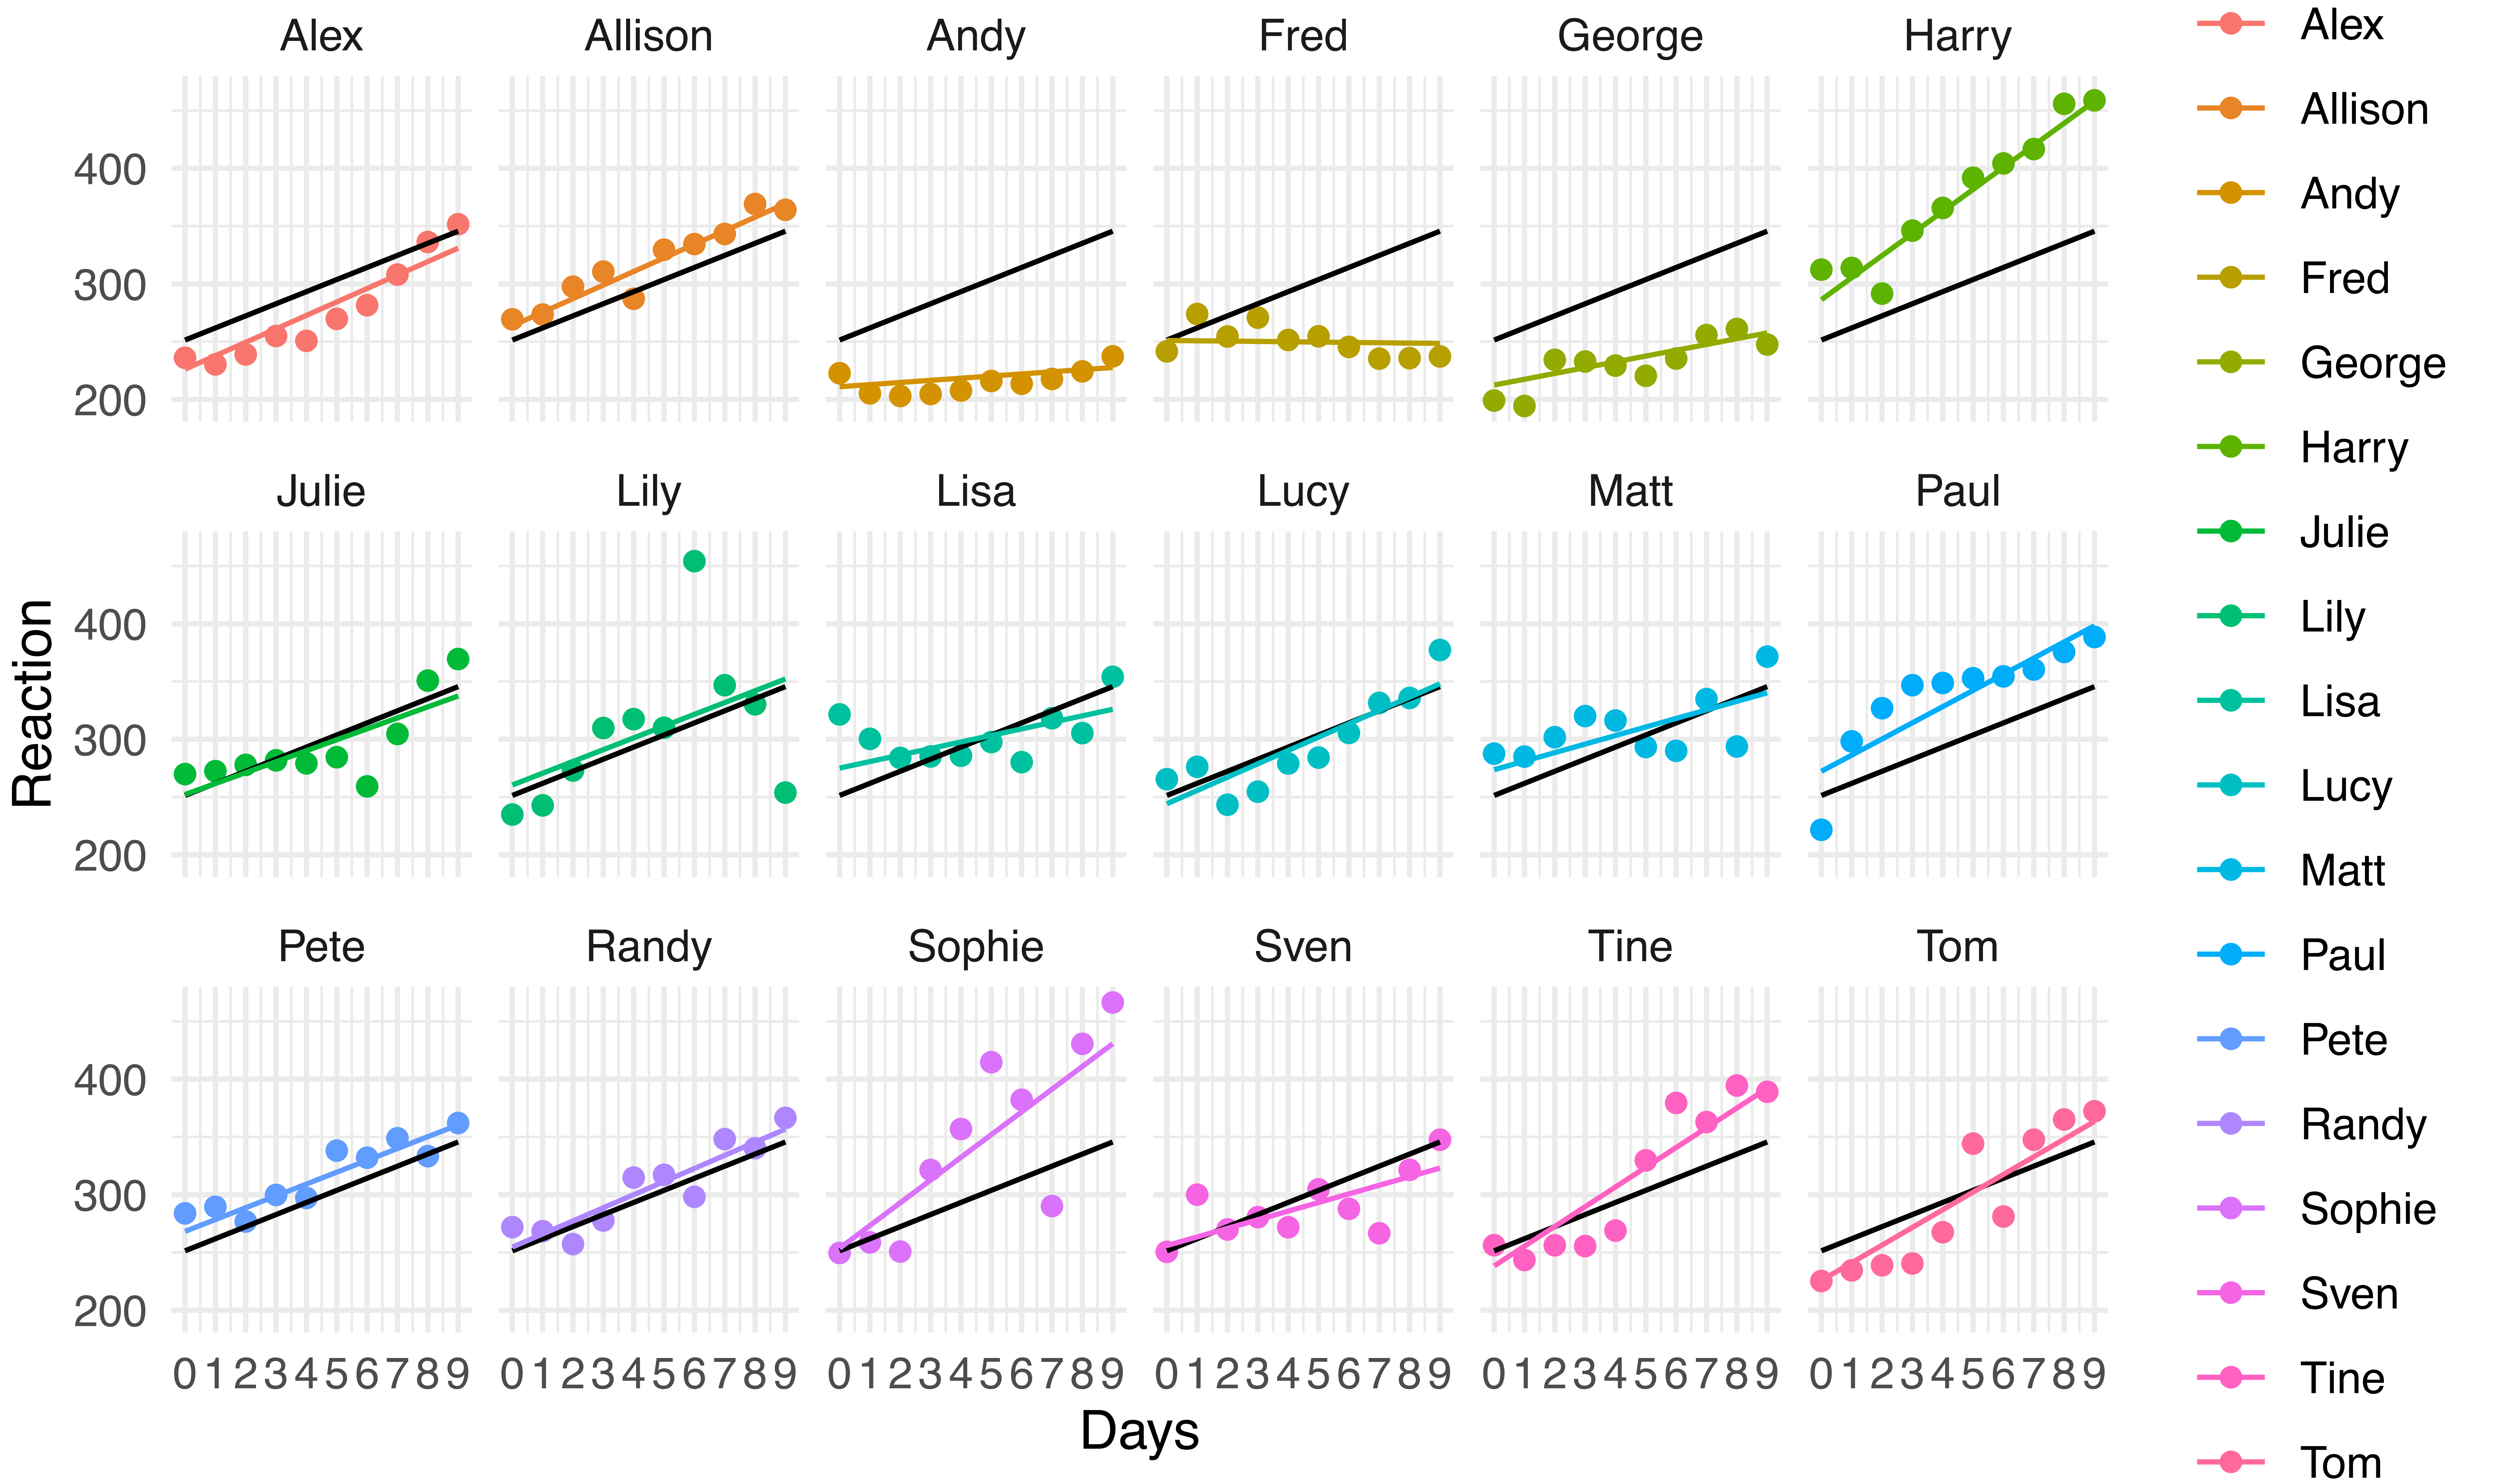

Reaction ~ 1 + Days: Linear model

Reaction ~ 1 + Days + (1 | Subject): Linear mixed effects model

Reaction ~ 1 + Days + (0 + Days | Subject): Linear mixed effects model

Reaction ~ 1 + Days + (1 + Days | Subject): Linear mixed effects model

Mixed effects (aka hierarchical) model: both fixed and random effects (intercept and/ or slope)

What is **generally** the effect of sleep deprivation on reaction time?

1) If we are interested in the population-level effect… What's the point?

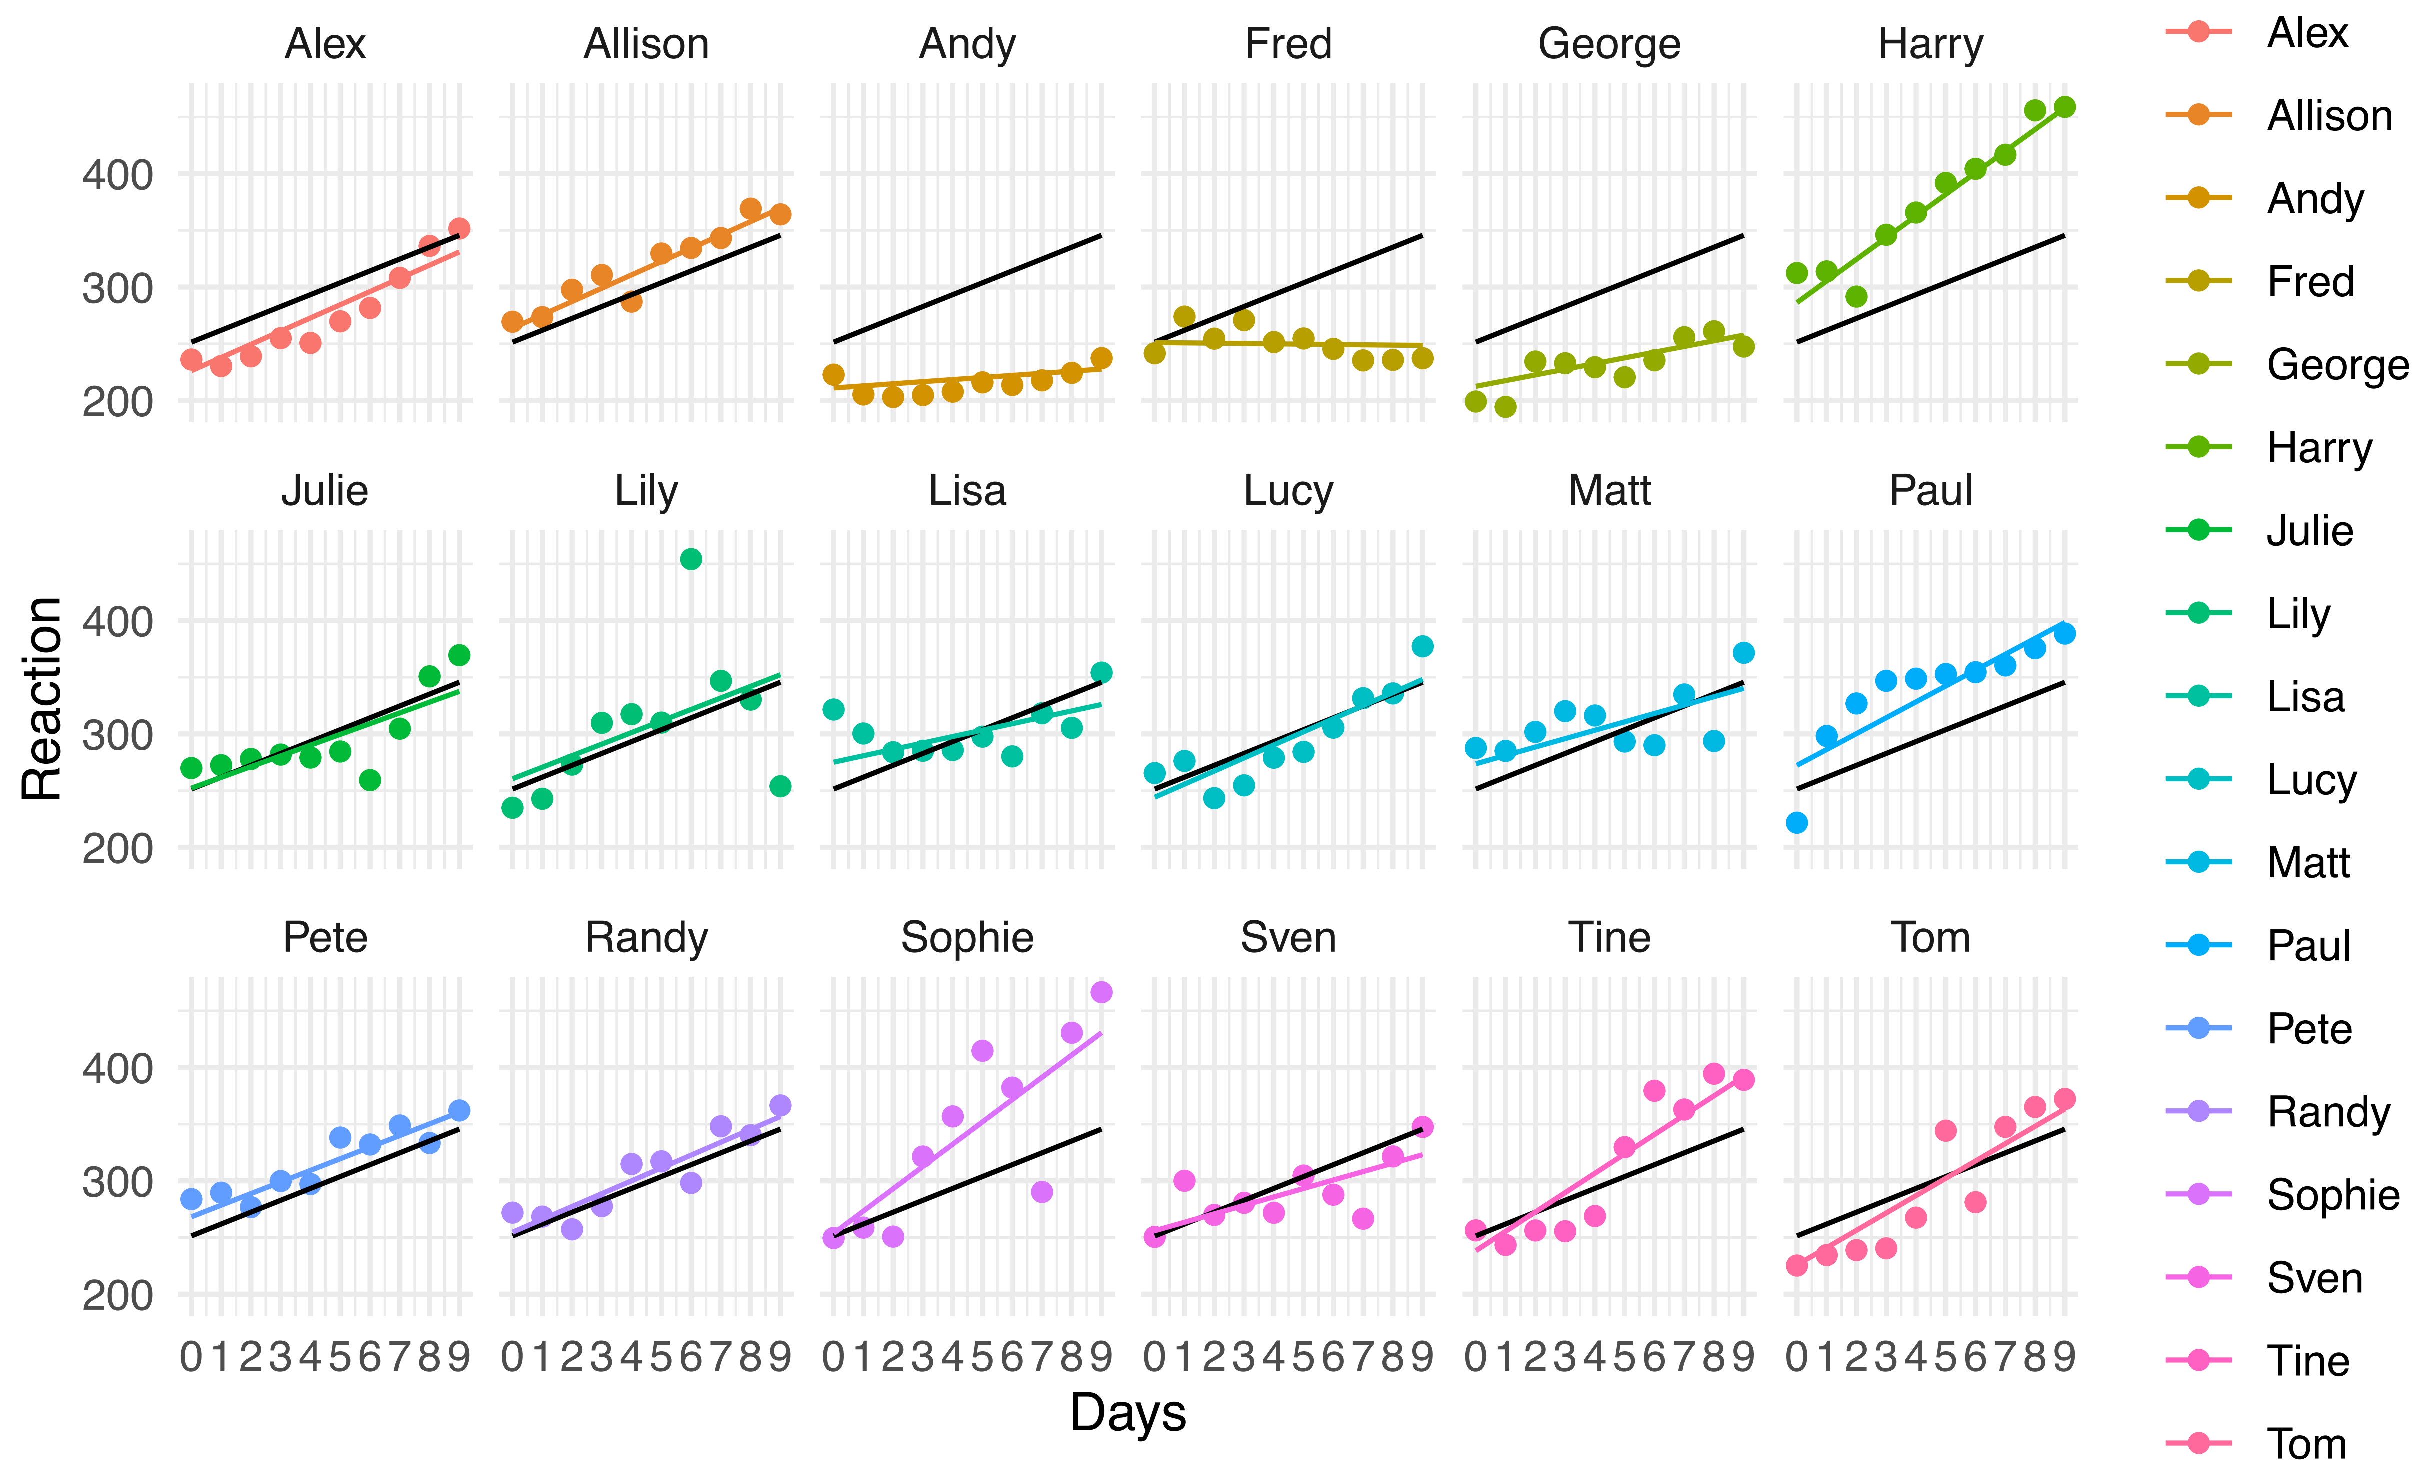

What is **generally** the effect of sleep deprivation on reaction time?

1) If we are interested in the population-level effect… What's the point?

We get an idea of how much people vary on average

summary(fit)

```
Random effects:
 Groups   Name        Variance Std.Dev.
 Subject  (Intercept) 612.10   24.741
          Days         35.07    5.922
 Residual             654.94   25.592
```

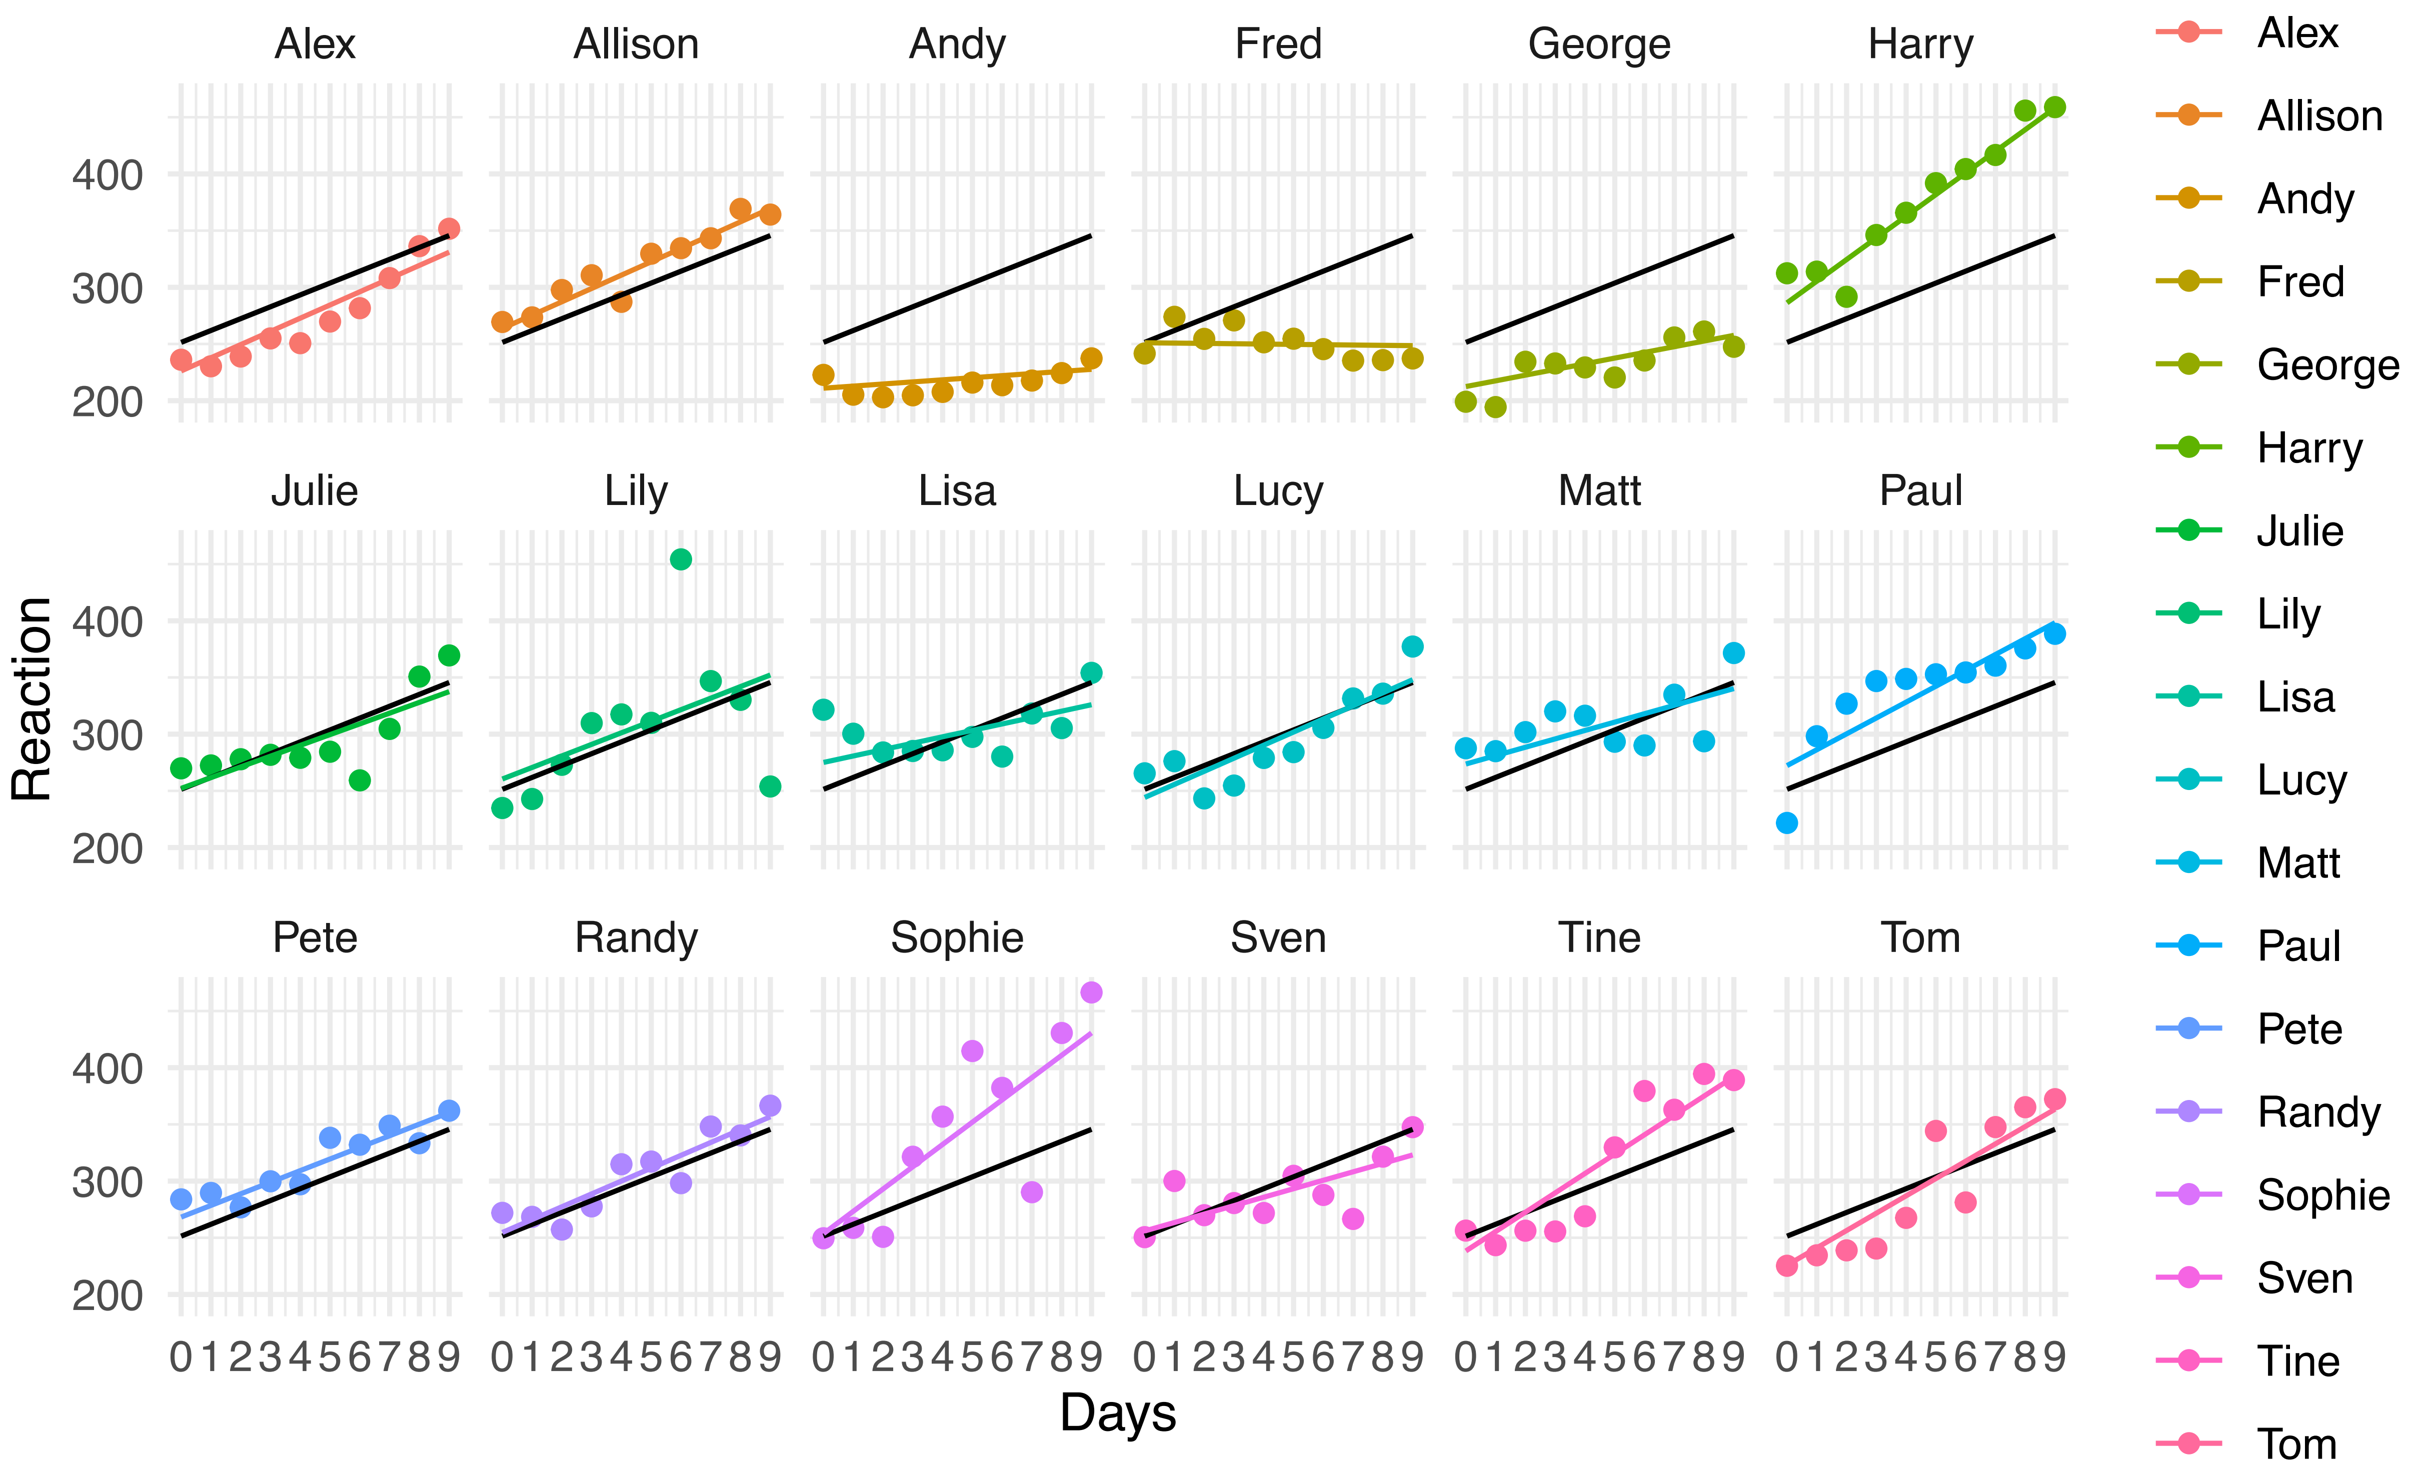

1) If we are interested in the population-level effect… What's the point?

We get an idea of how much people vary on average

summary(fit)

```
Random effects:
 Groups   Name        Variance Std.Dev.
 Subject  (Intercept) 612.10   24.741
          Days         35.07    5.922
 Residual             654.94   25.592
```

2) If we are interested in the subject-specific effects… Why not fixed interaction (instead of random effects)?

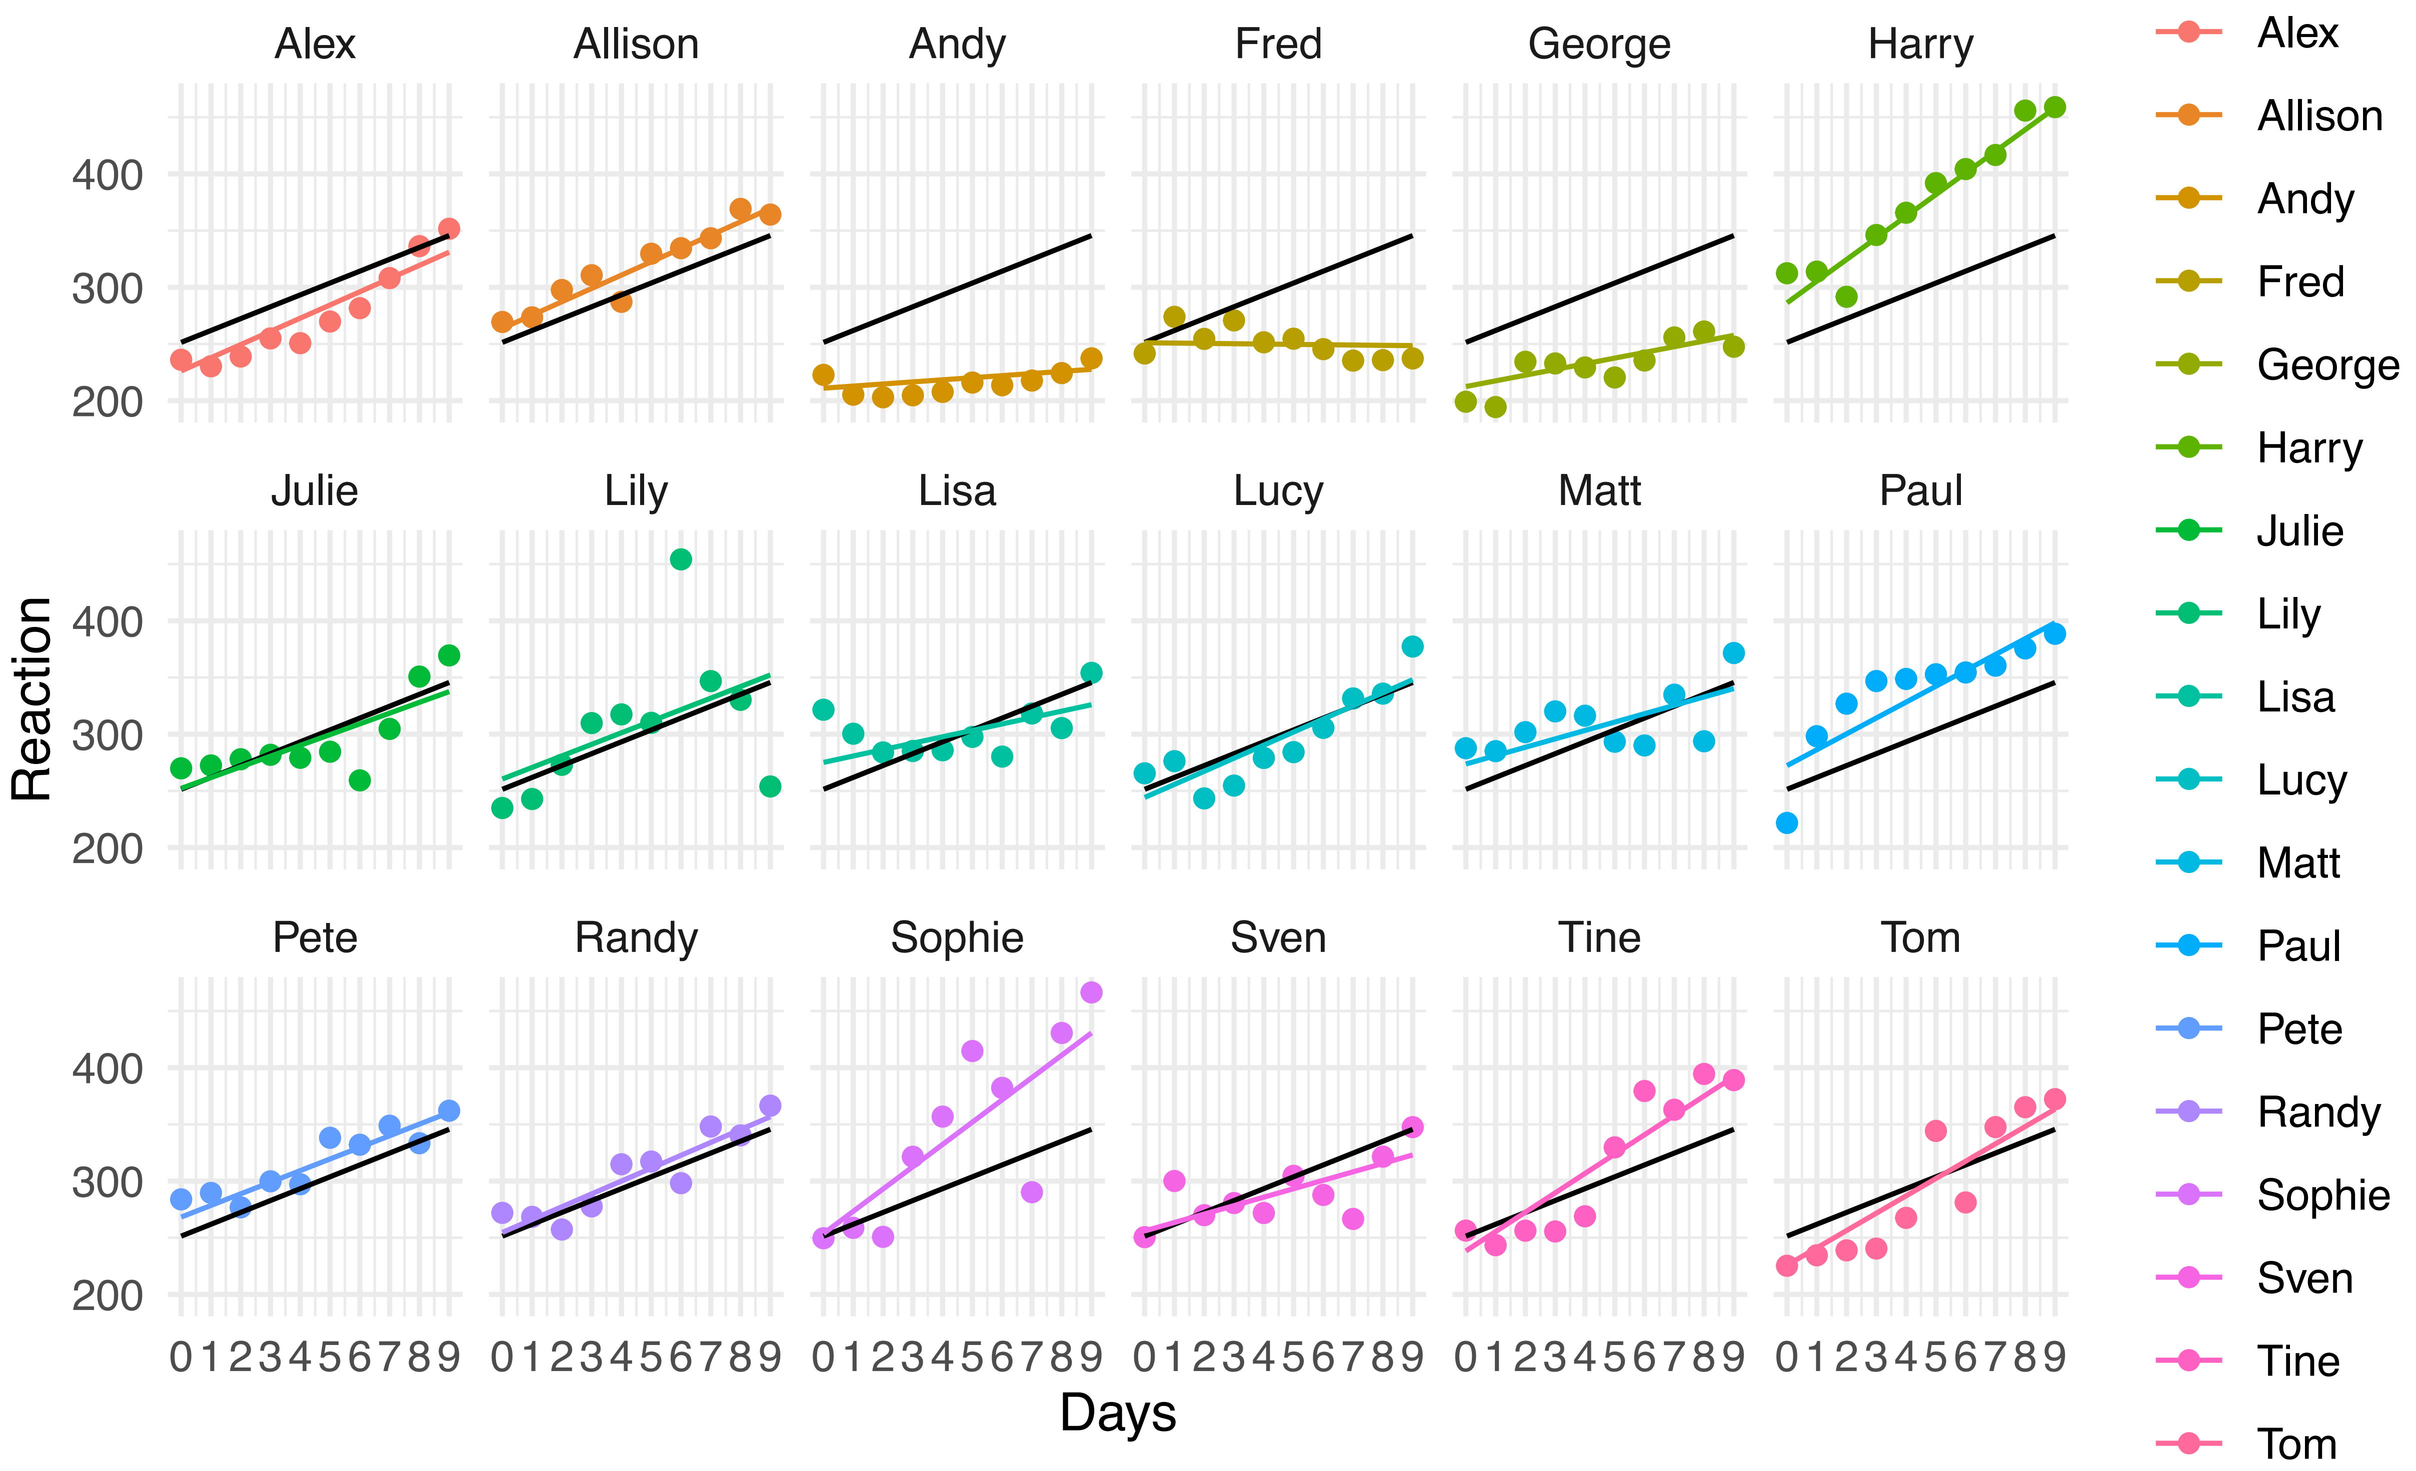

Fixed interaction of days and subject:
Model the effect of days on reaction time, separately for each subject

What's the difference?

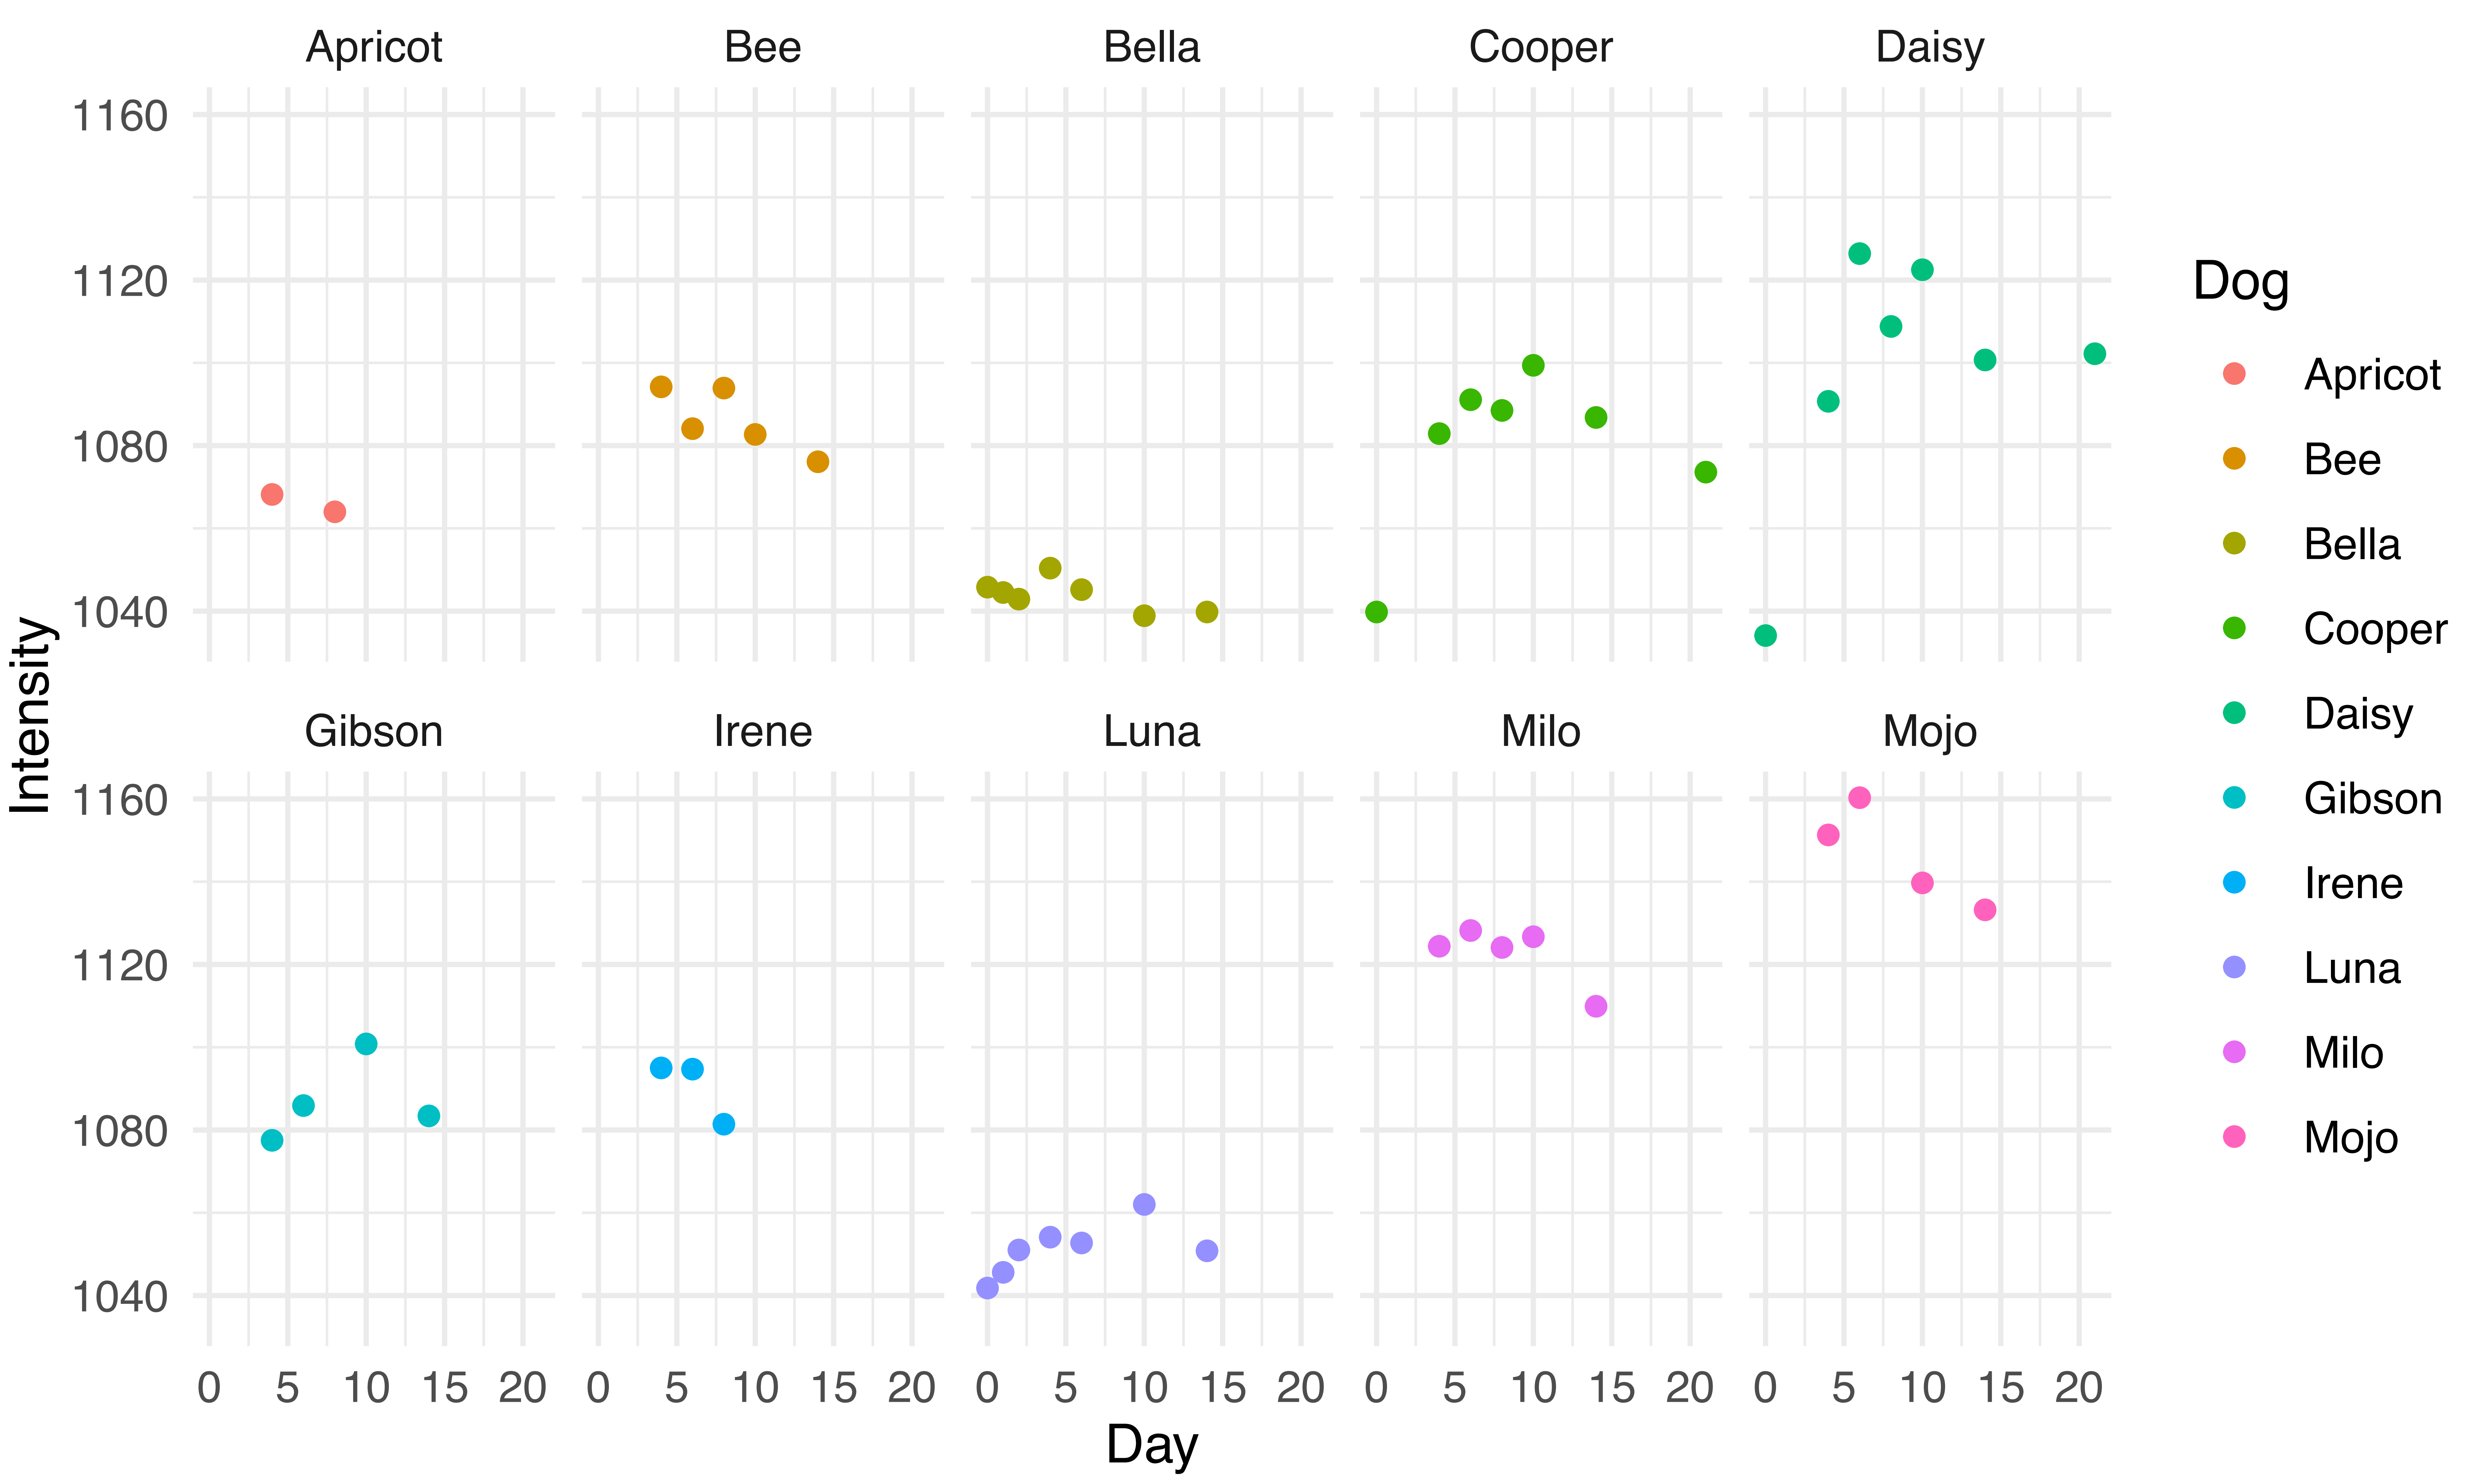

Pixel dataset from R package 'nlme'
Day= scanning day post injection of the contrast
Intensity= X-ray pixel intensity

$$y = \alpha + \beta_1 \cdot day + \beta_2 \cdot dog + \beta_3 \cdot day \cdot dog$$

```
fit_1 <- lm(Intensity ~ 1 + Day*Dog, data = dataR)
```

R package 'stats'

$$y = \alpha + \beta_1 \cdot day + \beta_2 \cdot dog + \beta_3 \cdot day \cdot dog$$

```
fit_1 <- lm(Intensity ~ 1 + Day*Dog, data = dataR)
```

Fixed interaction = Effect of day, separately for each dog

$$y = \alpha + \beta_1 \cdot day + \beta_2 \cdot dog + \beta_3 \cdot day \cdot dog$$

```
fit_1 <- lm(Intensity ~ 1 + Day*Dog, data = dataR)
```

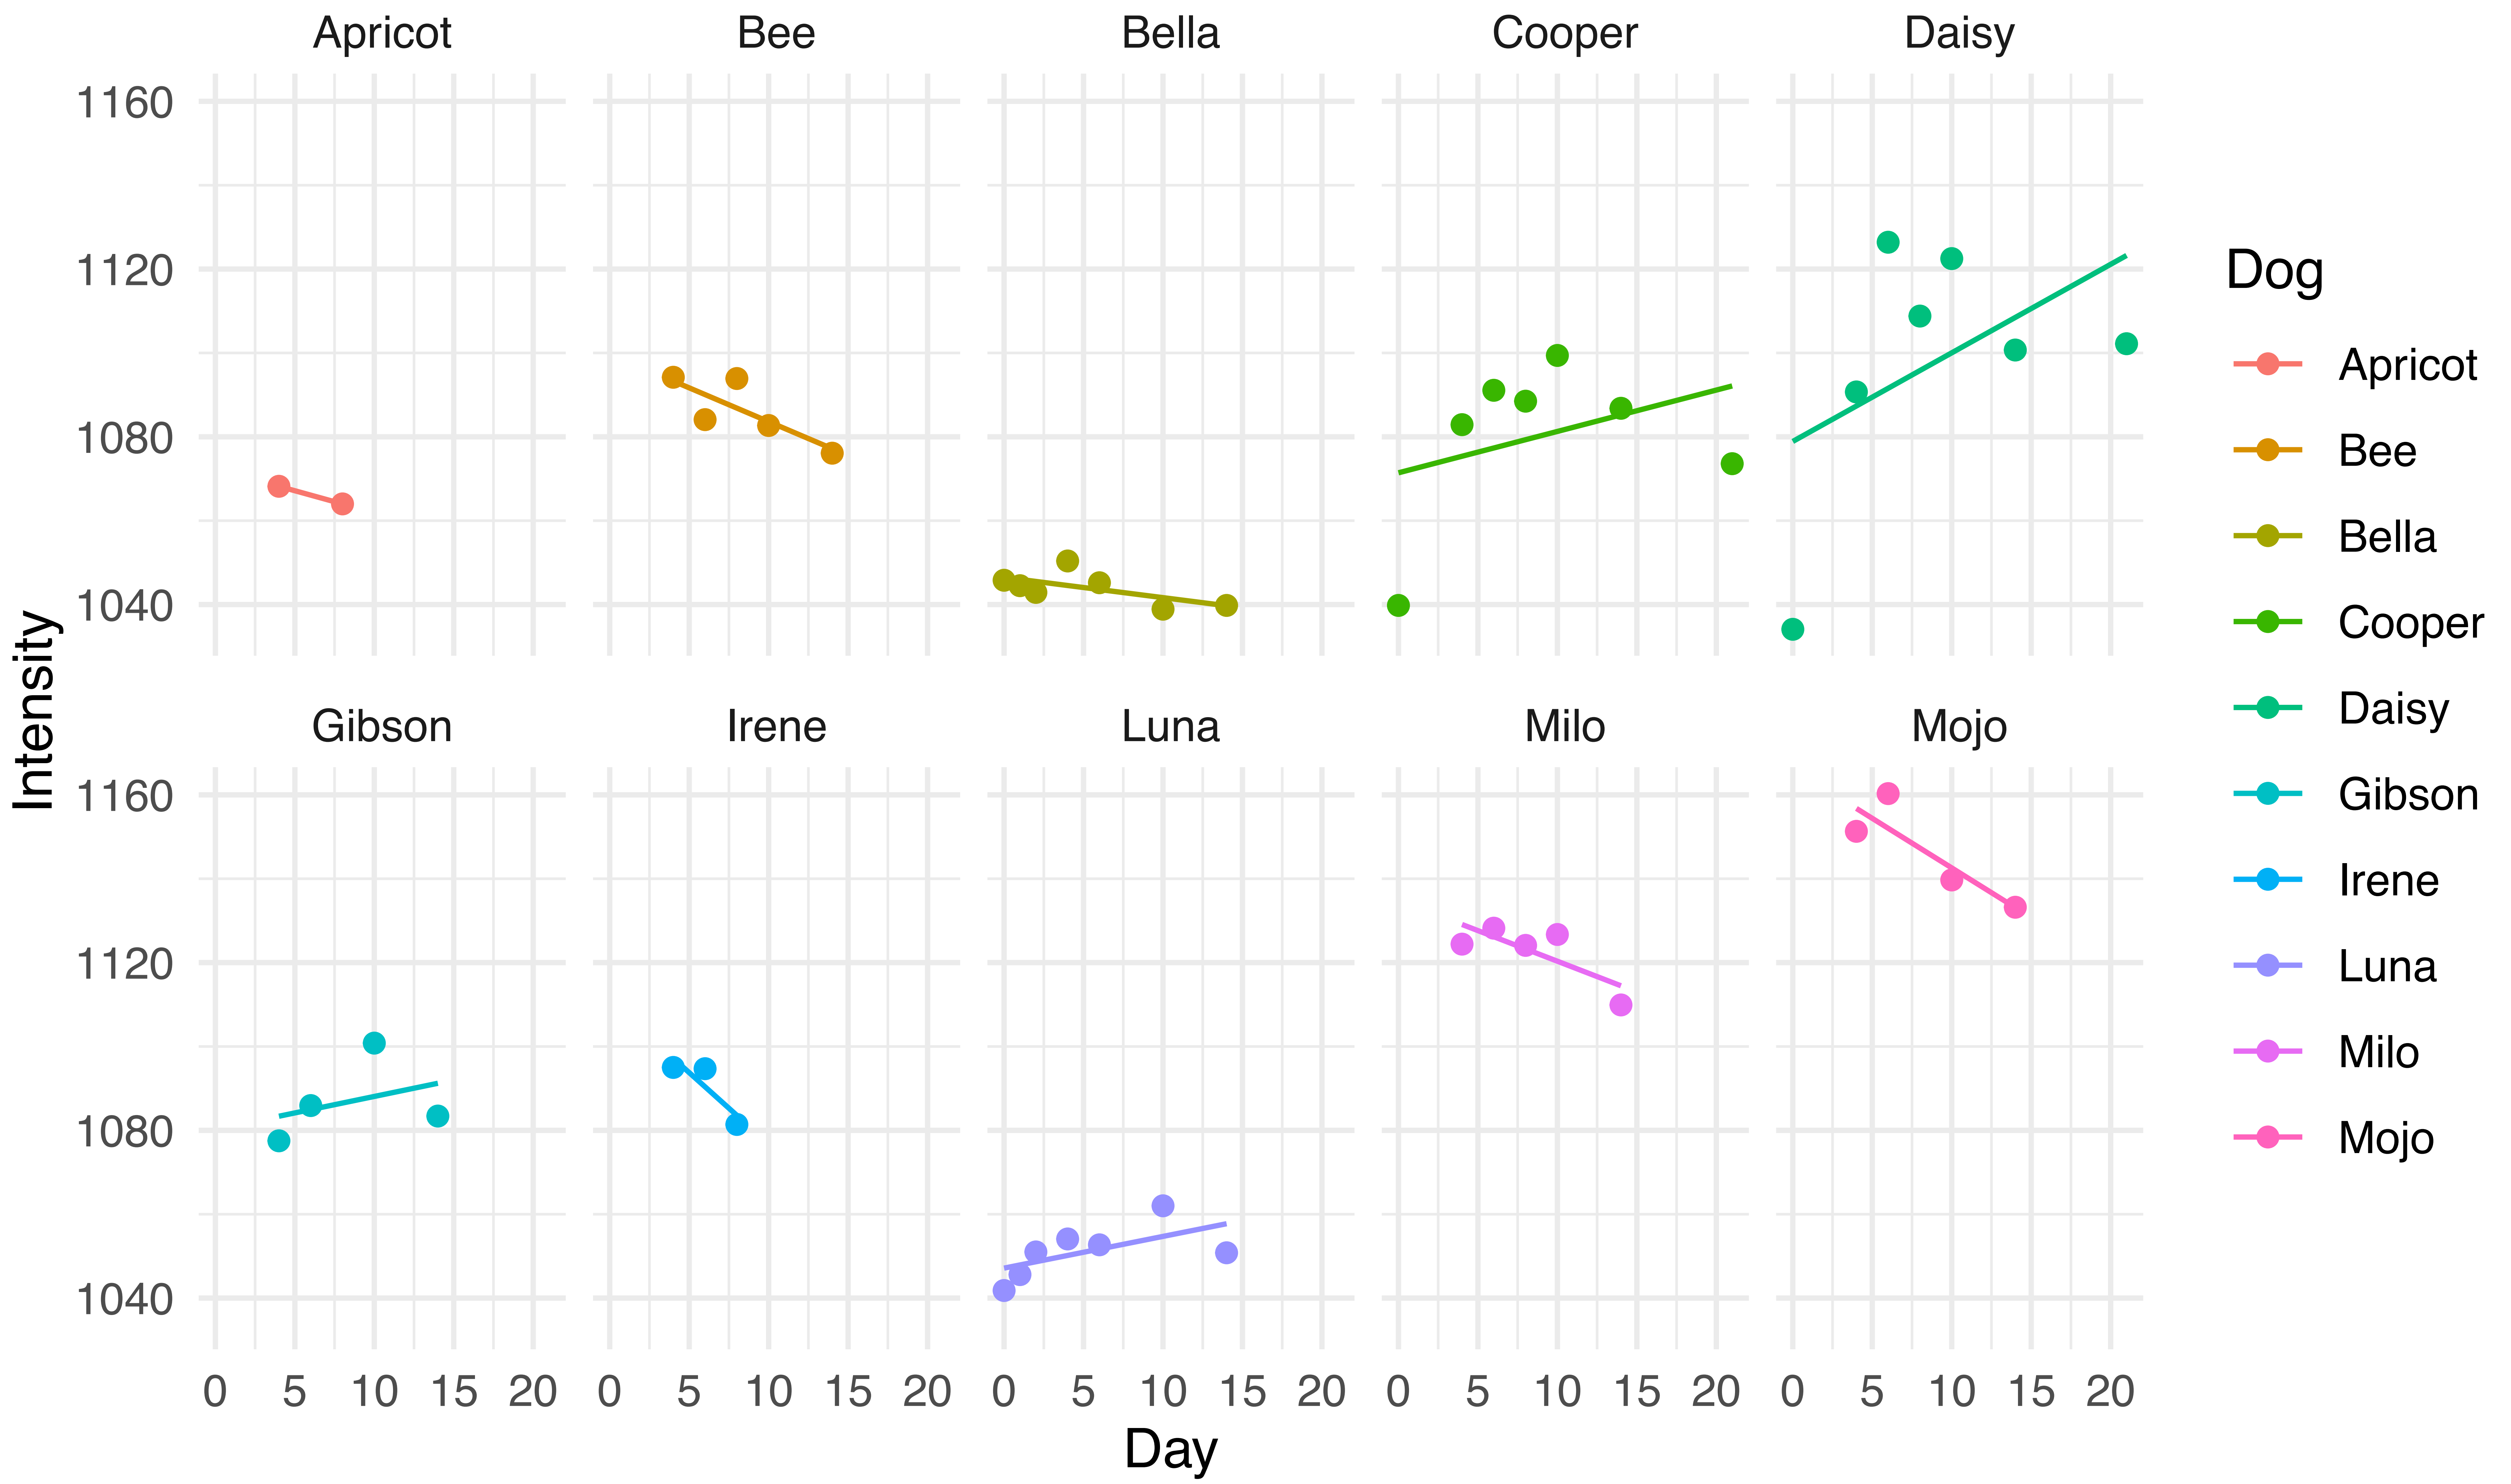

$$y = (\alpha + b_{\alpha,dog}) + (\beta + b_{\beta,dog}) \cdot day$$

```
fit_1 <- lm(Intensity ~ 1 + Day*Dog, data = dataR)
fit_2 <- lmer(Intensity ~ 1 + Day + (1 + Day | Dog) , data = dataR)
```

R package 'lme4'

$y = (\alpha + b_{\alpha,dog}) + (\beta + b_{\beta,dog}) \cdot day$

```
fit_1 <- lm(Intensity ~ 1 + Day*Dog, data = dataR)
fit_2 <- lmer(Intensity ~ 1 + Day + (1 + Day | Dog) , data = dataR)
```

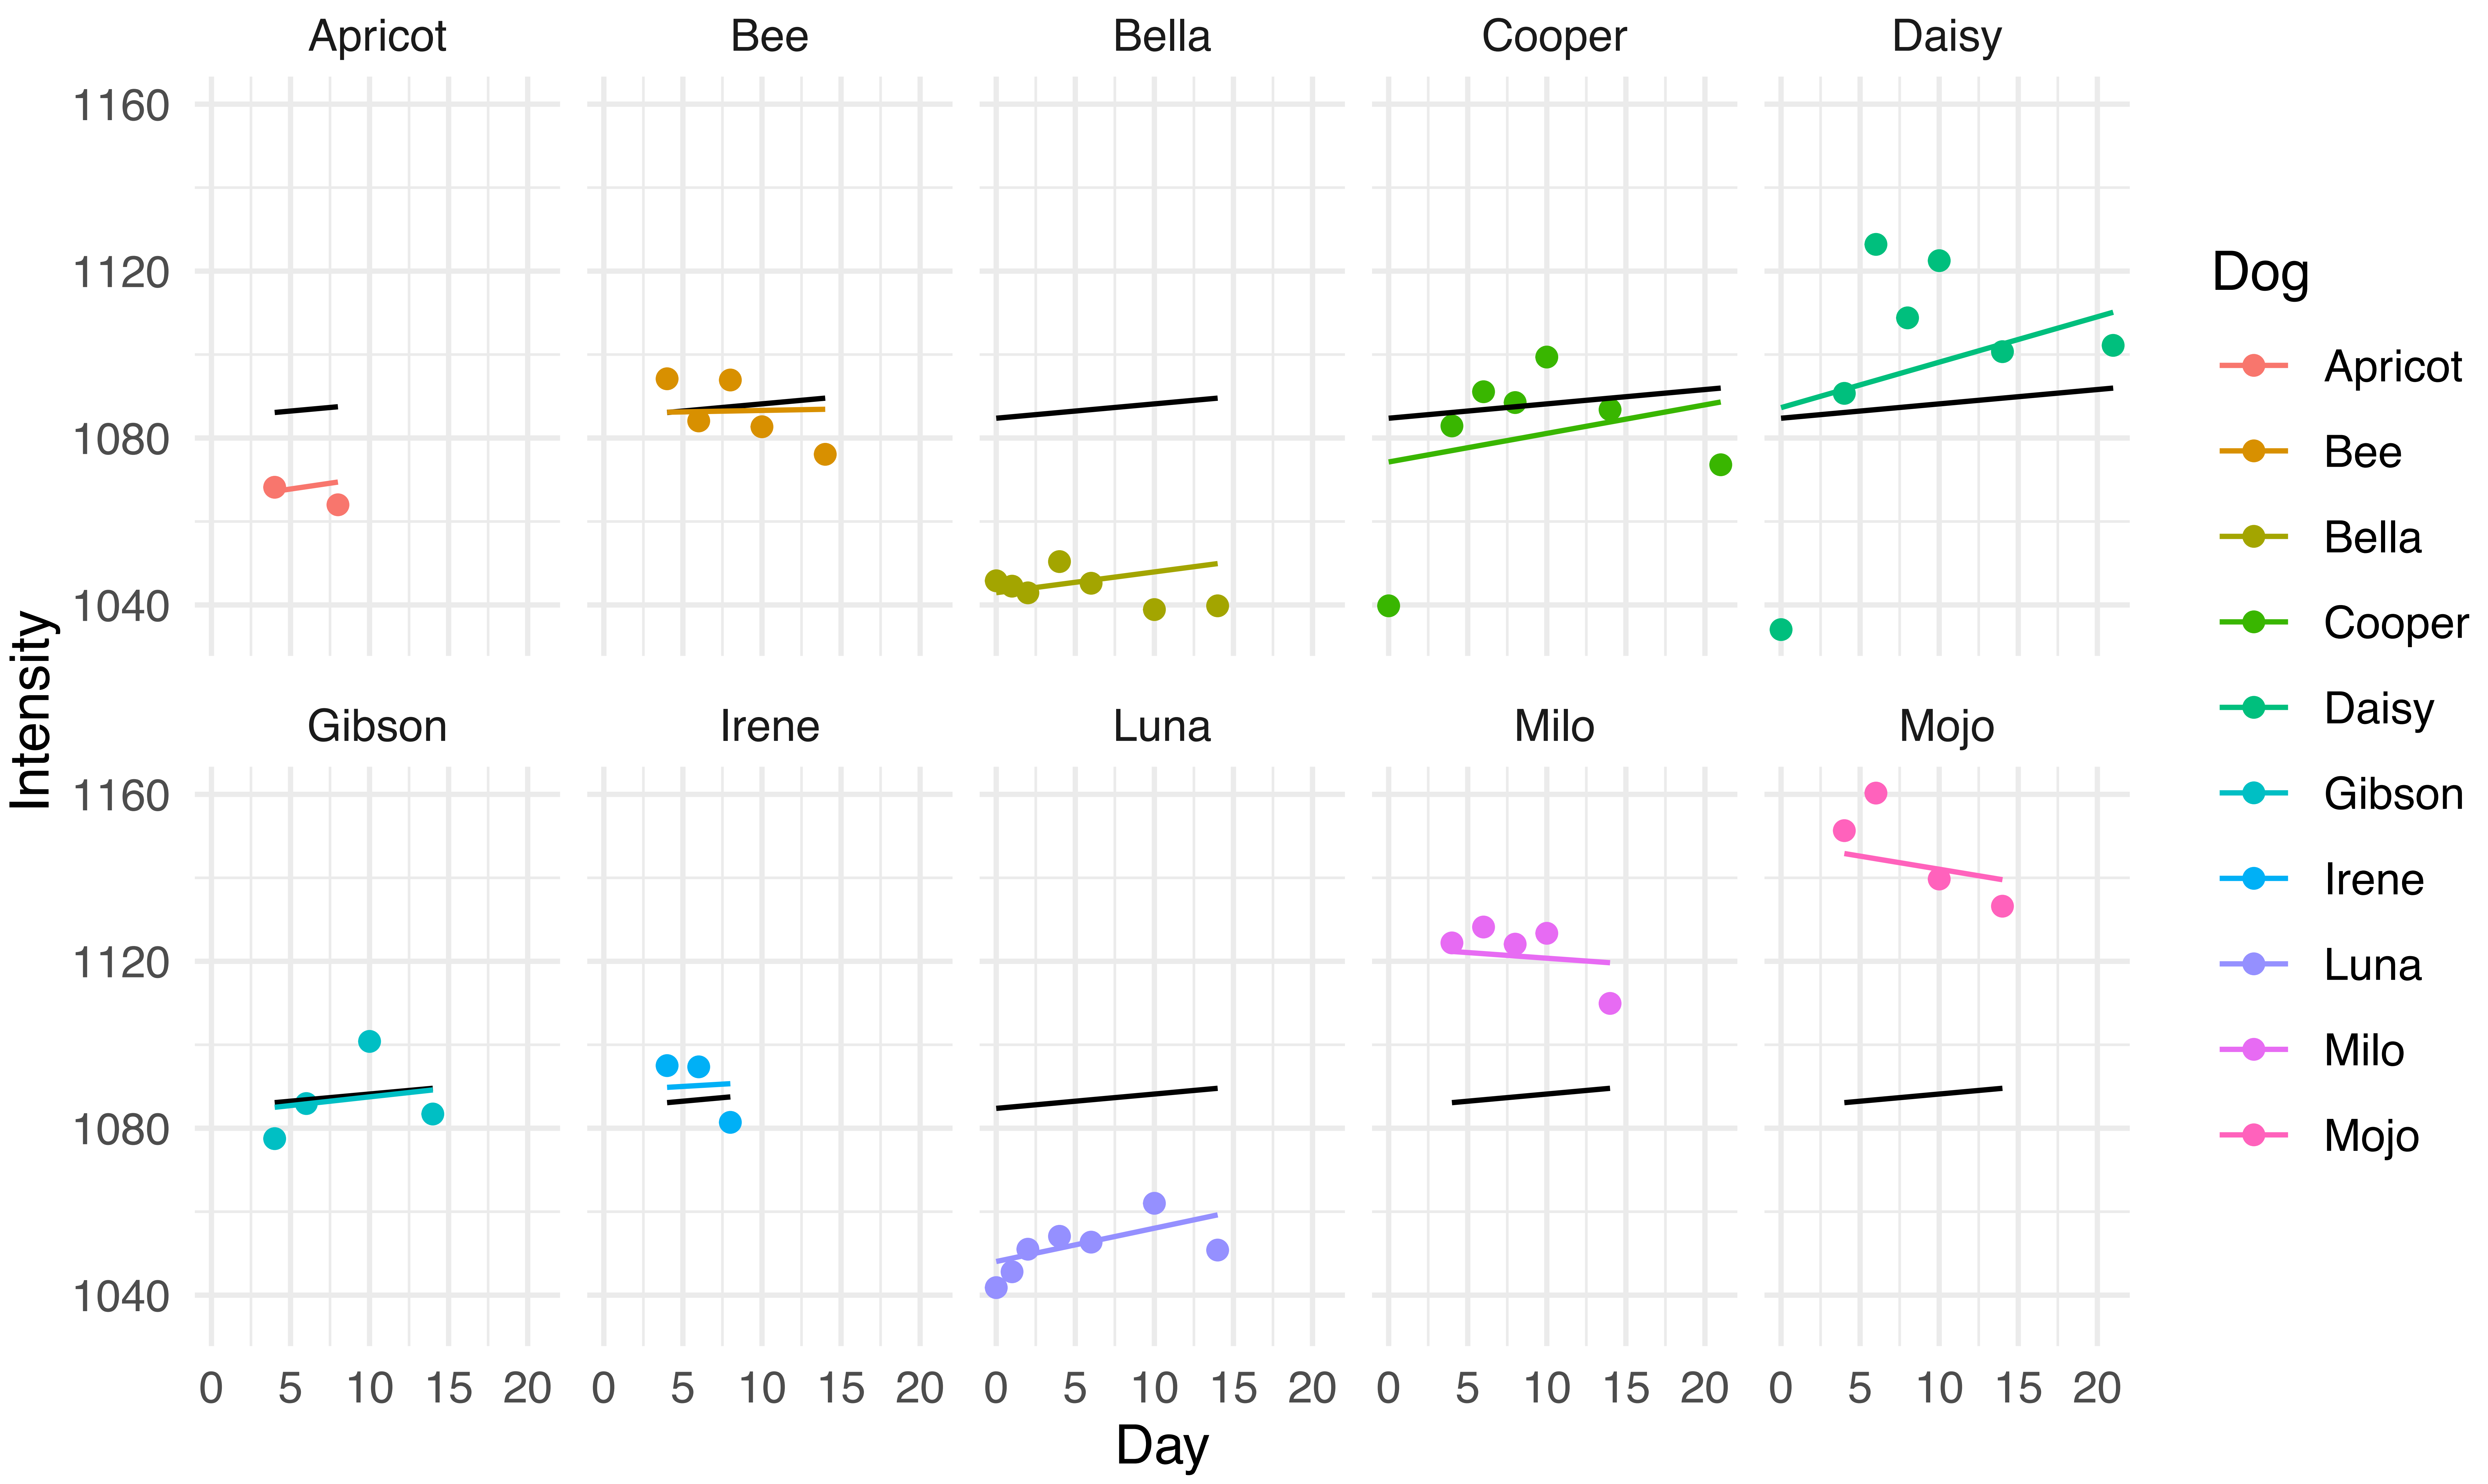

```
fit_1 <- lm(Intensity ~ 1 + Day*Dog, data = dataR)
fit_2 <- lmer(Intensity ~ 1 + Day + (1 + Day | Dog) , data = dataR)
```

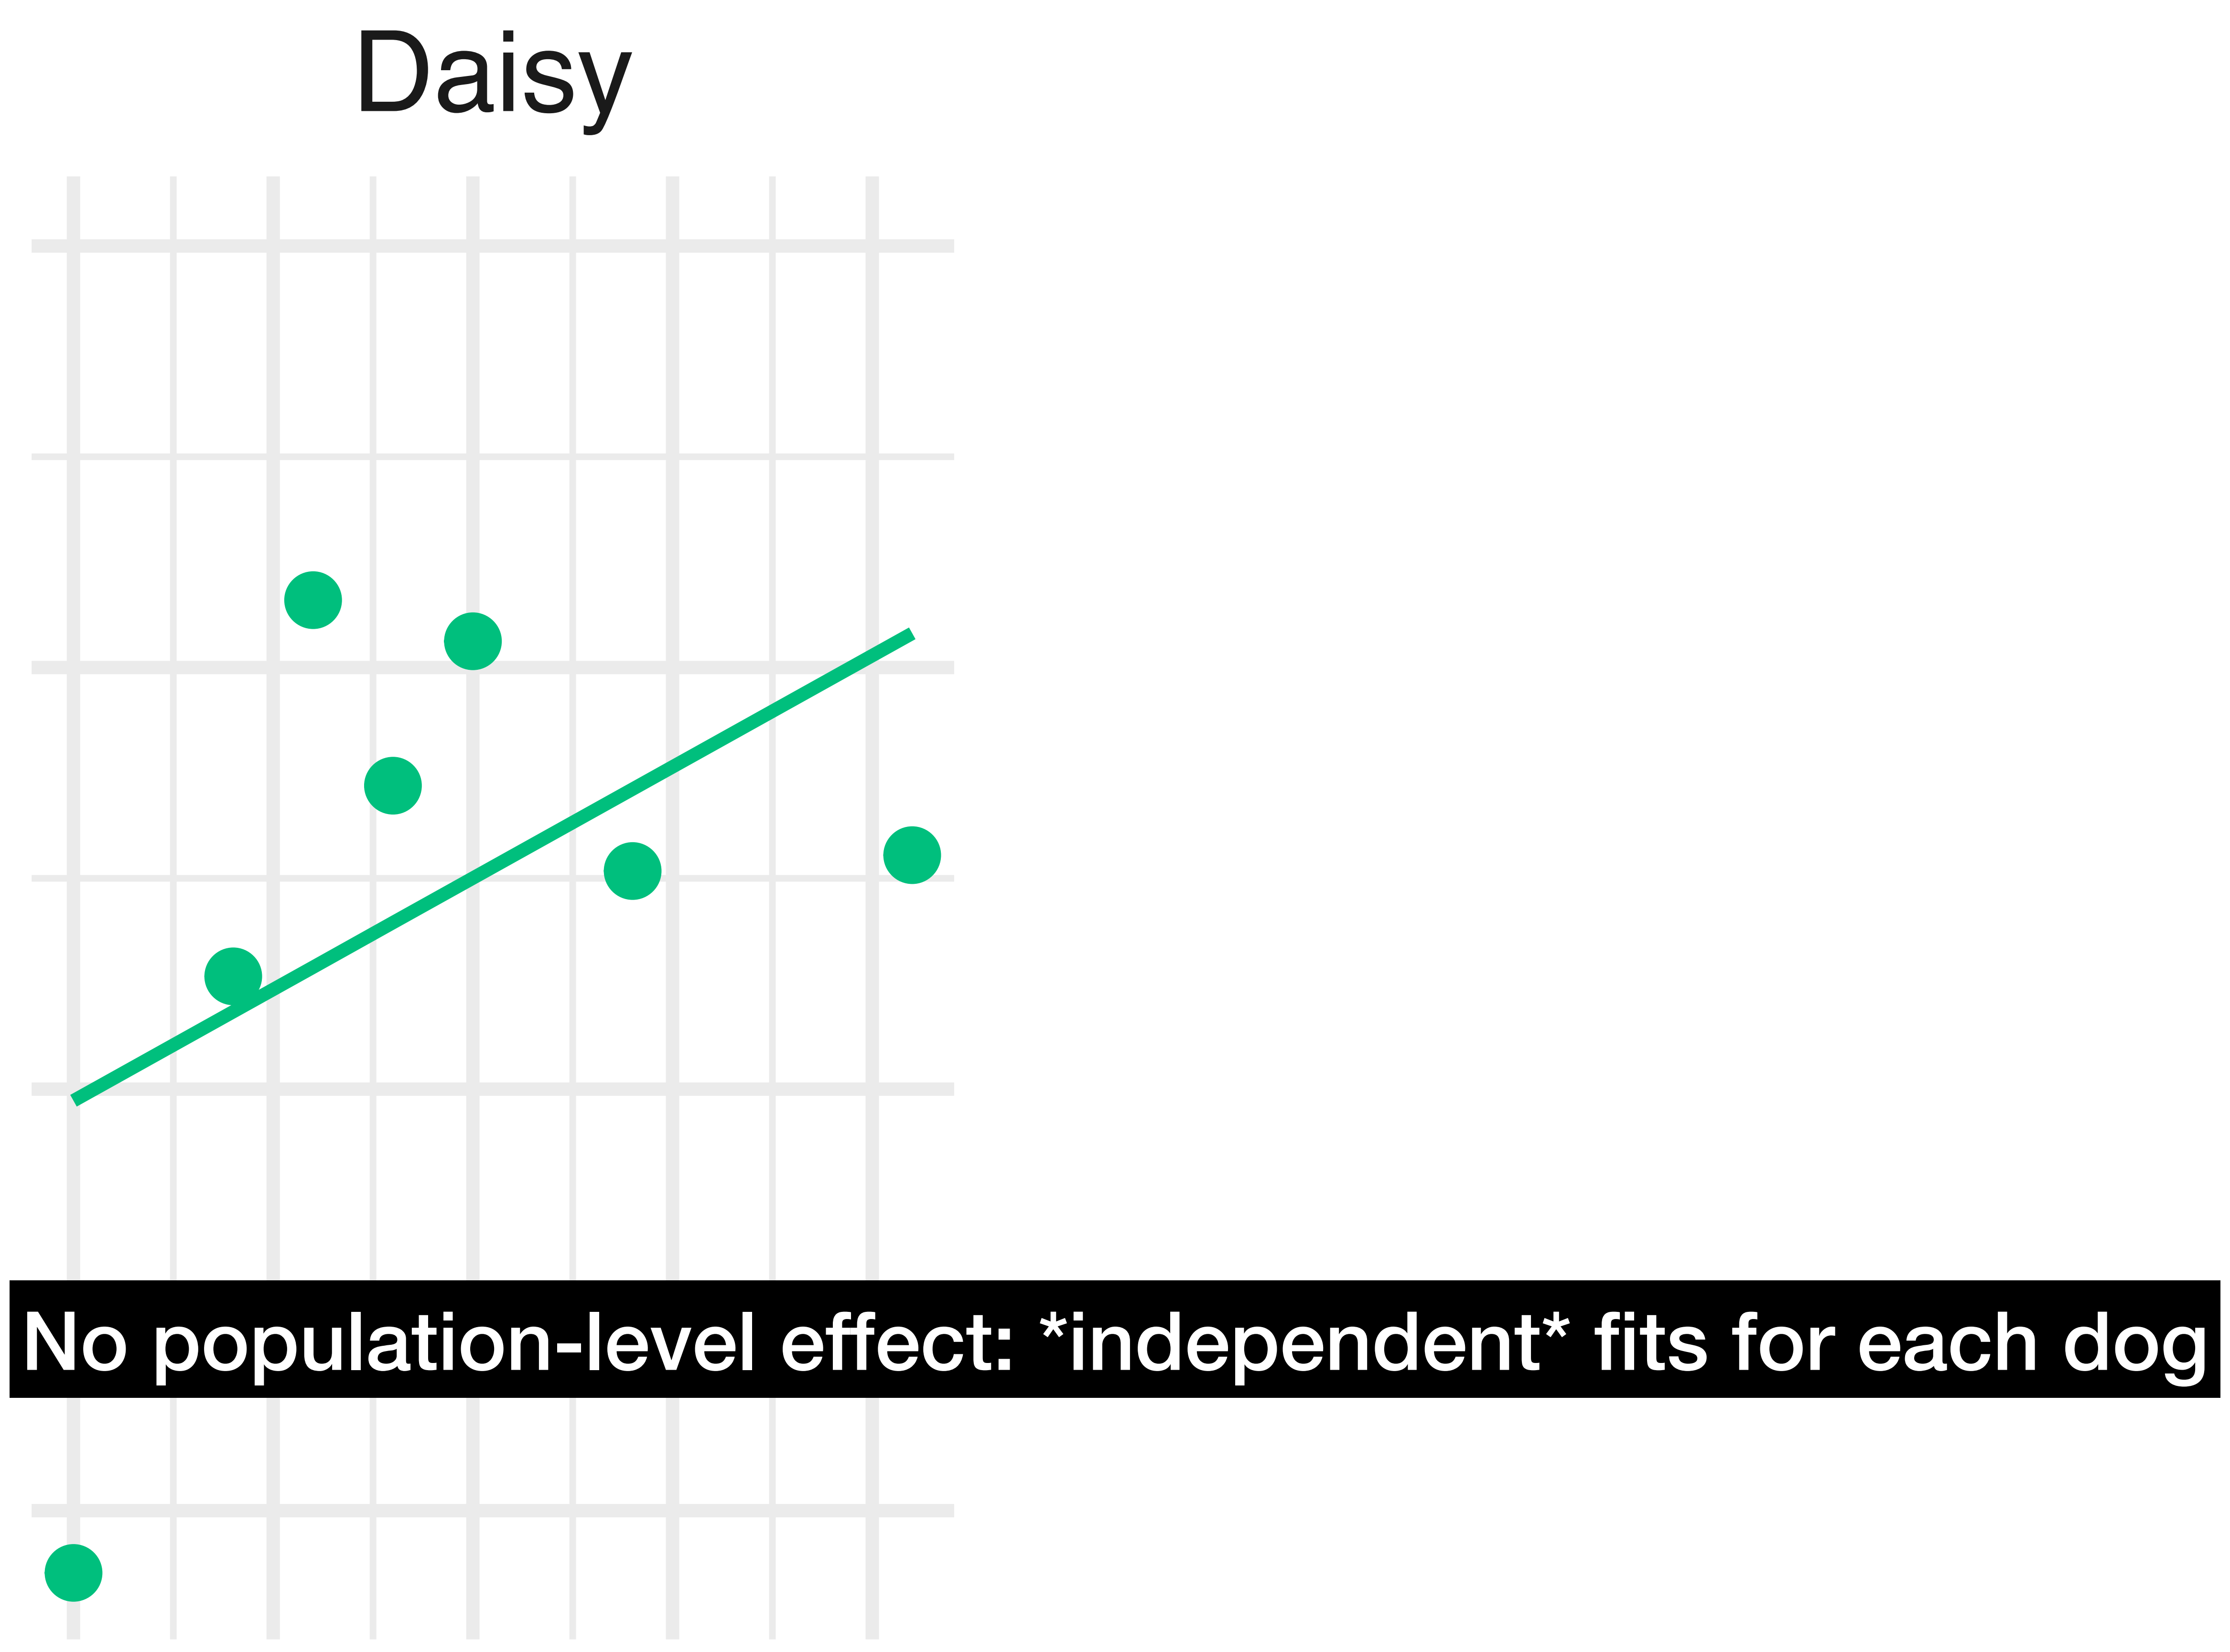

Fixed interaction

$$y = \alpha + \beta_1 \cdot day + \beta_2 \cdot dog + \beta_3 \cdot day \cdot dog$$

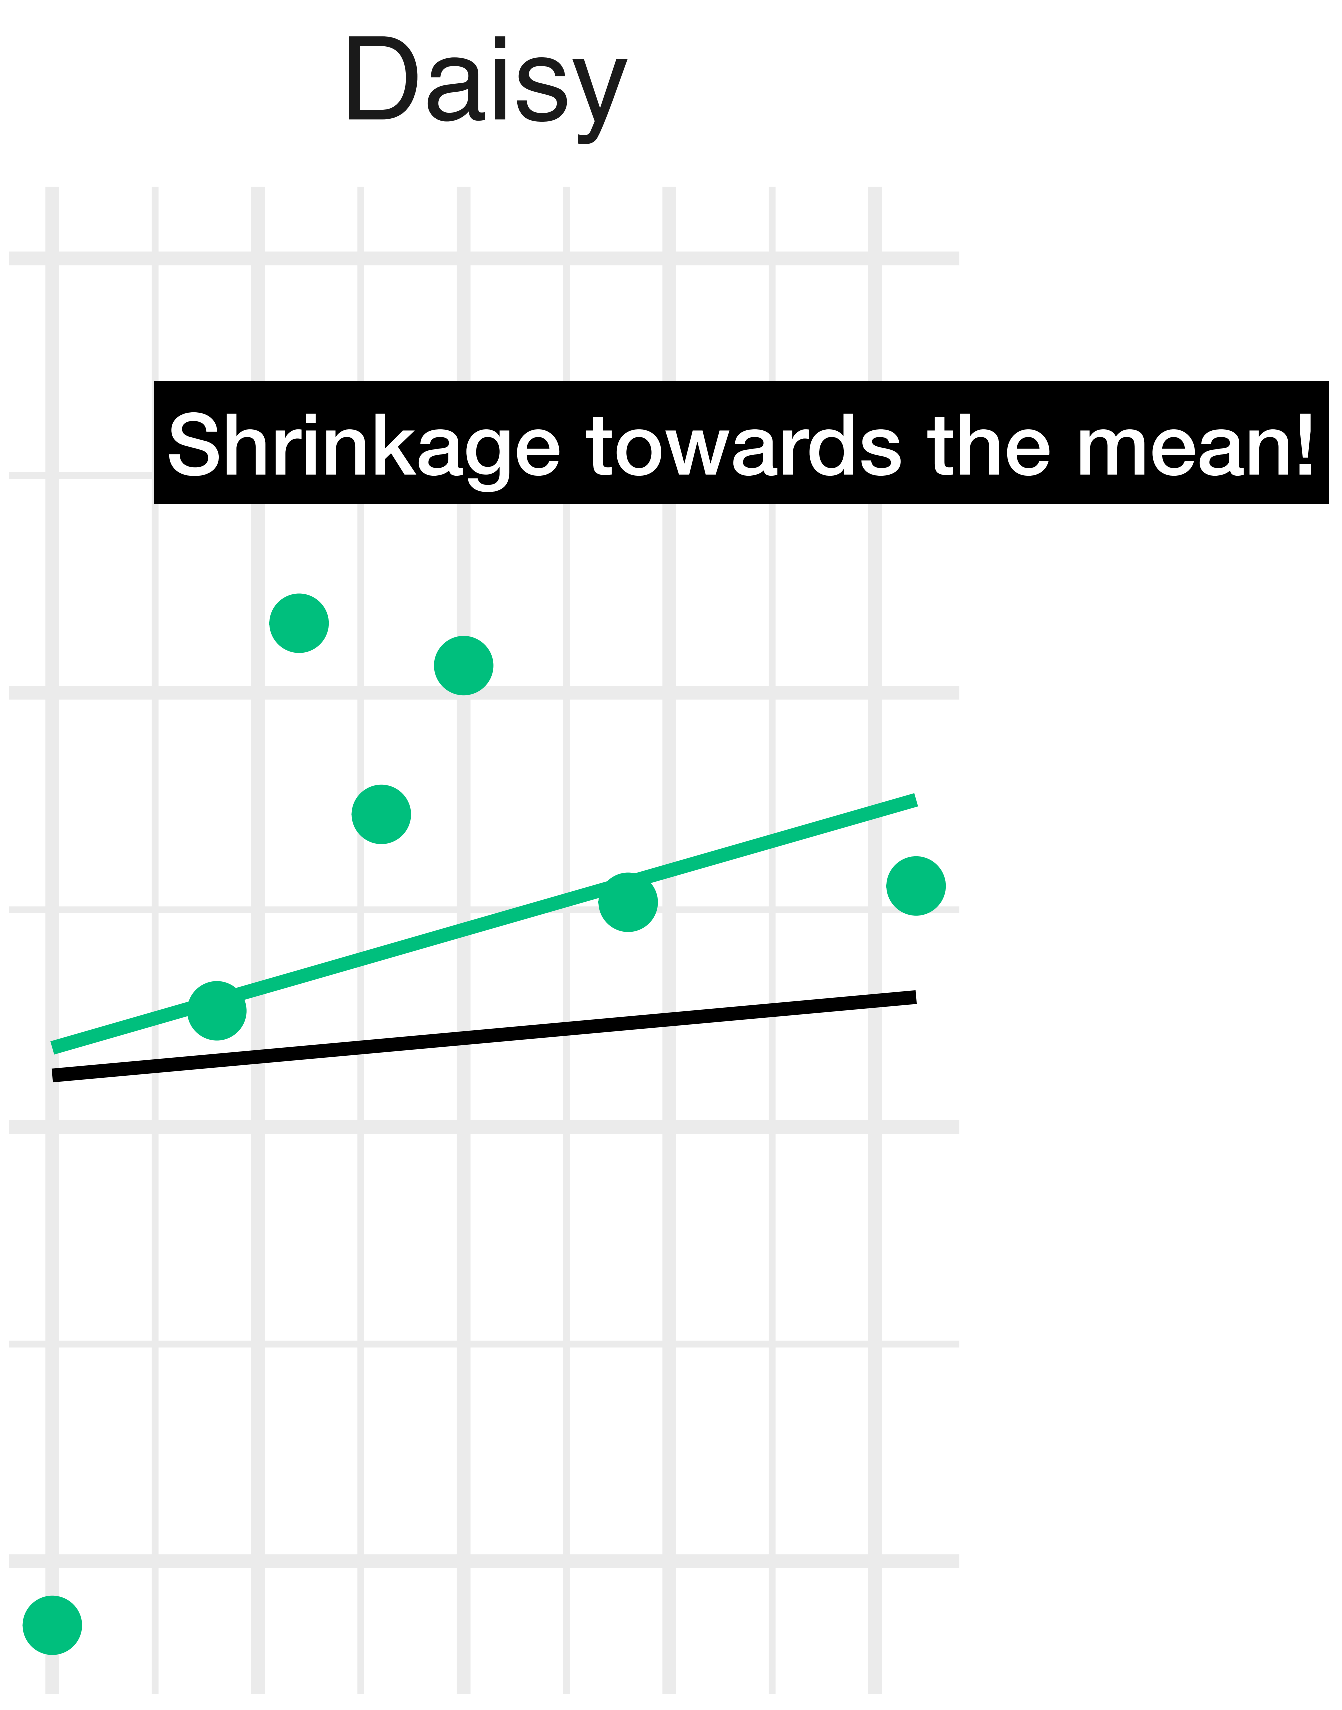

Random intercept and slope

$$y = (\alpha + b_{\alpha,dog}) + (\beta + b_{\beta,dog}) \cdot day$$

1) If we are interested in the population-level effect… What's the point?

We get an idea of how much people vary on average

summary(fit)

```
Random effects:
 Groups   Name        Variance Std.Dev.
 Subject  (Intercept) 612.10   24.741
          Days         35.07    5.922
 Residual             654.94   25.592
```

2) If we are interested in the subject-specific effects… Why not fixed interaction (instead of random effects)?

Extreme / rare observations might be unrealiable, utilize the information from others.
Effects are pooled towards the population-level effect (shrinkage).

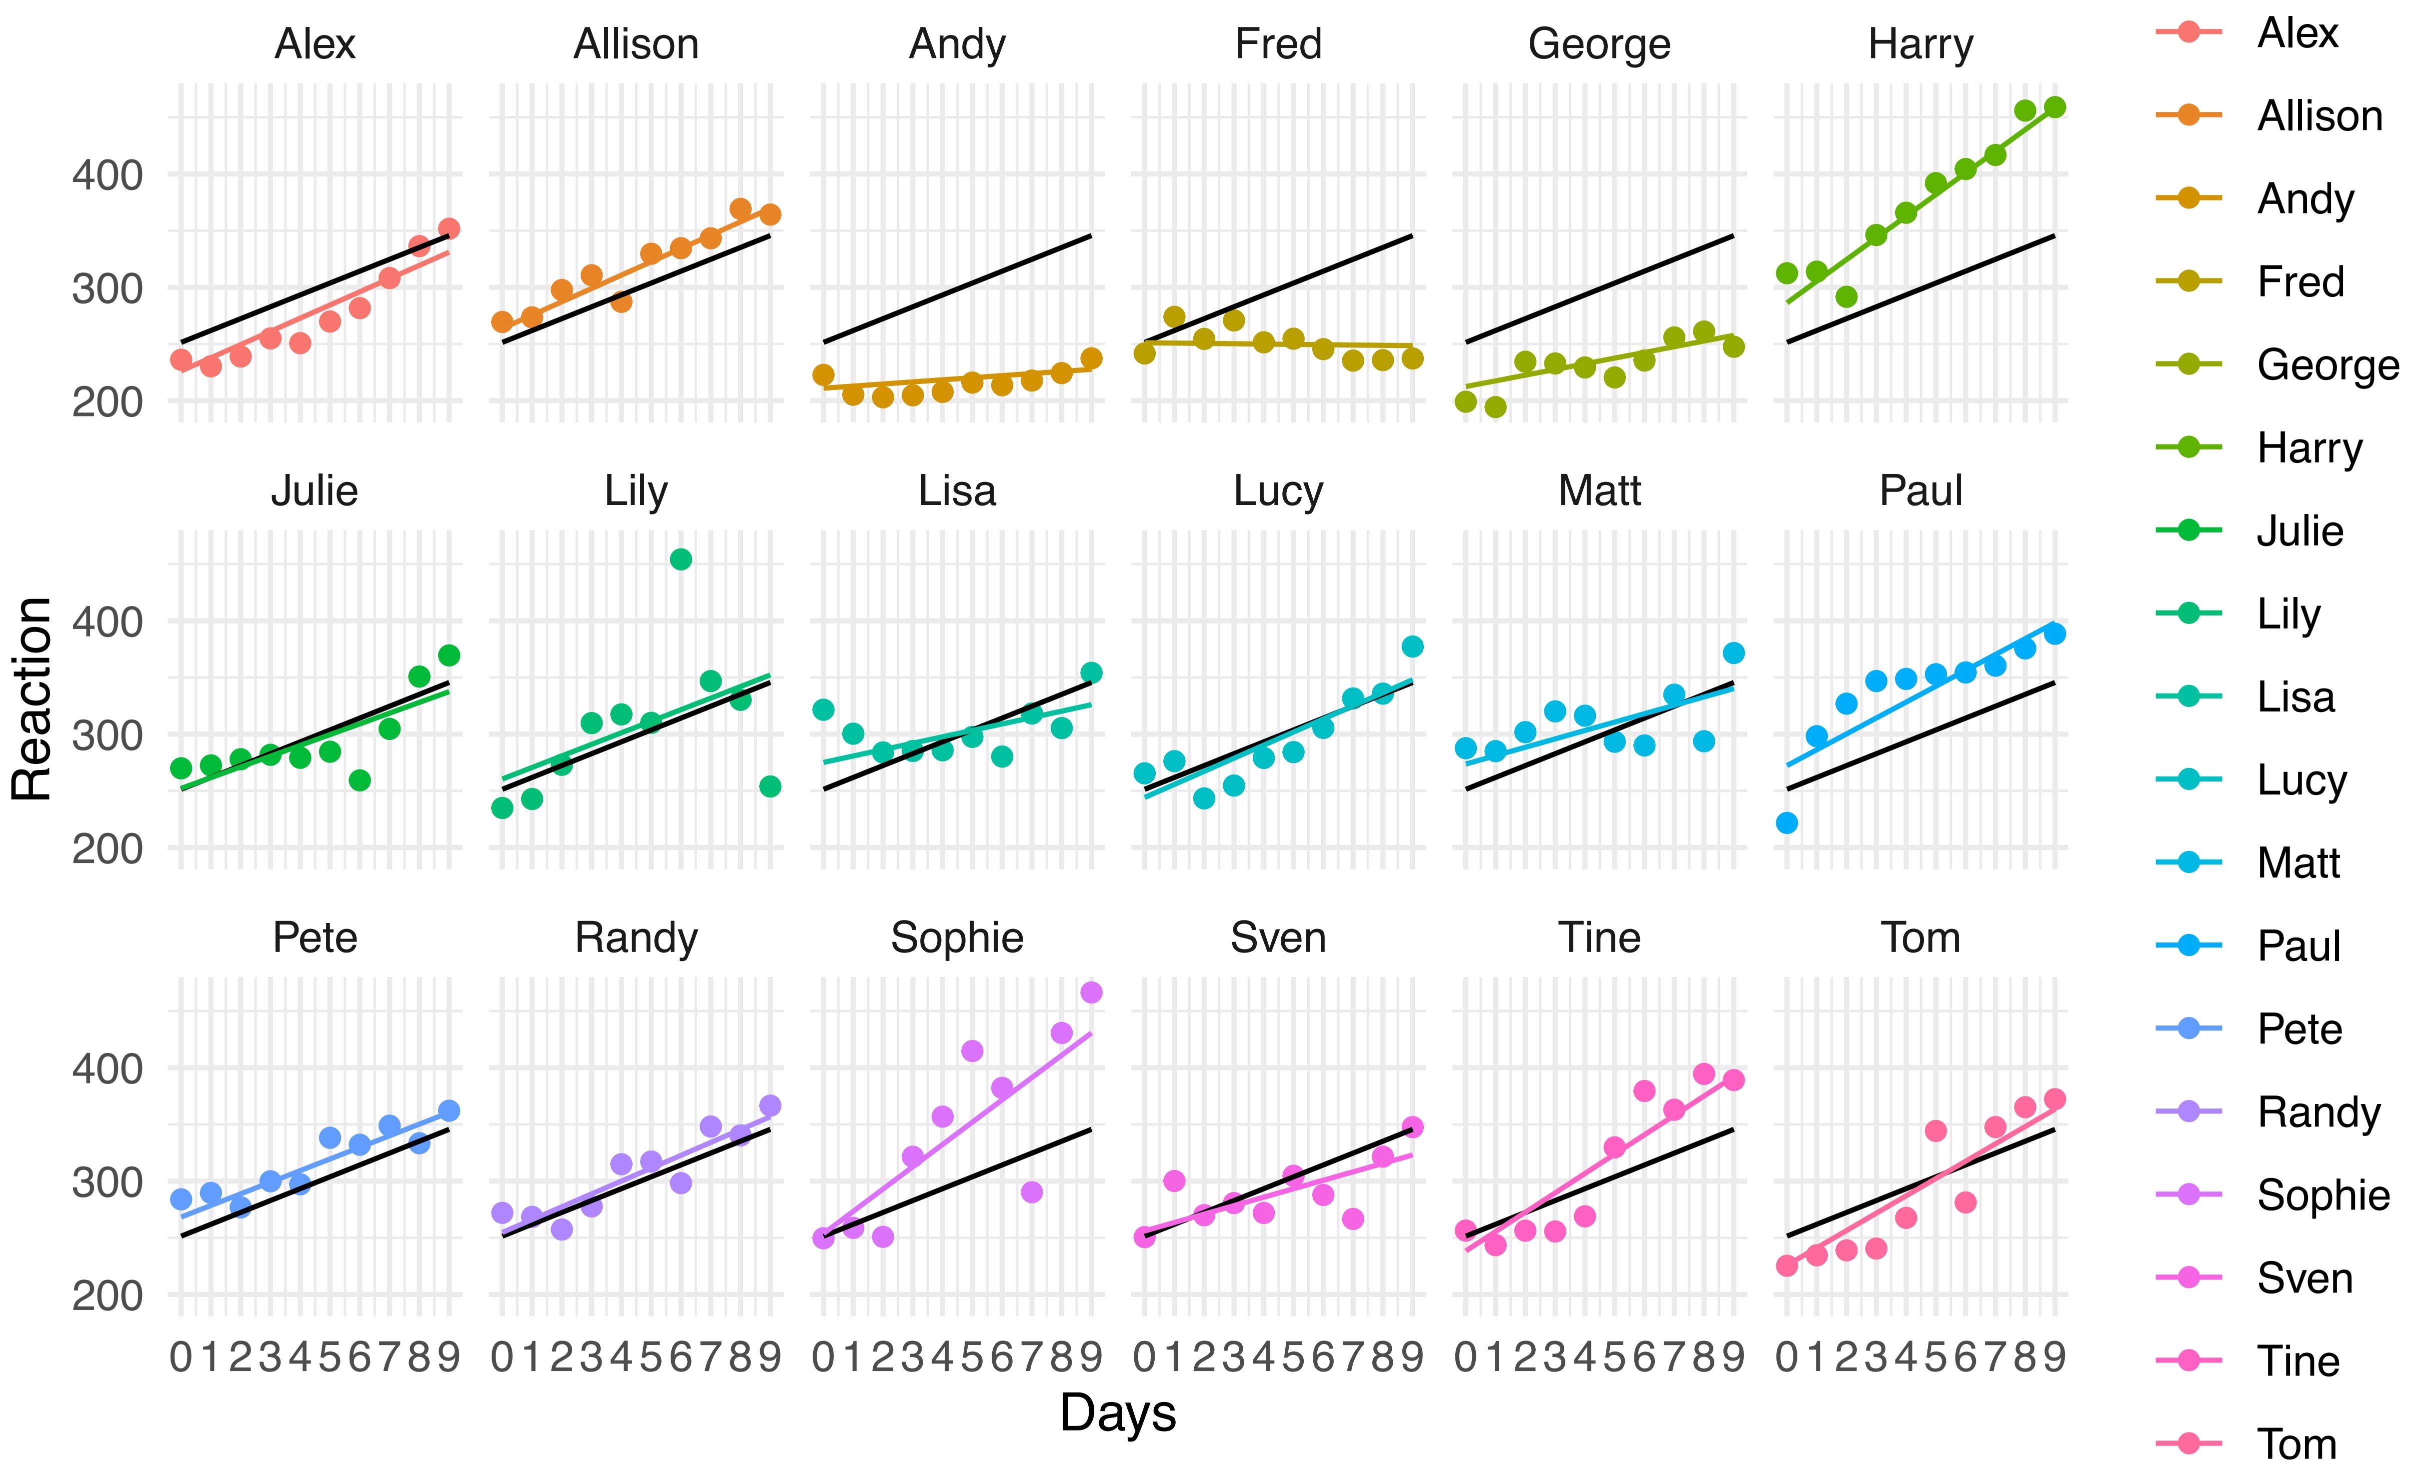

When a predictor should be modeled as fixed and when as random?

Turku PET Centre

Aarhus University Hospital

Yale, New Haven

University of British Columbia, Vancouver

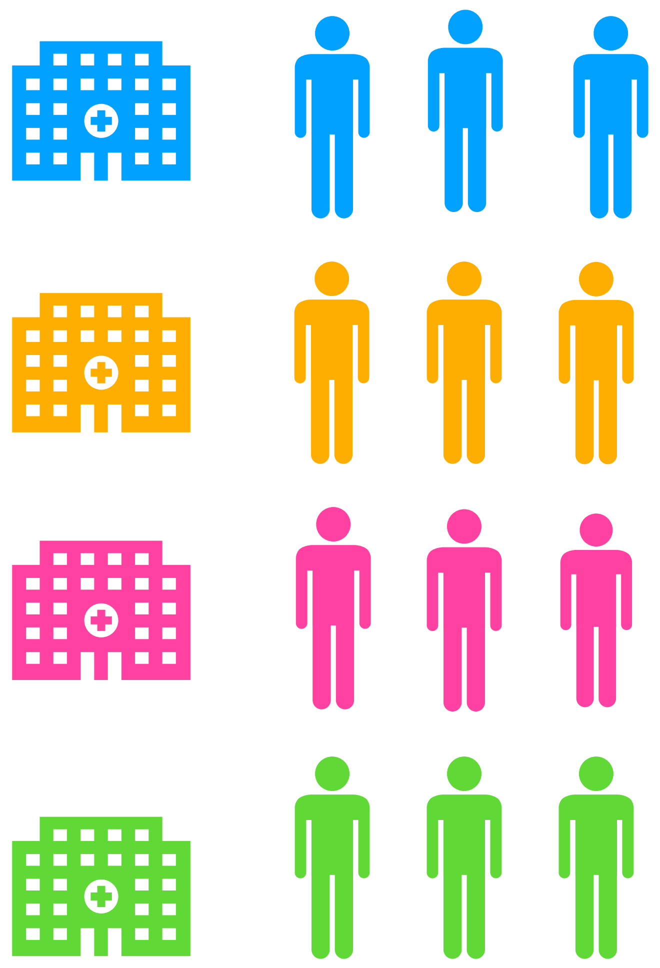

How should we model the site?

Fixed, if

- Effect of site on PET measure
- The difference between the sites are the key interest

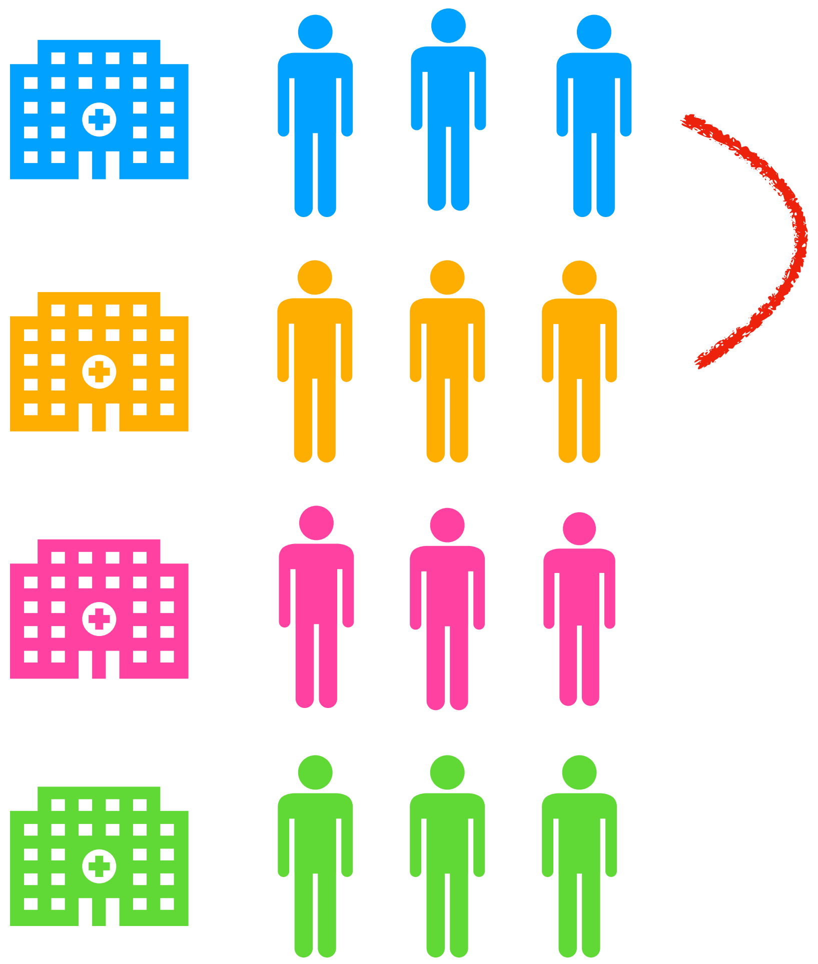

How should we model the site?

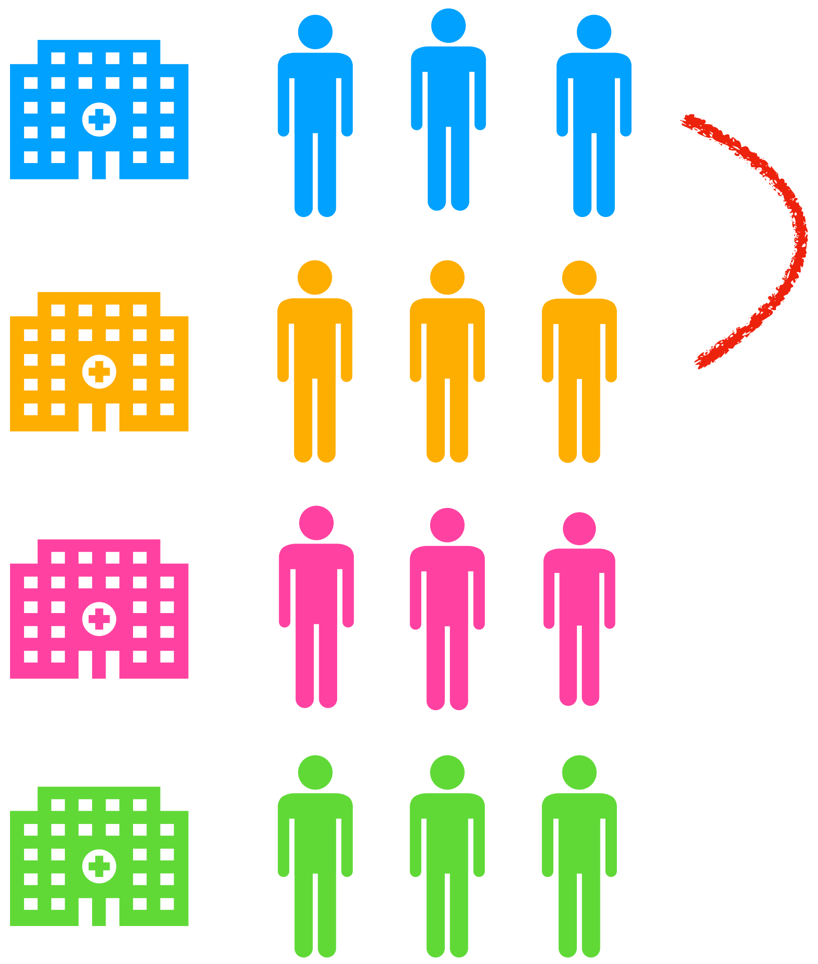

Fixed, if

- Effect of site on PET measure
- The difference between the sites are the key interest

Random, if

- Effect of something else, e.g. age on PET measure
- Random sample of sites (from all possible sites): might affect but the contrasts not interesting
- 4+ levels (here sites) recommended for random effects

Databases & registers!

Random effects in neuroimaging data

Site

Random effects in neuroimaging data

Site

Scanner

Random effects in neuroimaging data

Site

Scanner

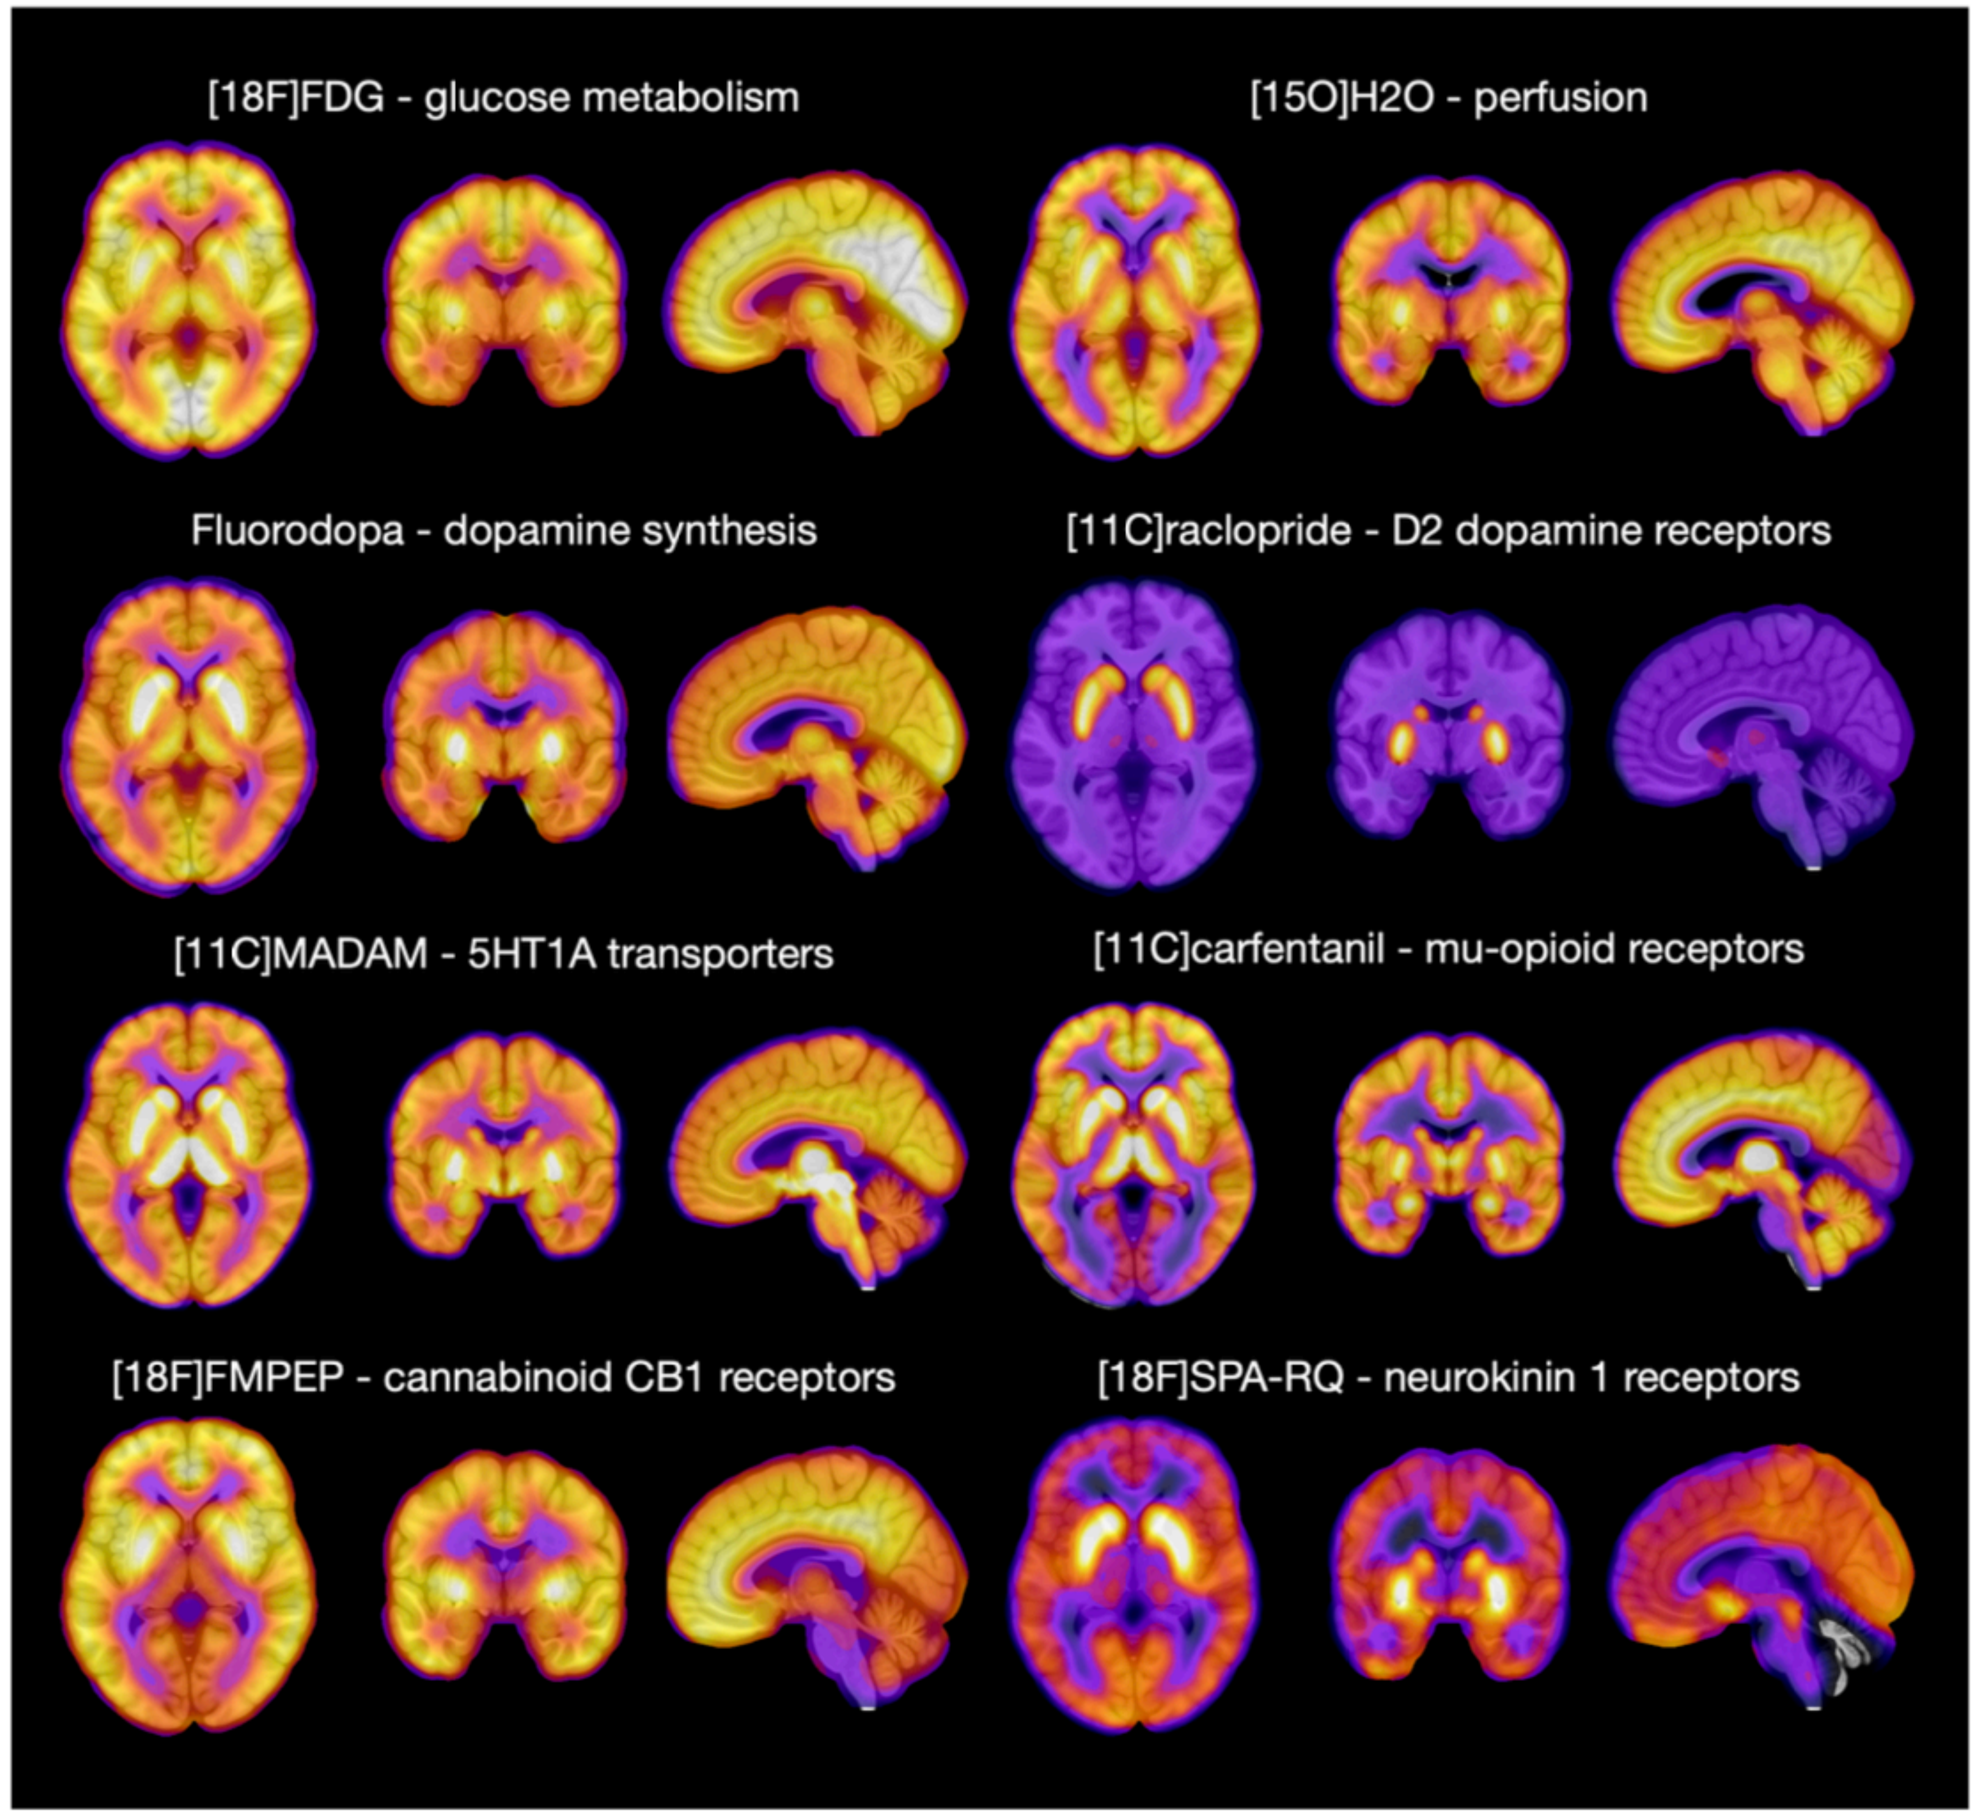

Brain region?

Figure from Nummenmaa, L., & Hirvonen, J. (preprint). Mapping the human emotion circuits with positron emission tomography. Human Emotion Systems Laboratory Preprints.

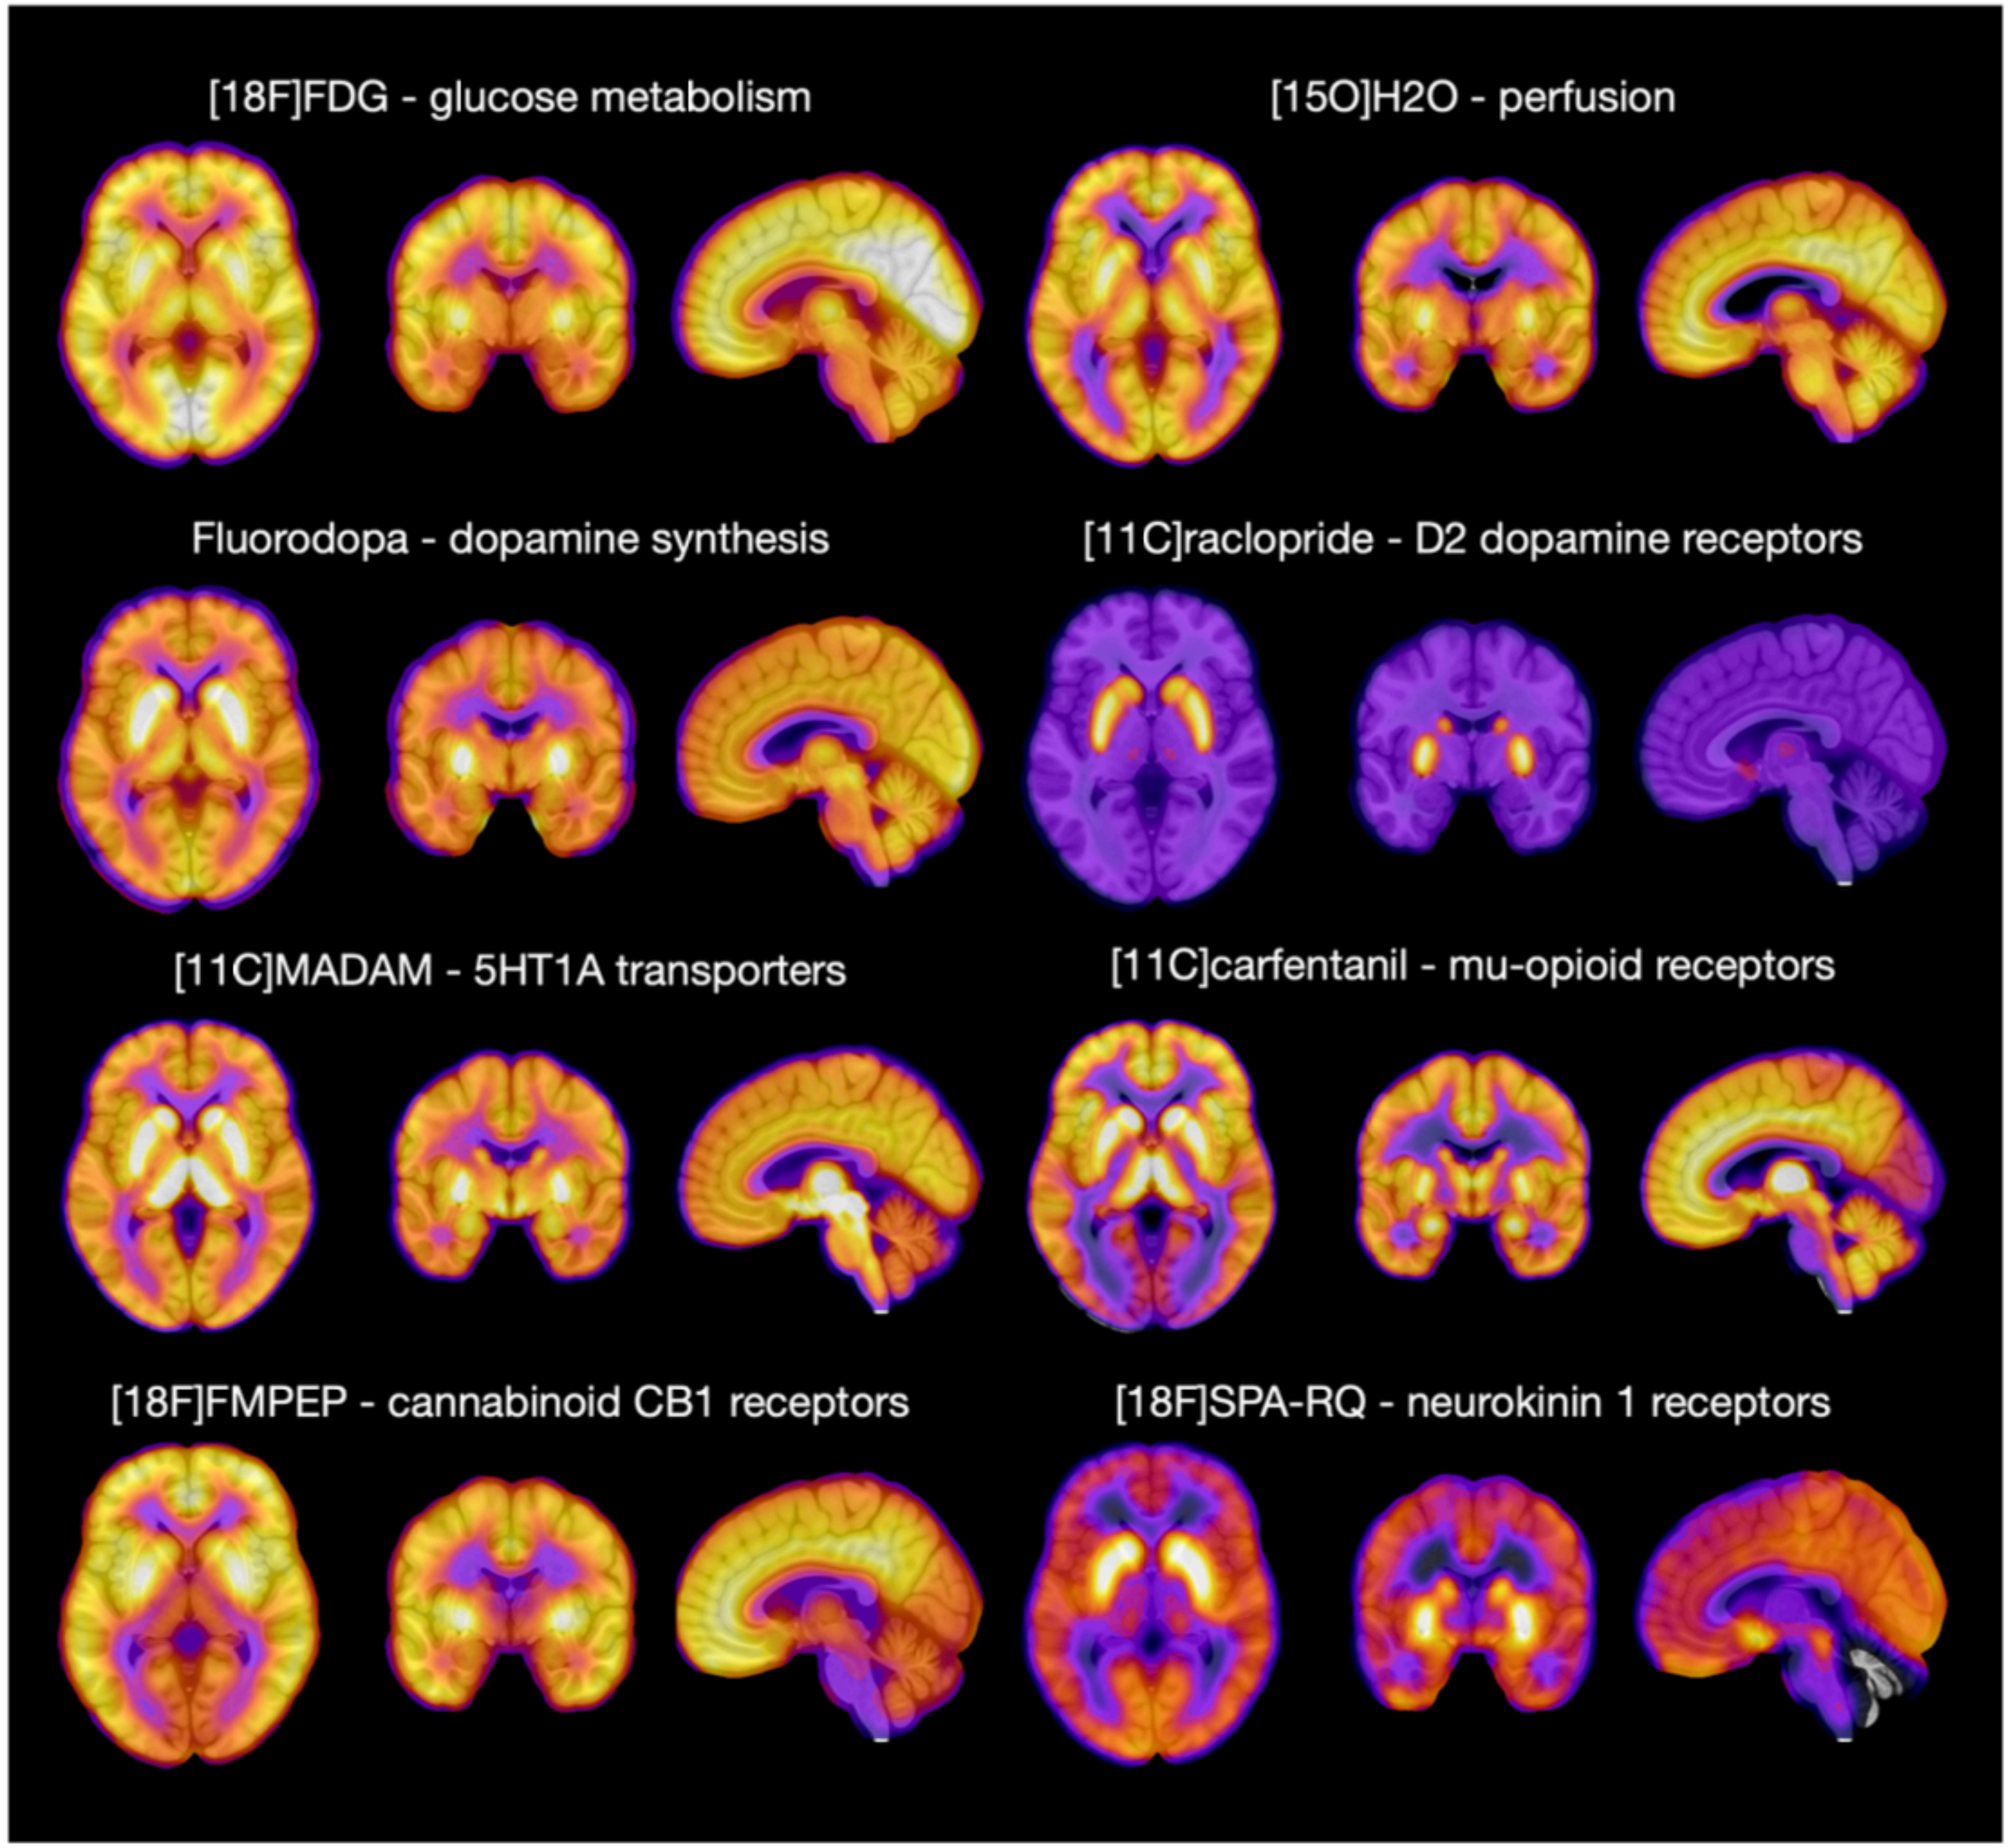

Brain region: fixed or random?

Dependent variable:
PET measure binding potential (receptor availability)

Examples:

Is binding higher in putamen than caudate?

What is the effect of age on binding?

Figure from Nummenmaa, L., & Hirvonen, J. (preprint). Mapping the human emotion circuits with positron emission tomography. Human Emotion Systems Laboratory Preprints.

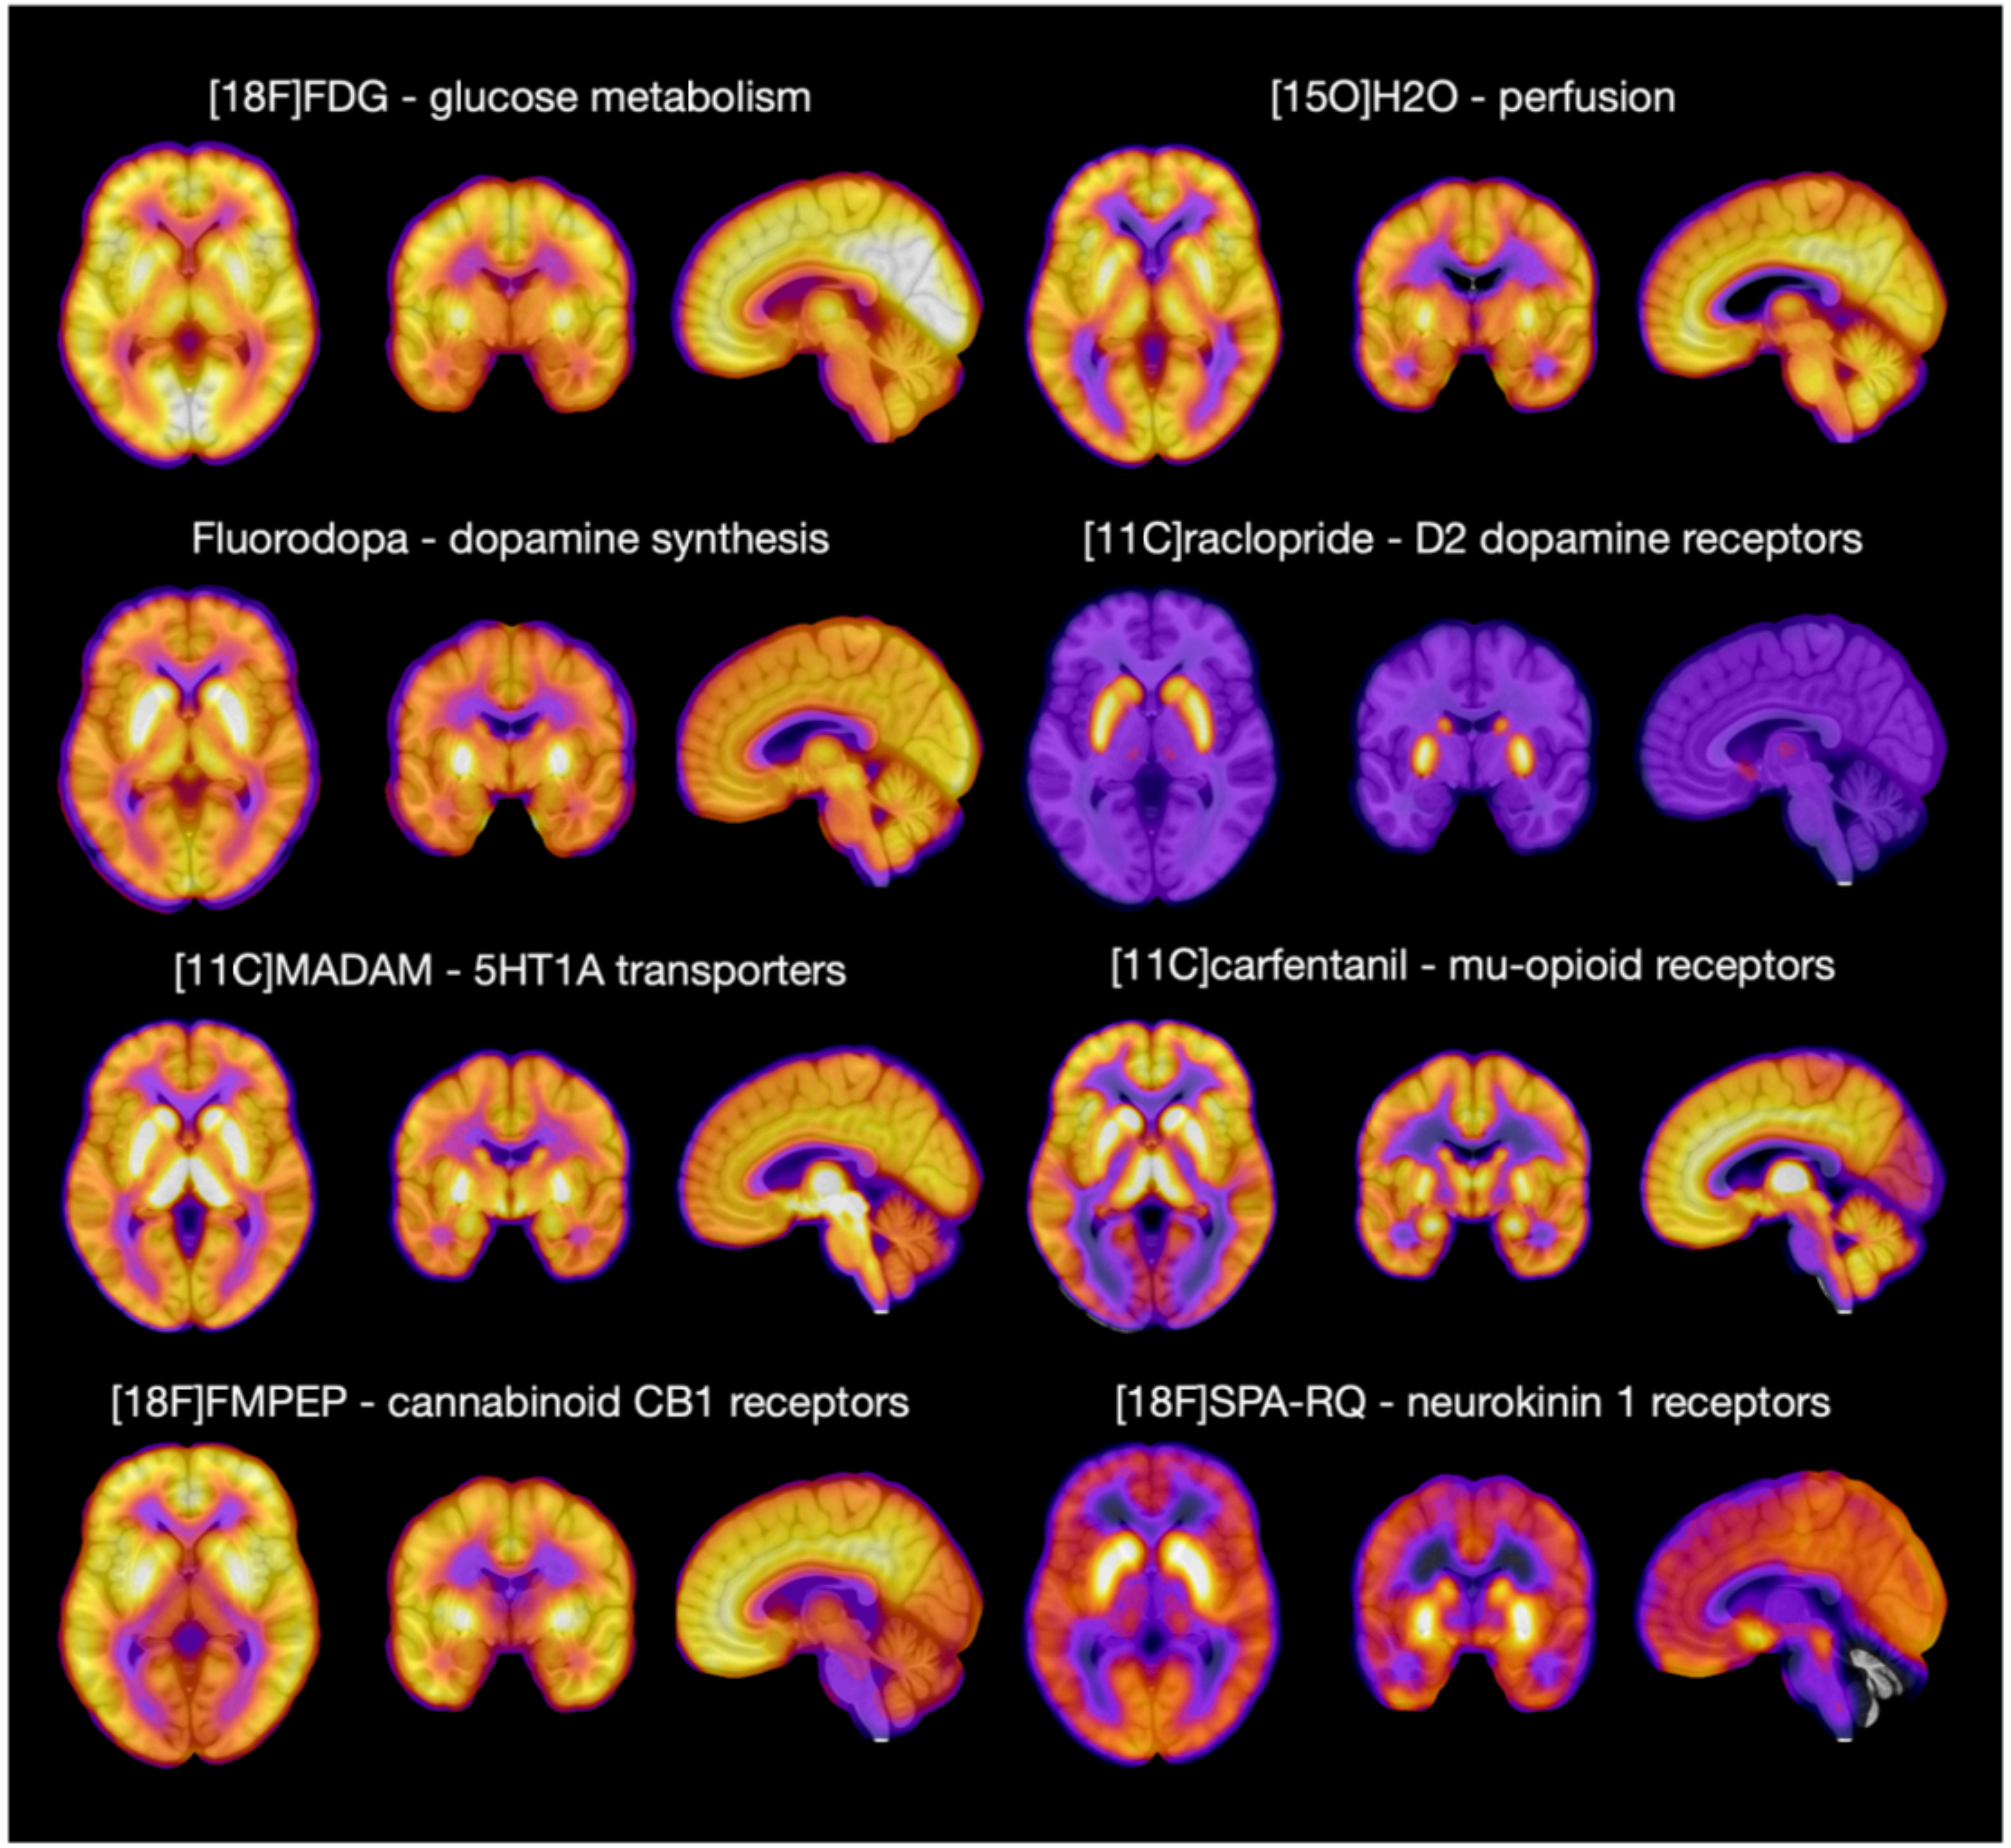

Brain region: fixed or random?

Can be either, depends on the goal

Some insights:

Figure from Nummenmaa, L., & Hirvonen, J. (preprint). Mapping the human emotion circuits with positron emission tomography. Human Emotion Systems Laboratory Preprints.

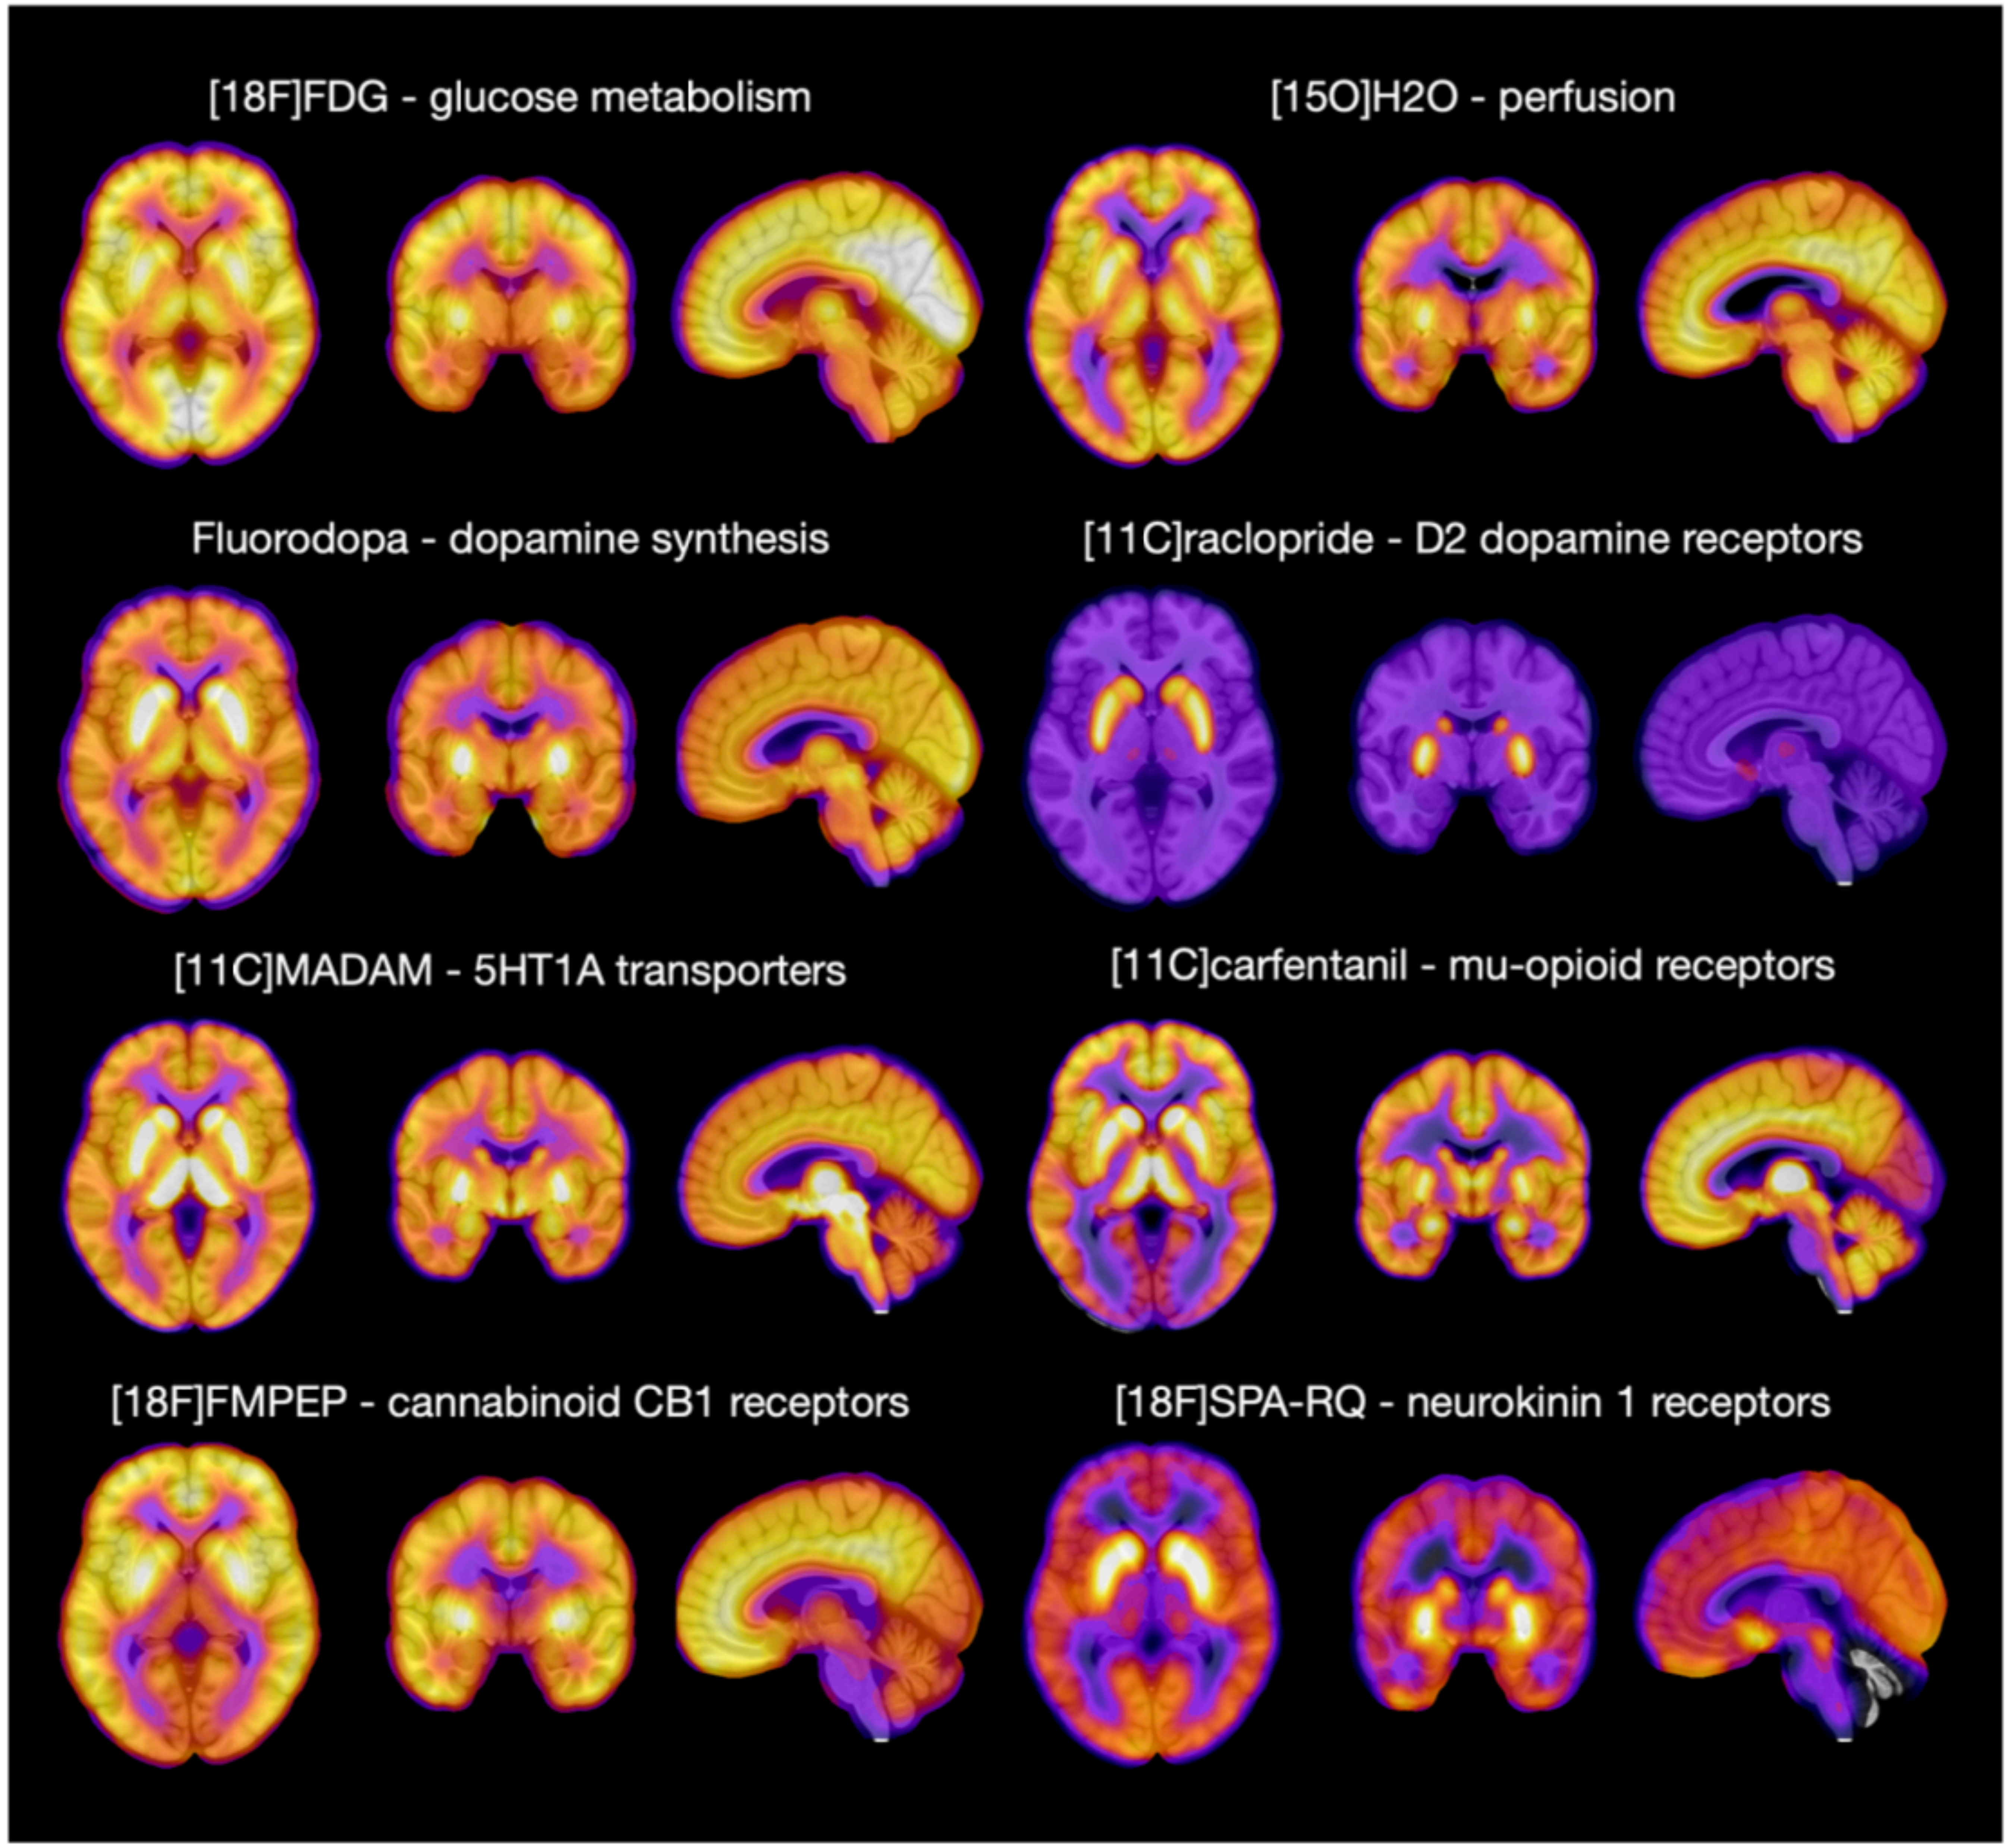

Brain region: fixed or random?

Can be either, depends on the goal

Some insights:

BP_{ND} ~ 1 + region Fixed

- Compare binding independently between regions

Figure from Nummenmaa, L., & Hirvonen, J. (preprint). Mapping the human emotion circuits with positron emission tomography. Human Emotion Systems Laboratory Preprints.

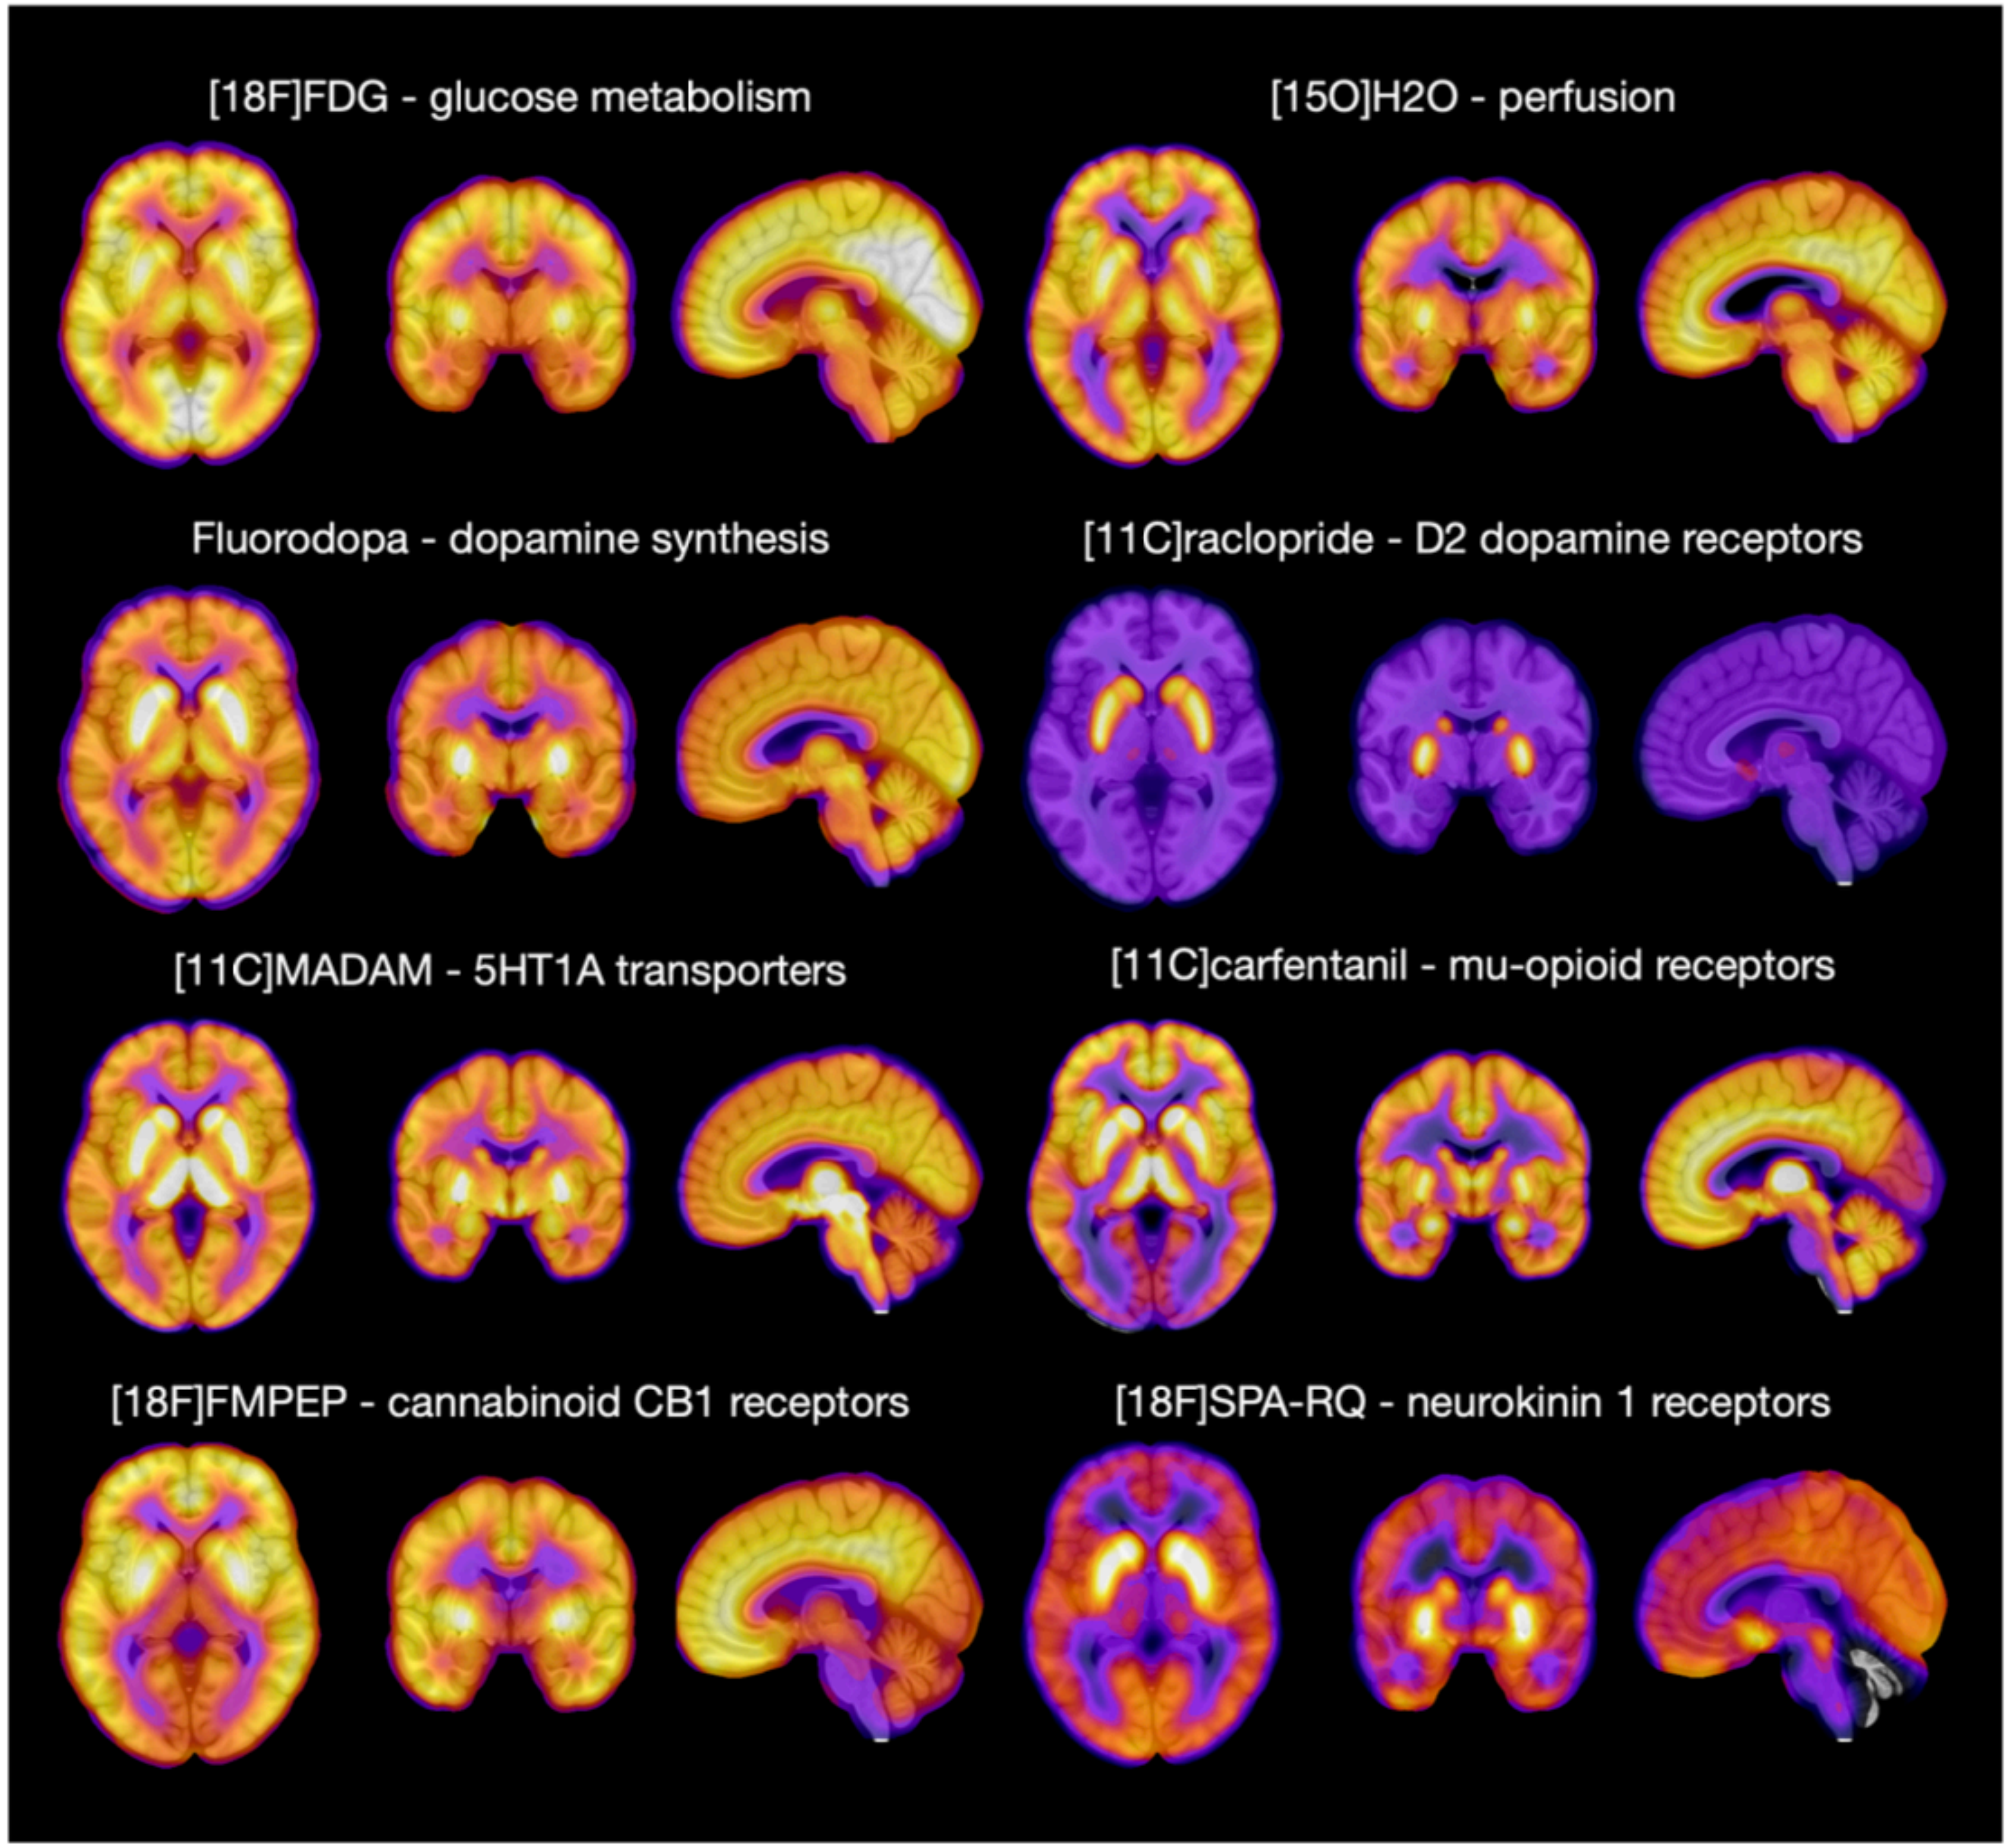

Brain region: fixed or random?

Can be either, depends on the goal

Some insights:

BP_{ND} ~ 1 + region Fixed

- Compare binding independently between regions

BP_{ND} ~ 1 + age + (1 + age | region) Random

- Intercept and age effect generally in the brain (across regions), how much regional variation on average

Figure from Nummenmaa, L., & Hirvonen, J. (preprint). Mapping the human emotion circuits with positron emission tomography. Human Emotion Systems Laboratory Preprints.

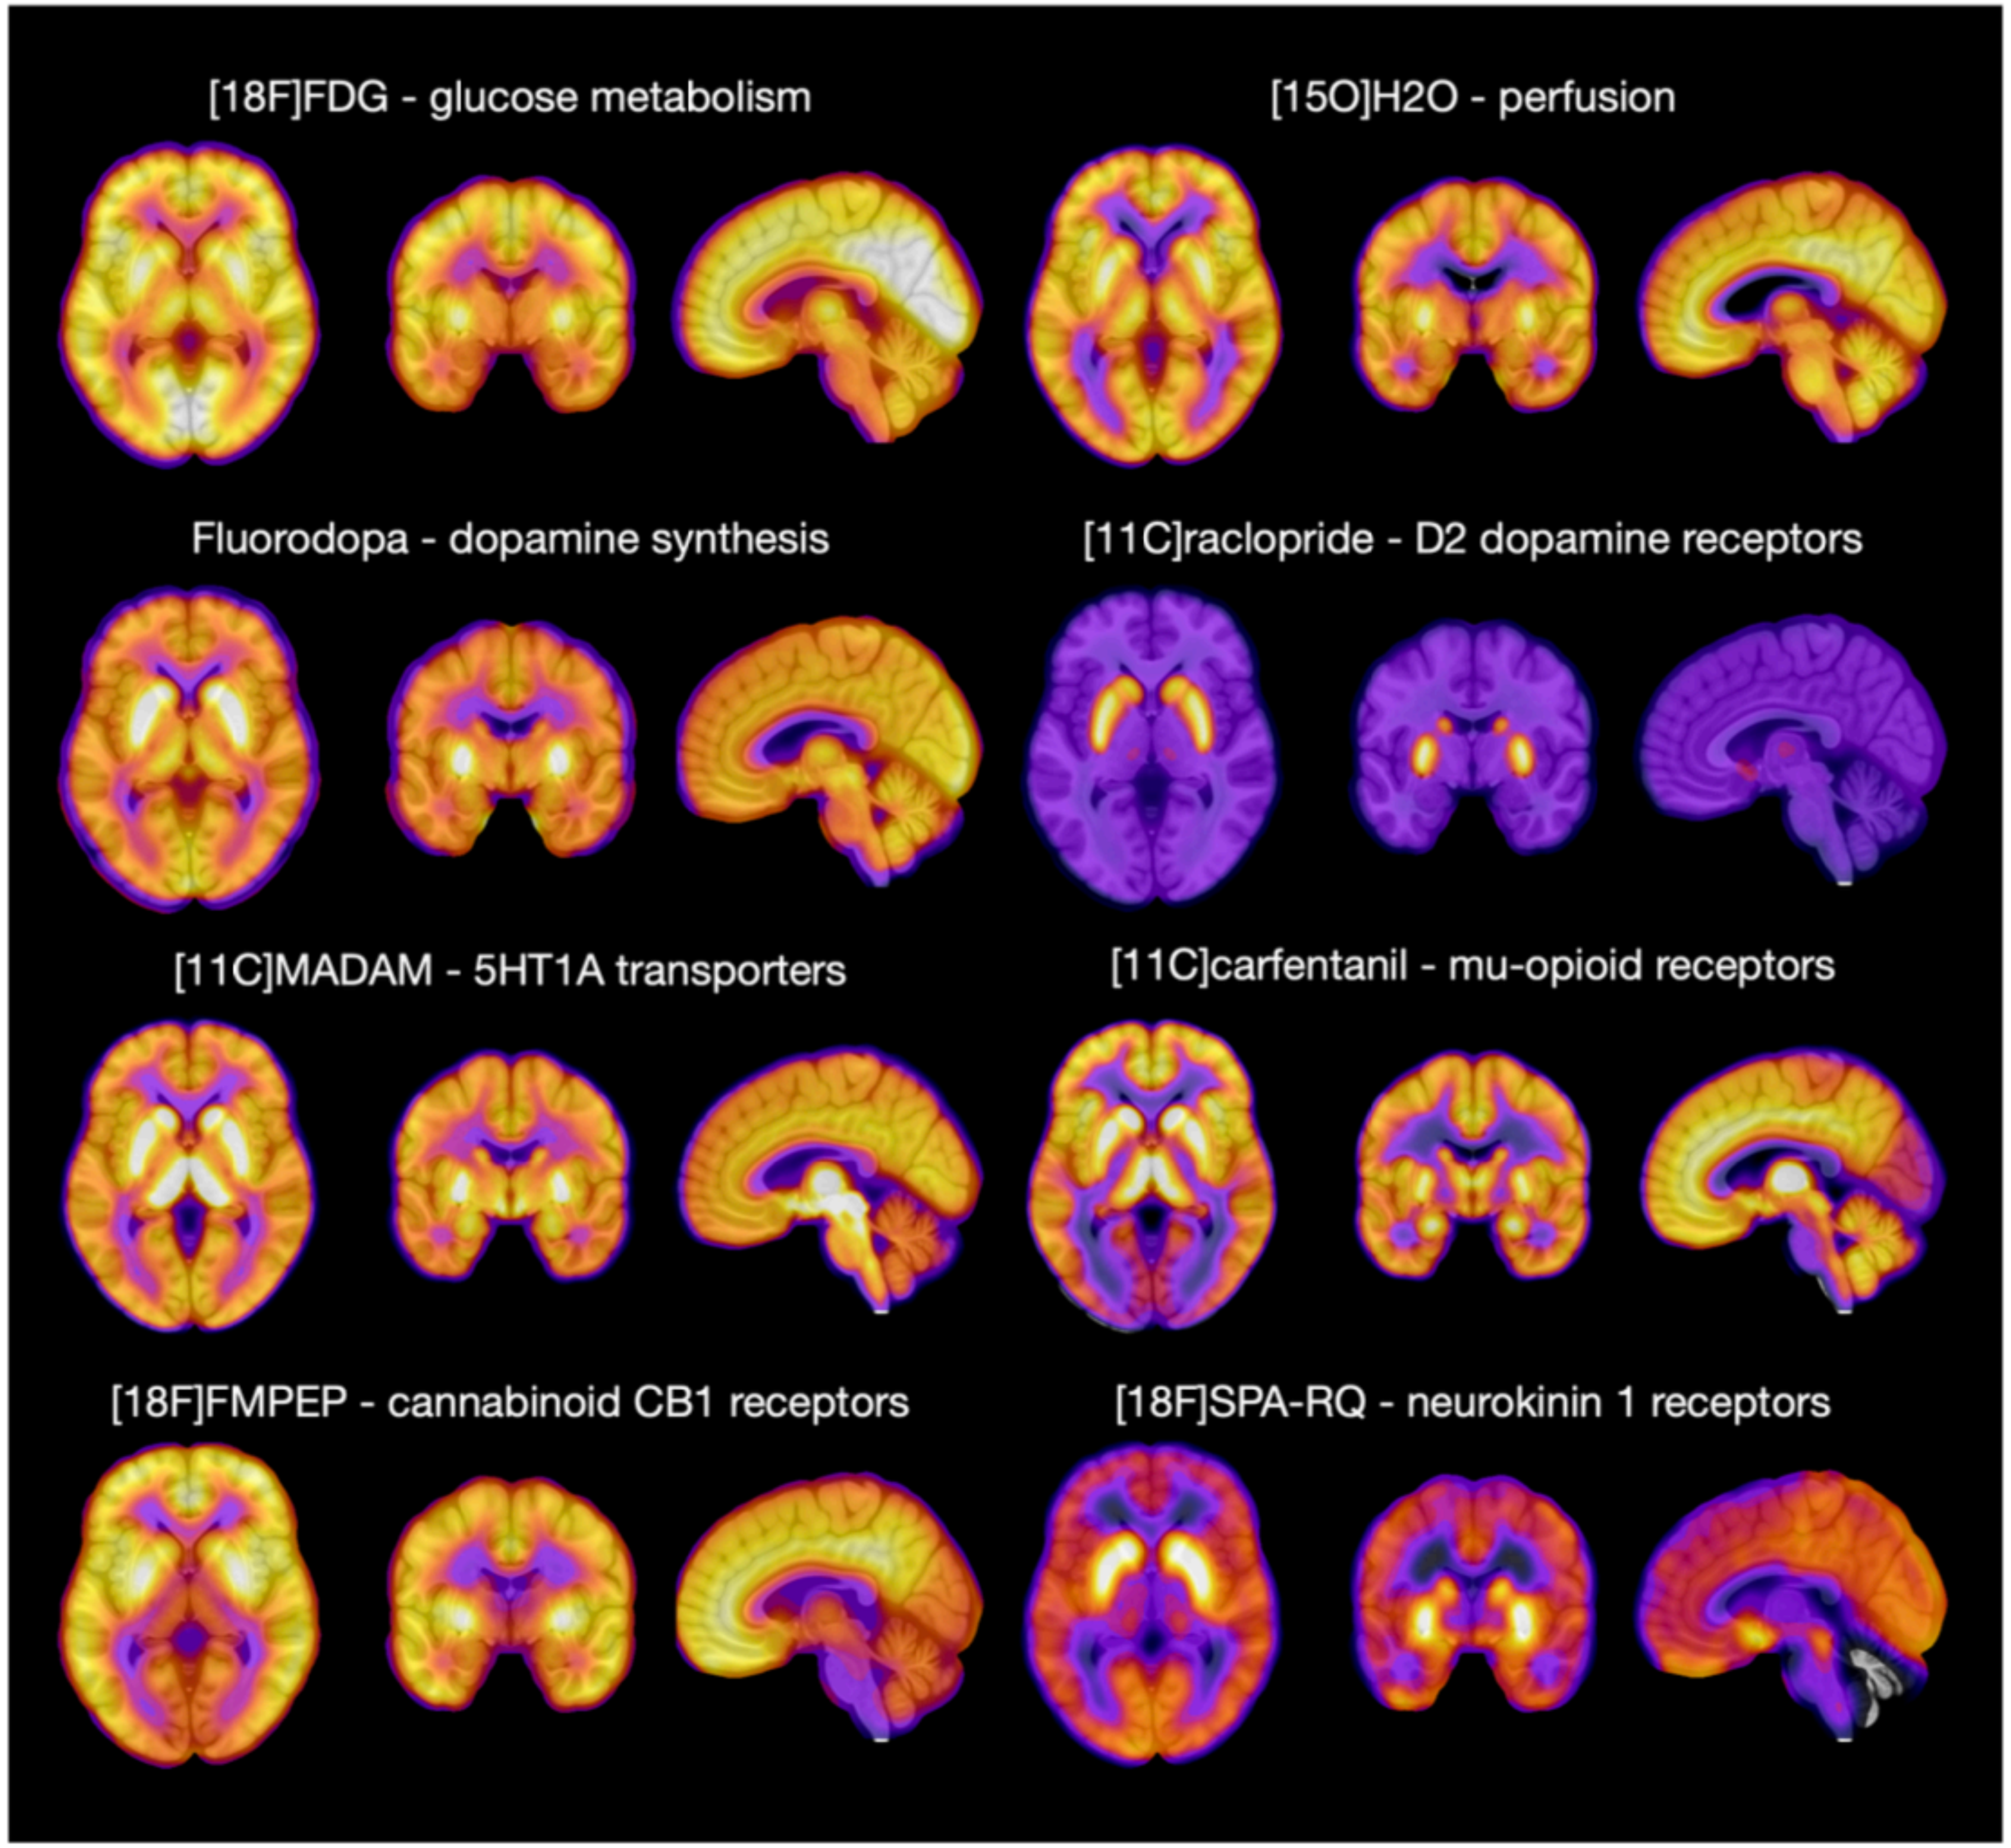

Brain region: fixed or random?

Can be either, depends on the goal

Some insights:

BP_{ND} ~ 1 + region Fixed

- Compare binding independently between regions

BP_{ND} ~ 1 + age + (1 + age | region) Random

- Intercept and age effect generally in the brain (across regions), how much regional variation on average
- Intercept and age effect separately for different regions, regularized by shrinkage (binding in certain region less reliable than others, utilize the estimates of the more reliable regions)

Figure from Nummenmaa, L., & Hirvonen, J. (preprint). Mapping the human emotion circuits with positron emission tomography. Human Emotion Systems Laboratory Preprints.

Diagnostics (the post-marathon run)

Residual size

Fixed effect

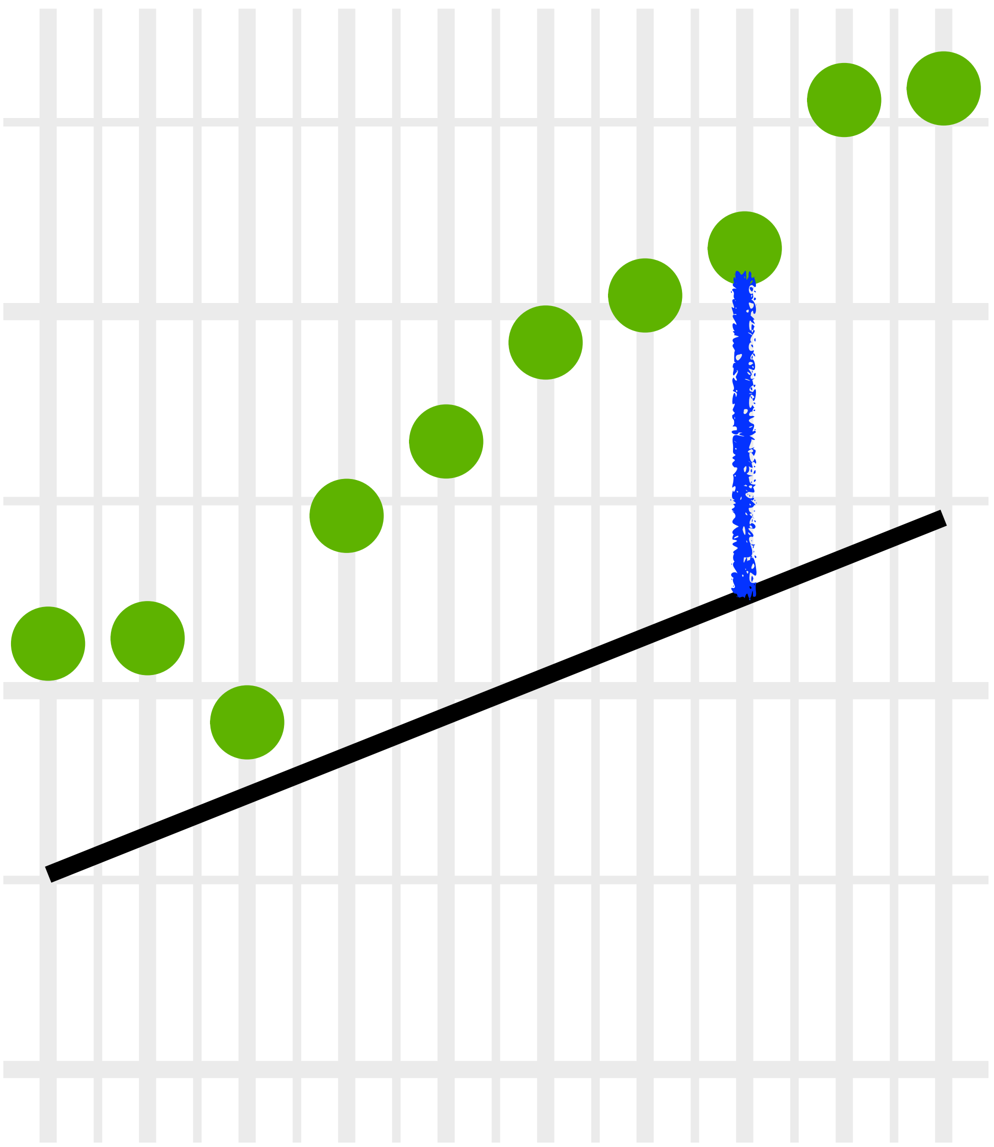

Reaction ~ 1 + Days

Compare models

Random intercept and slope

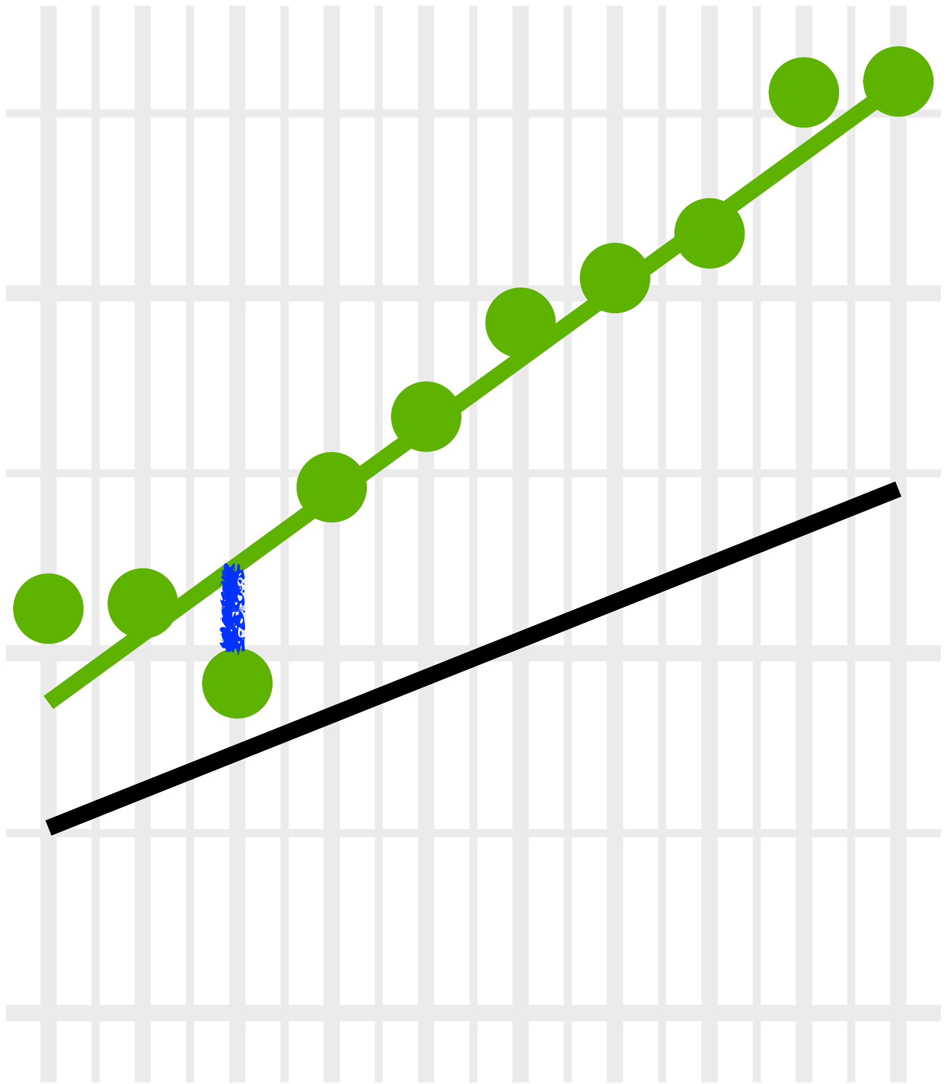

Reaction ~ 1 + Days + (1 + Days | Subject)

Residual distribution

No random effects:
Reaction ~ 1 + Days

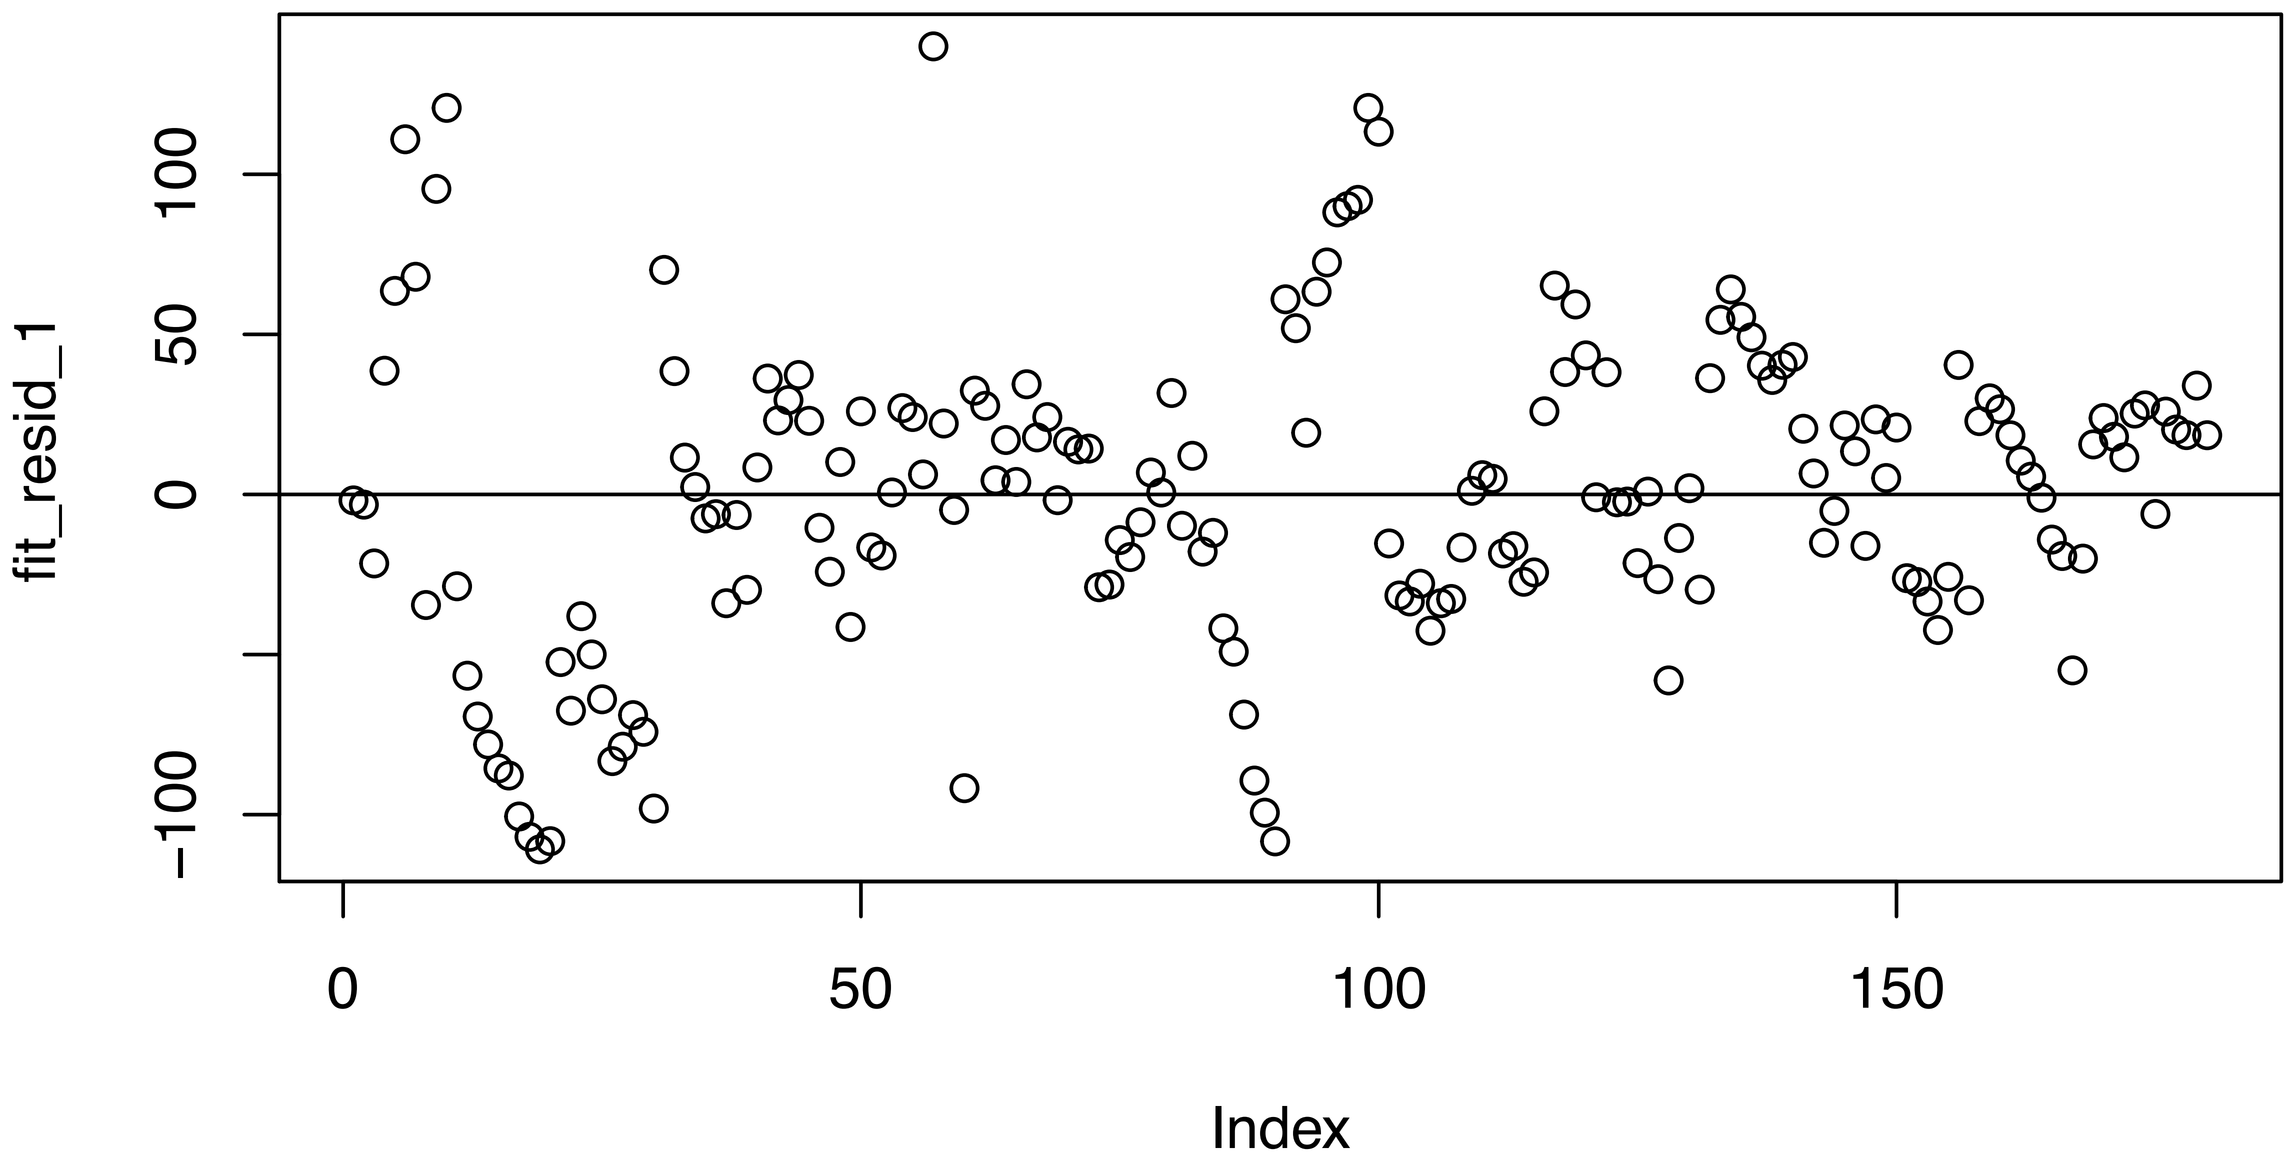

Subject as random intercept and slope:
Reaction ~ 1 + Days + (1 + Days | Subject)

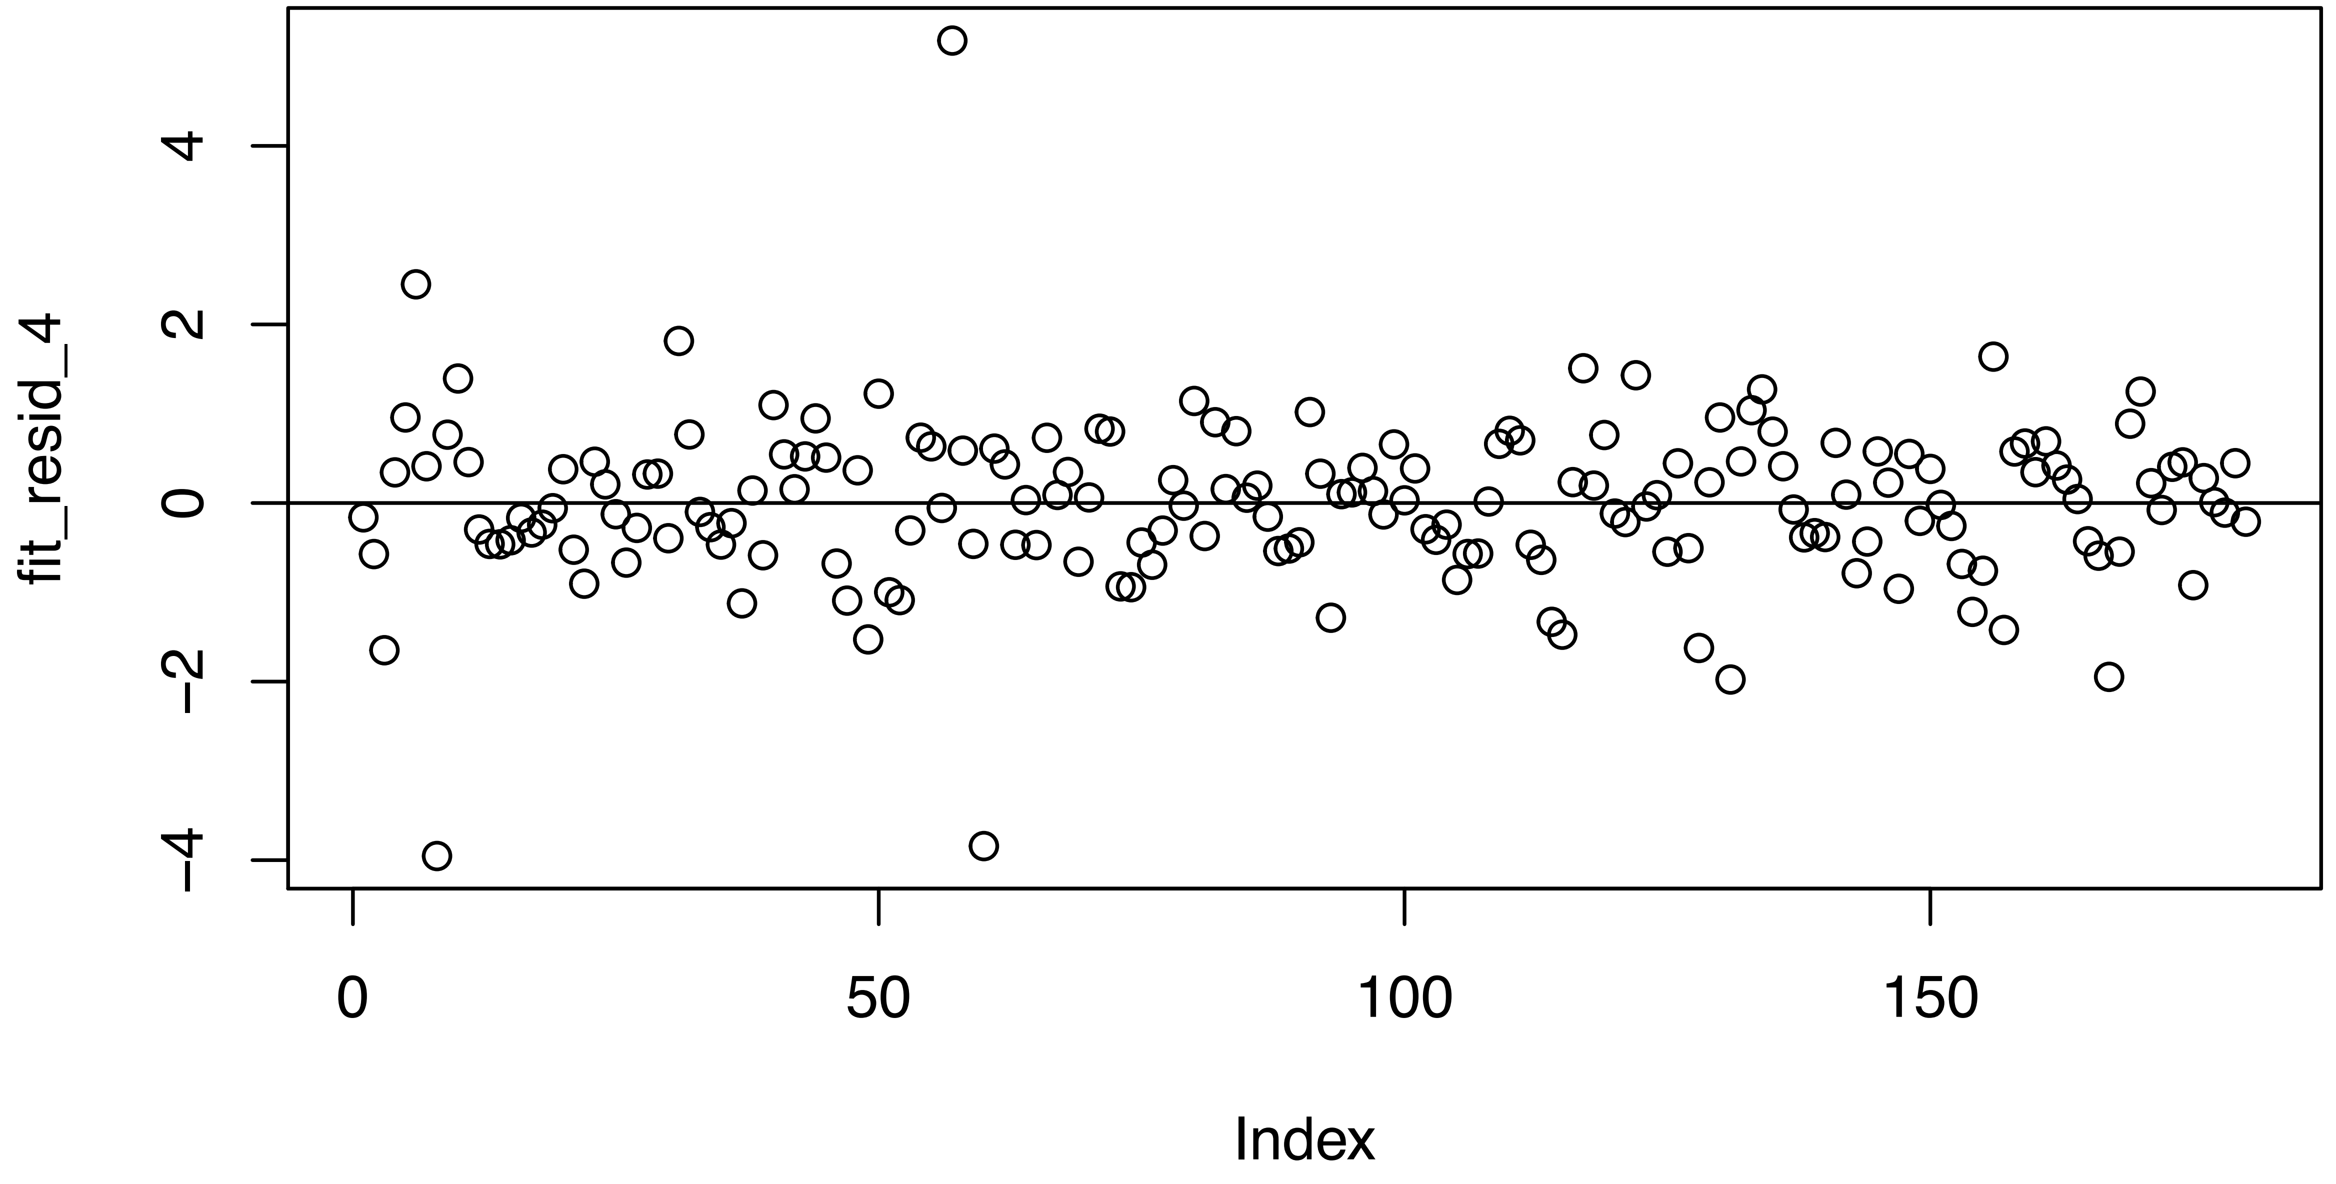


```
resid_1 <- resid(fit_1, scaled = TRUE)
plot(resid_1)
```

R package 'stats'

Residual distribution

No random effects:
Reaction ~ 1 + Days

Subject as random intercept and slope:
Reaction ~ 1 + Days + (1 + Days | Subject)

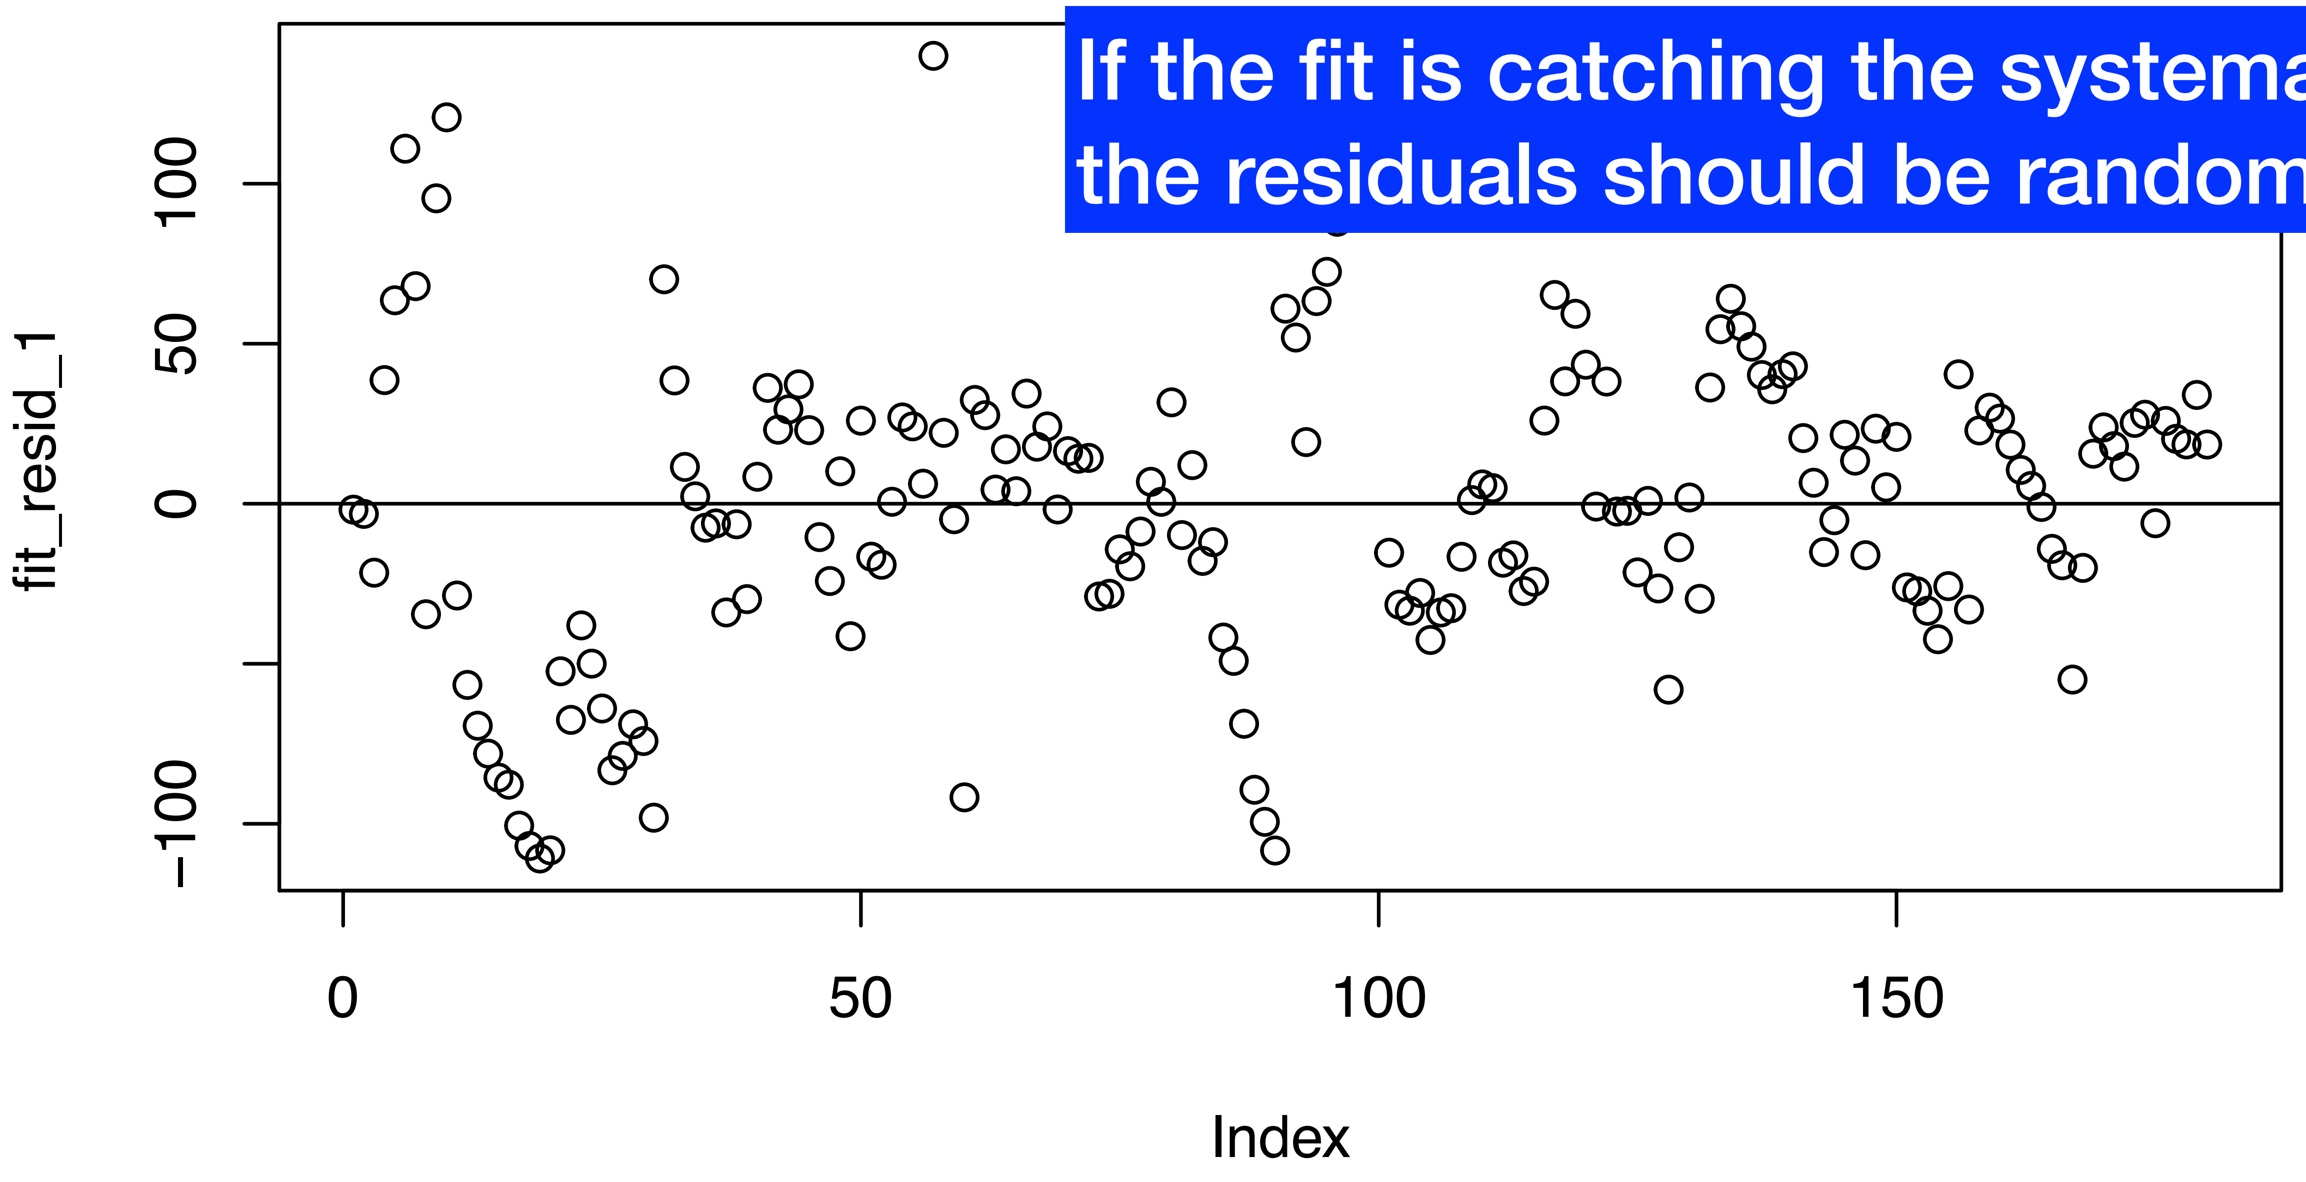

If the fit is catching the systematic variation in the data (fit is good),
the residuals should be randomly (not systematically) distributed around the fit

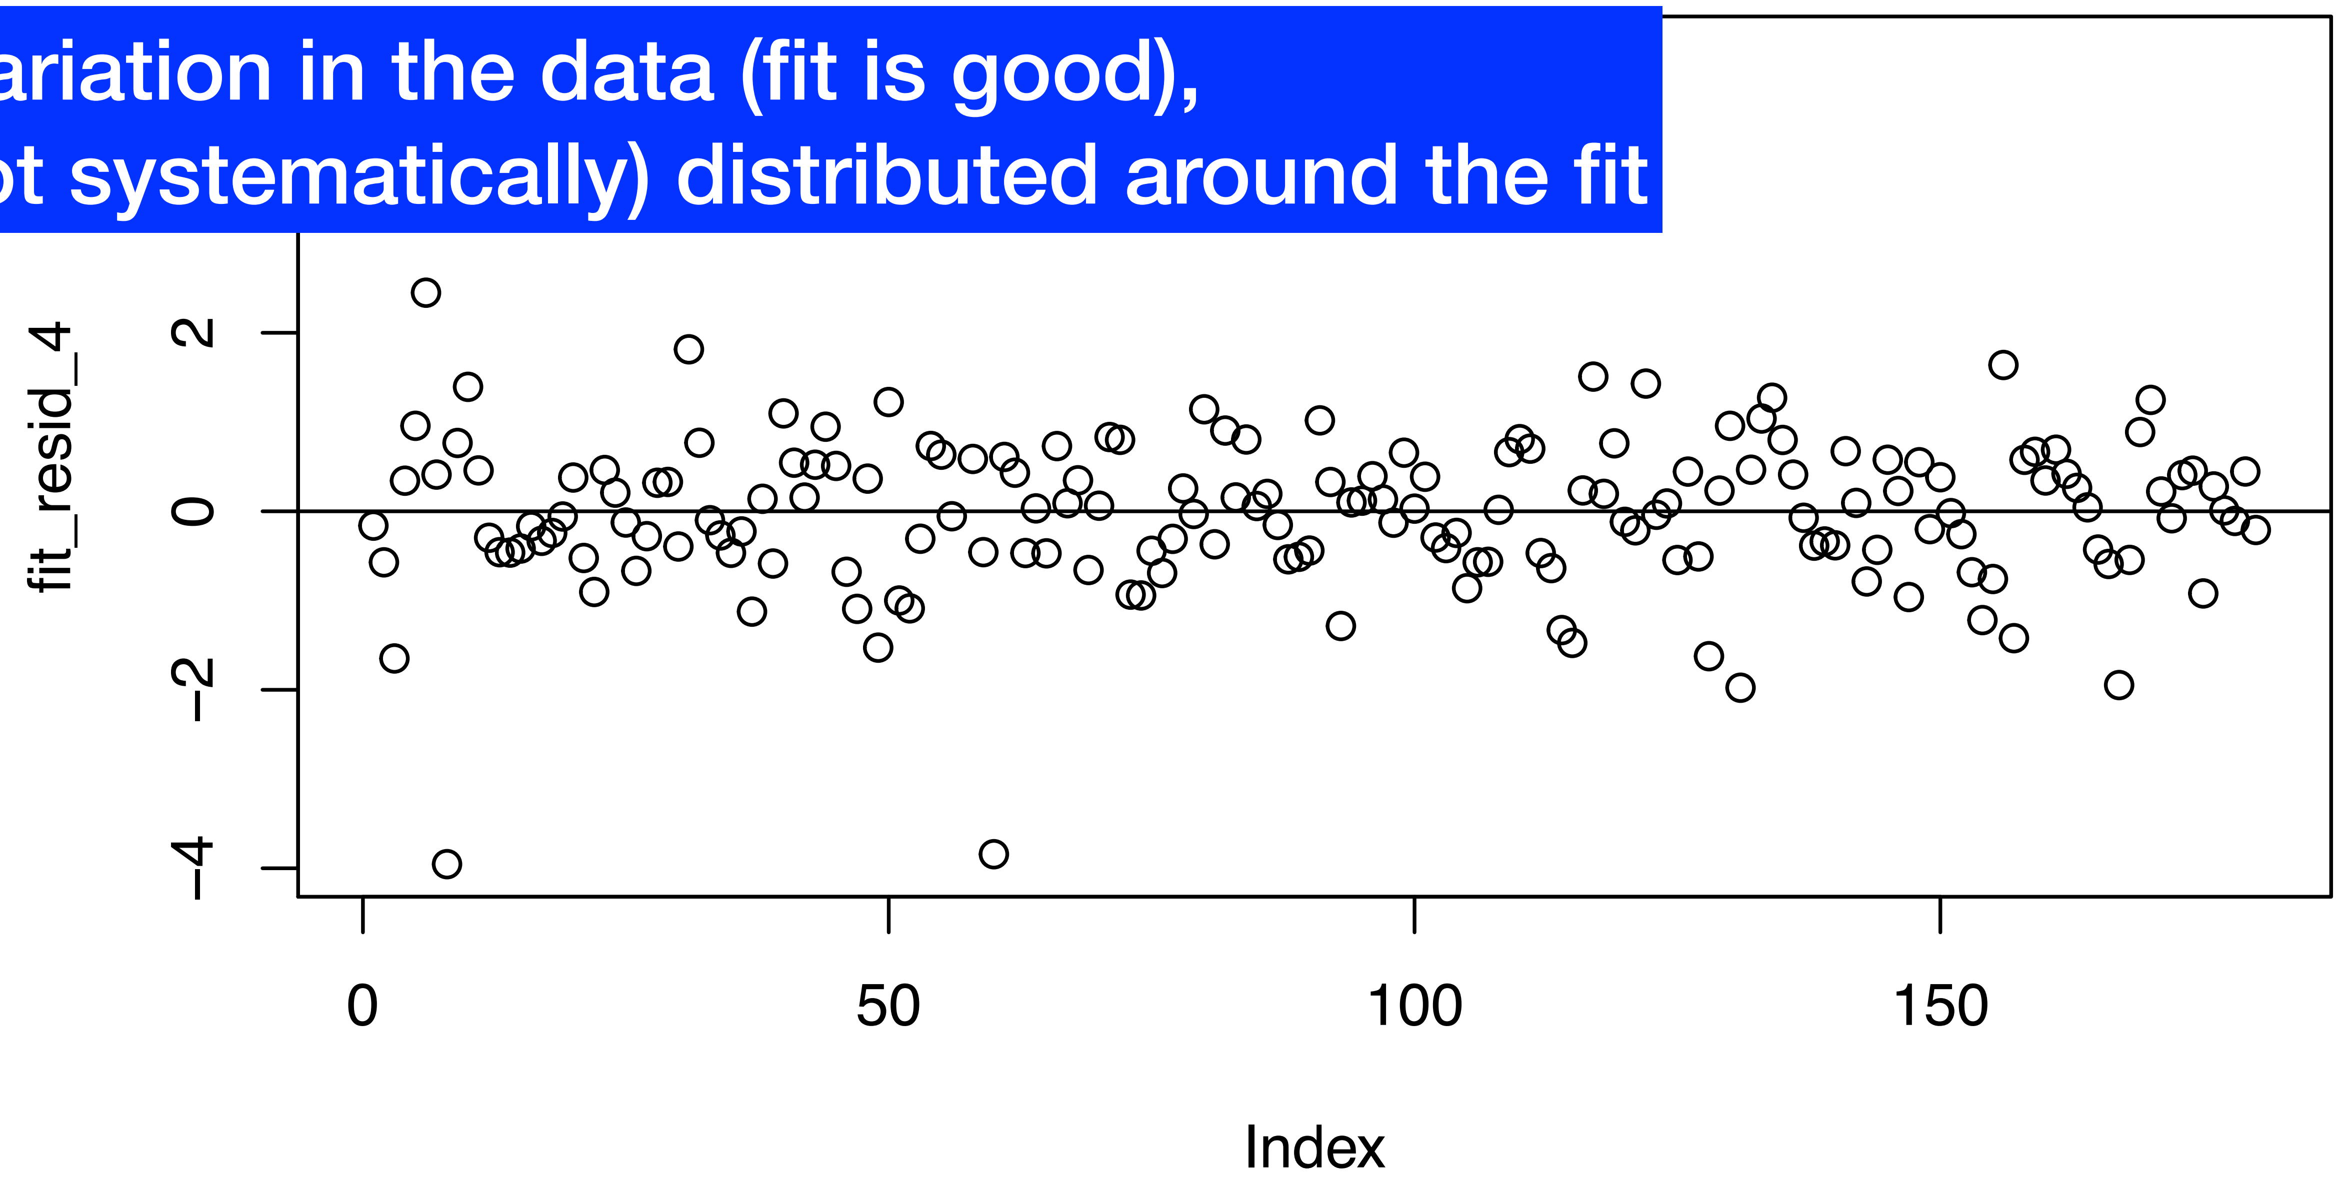

resid_1 <- resid(fit_1, scaled = TRUE)
plot(resid_1)

R package 'stats'

To take home

Linear mixed effects model is one type of linear model

What makes a linear model mixed: random effect(s) in addition to the fixed effect(s)

Mixed effects modeling is useful when we have grouping structure in the data (e.g. grouped by subject, scanner, site)

Thanks!

- **Thank you for your kind attention!**

- **Colleagues from Nummenmaa Lab, particularly Birgitta Paranko and Severi Santavira**

- Academy Research Fellow / Assistant Professor **Joni Virta**
  - UTU Course: Sekamallit (Mixed effects models)

References

- Harrison, X. A., Donaldson, L., Correa-Cano, M. E., Evans, J., Fisher, D. N., Goodwin, C. E., ... & Inger, R. (2018). A brief introduction to mixed effects modelling and multi-model inference in ecology. *PeerJ*, *6*, e4794.

- McElreath, R. (2020). *Statistical rethinking: A Bayesian course with examples in R and Stan*. CRC press.

Courses at the University of Turku, mainly:
- **Mixed effects (Sekamallit)** course in the University of Turku, Department of Mathematics and Statistics, by Academy Research Fellow / Assistant Professor **Joni Virta**
- Regression analysis and statistical learning (Regressioanalyysi ja tilastollinen oppiminen), by Professor **Henri Nyberg**

- R Core Team. (2021). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL https:/ / www.R-project.org/ .

- Douglas Bates, Martin Maechler, Ben Bolker, Steve Walker (2015). Fitting Linear Mixed-Effects Models Using lme4. Journal of Statistical Software, 67(1), 1-48. doi:10.18637/jss.v067.i01.

- Pinheiro J, Bates D, DebRoy S, Sarkar D, R Core Team (2021). _nlme: Linear and Nonlinear Mixed Effects Models_. R package version 3.1-152, <URL: https://CRAN.R-project.org/package=nlme>.

- Figures made with ggplot2: H. Wickham. ggplot2: Elegant Graphics for Data Analysis. Springer-Verlag New York, 2016.

Bonus material

Differences between the performances of machines A, B and C?
Tested with 6 workers (random sample), 3 observations from each

$$y_{ijk} = \mu + \beta_i + \boxed{b_j} + \varepsilon_{ijk}$$

y= fit

Fixed effects:

μ (intercept)

β_B, β_C (slopes)

Random intercepts for workers:

$b_j \sim N(0, \tau^2)$

Error terms (residuals):

$\varepsilon_{ijk} \sim N(0, \sigma^2)$

Parameters:

μ, β_B, β_C, τ^2, σ^2

$$\frac{\tau^2 = \text{variance between the workers}}{\sigma^2 = \text{residual variance (variance within workers)}} > 1$$

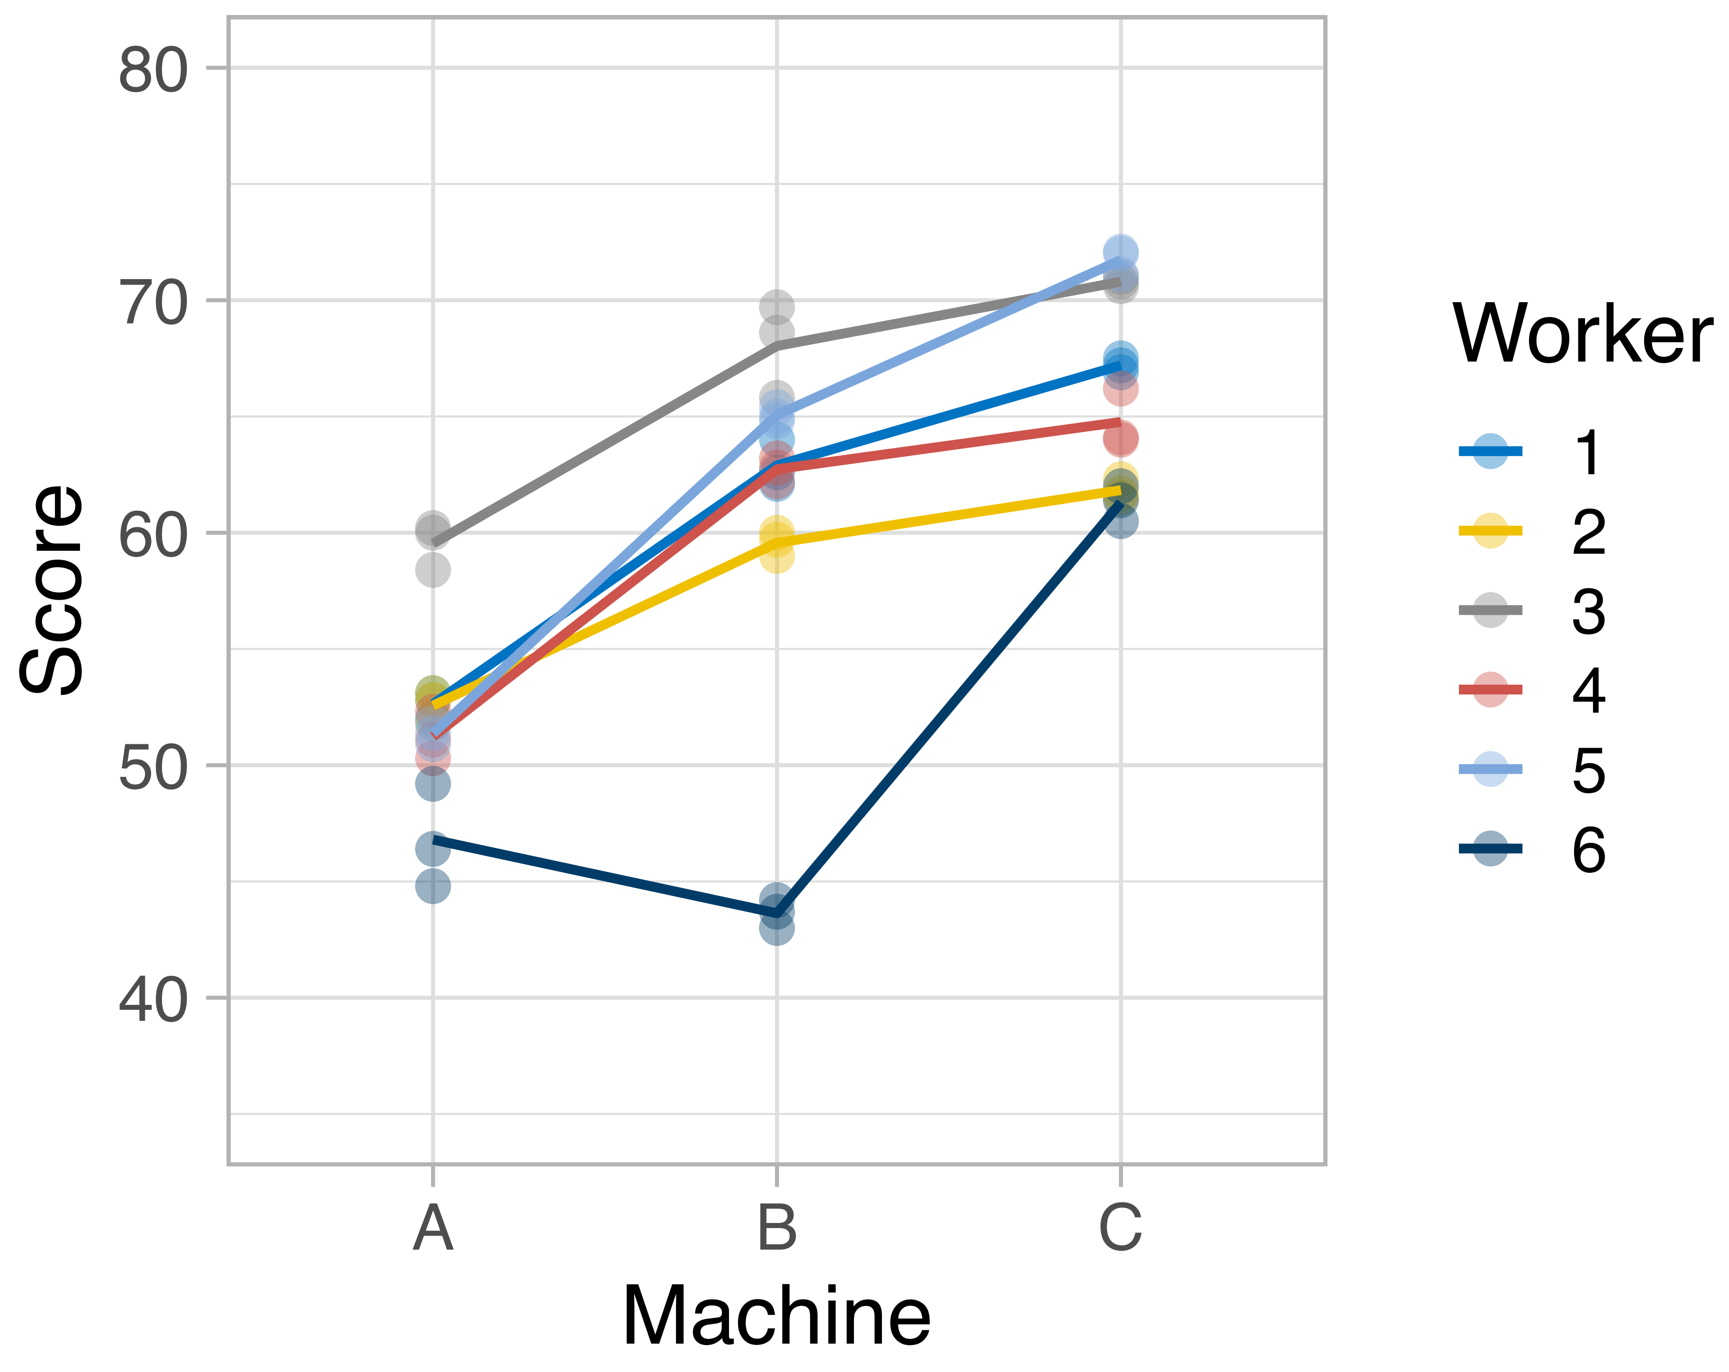

i= machines: A, B, C

j= workers: 1, 2, 3, 4, 5, 6

k= repeated measures: day1, day2, day3

Machines dataset from R package 'nlme'

Model with only fixed effects = General linear model

Population-level intercept

$$y = \alpha + \beta \cdot day$$

Population-level effect of days

Model with fixed and (1 or more) random effects= (Linear) mixed effects model

Random (subject-specific) effect of days

$$y = (\alpha + b_{\alpha sub}) + (\beta + b_{\beta sub}) \cdot day$$

Random (subject-specific) intercept