Lecture 5

Random Effects Models

2/01/2018

Random Effects Models

Sleep Study Data

The 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. The observations represent the average reaction time on a series of tests given each day to each subject.

```
sleep = lme4::sleepstudy %>% tbl_df()

sleep
## # A tibble: 180 x 3
##    Reaction  Days Subject
##       <dbl> <dbl> <fct>
##  1      250  0    308
##  2      259  1.00 308
##  3      251  2.00 308
##  4      321  3.00 308
##  5      357  4.00 308
##  6      415  5.00 308
##  7      382  6.00 308
##  8      290  7.00 308
##  9      431  8.00 308
## 10      466  9.00 308
## # ... with 170 more rows
```

EDA

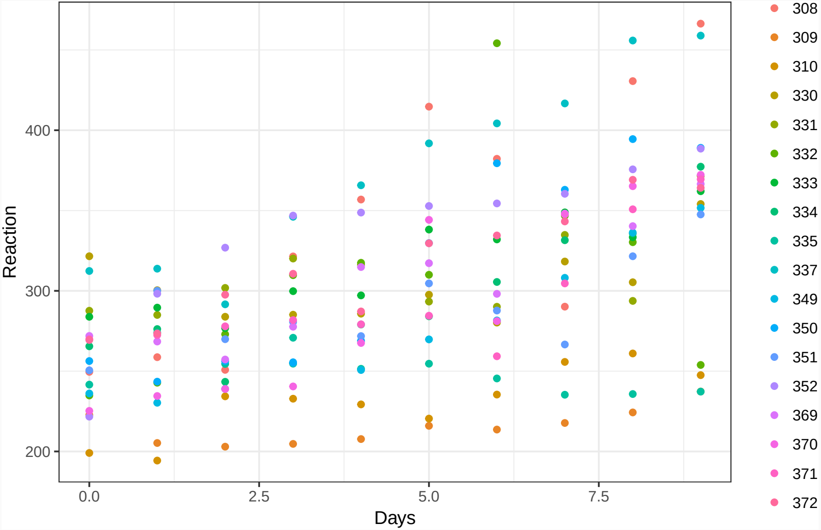

Reaction
400
300
200
0.0
2.5
5.0
7.5
Days
308
309
310
330
331
332
333
334
335
337
349
350
351
352
369
370
371
372

EDA (small multiples)

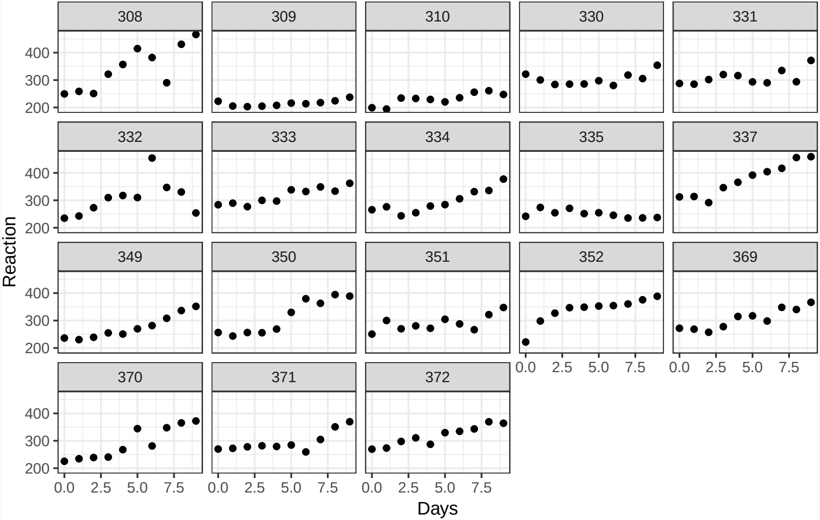

Bayesian Linear Model

```
sleep_lm = "model{
  # Likelihood
  for(i in 1:length(y)){
    y[i] ~ dnorm(mu[i], tau)
    mu[i] = beta[1] + beta[2]*x[i]

    y_pred[i] ~ dnorm(mu[i],tau)
  }

  # Prior for beta
  beta[1] ~ dnorm(0,1/10000)
  beta[2] ~ dnorm(0,1/10000)

  # Prior for sigma / tau
  sigma ~ dunif(0, 100)
  tau = 1/(sigma*sigma)
}"
```

MCMC Diagnostics

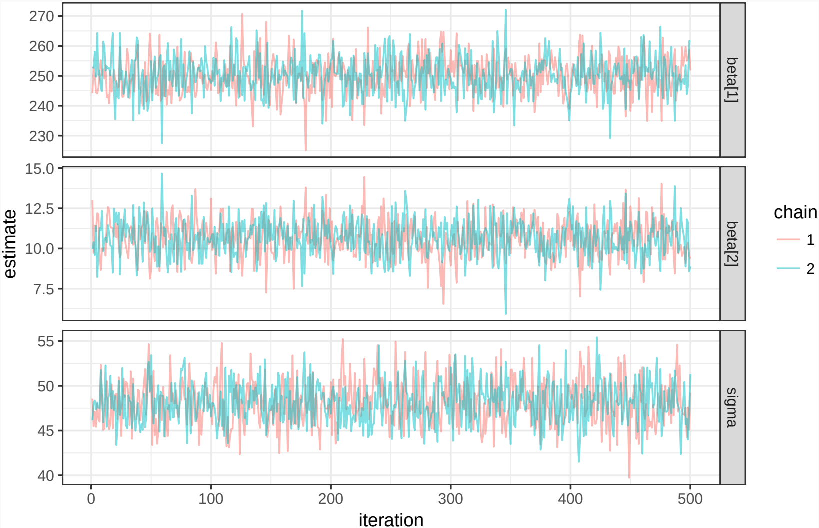

Model fit

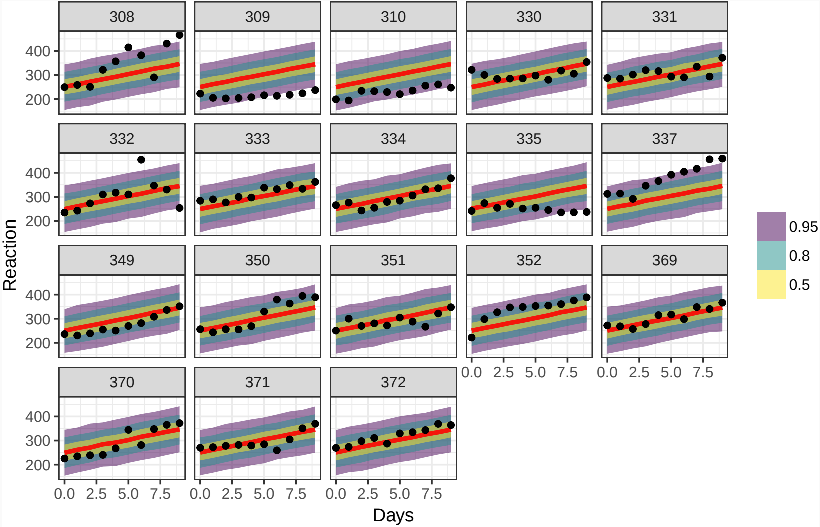
308
309
310
330
331
332
333
334
335
337
349
350
351
352
369
370
371
372
Reaction
400
300
200
0.0
2.5
5.0
7.5
Days
0.95
0.8
0.5

Residuals

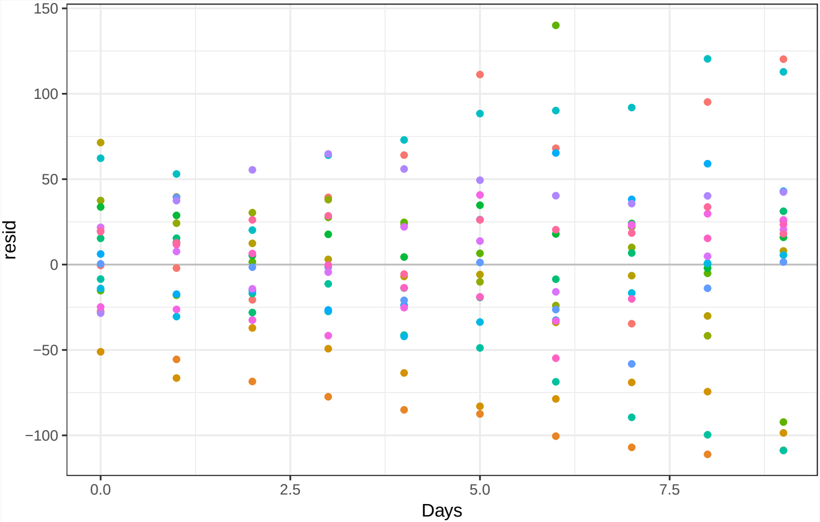

Residuals by subject

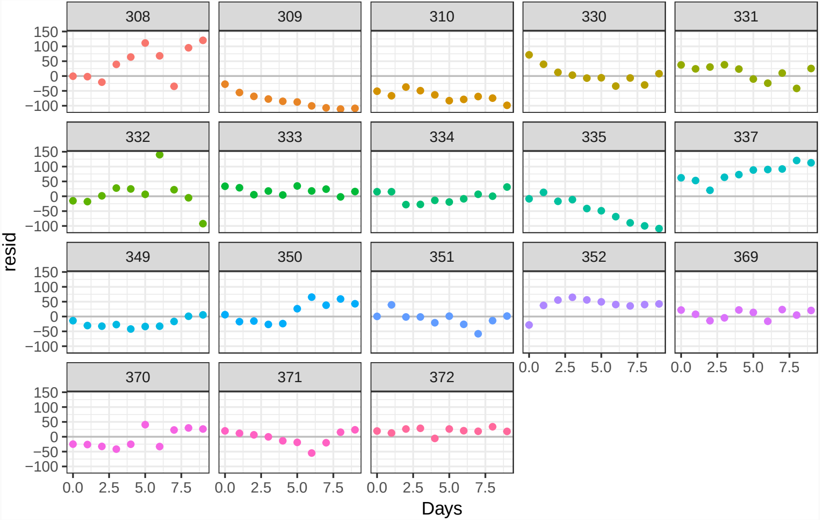

Random Intercept Model

Coding

```
sleep = sleep %>%
  mutate(Subject_index = as.integer(Subject))

sleep[c(1:2,11:12,21:22,31:32),]
## # A tibble: 8 x 4
##   Reaction  Days Subject Subject_index
##      <dbl> <dbl> <fct>           <int>
## 1      250  0    308                 1
## 2      259  1.00 308                 1
## 3      223  0    309                 2
## 4      205  1.00 309                 2
## 5      199  0    310                 3
## 6      194  1.00 310                 3
## 7      322  0    330                 4
## 8      300  1.00 330                 4
```

Let i represent each observation and $j(i)$ be subject in oberservation i then

$$y_i = \alpha_{j(i)} + \beta \times \text{Days} + \epsilon_i$$

$$\begin{aligned} \alpha_j &\sim \mathcal{N}(\beta_\alpha,\ \sigma_\alpha^2) \\ \epsilon_i &\sim \mathcal{N}(0,\ \sigma^2) \end{aligned}$$

$$\begin{aligned} \beta_\alpha &\sim \mathcal{N}(0, 10^4) \\ \beta &\sim \mathcal{N}(0, 10^4) \\ \sigma, \sigma_\alpha &\sim \text{Unif}(0, 10^2) \end{aligned}$$

Random Intercept Model - JAGS

```
sleep_ri = "model{
  for(i in 1:length(Reaction)) {
    Reaction[i] ~ dnorm(mu[i],tau)
    mu[i] = alpha[Subject_index[i]] + beta*Days[i]

    y_pred[i] ~ dnorm(mu[i],tau)
  }

  for(j in 1:18) {
    alpha[j] ~ dnorm(beta_alpha, tau_alpha)
  }

  beta_alpha ~ dnorm(0,1/10000)
  sigma_alpha ~ dunif(0, 100)
  tau_alpha = 1/(sigma_alpha*sigma_alpha)


  beta ~ dnorm(0,1/10000)
  sigma ~ dunif(0, 100)
  tau = 1/(sigma*sigma)
}"
```

MCMC Diagnostics

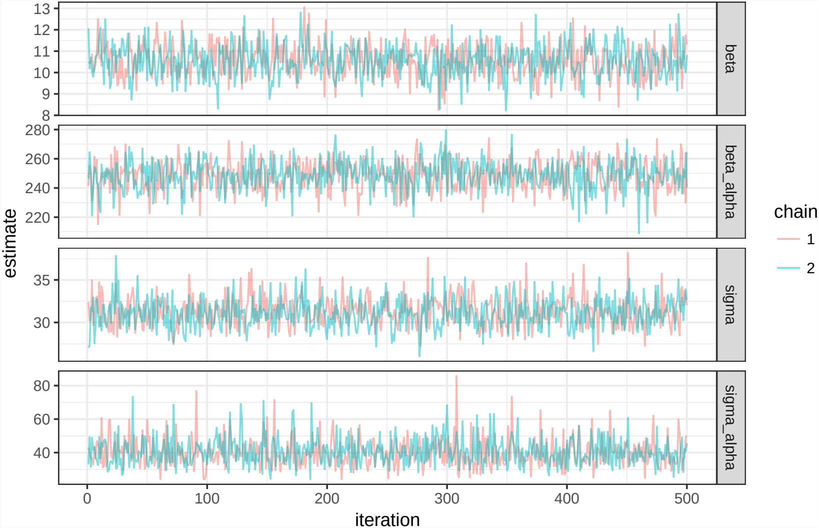

Model fit

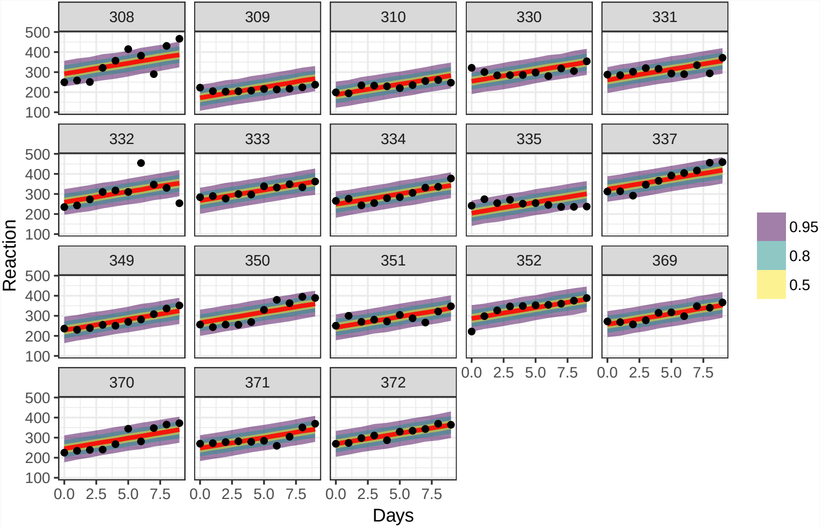
308
309
310
330
331
332
333
334
335
337
349
350
351
352
369
370
371
372
Reaction
Days
500
400
300
200
100
0.0
2.5
5.0
7.5
0.95
0.8
0.5

Model fit - lines

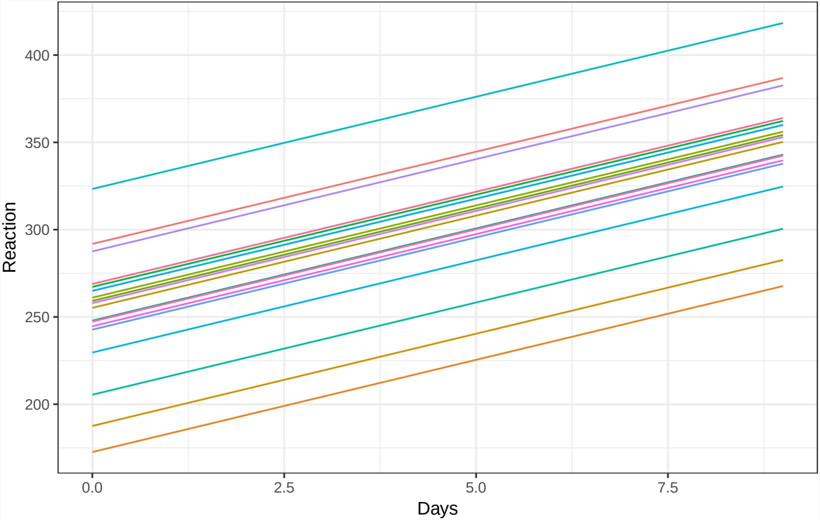

Random Effects

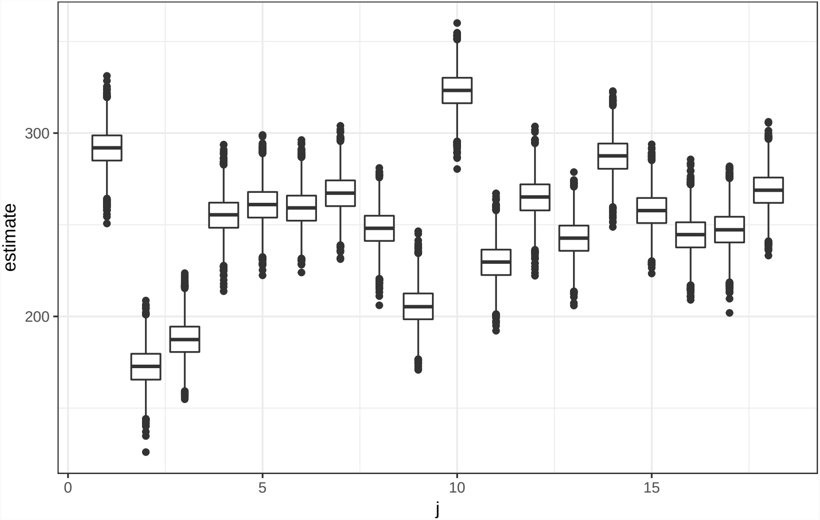

Residuals

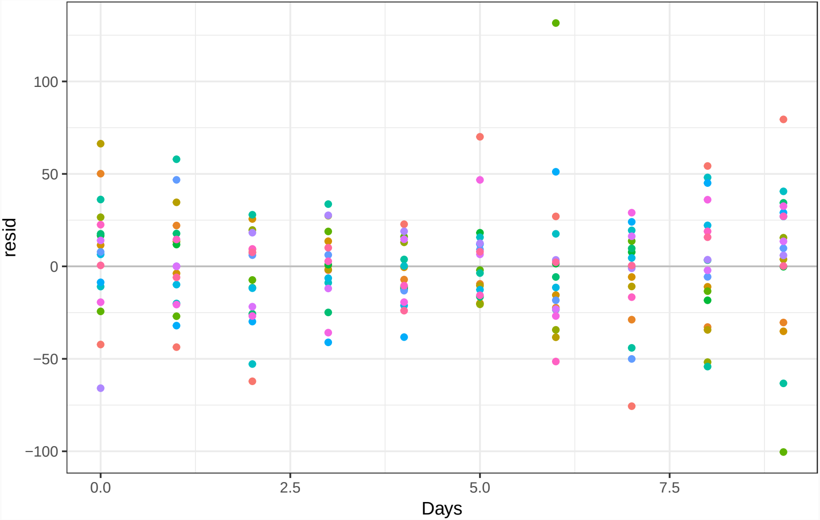

Residuals by subject

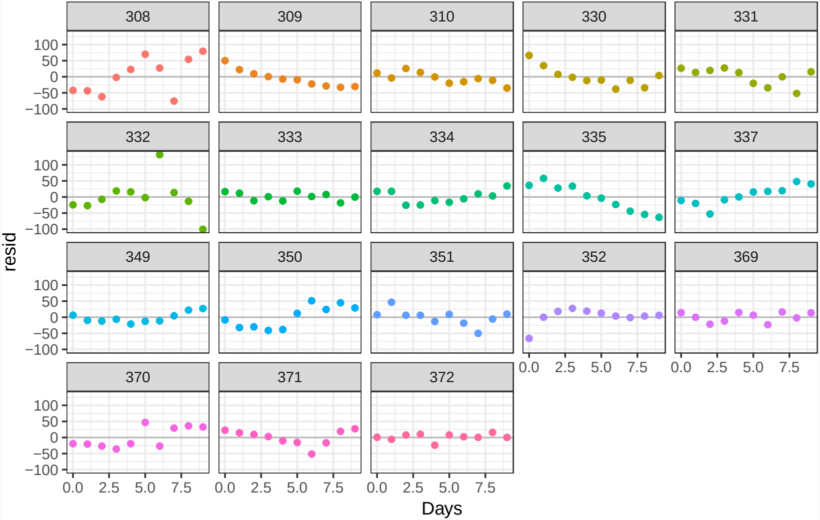

Why not a fixed effect for Subject?

We are not going to bother with the Bayesian model here to avoid the headache of dummy coding and the resulting βs.

Why not a fixed effect for Subject?

We are not going to bother with the Bayesian model here to avoid the headache of dummy coding and the resulting βs.

```
l = lm(Reaction ~ Days + Subject - 1, data=sleep)
summary(l)
##
## Call:
## lm(formula = Reaction ~ Days + Subject - 1, data = sleep)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -100.540  -16.389   -0.341   15.215  131.159
##
## Coefficients:
##            Estimate Std. Error t value Pr(>|t|)
## Days        10.4673     0.8042   13.02   <2e-16 ***
## Subject308 295.0310    10.4471   28.24   <2e-16 ***
## Subject309 168.1302    10.4471   16.09   <2e-16 ***
## Subject310 183.8985    10.4471   17.60   <2e-16 ***
## Subject330 256.1186    10.4471   24.52   <2e-16 ***
## Subject331 262.3333    10.4471   25.11   <2e-16 ***
## Subject332 260.1993    10.4471   24.91   <2e-16 ***
## Subject333 269.0555    10.4471   25.75   <2e-16 ***
## Subject334 248.1993    10.4471   23.76   <2e-16 ***
## Subject335 202.9673    10.4471   19.43   <2e-16 ***
## Subject337 328.6182    10.4471   31.45   <2e-16 ***
## Subject349 228.7317    10.4471   21.89   <2e-16 ***
## Subject350 266.4999    10.4471   25.51   <2e-16 ***
## Subject351 242.9950    10.4471   23.26   <2e-16 ***
## Subject352 290.3188    10.4471   27.79   <2e-16 ***
## Subject369 258.9319    10.4471   24.79   <2e-16 ***
## Subject370 244.5990    10.4471   23.41   <2e-16 ***
## Subject371 247.8813    10.4471   23.73   <2e-16 ***
## Subject372 270.7833    10.4471   25.92   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Comparing Model fit

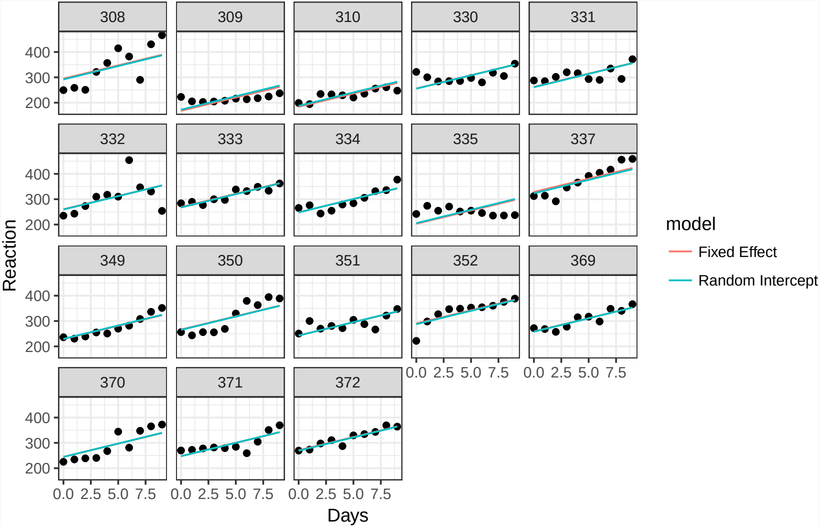

Random effects vs fixed effects

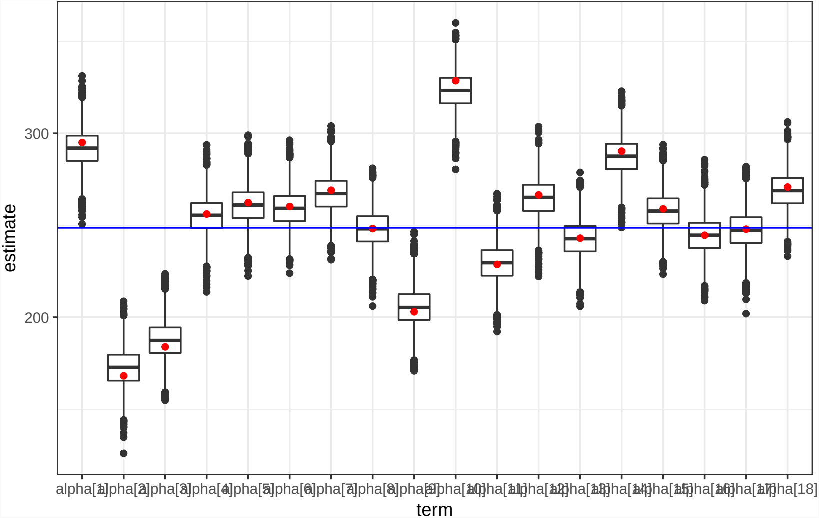

Random Intercept Model (strong prior for σ_α)

```
sleep_ri2 = "model{
  for(i in 1:length(Reaction)) {
    Reaction[i] ~ dnorm(mu[i],tau)
    mu[i] = alpha[Subject_index[i]] + beta*Days[i]

    y_pred[i] ~ dnorm(mu[i],tau)
  }

  for(j in 1:18) {
    alpha[j] ~ dnorm(beta_alpha, tau_alpha)
  }

  beta_alpha ~ dnorm(0,1/10000)
  sigma_alpha ~ dunif(0, 10)
  tau_alpha = 1/(sigma_alpha*sigma_alpha)


  beta ~ dnorm(0,1/10000)
  sigma ~ dunif(0, 100)
  tau = 1/(sigma*sigma)
}"
```

Comparing Model fit

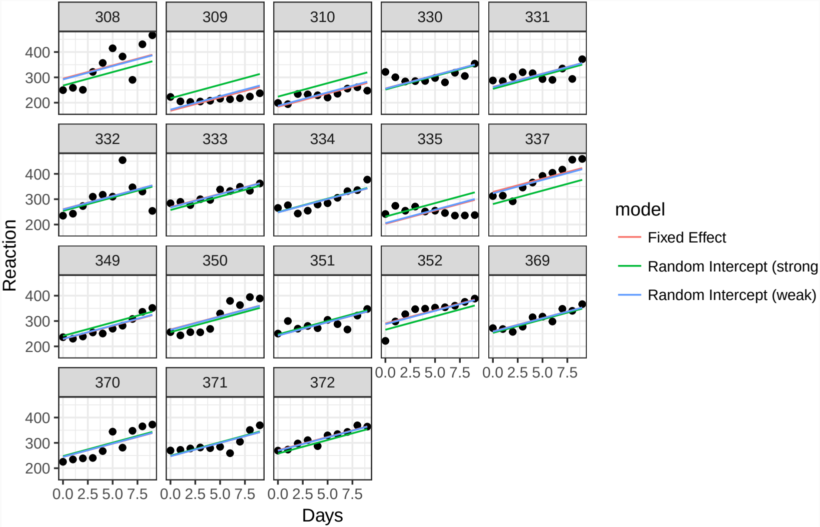

Prior Effect on α

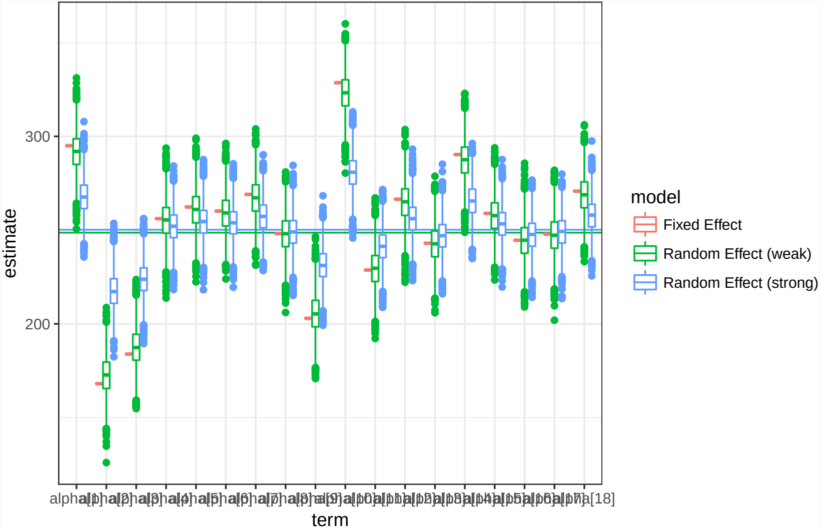

Some Distribution Theory

Random intercept and slope model

Let i represent each observation and $j(i)$ be the subject in oberservation i then

$$Y_i = \alpha_{j(i)} + \beta_{j(i)} \times \text{Days} + \epsilon_i$$

$$\begin{aligned} \alpha_j &\sim \mathcal{N}(\beta_0,\ \sigma^2_\alpha) \\ \beta_j &\sim \mathcal{N}(\beta_1,\ \sigma^2_\beta) \\ \epsilon_i &\sim \mathcal{N}(0,\ \sigma^2) \end{aligned}$$

$$\begin{aligned} \beta_\alpha, \beta_\beta &\sim \mathcal{N}(0, 10000) \\ \sigma, \sigma_\alpha, \sigma_\beta &\sim \text{Unif}(0, 100) \end{aligned}$$

Model - JAGS

```
sleep_ris = "model{
  for(i in 1:length(Reaction)) {
    Reaction[i] ~ dnorm(mu[i],tau)
    mu[i] = alpha[Subject_index[i]] + beta[Subject_index[i]] * Days[i]
    y_pred[i] ~ dnorm(mu[i], tau)
  }

  sigma ~ dunif(0, 100)
  tau = 1/(sigma*sigma)

  for(j in 1:18) {
    alpha[j] ~ dnorm(beta_alpha, tau_alpha)
    beta[j]  ~ dnorm(beta_beta, tau_beta)
  }

  beta_alpha ~ dnorm(0,1/10000)
  beta_beta ~ dnorm(0,1/10000)

  sigma_alpha ~ dunif(0, 100)
  tau_alpha = 1/(sigma_alpha*sigma_alpha)

  sigma_beta ~ dunif(0, 100)
  tau_beta = 1/(sigma_beta*sigma_beta)
}"
```

MCMC Diagnostics

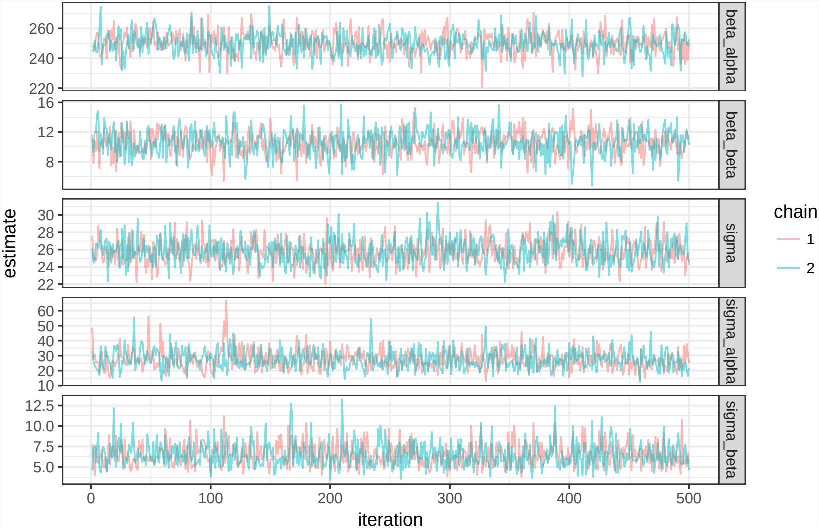

MCMC Diagnostics - α

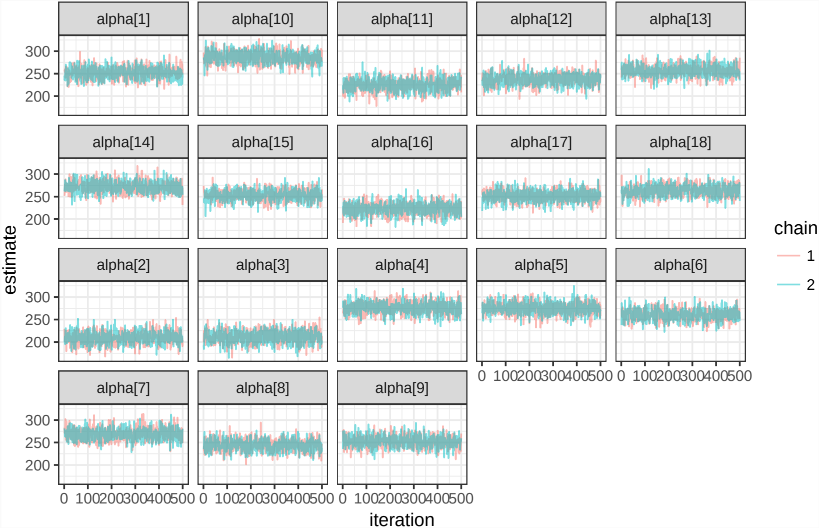

MCMC Diagnostics - β

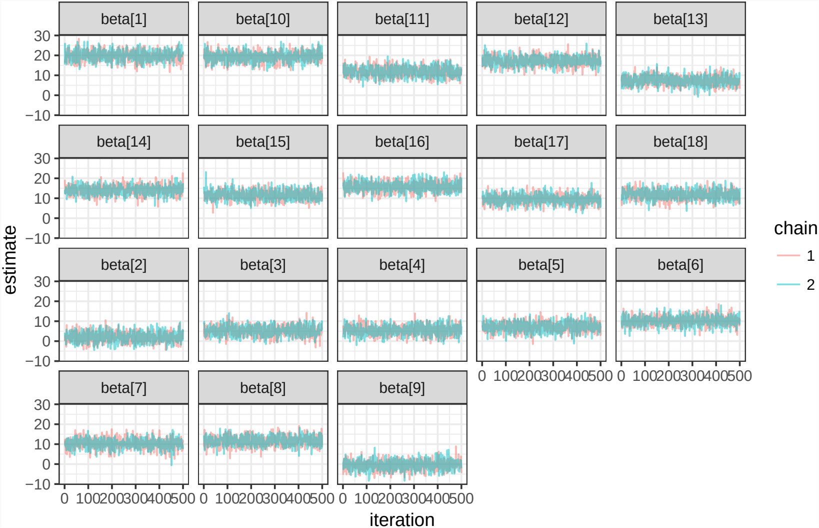

Model fit

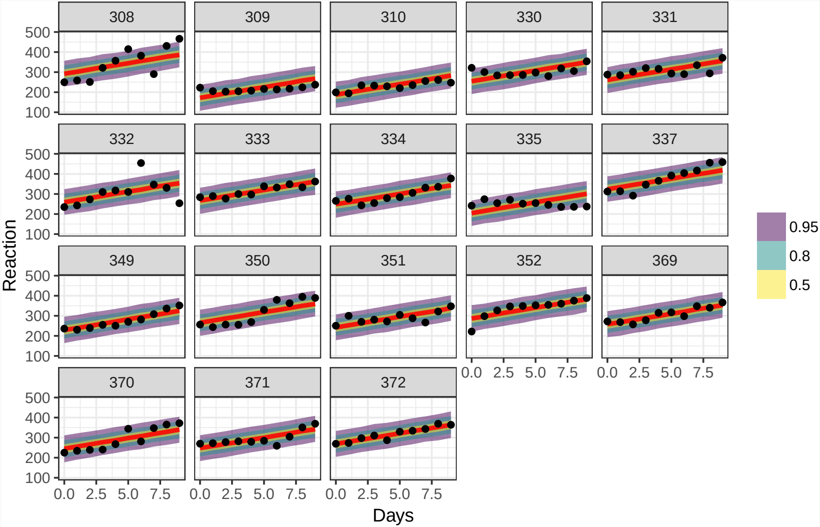

308
309
310
330
331
332
333
334
335
337
349
350
351
352
369
370
371
372
500
400
300
200
100
Reaction
0.0 2.5 5.0 7.5
Days
0.95
0.8
0.5

Model fit - lines

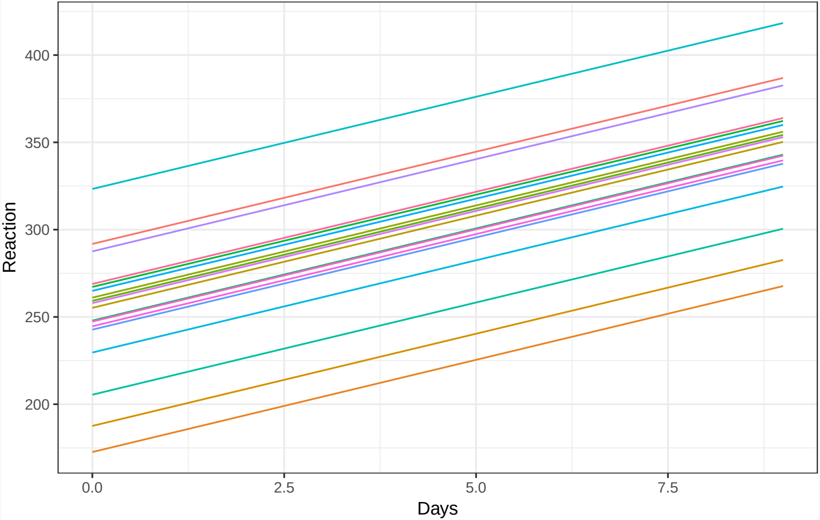

Random Effects

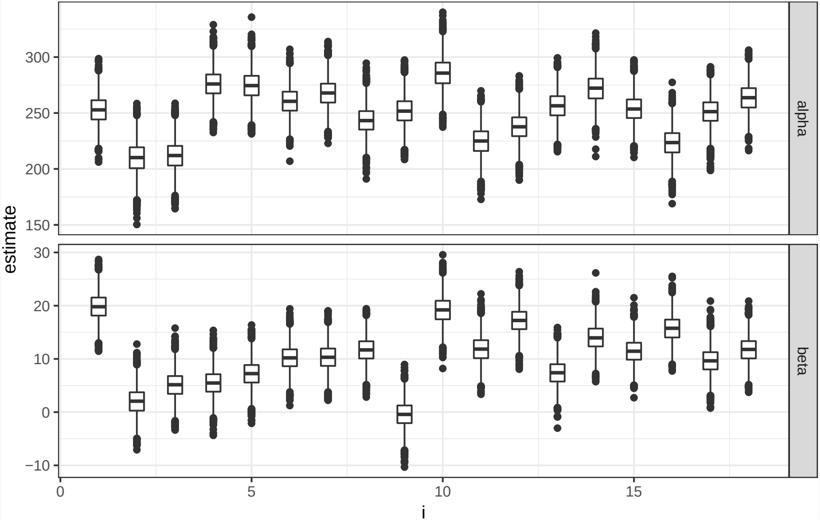

Residuals by subject