Statistics 222, Education 351A Autumn 2020

Statistical Methods for Longitudinal Research

Autumn 2020 Remote Asynchronous Instruction

David Rogosa Sequoia 224, rag{AT}stanford{DOT}edu
Course web page: http://rogosateaching.com/stat222//

To see full course materials from Autumn 2018 go here

Course Welcome and Logistics (first day stuff, to be posted in August, call it Week0)
Lecture slides, week 0 (pdf) **Audio companion, week 0**
For recreation of in-classroom experience, linked below are youtube versions of the music I play before starting lecture and after lecture concludes. Some may wish to reverse that ordering.

```
Registrar's information
STATS 222 (Same as EDUC 351A): Statistical Methods for Longitudinal Research   Units: 2
Grading Basis: Letter or Credit/No Credit

Course Description:
 STATS 222: Statistical Methods for Longitudinal Research (EDUC 351A)
Research designs and statistical procedures for time-ordered (repeated-measures) data.
The analysis of longitudinal panel data is central to empirical research on learning, development, aging, and the effects of interventions.
Topics include: measurement of change, growth curve models, analysis of durations including survival analysis,
experimental and non-experimental group comparisons, reciprocal effects, stability.
See http://rogosateaching.com/stat222/. Prerequisite: intermediate statistical methods
Terms: Aut | Units: 2 | Grading: Letter or Credit/No Credit
Instructors: Rogosa, D. (PI)
```

Preliminary Course Outline
Week 1. Course Overview, Longitudinal Research; Analyses of Individual Histories and Growth Trajectories
Week 2. Introduction to Data Analysis Methods for assessing Individual Change for Collections of Growth Curves (mixed-effects models)
Week 3. Analysis of Collections of growth curves: linear, generalized linear and non-linear mixed-effects models
Week 4. Special case of time-1, time-2 data; Traditional measurement of change for individuals and group comparisons
Week 5. Assessing Group Growth and Comparing Treatments: Traditional Repeated Measures Analysis of Variance and Linear Mixed-effects Models
Week 6. Comparing group growth continued: Power calculations, Cohort Designs, Cross-over Designs, Methods for missing data, Observational studies.
Week 7. Analysis of Durations: Introduction to Survival Analysis and Event History Analysis
Weeks 8-9. Further topics in analysis of durations: Diagnostics and model modification; Interval censoring, Time-dependence, Recurrent Events, Frailty Models, Behavioral Observations and Series of Events (renewal processes)
Dead Week. Assorted Special Topics (enrichment) and Overflow (weeks 1-8): Assessments of Stability (including Tracking), Reciprocal Effects, (mis)Applications of Structural Equation Models, Longitudinal Network Analysis

Texts and Resources for Course Content
1. Garrett M. Fitzmaurice Nan M. Laird James H. Ware Applied Longitudinal Analysis (Wiley Series in Probability and Statistics; 2nd ed 2011)
Text Website second edition website Text lecture slides
2. Judith D. Singer and John B. Willett . Applied Longitudinal Data Analysis: Modeling Change and Event Occurrence New York: Oxford University Press, March, 2003.
Text web page Text data examples at UCLA IDRE Powerpoint presentations good gentle intro to modelling collections of growth curves (and survival analysis) is Willett and Singer (1998)
3. Douglas M. Bates. lme4: Mixed-effects modeling with R February 17, 2010 Springer (chapters). A merged version of Bates book: lme4: Mixed-effects modeling with R January 11, 2010 has been refound
Manual for R-package lme4 and mlmRev, Bates-Pinheiro book datasets.
Additional Doug Bates materials. Collection of all Doug Bates lme4 talks Mixed models in R using the lme4 package Part 2: Longitudinal data, modeling interactions Douglas Bates 8th International Amsterdam Conference on Multilevel Analysis 2011-03-16 another version
Original Bates-Pinheiro text (2000). Mixed-Effects Models in S and S-PLUS (Stanford access). Appendix C has non-linear regression models.
Fitting linear mixed-effects models using lme4, *Journal of Statistical Software* Douglas Bates Martin Machler Ben Bolker. Technical topics: Mixed models in R using the lme4 package Part 4: Theory of linear mixed models
4. A handbook of statistical analyses using R (second edition). Brian Everitt, Torsten Hothorn CRC Press, Index of book chapters Stanford access
Longitudinal chapters: Chap11 Chap12 Chap13. Data sets etc Package 'HSAUR2' August 2014, Title A Handbook of Statistical Analyses Using R (2nd Edition)
There is now a third edition of HSAUR, but full text not yet available in crcnetbase.com. CRAN HSAUR3 page with Vignettes (chapter pieces) and data in reference manual
5. Peter Diggle , Patrick Heagerty, Kung-Yee Liang , Scott Zeger. Analysis of Longitudinal Data 2nd Ed, 2002
Amazon page Peter Diggle home page Book data sets
A Short Course in Longitudinal Data Analysis Peter J Diggle, Nicola Reeve, Michelle Stanton (School of Health and Medicine, Lancaster University), June 2011 earlier version associated exercises: Lab 1 Lab2 Lab3
6. Longitudinal and Panel Data: Analysis and Applications for the Social Sciences by Edward W. Frees (2004). Full book available and book data and

AIDS in Belgium example, (from Simon Wood) single trajectory, count data using glm. Rogosa R session for aids data
additional expositions of AIDS data, Poisson regression: Duke Kentucky
A *very* comprehensive introduction to analysis of count data Regression Models for Count Data in R Achim Zeileis Christian Kleiber Simon Jackman (Stanford University)
Non-linear models, esp logistic. From week 1, also week 3 Self-Starting Logistic model SSlogis help page, do `?SSlogis` post of annotated logistic curve with SSlogis arguments
Trend in Proportions: College fund raising example prop.trend.test help page `?prop.trend.test` in R-session.
Trend in proportions, group growth, Cochran-Armitage test. Expository paper: G. Salanti and K. Ulm (2003): Tests for Trend in Binary Response (SU access)

WEEK 1 Review Questions

1. For the straight-line (constant rate of change) fit example to subj 372 in the sleepstudy data. Obtain a confidence interval for the rate of change from the OLS fit. Now compare the OLS fit with day-to-day differences. Under the constant rate of change model these 9 day to day differences also estimate the rate of change. Obtain a estimate of the mean and a confidence interval for rate of change from these first differences. Compare with OLS results.

Solution for question 1

2. Revisit the Belgium Aids data example (counts of new cases by year). Use the parameter estimates for `am2` (quadratic in time glm fit) to compute by hand (or calculator) the values of the glm fit at year = 5 and year = 9. Compare those values with results from the model `am2` using `predict`

Solution for question 2

3. Paul Rosenbaum has a little data set on growth in vocabulary that I grabbed from his Wharton coursesite. Following the *chicks* class example, plot these data and try to fit a logistic growth curve to these data. What is the estimate of the final vocab level (asymptote)? Compare the data and the fits from the logistic growth curve.
For reference, Self-Starting Logistic model SSlogis help page, do `?SSlogis` post of annotated logistic curve with SSlogis arguments additional tools in the grofit package

Solution for question 3

4. More on autocorrelation[extension/enrichment]. In standard regression courses you may have seen in addition to Durbin-Watson test for AR(1) (dwtest()), versions of the Cochrane-Orcutt procedure for remediation. Uses a first difference transformation of the data with an estimate of the autocorrelation (therefore hopeless when you have 3,4 5 observations per unit). To illustrate the statements in class and the similarities to OLS result, the solution to this problem does the straight-line and polynomial examples from the Week 1 class handout using the R-package `orcutt`

Solution for question 4

WEEK 1 Exercises

1. Straight-line fits for NC Fem data: North Carolina Achievement Data (see Williamson, Applebaum, Epanchin, 1991). These education data are eight yearly observations on achievement test scores in math (Y), for 277 females each followed from grade 1 to grade 8, with a verbal ability background measure (W).
North Carolina, female math performance (also in Rogosa-Saner) North Carolina data (wide format); NC data (long)
a. Here we will use the 8 yearly observations on female ID `705810`, which you can obtain from either the long form or wide form of these data.
For that female, what is the rate of improvement over grades 1 through 8? Compare the observed improvement for grades 1 through 8 (the *difference score*) with the amount of improvement indicated by the model fit. Obtain a 95% confidence interval for each (if possible).
b. More on OLS and the difference score. Refer to an old publication: A growth curve approach to the measurement of change. Rogosa, David; Brandt, David; Zimowski, Michele Psychological Bulletin. 1982 Nov Vol 92(3) 726-748 APA record direct link; Equation 4, page 728, shows a useful form for the OLS slope. (actually reading the first three pages of that pub is a decent intro to the growth curve topic.) For equally spaced data, that Eq (4) gives a useful equivalence between difference scores (amounts of change) and OLS slopes (multiply rates of change by time interval). For the *part a* NC data show that the OLS slope can be expressed as a weighted sum of the four differences: { 8-1,7-2,6-3,5-4}. [to say that better {score at time 8 minus score at time 1; score at time 7 minus score at time 2; ...} and so forth]
Seperately, consider three observations at taken at equally spaced time intervals: What is a simple expression for the OLS slope (rate of change)?

2. Revisit the Berkeley Growth Data example from week 1 lecture. Consider the quadratic (polynomial degree 2) fit to these data, and also a (innapropriate?) constant-rate-of-change (straight-line) fit to these data. Then refer to Seigel, D. G. Several approaches for measuring average rates of change for a second degree polynomial. The American Statistician, 1975, 29, 36-37. JStor Link for equivalences for the slope of the straight-line fit to an *average rate of change* for the quadratic fit. Compare Seigel 'Approach 3" to 'Approach 1'.

Week 2. Analysis of collections of growth curves (Mixed-effects Models, lmer) Constant rate of change models

Lecture Topics. Analyses of collections of growth curves.
1. Plots, description and SFYS (smart first year student) analyses.
2. Mixed effects models using `lmer` . Growth modelling handout

Class Examples

1. Data frame `sleepstudy` available in lme4 package.
Music to accompany long-distance truck driver data: Flying Burrito Brothers "Six Days on the Road"
a. *Published Treatments, Sleepstudy example*
Source Publication: Belenky, G., Wesensten, N. J., Thorne, D. R., Thomas, M. L., Sing, H. C., Redmond, D. P., Russo, M., & Balkin, T. (2003). Patterns of performance degradation and restoration during sleep restriction and subsequent recovery: A sleep dose-response study. Journal of Sleep Research, 12(1), 1-12.
Sleepstudy data analysis from Doug Bates lme4 book lme4: Mixed-effects modeling with R February 17, 2010 (draft chapters) Chapter 4: Sleepstudy example or a set of Bates slides for the sleepstudy example
Why lmer (lme4) does not provide p-values for fixed effects : Doug Bates lmer, p-values and all that
A number of add-on packages seek to provide lmer p-values.(We show afex package in Review Question 1 with 2020 update).
b. *Class Materials, Sleepstudy example*
Individual plots (frame-by-frame) Plot of straight-line fits Initial descriptive analyses (SFYS)
Sleepstudy, Bates Ch 4, lme4 analyses, ascii Sleepstudy class handout, pdf scan
more Doug Bates Slides (pdf pages 8-28)

2. North Carolina, female math performance (also in Rogosa-Saner)

North Carolina data (wide format); NC data (long) plots for NC data

Data formatting: **wide to long** North Carolina data (wide format); making the "Long" version

The UCLA archive has a tutorial using built-in `reshape` function (rather than the reshape package).

North Carolina example. wide-form descriptives, background, plots Initial SFYS analyses of NC data, ascii

Model Comparisons for North Carolina, female math performance ascii version NC class handout, pdf scan model ncCon2 without redundent model term NC bootstrap results (SAS)

3. Brain Volume Data, in-class modeling exercise: analyses from "Variation in longitudinal trajectories of regional brain volumes of healthy men and women (ages 10 to 85 years) measured with atlas-based parcellation of MRI" cartoon plot of Lateral Ventricles data; actual data plot of Lateral Ventricles data; development of lmer (random effect) growth models

Background and Resources

Technical Formulation and extensions

Estimation in lmer.

Fitting linear mixed-effects models using lme4, Journal of Statistical Software Douglas Bates Martin Machler Ben Bolker also Rnews_2005 pp.27-30

Bates book, Chapter 5, Computational Methods. Bates talk slides: Mixed models in R using the lme4 package Part 4: Theory of linear mixed models

Extensions and Alternatives, lmer.

Plots and diagnostics: Package `Influence.me` RJournal intro Package `merTools` An Introduction to merTools Also, Prediction Intervals

Non-Gaussian modelling. Hierarchical Generalized Linear Models, Package `hglm` Hierarchical Generalized Linear Models, R Journal December 2010.

Extensions of lme4 modeling: Package `npmlreg` Nonparametric Maximum Likelihood (NPML) estimation;

Package `robustlmm`: An R Package for Robust Estimation of Linear Mixed-Effects Models

Package `RLRsim` Title Exact (Restricted) Likelihood Ratio Tests for Mixed and Additive Models

Data Examples

North Carolina Data also in (with full development of the modelling) Longitudinal Data Analysis Examples with Random Coefficient Models. David Rogosa; Hilary Saner . Journal of Educational and Behavioral Statistics, Vol. 20, No. 2, Special Issue: Hierarchical Linear Models: Problems and Prospects. (Summer, 1995), pp. 149-170. Jstor

Douglas Bates class resource item #3, Texts and Resources. Other Doug Bates materials: Three packages, "SASmixed", "mlmRev" and "MEMSS" with examples and data sets for mixed effect models

North Carolina Data also in (with full development of the modelling) Longitudinal Data Analysis Examples with Random Coefficient Models. David Rogosa; Hilary Saner . Journal of Educational and Behavioral Statistics, Vol. 20, No. 2, Special Issue: Hierarchical Linear Models: Problems and Prospects. (Summer, 1995), pp. 149-170. Jstor Data sets for Rogosa-Saner

Additional talk materials: An Assortment of Longitudinal Data Analysis Examples and Problems 1/97, Stanford biostat. Overview and Implementation for Basic Longitudinal Data Analysis CRESST Sept '97. Another version (short) of the expository material is from the Timepath '97 (old SAS progranms) site: Growth Curve models ; Data Analysis and Parameter Estimation ; Derived quantities for properties of collections of growth curves and bootstrap inference procedures

WEEK 2 Review Questions

1. More sleepstudy. Confidence interval and p-values. Add on, extension to class example.
I start by fitting the lmer model for the collection of growth curves: `sleeplmer = lmer(Reaction ~ Days + (1 + Days|Subject), sleepstudy).`
Then try out `confint` from lme4 (link to manual using likelihood profile or bootstrap methods.
Then look at the `pvalues` entry in the manual and try out add-on packages, esp for p-values for the fixed effects.

Solution for Review Question 1 2017 redo/update using 3.3.3 (barebones)

2. Ramus Data example. Example consists of 4 longitudinal observations on each of 20 cases. The measurement is the height of the mandibular ramus bone (in mm) for boys each measured at 8, 8.5, 9, 9.5 years of age. These data, which have been used by a number of authors (e.g., Elston and Grizzle 1962), can be found in Table 4.1 of Goldstein (1979). Ramus data example long form for Ramus data tutorial on creating long form data manipulation and 2017 redo/check of widetolong. Use lmList to obtain the 20 OLS fits, with the initial time set to 8 years of age, i.e. intercepts are fits for the time of initial measurement (not t=0). Fit the lmer model for the collection of growth curves (using initial time = 8); verify that fixed effects are the sample means (over persons) of the lmList intercepts and slopes. Verify that the random effects variance for "age" (i.e. slopes) is the method-of-moments estimate for Var(theta). Compare the random effect estimates (ranef) which borrow strength for each subject with the OLS estimates from lmList (c.f. Bates Chap 4 discussion of sleepstudy data)

Solution for Review Question 2

3. Artificial data example (used in Myths chapter to illustrate time-1,time-2 data analysis) Two part artificial data example. The bottom frame (the X's) is 40 subjects each with three equally spaced time observations (here in wide form).For these the fallible "X" measurements (constructed by adding noise to the Xi measurements). Follow the class examples 'wide-to-long' and obtain the plot showing each subject's data and straight-line fit. Use lmList to obtain the 40 slopes for the straight-line fits.

Solution for Review Question 3

4. More with North Carolina data

a. identify the fastest and slowest growth among the 277 females. Compare medians of growth rates for females with verbal ability (Z) at or above 106 with that for females with verbal ability below 106. Show side-by-side boxplots.

b. In the class handout version of the NC analyses (and other postings, but not all) the first thing to do was make the 'time' variable have intitial value = 0 (making the intercept of a straight line fit correspond to level at initial time): i.e. 1 to 8 becomes 0 to 7. Obtain lmList results and fit the ncUnc lmer model (straight-line growth, no Z) using time 1 to 8. Comment on differences of these analyses with those using timeInt in the class handout. In particular, look at the correlation of change and initial status. The correlation between observed change and observed initial status using timeInt was .279 from lmer (`Correlation of Fixed Effects`) and also from lmList (you should confirm that). What is the result you obtain using `time` rather than `timeInt`? The mle of of the correlation of 'true' change and 'true' initial status is .651 using `timeInt`. What do you obtain using `time`?.

Solution for Review Question 4

5. xyplot with large sample sizes.

North Carolina data has 277 subjects, a frame-by-frame display of individuals requires subsampling. Construct a plot for 24 (arbitrary) individuals data trajectories.

Solution for Review Question 5

WEEK 2 Exercises

1. *Tolerance data* [note: 10/12/17 data location updated]
A subsample of data from the National Youth Survey is obtained in long-form by
`read.table("https://stats.idre.ucla.edu/wp-content/uploads/2016/02/tolerance1_pp.txt", sep=",", header=T)`
and in wide form by
`read.table("https://stats.idre.ucla.edu/wp-content/uploads/2016/02/tolerance1.txt", sep=",", header=T)`
Yearly observations from ages 11 to 15 on the tolerance measure (tolerance to deviant behavior e.g. cheat, drug, steal, beat; larger values indicates more tolerance on a 1to4 scale). Also in this data set are gender (is_male) and an `exposure` measure obtained at age 11 (self report of close friends involvement in deviant behaviors). note: the `time` measure is age - 11.
i. obtain individual OLS fits (tolerance over time) and plot the collection of those straight-lines. Provide descriptive statistic summaries for the rate of change in tolerance and initial level.
ii. fit a mixed effects model for tolerance over time (unconditional) for this collection of individuals. Obtain interval estimates for the fixed and random effects. Show that the fixed effects estimates correspond to quantities obtained in part i. Explain.
iii. Investigate whether the `exposure` measure is a useful predictor of level or rate of change in `tolerance`. What appears to be the best fitting mixed model for these data using these measures? Show specifics.

2. *lmer/lme vs lm*
Consider the sleepstudy and Ramus examples, collections of growth trajectories with no exogenous variable. Ramus Data example. Example consists of 4 longitudinal observations on each of 20 cases. The measurement is the height of the mandibular ramus bone (in mm) for boys each measured at 8, 8.5, 9, 9.5 years of age. These data, which have been used by a number of authors (e.g., Elston and Grizzle 1962), can be found in Table 4.1 of Goldstein (1979).
Ramus data example long form for Ramus data tutorial on creating long form data manipulation .
Fitting the lmer models with `Formula: Reaction ~ 1 + Days + (1 + Days | Subject)` or `Formula: ramus ~ I(age - 8) + (age | subj)` has motivated the student question, what is going on here beyond what lm would do? So let's look at what `lm` would do in these examples. Verify (or disprove) the assertion that the fixed effects from lmer, which we have seen are the averages of the individual fit parameter estimates (i.e. lmList), and therefore the coefficients of the average growth curve are *identical* to the fit from `lm` (which ignores the existence of individual trajectories). Compare the results of the lm and lmer analyses for these two data sets.

3. *Early Education data* (From Bates and Willett-Singer).
Data on early childhood cognitive development described in Doug Bates talk materials (pdf pages 49-52). Obtain these data from the R-package "mlmRev" or the Willett-Singer book site (in our week 1 intro links). Data are in long form and consist of 3 observations 58 treatment and 45 control children; see the Early entry in the mlmRev package docs. Produce the plot of individual trajectories shown pdf p.49, Bates talk. (note:Bates does connect-the-dots, we have done straight-line fit, your choice). Show five-number summaries of rates of impovement in cognitive scores for treatment and control groups. Develop and fit the `fm12` lmer model shown in Bates pdf p.50 (note `fm12` allows trt to effect rates of improvement but not level;). Interpret results. Note: this moves us into the comparing groups topics, where the individual attribute is group membership.

4. *Standardizing is always a bad idea* is a good motto for life, especially with longitudinal data.
Artificial data example from Review Question 3 (used in Myths chapter to illustrate time-1,time-2 data analysis) Start out with the "X" data, and standardize (i.e. transform to mean 0, var 1) at each of the 3 time points. Note "scale" will do this for you (in wide form). For the standardized data obtain the plot showing each subject's data and straight-line fit. What do you have here? Compare the results the mixed-effects models fitting the collection of straight-line growth curves for the measured and standardized data.

Analyses of Collections of Growth Curves

Stat 222
Week 2

Descriptive Analyses (smart first year student)

SFYS Fit Y on t regressions (lmList)

describe $\hat{\alpha}_0$ $\hat{\alpha}_1$ plots etc $\hat{\alpha}_1 \,\big|\, \hat{\alpha}_0$ initial status growth

systematic indiv differences (z) in change

cor($\hat{\alpha}_1$, z) plots $\hat{\alpha}_1 \,\big|\, z$ etc

Mixed effects models (lmer)

A. Unconditional

sleep N/C

Level 1 $Y = \alpha_0 + \alpha_1 t + \varepsilon$ straight-line growth each unit

level at t=0 (define init status)

gradient

Level 2 $\alpha_0 = \gamma_{00} + u_0$ "null"

$\alpha_1 = \gamma_{10} + u_1$

Combined (estimation) model $Y = \underbrace{\gamma_{00} + \gamma_{10} t}_{\text{fixed}} + \underbrace{[\varepsilon + u_0 + u_1]}_{\text{random}}$

B. Conditional Model

exog var z, w systematic indiv diffs in growth

N/C ncCon2 (full)

Level 1 $Y = \alpha_0 + \alpha_1 t + \varepsilon$

Level 2

$\alpha_0 = \gamma_{00} + \gamma_{01} z + u_0$ diff random terms

$\alpha_1 = \gamma_{10} + \gamma_{11} z + u_1$

Combined model $Y = \gamma_{00} + \gamma_{01} z + \gamma_{10} t + \gamma_{11} z \cdot t + [\varepsilon + u_0 + u_1]$

model form $Y \sim z * t$

see note on redundant time Int term in NC handout lmer

LONGITUDINAL PANEL DATA

OBSERVATIONS X_{ip}

TAKEN AT TIME t_i $(i = 1, \ldots T)$

FOR INDIVIDUAL p $(p = 1, \ldots, n)$

T "WAVES" OF DATA

MEASUREMENT MODEL

$$X_{ip} = \xi_{ip} + \epsilon_{ip}$$

"TRUE SCORE" ξ_{ip}

reliability coeff
Var(Xi)/Var(X)

GROWTH MODELS

$$\xi_p(t) = f(\xi, t)$$

Collection of Growth Curves

For individual p, growth curve for single measure $\xi_p(t)$

Parameters of growth curve vary over p

Examples: Straight-line growth

$$\xi_p(t) = \xi_p(0) + \theta_p t$$

Exponential growth

$$\xi_p(t) = \lambda_p - (\lambda_p - \xi_p(0))e^{-\gamma_p t}$$

Alternative models

Autoregressive process/

Simplex models

Dictum:

Individual trajectory is the key starting point for conceptualization, modelling, data analysis.

In general model processes that generate the (individual's) data.

Models for Collections of Growth Curves

Straight-line Growth Curve Formulation.

attribute η, which exhibits systematic change over time. For individual p, growth curve in η is $\eta_p(t)$.

$$\eta_p(t) = \eta_p(0) + \theta_p t$$

Note: Rewrite using the centering parameter t^o ; θ and $\eta(t^o)$ are uncorrelated over the population of individuals $t^o = -\sigma_{\eta(0)\theta}/\sigma^2_\theta$

$$\eta_p(t) = \eta_p(t^o) + \theta_p(t - t^o) \quad .$$

Constant rate of change θ_p -- first two moments μ_θ σ^2_θ
For systematic individual differences in growth (i.e. correlates of change) exogenous characteristic W. Conditional expectation

$E(\theta|W) = \mu_\theta + \gamma\,(W - \mu_W)$, With no measured exogenous variable, this between-unit model is $E(\theta|W) = \mu_\theta$.

Shown below 15 straight-line growth curves corresponding to pop. parameters $t^o = 2$; $\sigma^2_\theta = 5.333$; $\sigma^2_{\eta(t^o)} = 48$; $\theta \sim U[1, 9]$, $\eta(t^o) \sim U[38, 62]$. correlations among $\eta(t_i)$ for observation times $\rho_{\eta(1)\eta(4)} = .614$, $\rho_{\eta(1)\eta(6)} = .316$, $\rho_{\eta(4)\eta(6)} = .943$. For Y, $\text{var}(\epsilon) = 5$, the pop. correlations are $\rho_{Y(1)Y(4)} = .567$, $\rho_{Y(1)Y(6)} = .297$, $\rho_{Y(4)Y(6)} = .894$.

Alternative: exponential growth to an asymptote

Exponential growth curve with asymptote λ_p and curvature δ

Straight-line Growth

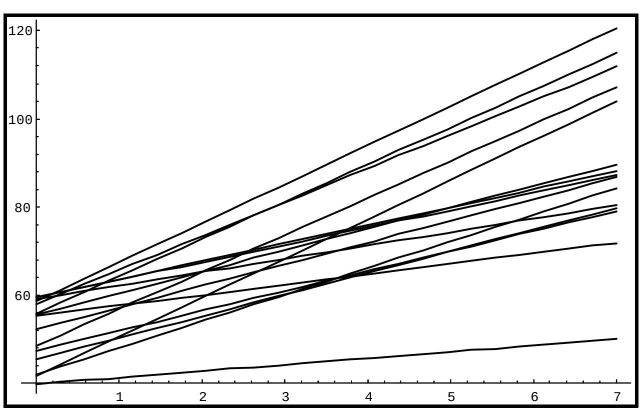

Exponential Growth

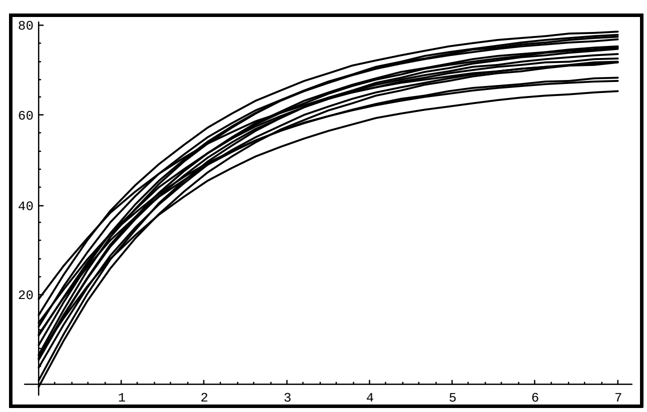

Analyses of Collections of Growth Curves

Descriptive Analyses (smart first year student)

SFYS Fit Y on t regressions (lmList)

describe $\hat{\alpha}_0$ $\hat{\alpha}_1$ plots etc $\hat{\alpha}_1$ $\hat{\alpha}_0$ initial status growth

systematic indiv differences (z) in change

cor $(\hat{\alpha}_1, z)$ plots $\hat{\alpha}_1$ z etc

Mixed effects models (lmer)

A. Unconditional

sleep NC

level at t=0 (define init status

Level 1 $\quad Y = \alpha_0 + \alpha_1 t + \varepsilon$ straight-line growth each unit

gradient

Level 2 $\quad \alpha_0 = \gamma_{00} + u_0$ "null"

$\alpha_1 = \gamma_{10} + u_1$

Combined (estimation) model $\quad Y = \gamma_{00} + \gamma_{10} t + [\varepsilon + u_0 + u_1]$

fixed random

B. Conditional Model

exog var z, w
systematic indiv diffs in growth

NC ncCon2 (full)

Level 1 $\quad Y = \alpha_0 + \alpha_1 t + \varepsilon$

Level 2

$\alpha_0 = \gamma_{00} + \gamma_{01} z + u_0$ diff random terms

$\alpha_1 = \gamma_{10} + \gamma_{11} z + u_1$

Combined model $\quad Y = \gamma_{00} + \gamma_{01} z + \gamma_{10} t + \gamma_{11} z \cdot t$

$+ [\varepsilon + u_0 + u_1]$

model form Y ~ z*t

see note on redundant time Int term in NC handout lmer

Two-Stage Analysis - "NIH Method"

One classic approach to the analysis of such data is known as two-stage or two-step analysis.

It is sometimes called the "*NIH Method*" because it was popularized by statisticians working at NIH.

In the two-stage method, we simply fit a straight line (or curve) to the response data for each subject (first stage), and then regress the estimates of the individual intercepts and slopes on subject-specific covariates (second stage).

One of the attractions of this method is that it is very easy to perform using existing statistical software for linear regression.

We can illustrate the method by considering a two-stage analysis of Feldman's clearance data.

Gelman secret weapon; Diggle clever ostrich, SFYS from Timepath site

Section 1– Descriptive analyses of growth rates.

Individual OLS Fits

The most basic step in the analysis is the fitting of a straight-line growth curve (the regression of Y on t for each p) by ordinary least-squares.

Individual OLS Fits displays for each individual (rows) the (columns):

`ID`

`RATE` (empirical rate; OLS estimate of θ_p, slope Y on t fit)

`INIT_LVL` (level for Y on t fit evaluated at the first anchor time point)

`MSR` (residual variance for Y on t fit)

`RSQ` (squared multiple correlation for Y on t fit)

`W` (value of exogenous variable)

Cross-sectional Description:

For *synchronous* (same times of observation for each individual) data sets (in which it makes sense to talk about time-1 etc observations) cross-sectional descriptive summaries are provided.

(In data sets such as smearmiss, this output is automatically not computed).

Cross-sectional means and between-wave correlations for the Y(t_i) (adding $\hat{\theta}$ and W)

OLS Fits: Descriptive Statistics

Estimation of the straight-line growth model allows comparisons of rates of change across individuals. Stem-and-leaf diagrams, boxplots, and the five-number summaries of the empirical rates are useful ways to describe both typical rates of learning and the heterogeneity across individuals. Displayed in this section are descriptive statistics for the quantities listed under Individual OLS Fits (`RATE` `INIT_LVL` `MSR` `RSQ` `W`) plus individual values for the Foulkes-Davis tracking index (`Gamma`). (Foulkes-Davis index of tracking estimated from a count of the number of intersections that each individual trajectory has with the other individuals; for each individual $\hat{\gamma}_p$ is one minus the number of intersections over $n - 1$. Individuals with a low value of $\hat{\gamma}_p$ are those whose relative standing changes considerably over the time period.)

Stem-and-leaf diagrams and accompanying boxplots are displayed for `RATE` (Empirical rate) `INIT_LVL` (Fitted Initial Level) `W` (Exogenous Variable).

OLS Fitted Values for Anchor Times

Regardless of whether the Cross-sectional Description for (*synchronous* data) is printed above, descriptive (cross-sectional) statistics and between-wave correlations are provided for the following measures: `TA_FITi` values of the individual fits for each of the specified anchor time points, plus `RATE W`

Rate and Fitted Initial Level Scatter plot of RATE vs INIT_LVL

To provide some descriptive augmentation to the parameter estimate for the correlation $\rho_{\eta(t_I)\theta}$ (correlation between change θ_p and true initial status $\eta_p(t_I)$, where t_I indicates a time of initial status designated by the first time anchor point).

OLS Theta-hat on W Regression When W is present, OLS regression and corresponding scatterplot is given–provides graphical diagnostic and shows correspondence (both point estimate and standard error) with fixed-effects estimates from mixed-model estimation (exact match for complete synchronous data)

OLS Fitted Initial Level on W Regression When W is present, OLS regression is given–provides graphical diagnostic and shows correspondence (both point estimate and standard error) with fixed-effects estimates from mixed-model estimation (exact match for complete synchronous data)

Chapter 4
Models for Longitudinal Data

Longitudinal data consist of repeated measurements on the same subject (or some other "experimental unit") taken over time. Generally we wish to characterize the time trends within subjects and between subjects. The data will always include the response, the time covariate and the indicator of the subject on which the measurement has been made. If other covariates are recorded, say whether the subject is in the treatment group or the control group, we may wish to relate the within- and between-subject trends to such covariates.

In this chapter we introduce graphical and statistical techniques for the analysis of longitudinal data by applying them to a simple example.

4.1 The `sleepstudy` Data

Belenky et al. [2003] report on a study of the effects of sleep deprivation on reaction time for a number of subjects chosen from a population of long-distance truck drivers. These subjects were divided into groups that were allowed only a limited amount of sleep each night. We consider here the group of 18 subjects who were restricted to three hours of sleep per night for the first ten days of the trial. Each subject's reaction time was measured several times on each day of the trial.

```
> str(sleepstudy)
```

```
'data.frame':        180 obs. of  3 variables:
 $ Reaction: num  250 259 251 321 357 ...
 $ Days    : num  0 1 2 3 4 5 6 7 8 9 ...
 $ Subject : Factor w/ 18 levels "308","309","310",..: 1 1 1 1 1 1 1 1 1..
```

In this data frame, the response variable `Reaction`, is the average of the reaction time measurements on a given subject for a given day. The two covariates are `Days`, the number of days of sleep deprivation, and `Subject`, the identifier of the subject on which the observation was made.

or opportunity for sleep except as required by the periodic sleep latency tests (described below).

Test instruments and measures

Psychomotor vigilance test

The PVT measures simple reaction time to a visual stimulus, presented approximately 10 times/minute (interstimulus interval varied from 2 to 10 s in 2-s increments) for 10 min and implemented in a thumb-operated, hand-held device (Dinges and Powell 1985). Subjects attended to the LED timer display on the device and pressed the response button with the preferred thumb as quickly as possible after the appearance of the visual stimulus. The visual stimulus was the LED timer turning on and incrementing from 0 at 1-ms intervals. In response to the subject's button press, the LED timer display stopped incrementing and displayed the subject's response latency for 0.5 s, providing trial-by-trial performance feedback. At the end of this 0.5-s interval the display turned off for the remainder of the foreperiod preceding the next stimulus. Foreperiods varied randomly from 2 to 10 s. Dependent measures, averaged or summed across the 10-min PVT session, included mean speed (reciprocal of average response latency), number of lapses (lapse = response latency exceeding 500 ms), and mean speed for the fastest 10% of all responses.

Polysomnography

Polysomnographic (PSG) measures [EEG (C3 and C4); EOG (outer canthi of each eye), EMG (mental/submental)], and EKG (from just below left and right clavicle) were recorded continuously throughout the study using Medilog 9000-II magnetic cassette recorders (Oxford Instruments, Largo, FL, USA). Raw data were digitized, and both night-time sleep and sleep latency tests (described below) were scored in accordance with Rechtschaffen and Kales (1968) criteria using Eclipse software (Stellate Systems, Westmont, Quebec, Canada).

Night-time sleep

Six technicians, whose inter-rater reliabilities were at least 85% compared with the scoring of a diplomate of the American Board of Sleep Medicine (TJB), scored night-time sleep periods (defined as lights out to lights on). Dependent measures included minutes of individual sleep stages [1, 2, slow wave sleep (SWS) and REM] and minutes of total sleep time (TST) (sum of minutes spent in all sleep stages).

Sleep latency test

Subjects were placed in bed in a quiet, darkened room and instructed to close their eyes and not resist the urge to fall asleep. Staff monitored PSG signals from outside of the bedroom using Oxford Mentor systems. To decrease the likelihood of premature test termination during ambiguous stage 1, the sleep latency test (SLT) was terminated immediately after the onset of stage 2 sleep (or after 20 min without sleep onset) by a staff member opening the bedroom door, turning on the lights, and announcing that the test was over. As stage 1 sleep does not appear to confer recuperative benefit in otherwise normal, healthy adults (Wesensten *et al.* 1999), the potential accumulation of up to 40 min stage 1 sleep daily (20 min per SLT $\times$ 2 SLTs per day) was not expected to affect performance. SLT schedules were staggered by 25 min for subject roommates, so that each could be tested in the bedrooms individually. For purposes of analyses, sleep latency was re-scored off-line from lights out to the first 30 s of stage 1 sleep.

Subjective alertness/sleepiness

The Stanford Sleepiness Scale (SSS; Hoddes *et al.* 1973) assessed subjective sleepiness on a single-item scale ranging from 1 ('feeling active and vital; alert; wide awake') to 7 ('almost in reverie; sleep onset soon; losing struggle to remain awake'). The dependent measure was the subject's sleepiness rating.

Procedure

Subjects reported to the Division of Neuropsychiatry, Walter Reed Army Institute of Research at 10:00 h on the day prior to T1. After being provided with verbal and written descriptions of study procedures and rules, subjects were individually informed of the sleep schedule to which their group (of two to four subjects) was being assigned and electrodes for ambulatory PSG (Oxford Medilog 9200-II), including EOG, EMG, C3, C4, O1, O2, and EKG were applied. Subjects then underwent training on the various performance tasks. At 18:00 h, they were transported to the Johns Hopkins Bayview General Clinical Research Center (GCRC, Baltimore, MD, USA) where they resided until the end of the study. Throughout the study, meals were served at 08:30, 12:30, and 17:30 h, with snacks and beverages available *ad libitum* between performance tests. Vital signs (blood pressure, pulse, and tympanic temperature) were recorded periodically for purposes of checking general health status. Subjects did not use/consume nicotine or caffeine-containing products during the study; random urine drug screens verified compliance. Use of medications during the study (e.g. acetaminophen for headache) was allowed at the discretion of the attending physician. For all women enrolled in the study, serum pregnancy tests performed at the beginning of the study were negative.

Same-sex subject pairs were assigned to share 2-person hospital-style bedrooms. T1 and T2 were devoted to training on the performance tests and familiarization with study procedures. Baseline testing commenced on the morning of the third day (B) and testing continued for the duration of the study (E1–E7, R1–R3). On the morning of R4 electrodes were removed shortly after awakening, and subjects were debriefed and released from the study. No testing occurred on R4.

Sleep deprivation data

- This laboratory experiment measured the effect of sleep deprivation on cognitive performance.
- There were 18 subjects, chosen from the population of interest (long-distance truck drivers), in the 10 day trial. These subjects were restricted to 3 hours sleep per night during the trial.
- On each day of the trial each subject's reaction time was measured. The reaction time shown here is the average of several measurements.
- These data are *balanced* in that each subject is measured the same number of times and on the same occasions.

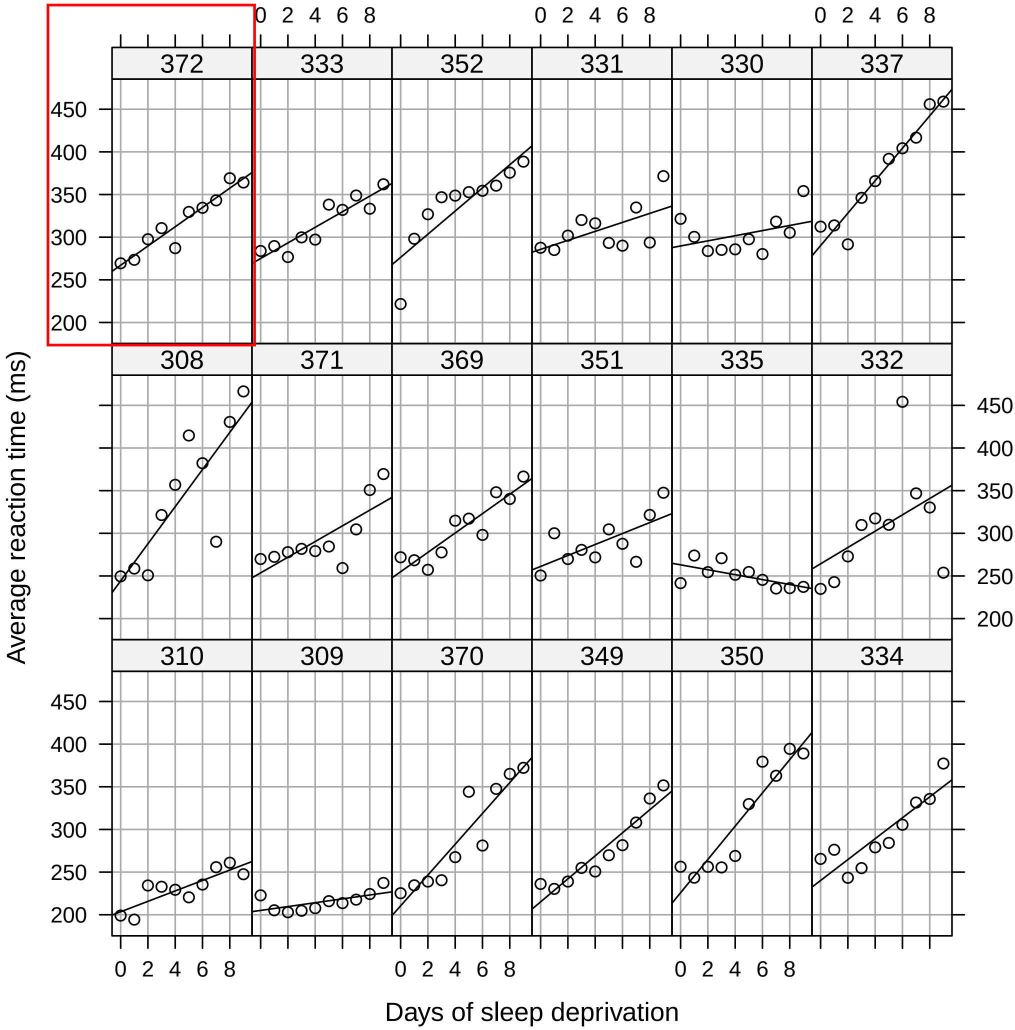

Fig. 4.1 A lattice plot of the average reaction time versus number of days of sleep deprivation by subject for the `sleepstudy` data. Each subject's data are shown in a separate panel, along with a simple linear regression line fit to the data in that panel. The panels are ordered, from left to right along rows starting at the bottom row, by increasing intercept of these per-subject linear regression lines. The subject number is given in the strip above the panel.

As recommended for any statistical analysis, we begin by plotting the data. The most important relationship to plot for longitudinal data on multiple subjects is the trend of the response over time by subject, as shown in Fig. 4.1. This plot, in which the data for different subjects are shown in separate panels with the axes held constant for all the panels, allows for examination of the time-trends within subjects and for comparison of these patterns between

Assessing the linear fits

- In most cases a simple linear regression provides an adequate fit to the within-subject data.
- Patterns for some subjects (e.g. 350, 352 and 371) deviate from linearity but the deviations are neither widespread nor consistent in form.
- There is considerable variation in the intercept (estimated reaction time without sleep deprivation) across subjects – 200 ms. up to 300 ms. – and in the slope (increase in reaction time per day of sleep deprivation) – 0 ms./day up to 20 ms./day.
- We can examine this variation further by plotting confidence intervals for these intercepts and slopes. Because we use a pooled variance estimate and have balanced data, the intervals have identical widths.
- We again order the subjects by increasing intercept so we can check for relationships between slopes and intercepts.

Individual Plots

```
>  xyplot(Reaction ~ Days | Subject, sleepstudy, type = c("g","p","r"),
+  index = function(x,y) coef(lm(y ~ x))[1],
+  xlab = "Days of sleep deprivation",
+  ylab = "Average reaction time (ms)", aspect = "xy")
```

```
>  xyplot(Reaction ~ Days | Subject, sleepstudy, type = c("g","b"),
+  index = function(x,y) coef(lm(y ~ x))[1],
+  xlab = "Days of sleep deprivation",
+  ylab = "Average reaction time (ms)", aspect = "xy")
```

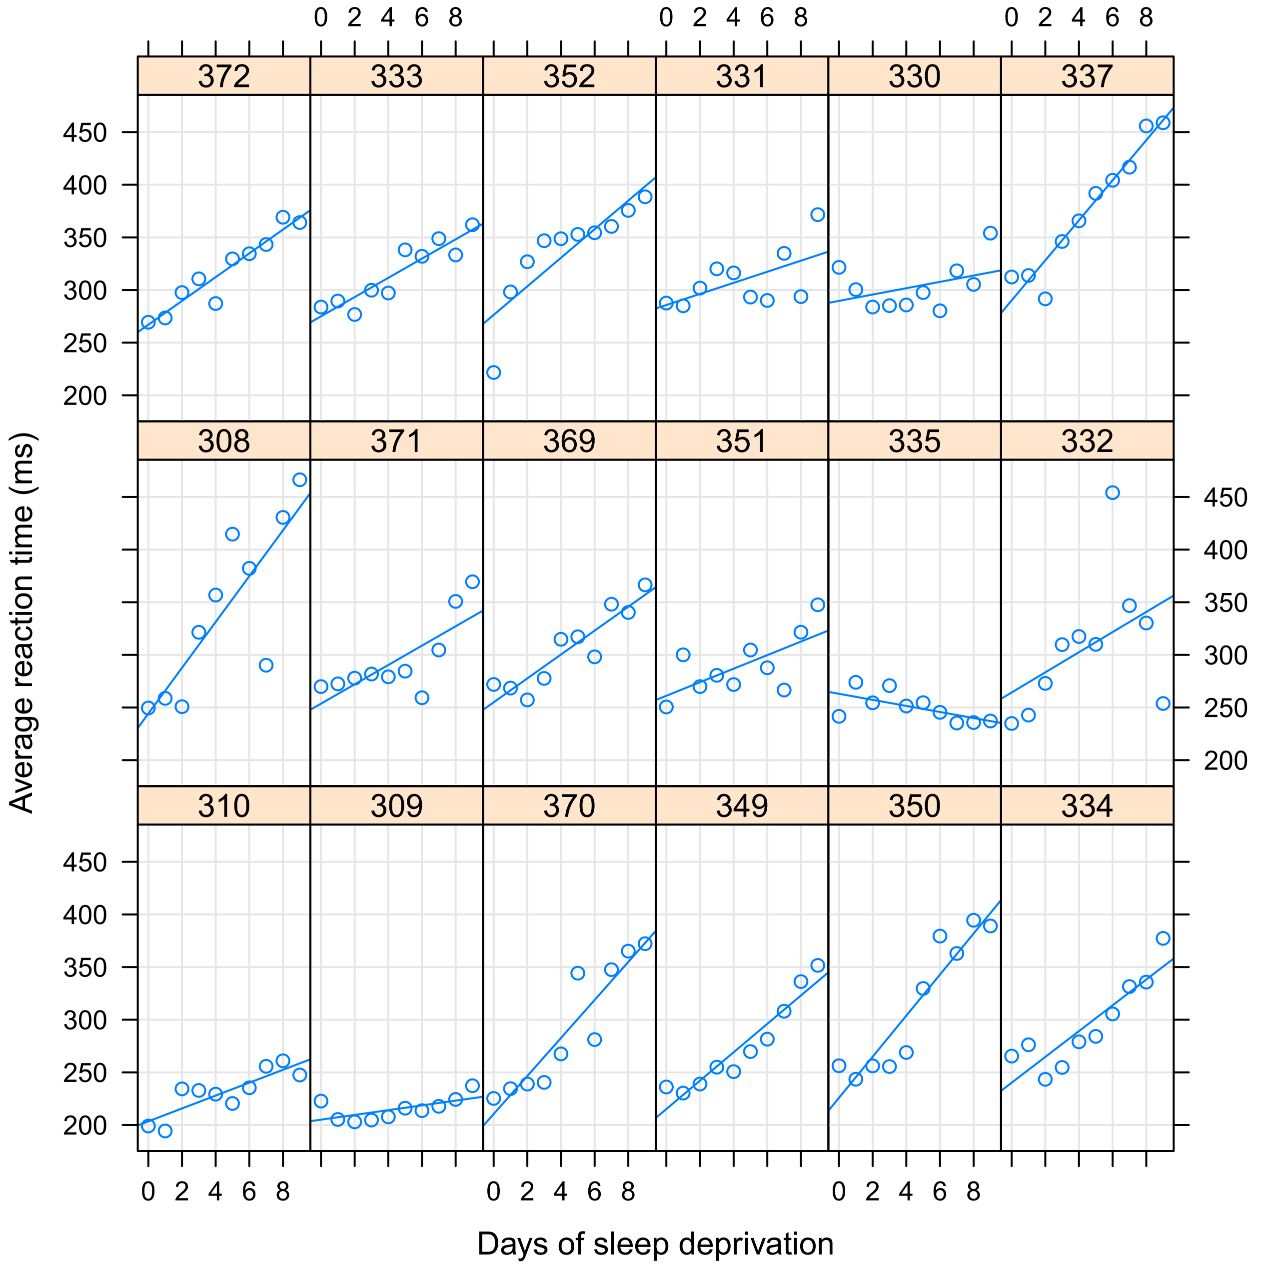

372
333
352
331
330
337
308
371
369
351
335
332
310
309
370
349
350
334
Average reaction time (ms)
Days of sleep deprivation

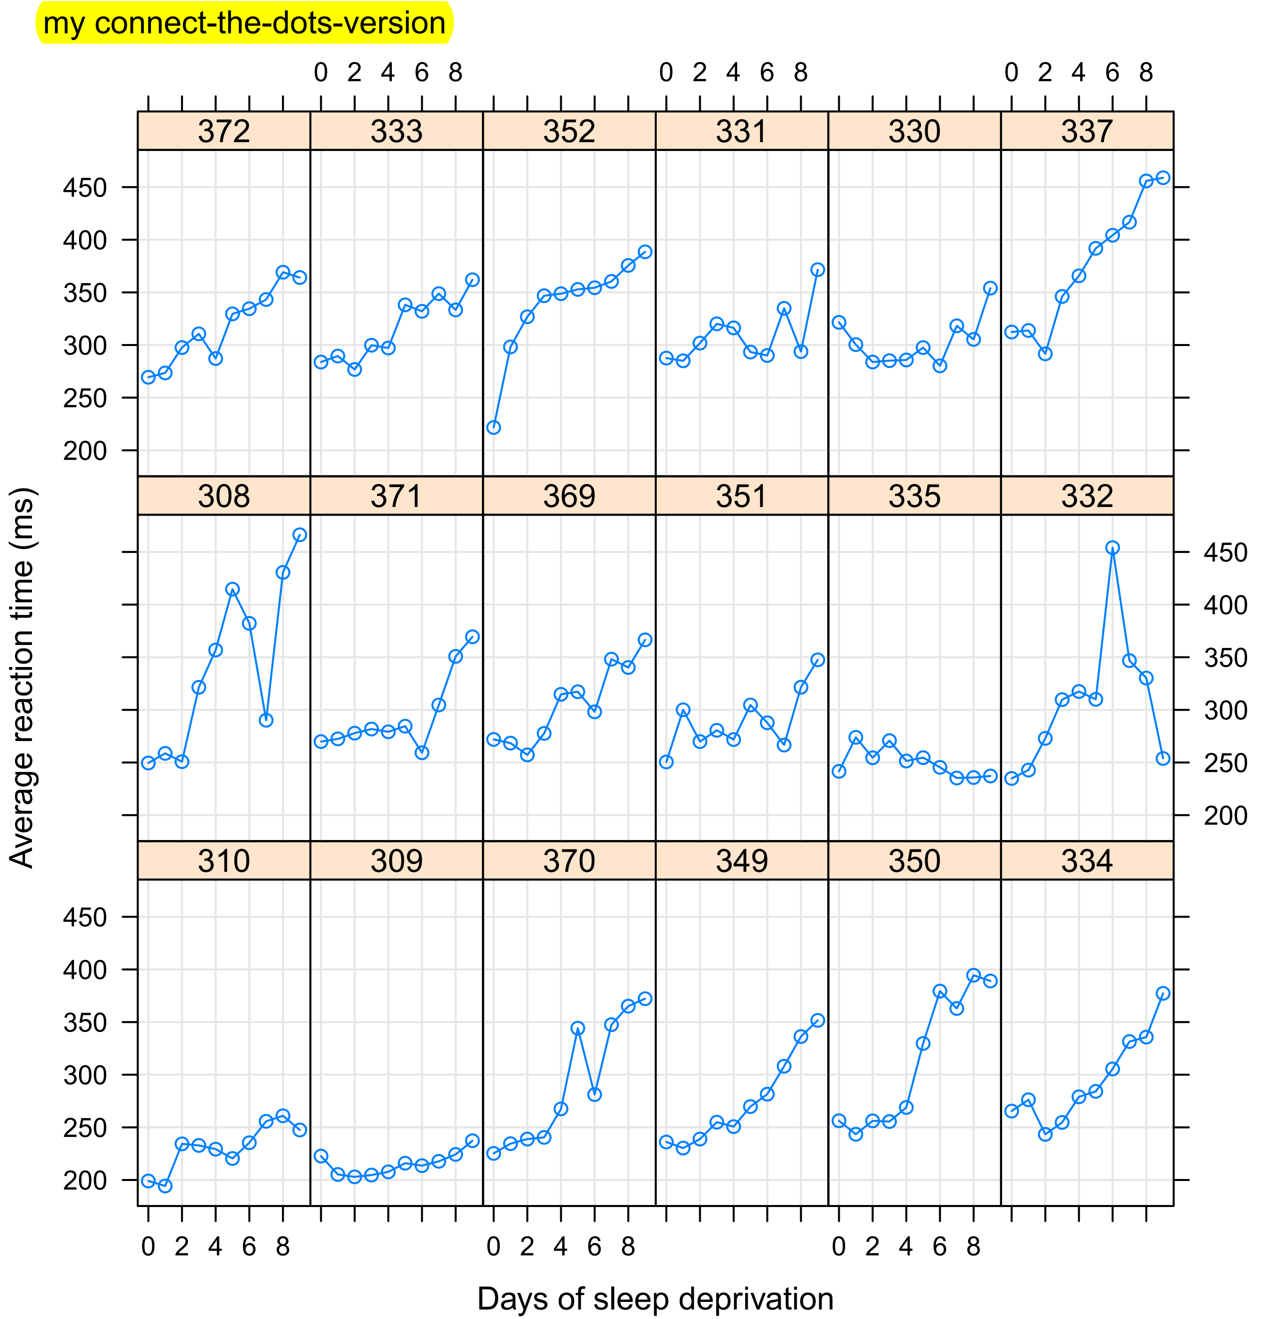

my connect-the-dots-version
372
333
352
331
330
337
308
371
369
351
335
332
310
309
370
349
350
334
Average reaction time (ms)
Days of sleep deprivation

4.1.1 *Characteristics of the* sleepstudy *Data Plot*

The principles of "Trellis graphics", developed by Bill Cleveland and his coworkers at Bell Labs and implemented in the `lattice` package for R by Deepayan Sarkar, have been incorporated in this plot. As stated above, all the panels have the same vertical and horizontal scales, allowing us to evaluate the pattern over time for each subject and also to compare patterns between subjects. The line drawn in each panel is a simple least squares line fit to the data in that panel only. It is provided to enhance our ability to discern patterns in both the slope (the typical change in reaction time per day of sleep deprivation for that particular subject) and the intercept (the average response time for the subject when on their usual sleep pattern).

The aspect ratio of the panels (ratio of the height to the width) has been chosen, according to an algorithm described in Cleveland [1993], to facilitate comparison of slopes. The effect of choosing the aspect ratio in this way is to have the slopes of the lines on the page distributed around $\pm 45°$, thereby making it easier to detect systematic changes in slopes.

The panels have been ordered (from left to right starting at the bottom row) by increasing intercept. Because the subject identifiers, shown in the strip above each panel, are unrelated to the response it would not be helpful to use the default ordering of the panels, which is by increasing subject number. If we did so our perception of patterns in the data would be confused by the, essentially random, ordering of the panels. Instead we use a characteristic of the data to determine the ordering of the panels, thereby enhancing our ability to compare across panels. For example, a question of interest to the experimenters is whether a subject's rate of change in reaction time is related to the subject's initial reaction time. If this were the case we would expect that the slopes would show an increasing trend (or, less likely, a decreasing trend) in the left to right, bottom to top ordering.

There is little evidence in Fig. 4.1 of such a systematic relationship between the subject's initial reaction time and their rate of change in reaction time per day of sleep deprivation. We do see that for each subject, except 335, reaction time increases, more-or-less linearly, with days of sleep deprivation. However, there is considerable variation both in the initial reaction time and in the daily rate of increase in reaction time. We can also see that these data are balanced, both with respect to the number of observations on each subject, and with respect to the times at which these observations were taken. This can be confirmed by cross-tabulating `Subject` and `Days`.

```
> xtabs(~ Subject + Days, sleepstudy)

       Days
Subject 0 1 2 3 4 5 6 7 8 9
    308 1 1 1 1 1 1 1 1 1 1
    309 1 1 1 1 1 1 1 1 1 1
    310 1 1 1 1 1 1 1 1 1 1
    330 1 1 1 1 1 1 1 1 1 1
```

CODE???

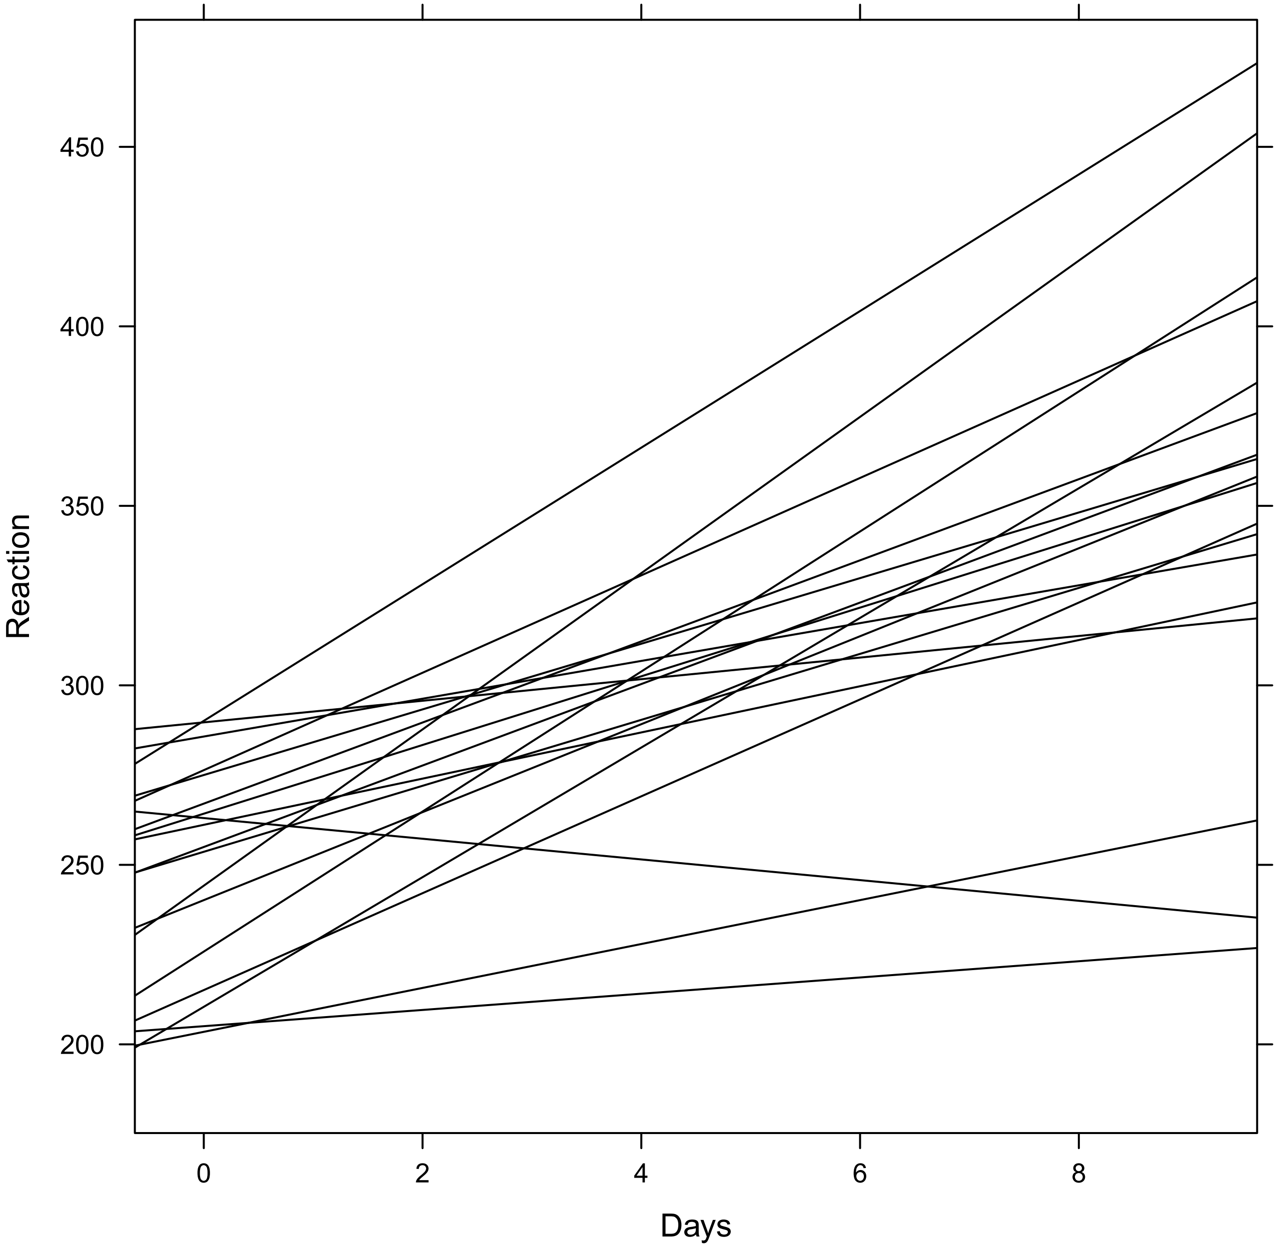

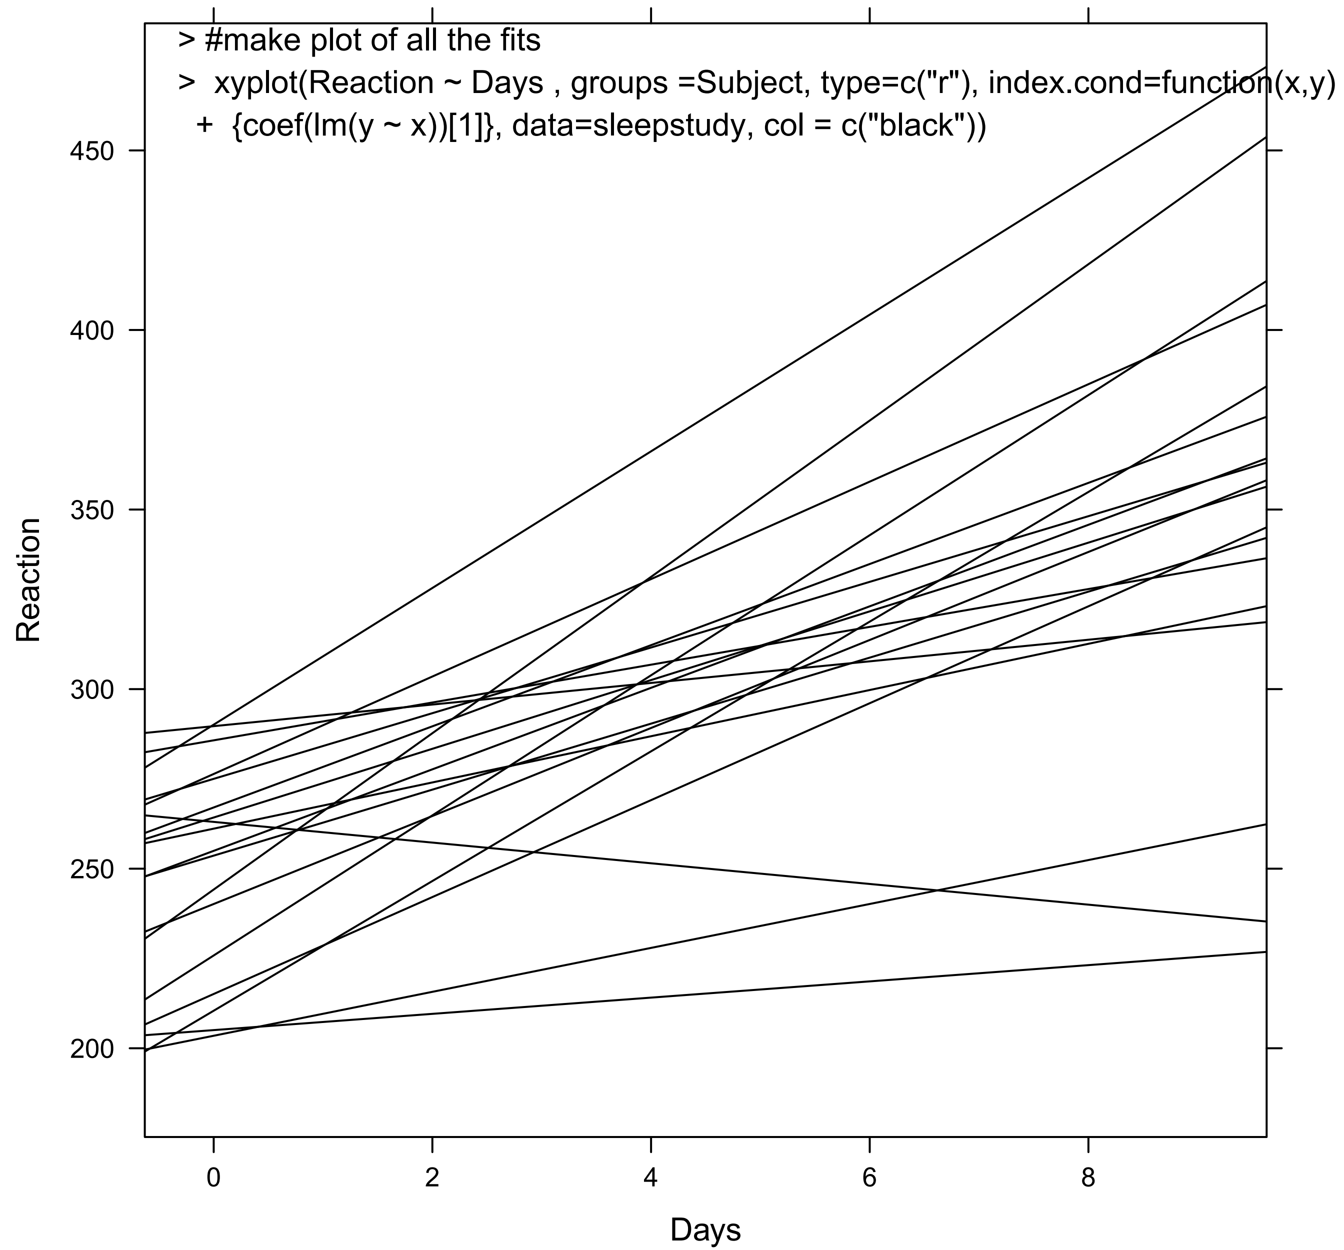

> #make plot of all the fits
> xyplot(Reaction ~ Days , groups =Subject, type=c("r"), index.cond=function(x,y)
+ {coef(lm(y ~ x))[1]}, data=sleepstudy, col = c("black"))
Reaction
450
400
350
300
250
200
0
2
4
6
8
Days

```
R version 3.4.4 (2018-03-15) -- "Someone to Lean On"
Copyright (C) 2018 The R Foundation for Statistical Computing
> library(lme4)
Loading required package: Matrix
### Sleepstudy descriptives SFYS
> data(sleepstudy)  # sleep deprivation 3hrs/night truckers
> dim(sleepstudy)
[1] 180   3
> head(sleepstudy)
  Reaction Days Subject
1 249.5600    0     308
2 258.7047    1     308
3 250.8006    2     308
4 321.4398    3     308
5 356.8519    4     308
6 414.6901    5     308
## note initial observation time = 0
> attach(sleepstudy)
> table(Subject)
Subject
308 309 310 330 331 332 333 334 335 337 349 350 351 352 369 370 371 372
 10  10  10  10  10  10  10  10  10  10  10  10  10  10  10  10  10  10
> #yes 18 subjects (3hrs sleep) 10 observations on each, page 64 of bates ch4
>
>
>   #lmList has the old nlme syntax, even in lme4
>  sleeplmList = lmList(Reaction ~ Days |Subject, data = sleepstudy)
>  sleeplmList
Call: lmList(formula = Reaction ~ Days | Subject, data = sleepstudy)
Coefficients:
    (Intercept)      Days
308    244.1927 21.764702
309    205.0549  2.261785
310    203.4842  6.114899
330    289.6851  3.008073
331    285.7390  5.266019
332    264.2516  9.566768
333    275.0191  9.142045
334    240.1629 12.253141
335    263.0347 -2.881034
337    290.1041 19.025974
349    215.1118 13.493933
350    225.8346 19.504017
351    261.1470  6.433498
352    276.3721 13.566549
369    254.9681 11.348109
370    210.4491 18.056151
371    253.6360  9.188445
372    267.0448 11.298073
```

```
Degrees of freedom: 180 total; 144 residual
Residual standard error: 25.59182
> # note this matches lmer Residual (random eff)
>  mean(coef(sleeplmList)[,1])
[1] 251.4051
>  mean(coef(sleeplmList)[,2])
[1] 10.46729
> #mean int and slope match lmer Fixed effects results
>
> #### Random -> varies over units (subj); Fixed -> not vary

> # heterogeneity across units
>  var(coef(sleeplmList)[,1])
[1] 838.3423
>  var(coef(sleeplmList)[,2])
[1] 43.01034

>  #these are too big as they should be, compare with variance random effects
>  # mle subtracts off wobble in estimated indiv regressions Y on t
>  # method of moments - SSR/SST, see Review Question
>
> ## more useful descriptives using lmList results
>  quantile(coef(sleeplmList)[,1])
      0%      25%      50%      75%     100% 
203.4842 229.4167 258.0576 273.0255 290.1041 
>
>  quantile(coef(sleeplmList)[,2])
       0%       25%       50%       75%      100% 
-2.881034  6.194548 10.432421 13.548395 21.764702 
>
>  stem(coef(sleeplmList)[,2])

  The decimal point is 1 digit(s) to the right of the |

  -0 | 3
   0 | 23
   0 | 56699
   1 | 011234
   1 | 89
   2 | 02

> ## can also rename quantities
>  rate = coef(sleeplmList)[,2]
> quantile(rate)
       0%       25%       50%       75%      100% 
-2.881034  6.194548 10.432421 13.548395 21.764702 
>
```

STAT 222 Sleepstudy (Bates) 1

```
R version 2.10.1 (2009-12-14)                    Stat222, Week2
> library(lme4)  # the "new" random effects (vs nlme) see Bates book 2010
Loading required package: Matrix   Loading required package: lattice
> data(sleepstudy)  # sleep deprivation 3hrs/night truckers
> dim(sleepstudy)    [1] 180    3
> head(sleepstudy)
  Reaction Days Subject
1 249.5600    0     308
2 258.7047    1     308
3 250.8006    2     308
4 321.4398    3     308
5 356.8519    4     308
6 414.6901    5     308
```

data source lme4 package

note initial observation days=0

```
> attach(sleepstudy) > table(Subject)   #"balanced data"
Subject
308 309 310 330 331 332 333 334 335 337 349 350 351 352 369 370 371 372
 10  10  10  10  10  10  10  10  10  10  10  10  10  10  10  10  10  10
> #yes 18 subjects (3hrs sleep) 10 observations on each, page 64 of bates ch4
  #lmList has the old nlme syntax, even in lme4  #Start "Smart First Year Student" description
> sleeplmList = lmList(Reaction ~ Days |Subject, data = sleepstudy)
> sleeplmList      Call: lmList(formula = Reaction ~ Days | Subject, data = sleepstudy)
Coefficients:
    (Intercept)       Days
308    244.1927 21.764702
309    205.0549  2.261785
310    203.4842  6.114899
330    289.6851  3.008073
331    285.7390  5.266019
332    264.2516  9.566768
333    275.0191  9.142045
334    240.1629 12.253141
335    263.0347 -2.881034
337    290.1041 19.025974
349    215.1118 13.493933
350    225.8346 19.504017
351    261.1470  6.433498
352    276.3721 13.566549
369    254.9681 11.348109
370    210.4491 18.056151
371    253.6360  9.188445
372    267.0448 11.298073
Degrees of freedom: 180 total; 144 residual
Residual standard error: 25.59182  # note this matches Bates lmer Residual (random eff)
> mean(coef(sleeplmList)[,1])    [1] 251.4051
> mean(coef(sleeplmList)[,2])    [1] 10.46729
> #mean int and slope match lmer Fixed effects results p.67
> var(coef(sleeplmList)[,1])    [1] 838.3423
> var(coef(sleeplmList)[,2])    [1] 43.01034
> #these are too big as they should be, compare with variance random effects p.67
> # mle subtracts off wobble in estimated indiv regressions Y on t
> # - SSR/SST
```

10 obs, 18 subjects

SFYS

lmList

[show data and fit plots linked]

```
> quantile(coef(sleeplmList)[,1])
      0%      25%      50%      75%     100%
203.4842 229.4167 258.0576 273.0255 290.1041
> quantile(coef(sleeplmList)[,2])
       0%       25%       50%       75%      100%
-2.881034  6.194548 10.432421 13.548395 21.764702
```

level (day0) fit

rate

```
> stem(coef(sleeplmList)[,2])
  The decimal point is 1 digit(s) to the right of the |
 -0 | 3
  0 | 23
  0 | 56699
  1 | 011234
  1 | 89
  2 | 02
   # series of (redundent) lmer analyses for expository
  # first 2 do REML, second pair match Bates p.67 in doing MLE (REML FALSE)
> sleeplmer = lmer(Reaction ~ Days + (1 + Days|Subject), sleepstudy)
> summary(sleeplmer)
Linear mixed model fit by REML
Formula: Reaction ~ Days + (1 + Days | Subject)    Data: sleepstudy
  AIC  BIC logLik deviance REMLdev
 1756 1775 -871.8     1752    1744
```

rate

growth curve

mixed effects

Level 1 $R = \alpha_0 + \alpha_1 D + \varepsilon$

Level 2 $\alpha_0 = \gamma_{00} + u_0$

$\alpha_1 = \gamma_{10} + u_1$ fixed effects

lmer replaces lme in nlme

Stat222 Week2

```
R version 3.4.4 (2018-03-15) -- "Someone to Lean On"

>     # series of (redundent) lmer analyses for expository
>    # first 2 do REML, second pair match Bates p.67 in doing MLE (REML FALSE)
> # refer to unconditional model "A" in handout
> ## note dataframe has initial time = 0;
          that's assumed in the "intercept" setting in lmer

>  sleeplmer = lmer(Reaction ~ Days + (1 + Days|Subject), sleepstudy)
>  summary(sleeplmer)
Linear mixed model fit by REML ['lmerMod']
Formula: Reaction ~ Days + (1 + Days | Subject)
   Data: sleepstudy

REML criterion at convergence: 1743.6

Scaled residuals:
    Min      1Q  Median      3Q     Max
-3.9536 -0.4634  0.0231  0.4634  5.1793

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 Subject  (Intercept) 612.09   24.740
          Days         35.07    5.922   0.07
 Residual             654.94   25.592
Number of obs: 180, groups:  Subject, 18

Fixed effects:
            Estimate Std. Error t value
(Intercept)  251.405      6.825  36.838
Days          10.467      1.546   6.771

Correlation of Fixed Effects:
     (Intr)
Days -0.138
>  ## also note Corr .07 from random effects vs Corr -.138 from fixed effects
                 see RQ
> cor(coef(sleeplmList)[,1],coef(sleeplmList)[,2])
[1] -0.1375534
> confint(sleeplmer)  # best inference tool, see RQ for bootstrap version
Computing profile confidence intervals ...
                 2.5 %     97.5 %
.sig01       14.3815822  37.715996
.sig02       -0.4815007   0.684986
.sig03        3.8011641   8.753383
.sigma       22.8982669  28.857997
(Intercept) 237.6806955 265.129515
Days          7.3586533  13.575919
```

```
> confint(sleeplmer, oldNames = FALSE)  # to get good labels
Computing profile confidence intervals ...
                                  2.5 %     97.5 %
sd_(Intercept)|Subject        14.3815822  37.715996
cor_Days.(Intercept)|Subject  -0.4815007   0.684986
sd_Days|Subject                3.8011641   8.753383
sigma                         22.8982669  28.857997
(Intercept)                  237.6806955 265.129515
Days                           7.3586533  13.575919

# useful residual plots in RQ, week 3
#########  end of analysis ########################
#For comparison

> sleeplmer2 = lmer(Reaction ~ 1 + Days + (1 + Days|Subject), sleepstudy)
>  summary(sleeplmer2)
Linear mixed model fit by REML ['lmerMod']
Formula: Reaction ~ 1 + Days + (1 + Days | Subject)
   Data: sleepstudy

REML criterion at convergence: 1743.6

Scaled residuals:
    Min      1Q  Median      3Q     Max
-3.9536 -0.4634  0.0231  0.4634  5.1793

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 Subject  (Intercept) 612.09   24.740
          Days         35.07    5.922   0.07
 Residual             654.94   25.592
Number of obs: 180, groups:  Subject, 18

Fixed effects:
            Estimate Std. Error t value
(Intercept)  251.405      6.825  36.838
Days          10.467      1.546   6.771

Correlation of Fixed Effects:
     (Intr)
Days -0.138
### Maximum liklelihood instead of REML
> sleeplmer3 = lmer(Reaction ~ 1 + Days + (1 + Days|Subject),
                                   sleepstudy, REML = FALSE)
>  summary(sleeplmer3)
Linear mixed model fit by maximum likelihood  ['lmerMod']
Formula: Reaction ~ 1 + Days + (1 + Days | Subject)
   Data: sleepstudy

     AIC      BIC   logLik deviance df.resid
```

```
  1763.9   1783.1   -876.0   1751.9      174

Scaled residuals: 
    Min      1Q  Median      3Q     Max 
-3.9416 -0.4656  0.0289  0.4636  5.1793 

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 Subject  (Intercept) 565.52   23.781       
          Days         32.68    5.717   0.08
 Residual             654.94   25.592       
Number of obs: 180, groups:  Subject, 18

Fixed effects:
            Estimate Std. Error t value
(Intercept)  251.405      6.632  37.906
Days          10.467      1.502   6.968

Correlation of Fixed Effects:
     (Intr)
Days -0.138

> sleeplmer4 = lmer(Reaction ~ Days + (1 + Days|Subject),
                              sleepstudy, REML = FALSE)
>  summary(sleeplmer4)
Linear mixed model fit by maximum likelihood  ['lmerMod']
Formula: Reaction ~ Days + (1 + Days | Subject)
   Data: sleepstudy

     AIC      BIC   logLik deviance df.resid 
  1763.9   1783.1   -876.0   1751.9      174 

Scaled residuals: 
    Min      1Q  Median      3Q     Max 
-3.9416 -0.4656  0.0289  0.4636  5.1793 

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 Subject  (Intercept) 565.52   23.781       
          Days         32.68    5.717   0.08
 Residual             654.94   25.592       
Number of obs: 180, groups:  Subject, 18

Fixed effects:
            Estimate Std. Error t value
(Intercept)  251.405      6.632  37.906
Days          10.467      1.502   6.968

Correlation of Fixed Effects:
     (Intr)
```

```
Days -0.138
>
```

②

```
Random effects:
 Groups   Name        Variance Std.Dev. Corr
 Subject  (Intercept) 612.092  24.7405
          Days         35.072   5.9221  0.066
 Residual             654.941  25.5918
Number of obs: 180, groups: Subject, 18
Fixed effects:
            Estimate Std. Error t value
(Intercept)  251.405      6.825   36.84
Days          10.467      1.546    6.77
Correlation of Fixed Effects:
     (Intr)
Days -0.138
> cor(coef(sleeplmList)[,1],coef(sleeplmList)[,2])  [1] -0.1375534
```

$\hat{\gamma}_{00}$ $\hat{\gamma}_{10}$ means from SFYS

random: varies over units (subj)

fixed: constant over units (subj)

confint in revQ#1

```
> sleeplmer2 = lmer(Reaction ~ 1 + Days + (1 + Days|Subject), sleepstudy)
> summary(sleeplmer2)   #same as above, intercept implicit (else set to 0)
Linear mixed model fit by REML
Formula: Reaction ~ 1 + Days + (1 + Days | Subject)    Data: sleepstudy
  AIC  BIC logLik deviance REMLdev
 1756 1775 -871.8     1752    1744
Random effects:
 Groups   Name        Variance Std.Dev. Corr
 Subject  (Intercept) 612.092  24.7405
          Days         35.072   5.9221  0.066
 Residual             654.941  25.5918
Number of obs: 180, groups: Subject, 18
Fixed effects:
            Estimate Std. Error t value
(Intercept)  251.405      6.825   36.84
Days          10.467      1.546    6.77
Correlation of Fixed Effects:
     (Intr)
Days -0.138
```

same model

Laird-Ware do REML

```
> sleeplmer3 = lmer(Reaction ~ 1 + Days + (1 + Days|Subject), sleepstudy, REML = FALSE)
> summary(sleeplmer3)  #mle results a little different than REML, as it should
Linear mixed model fit by maximum likelihood
Formula: Reaction ~ 1 + Days + (1 + Days | Subject)    Data: sleepstudy
  AIC  BIC logLik deviance REMLdev
 1764 1783   -876     1752    1744
Random effects:
 Groups   Name        Variance Std.Dev. Corr
 Subject  (Intercept) 565.518  23.7806
          Days         32.682   5.7168  0.081
 Residual             654.941  25.5918
Number of obs: 180, groups: Subject, 18
Fixed effects:
            Estimate Std. Error t value
(Intercept)  251.405      6.632   37.91
Days          10.467      1.502    6.97
Correlation of Fixed Effects:
     (Intr)
Days -0.138
```

mle versions in Bates redundants

```
> sleeplmer4 = lmer(Reaction ~ Days + (1 + Days|Subject), sleepstudy, REML = FALSE)
> summary(sleeplmer4)
Linear mixed model fit by maximum likelihood
Formula: Reaction ~ Days + (1 + Days | Subject)    Data: sleepstudy
  AIC  BIC logLik deviance REMLdev
 1764 1783   -876     1752    1744
Random effects:
 Groups   Name        Variance Std.Dev. Corr
 Subject  (Intercept) 565.518  23.7806
          Days         32.682   5.7168  0.081
 Residual             654.941  25.5918
Number of obs: 180, groups: Subject, 18
Fixed effects:
            Estimate Std. Error t value
(Intercept)  251.405      6.632   37.91
Days          10.467      1.502    6.97
Correlation of Fixed Effects:
     (Intr)
Days -0.138
```

Review Question #1

```
R version 3.0.3 (2014-03-06) -- "Warm Puppy"
Copyright (C) 2014 The R Foundation for Statistical Computing
Platform: x86_64-w64-mingw32/x64 (64-bit)


> library(lme4)
Loading required package: lattice
Loading required package: Matrix
> data(sleepstudy)
> ?data
starting httpd help server ... done
> try(data(package = "lme4")) # to list datasets in lme4
> ?sleepstudy

> sleeplmer = lmer(Reaction ~ Days + (1 + Days|Subject), sleepstudy)

> summary(sleeplmer) # as we did in class exs
Linear mixed model fit by REML ['lmerMod']
Formula: Reaction ~ Days + (1 + Days | Subject)
   Data: sleepstudy

REML criterion at convergence: 1743.628

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 Subject  (Intercept) 612.10   24.741
          Days         35.07    5.922   0.07
 Residual             654.94   25.592
Number of obs: 180, groups: Subject, 18

Fixed effects:
            Estimate Std. Error t value
(Intercept)  251.405      6.825   36.84
Days          10.467      1.546    6.77

Correlation of Fixed Effects:
     (Intr)
Days -0.138

> confint(sleeplmer) # basic set of 95% CI (profile method)
Computing profile confidence intervals ...
                  2.5 %     97.5 %
.sig01       14.3815734  37.715996   # parameter sqrt(variance component for intercept parameter) i.e sqrt(Var(alpha_0)
.sig02       -0.4815007   0.684986   # parameter Cor(alpha_0, alpha_1) , rather wide
.sig03        3.8011641   8.753383   # parameter sqrt(variance component for rate (slope) parameter) i.e sqrt(Var(alpha_1)
.sigma       22.8982669  28.857997   # parameter sqrt(residual variance individual regressions) i.e. sqrt(Var(epsilon)
(Intercept) 237.6806955 265.129515   # parameter fixed effect gamma_00, i.e. mean(alpha_0)
Days          7.3586533  13.575919   # parameter fixed effect gamma_10, i.e. mean(alpha_1)
### model defined in class handout sleepstudy

> confint(sleeplmer, level = .99)  # try higher confidence level, wider intervals
Computing profile confidence intervals ...
                 0.5 %      99.5 %
.sig01       11.697852  43.9249121
.sig02       -0.611183   0.8620229
.sig03        3.313161  10.1306220
.sigma       22.149064  30.0282535
(Intercept) 232.619539 270.1906680
Days          6.212281  14.7222899

> confint(sleeplmer, method = "boot", nsim = 1000, boot.type = "perc") # try bootstrap, just percentile method, CI; more useful labels here
Computing bootstrap confidence intervals ...
                                 2.5 %      97.5 %
sd_(Intercept)|Subject        13.0598528  35.1257317
cor_Days.(Intercept)|Subject  -0.5137308   0.9294877
sd_Days|Subject                3.6573704   8.3907217
sigma                         22.8591370  28.4592246
(Intercept)                  239.0486723 265.1112312
Days                           7.1933705  13.6621452
Warning messages:
1: In cov2cor(m) :
  diag(.) had 0 or NA entries; non-finite result is doubtful
2: In bootMer(object, bootFun, nsim = nsim, ...) :
  some bootstrap runs failed (1/1000)
> confint(sleeplmer, method = "boot", nsim = 1000, boot.type = "perc") # try again as we had one bad resample in last one, not important
Computing bootstrap confidence intervals ...
                                 2.5 %      97.5 %
sd_(Intercept)|Subject        12.5289697  35.3488520
cor_Days.(Intercept)|Subject  -0.4756043   0.9096473
sd_Days|Subject                3.2938939   8.3535371
sigma                         22.6564232  28.4105186
(Intercept)                  237.3140641 265.0539370
Days                           7.5707244  13.7008761
> confint(sleeplmer, method = "boot", nsim = 3000, boot.type = "perc") # with a higher number resamples
Computing bootstrap confidence intervals ...
                                 2.5 %     97.5 %
sd_(Intercept)|Subject        12.6096568  35.420349
cor_Days.(Intercept)|Subject  -0.5146076   0.961081
sd_Days|Subject                3.3717408   8.364048
sigma                         22.6353905  28.474294
(Intercept)                  238.1725151 264.494443
Days                           7.4794324  13.426824

# bootstrap can take a couple minutes on your machines

##############################################################################################
Now to p-values, just because I promised.
I don't have much concern about the abscence of p-values in lme4
See the p-values section of lme4 (p.77)
```

We try the mixed function from package afex (among those discussed in lme4 manual)

```
> install.packages("afex")
Installing package into ‘C:/Users/rag/Documents/R/win-library/3.0’
(as ‘lib’ is unspecified)
```

see 2020 update for RQ1; afex is easier and better output if you want it

```
The downloaded binary packages are in
        C:\Users\rag\AppData\Local\Temp\Rtmp88hftJ\downloaded_packages

> library(afex)
Loading required package: car
Loading required package: reshape2
************
Welcome to afex. Important notes:

Due to popular demand, afex doesn't change the contrasts globally anymore.
To set contrasts globally to contr.sum run set_sum_contrasts().
To set contrasts globally to the default (treatment) contrasts run set_default_contrasts().

All afex functions are unaffected by global contrasts and use contr.sum as long as check.contr = TRUE (which is the default).
************
> ?mixed
> #Obtain p-values for a mixed-model from lmer().

> mixed(Reaction ~ Days + (1 + Days|Subject), sleepstudy)
Contrasts set to contr.sum for the following variables: Subject
Numerical variables NOT centered on 0 (i.e., interpretation of all main effects might be difficult if in interactions): Days
Fitting 2 (g)lmer() models:
[..]
Obtaining 1 p-values:
[Note: method with signature ‘sparseMatrix#ANY’ chosen for function ‘kronecker’,
 target signature ‘dgCMatrix#ngCMatrix’.
 "ANY#sparseMatrix" would also be valid
.]
  Effect     F ndf   ddf F.scaling p.value
1   Days 45.85   1 17.00      1.00  <.0001  # it's good mixed only gives you the p-value is small rather than some rediculous e^-10

> mixed(Reaction ~ 1 + Days + (1 + Days|Subject), sleepstudy) # no diff in way maodel is specified
Contrasts set to contr.sum for the following variables: Subject
Numerical variables NOT centered on 0 (i.e., interpretation of all main effects might be difficult if in interactions): Days
Fitting 2 (g)lmer() models:
[..]
Obtaining 1 p-values:
[.]
  Effect     F ndf   ddf F.scaling p.value
1   Days 45.85   1 17.00      1.00  <.0001
>
> sqrt(45.85)
[1] 6.771263
```

Fitting the model

```
> (fm1 <- lmer(Reaction ~ Days + (Days|Subject), sleepstudy))
```

```
Linear mixed model fit by REML ['merMod']
Formula: Reaction ~ Days + (Days | Subject)
   Data: sleepstudy
REML criterion at convergence: 1743.628
Random effects:
 Groups   Name        Variance Std.Dev. Corr
 Subject  (Intercept) 612.09   24.740
          Days         35.07    5.922   0.066
 Residual             654.94   25.592
Number of obs: 180, groups: Subject, 18

Fixed effects:
            Estimate Std. Error t value
(Intercept)  251.405      6.825   36.84
Days          10.467      1.546    6.77

Correlation of Fixed Effects:
     (Intr)
Days -0.138
```

```
> (fm8 <- lmer(Reaction ~ 1 + Days + (1 + Days|Subject), sleepstudy,
+              REML = 0))
```

```
Linear mixed model fit by maximum likelihood
Formula: Reaction ~ 1 + Days + (1 + Days | Subject)
   Data: sleepstudy
  AIC  BIC logLik deviance
 1764 1783   -876     1752
Random effects:
 Groups   Name        Variance Std.Dev. Corr
 Subject  (Intercept) 565.516  23.7806
          Days         32.682   5.7168  0.081
 Residual             654.941  25.5918
Number of obs: 180, groups: Subject, 18

Fixed effects:
            Estimate Std. Error t value
(Intercept)  251.405      6.632   37.91
Days          10.467      1.502    6.97

Correlation of Fixed Effects:
     (Intr)
Days -0.138
```

2x2 cov matrix of random effects

From the display we see that this model incorporates both an intercept and a slope (with respect to `Days`) in the fixed effects and in the random effects. Extracting the conditional modes of the random effects

```
> head(ranef(fm8)[["Subject"]])
```

```
    (Intercept)       Days
308    2.815683  9.0755340
309  -40.048490 -8.6440671
310  -38.433156 -5.5133785
330   22.832297 -4.6587506
331   21.549991 -2.9445203
332    8.815587 -0.2352093
```

confirms that these are *vector-valued* random effects. There are a total of $q = 36$ random effects, two for each of the 18 subjects.

The random effects section of the model display,

```
 Groups   Name        Variance Std.Dev. Corr
 Subject  (Intercept) 565.516  23.7806
          Days         32.682   5.7168  0.081
 Residual             654.941  25.5918
```

indicates that there will be a random effect for the intercept and a random effect for the slope with respect to `Days` at each level of `Subject` and, furthermore, the unconditional distribution of these random effects, $\mathscr{B} \sim \mathscr{N}(\mathbf{0}, \Sigma)$, allows for correlation of the random effects for the same subject.

We can confirm the potential for correlation of random effects within subject in the images of Λ, Σ and $\mathbf{L}$ for this model (Fig. 4.2). The matrix Λ has

Data Analysis and Parameter Estimation

Precursor: Descriptive Growth Curve Analyses

SFYS: fit Y on t regressions, describe resulting $\hat{\theta}_p$, fit $\hat{\theta}_p$ on W regr,
Examples: WISC, frames 1-4; Ramus, frames 1-3; SmearMiss, frames 1-3.
Even non-synchronous data, get variance comps and derived quants by approx method-of-moments (Rogosa-Saner 1995); works surprisingly well.

Maximum Likelihood estimation for parameters

Special, simple case; Complete, Synchronous Data.
ml estimation equations for full data in closed form (Blomqvist 1977)
example estimation of var(theta) σ^2_θ
MSR_p mean squared residual for OLS fit individual p; $\hat{\sigma}^2$ is Ave(MSR_p).
estimate for σ^2_θ : $\hat{\sigma}^2_\theta$ = SS($\hat{\theta}_p$)/"n" - $\hat{\sigma}^2/SSt$,
reliability estimate for $\hat{\theta}_p$: $\hat{\rho}(\hat{\theta})$ = $\hat{\sigma}^2_\theta$ /SS($\hat{\theta}_p$)/"n"

General strategy: get elements of 2x2 est. covariance matrix of θ and $\eta(0)$ for full or incomplete data. Common to **All programs** (**LISREL** HLM Tp) Tp: further substitute for derived quantities.
Also, fixed effects from separate run with W (when exists)–OLS equiv

properties of mle: bias, precision: Is reml best?
bias and mean-square-error : compare ML and REML
mle and reml simulation (50,000); complete synchronous data

Estimation of σ^2_θ var(theta) = 5.0

n	ML		REML	
10	4.37	[7.39]	4.99	[8.61]
15	4.58	[5.06]	4.99	[5.59]

MAJOR MESSAGES
1. OLS equivalences for fixed effects; Method-of-moments match for random effects
2. 2x2 covariance matrix ($\eta_p(0)$ θ_p)-- elements σ^2_θ $\sigma^2_{\eta(0)}$ $\sigma_{\eta(0)\theta}$ –starting point for growth statistics
3. uncertainty, via s.e. and CI, reporting essential–for small (or medium) n, BCa intervals vs standard

Properties (Moments of Observables) of Collections of Growth Curves

STAT 222
[redacted]
D Rogosa

for indiv p $\quad \xi_p(t) = \xi_p(0) + \theta_p t$ $\qquad t_i \; (i = 1, \ldots T)$
$p \; (p = 1, \ldots n)$

centering, scale
$t^o = -\sigma_{\xi(0)\theta}/\sigma_\theta^2$
$\rho_{\xi(t^o)\theta} = 0$, min $Var(\xi)$
Scale $\kappa = \sigma_{\xi(t^o)}/\sigma_\theta$
time metric

"centering"

$$\xi_p(t) = \xi_p(t^o) + \theta_p(t - t^o)$$

Moments

Covariance $\sigma_{\xi(t_1)\xi(t_2)} =$
$$(t_1 - t^o)(t_2 - t^o)\sigma_\theta^2 + \sigma^2_{\xi(t^o)}$$

Variance $\quad \sigma^2_{\xi(t)} = \sigma^2_{\xi(t^o)} + ((t - t^o)/\kappa)^2 \sigma^2_{\xi(t^o)}$
$$\sigma^2_{\xi(t)}/\sigma^2_{\xi(t^o)} = 1 + \left(\frac{t - t^o}{\kappa}\right)^2$$

Correl change, initial status $\quad \rho_{\theta\xi(t)} = \dfrac{t - t^o}{[\kappa^2 + (t - t^o)^2]^{1/2}}$

exogenous var W

$$\rho_{W\xi(t)} = \frac{(t - t^o)\rho_{W\theta} + \kappa \rho_{W\xi(t^o)}}{[\kappa^2 + (t - t^o)^2]^{1/2}}$$

where $\quad t^u = t^o + \kappa\left(\dfrac{\rho_{W\theta}}{\rho_{W\xi(t^o)}}\right) \qquad t^l = t^o - \kappa\left(\dfrac{\rho_{W\xi(t^o)}}{\rho_{W\theta}}\right)$

Myths Week 1 example: (p.64) $\theta \sim U[1, 9]$, $\xi(t^o) \sim U[38, 62]$
$t^o = 2 \quad \sigma_\theta^2 = 5.333 \quad \sigma^2_{\xi(t^o)} = 48 \quad \rho_{W\theta} = 0 \quad \rho_{W\xi(t^o)} =$

at time t_i $\quad X_{ip} = \xi_{ip} + \epsilon \qquad \epsilon \sim (0, \sigma_\epsilon^2)$ $\quad$ errors in variables
week 1 ex $\sigma_\epsilon^2 = 10$

WISC

4 observations, Wechsler Intelligence Scale for Children,Performance Scale, 86 children (times: begin first, end first, third, fifth grades). Gender is W

NC Fem

North Carolina Achievement Data (see Williamson, Applebaum, Epanchin, 1991). These education data are eight yearly observations on achievement test scores in math (Y), for 277 females each followed from grade 1 to grade 8, with a verbal ability background measure (W) Each individual has a row of data; the first column contains the verbal ability score, which is used as the exogenous background measure, W. The multiple longitudinal observations follow: 8 waves of achievement test scores in math (grades 1-8).

	Observation time								
ID	1	2	3	4	5	6	7	8	W
1	380	377	460	472	495	566	637	628	120
2	362	382	392	475	475	543	601	576	95
3	387	405	438	418	484	533	570	589	99
4	342	368	408	422	470	543	493	589	101
5	335	372	450	424	500	510	540	583	109
6	362	444	473	482	567	597	651	655	115
7	354	409	410	445	460	540	567	620	115
8	365	381	455	482	533	554	591	602	109
9	359	371	438	452	497	591	573	593	107

Ramus

Ramus Data. 4 longitudinal observations on each of 20 cases. The measurement is the height of the mandibular ramus bone (in mm) for boys each measured at 8, 8.5, 9, 9.5 years of age. These data, used by a number of authors, can be found in Table 4.1 of Goldstein (1979).

ID	T-> 8.00	8.50	9.00	9.50
1	47.80	48.80	49.00	49.70
2	46.40	47.30	47.70	48.40
3	46.30	46.80	47.80	48.50
4	45.10	45.30	46.10	47.20
5	47.60	48.50	48.90	49.30
6	52.50	53.20	53.30	53.70
7	51.20	53.00	54.30	54.50
8	49.80	50.00	50.30	52.70
9	48.10	50.80	52.30	54.40
10	45.00	47.00	47.30	48.30
11	51.20	51.40	51.60	51.90
12	48.50	49.20	53.00	55.50
13	52.10	52.80	53.70	55.00
14	48.20	48.90	49.30	49.80
15	49.60	50.40	51.20	51.80
16	50.70	51.70	52.70	53.30
17	47.20	47.70	48.40	49.50
18	53.30	54.60	55.10	55.30
19	46.20	47.50	48.10	48.40
20	46.30	47.60	51.30	51.80

Smearmiss.

```
R version 2.14.1 (2011-12-22)
Copyright (C) 2011 The R Foundation for Statistical Computing
ISBN 3-900051-07-0
Platform: x86_64-pc-mingw32/x64 (64-bit)

R is free software and comes with ABSOLUTELY NO WARRANTY.
You are welcome to redistribute it under certain conditions.
Type 'license()' or 'licence()' for distribution details.

  Natural language support but running in an English locale

R is a collaborative project with many contributors.
Type 'contributors()' for more information and
'citation()' on how to cite R or R packages in publications.

Type 'demo()' for some demos, 'help()' for on-line help, or
'help.start()' for an HTML browser interface to help.
Type 'q()' to quit R.

> week3NC = read.table(file="http://www-stat.stanford.edu/~rag/stat222/nc8wide.dat", header = T)
> summary(week3NC) # this is the wide data
       ID                Y.1             Y.2             Y.3             Y.4            Y.5             Y.6
 Min.   :  705810   Min.   :270.0   Min.   :302.0   Min.   :343.0   Min.   :353   Min.   :348.0   Min.   :394.
 1st Qu.:  847813   1st Qu.:315.0   1st Qu.:364.0   1st Qu.:398.0   1st Qu.:426   1st Qu.:467.0   1st Qu.:491.
 Median : 1046817   Median :331.0   Median :384.0   Median :420.0   Median :452   Median :492.0   Median :526.
 Mean   : 1461655   Mean   :332.8   Mean   :384.1   Mean   :416.7   Mean   :454   Mean   :493.2   Mean   :525.
 3rd Qu.: 1290819   3rd Qu.:349.0   3rd Qu.:409.0   3rd Qu.:438.0   3rd Qu.:478   3rd Qu.:523.0   3rd Qu.:558.
 Max.   :11090821   Max.   :392.0   Max.   :456.0   Max.   :490.0   Max.   :569   Max.   :615.0   Max.   :662.
      Y.7             Y.8               Z
 Min.   :352.0   Min.   :382.0   Min.   : 64.0
 1st Qu.:517.0   1st Qu.:555.0   1st Qu.: 97.0
 Median :558.0   Median :593.0   Median :106.0
 Mean   :560.2   Mean   :592.1   Mean   :106.1
 3rd Qu.:601.0   3rd Qu.:628.0   3rd Qu.:115.0
 Max.   :724.0   Max.   :762.0   Max.   :145.0
> cor(week3NC)
              ID         Y.1         Y.2         Y.3         Y.4         Y.5        Y.6         Y.7          Y.
ID   1.000000000 -0.02805969 0.009225629 -0.03780944 -0.02499697 -0.03051227 -0.0306747 -0.00460692 -0.0142784
Y.1 -0.028059693  1.00000000 0.515791280  0.52307596  0.43523228  0.46388027  0.4455050  0.49203106  0.4466890
Y.2  0.009225629  0.51579128 1.000000000  0.72680549  0.73257228  0.70828888  0.6862661  0.73810725  0.6837425
Y.3 -0.037809439  0.52307596 0.726805489  1.00000000  0.78289485  0.83265946  0.8035561  0.78578095  0.7676681
Y.4 -0.024996972  0.43523228 0.732572280  0.78289485  1.00000000  0.83173846  0.8125849  0.82101419  0.8036211
Y.5 -0.030512274  0.46388027 0.708288880  0.83265946  0.83173846  1.00000000  0.8621471  0.83695336  0.8338277
Y.6 -0.030674705  0.44550498 0.686266086  0.80355612  0.81258495  0.86214713  1.0000000  0.87668055  0.8342427
Y.7 -0.004606920  0.49203106 0.738107252  0.78578095  0.82101419  0.83695336  0.8766805  1.00000000  0.8719083
Y.8 -0.014278455  0.44668909 0.683742594  0.76766816  0.80362111  0.83382777  0.8342427  0.87190838  1.0000000
Z   -0.067210733  0.42052594 0.524921660  0.67013710  0.61108627  0.66648797  0.6824815  0.69833019  0.6634602

> pairs(week3NC, pch = 20) # prob need to go to a "dot"
> dif = week3NC$Y.8 - week3NC$Y.1 # the observed amount of change, years 1-8
> cor(dif, week3NC$Y.1)
[1] 0.06579661
> cor(dif, week3NC$Z)
[1] 0.5581162
> cor.test(dif, week3NC$Y.1) # not significant correlation obs change, obs initial status

        Pearson's product-moment correlation

data:  dif and week3NC$Y.1
t = 1.0935, df = 275, p-value = 0.2751
alternative hypothesis: true correlation is not equal to 0
95 percent confidence interval:
 -0.05246573  0.18223893
sample estimates:
       cor
0.06579661

> week3NC = read.table(file="http://www-stat.stanford.edu/~rag/stat222/ncLong.dat", header = T)
> # oops, I overwrote the wide file (I may not need it again) for the session
> #plots
> library(lattice)
> head(week3NC)
      ID time   Y   Z
1 705810    1 380 120
```

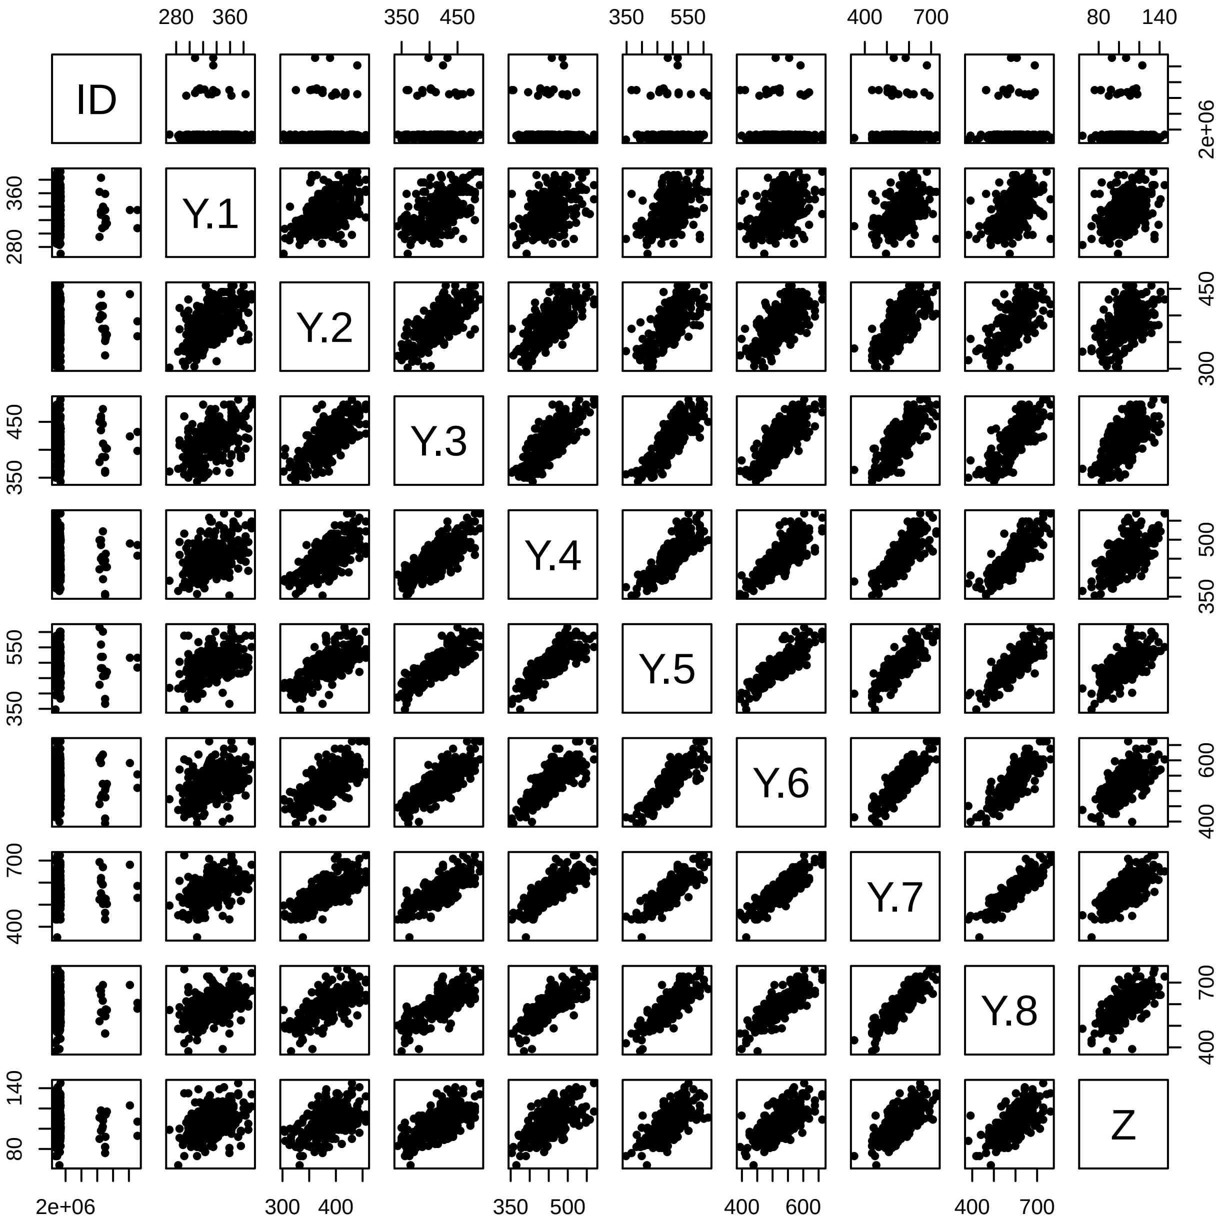

ID
Y.1
Y.2
Y.3
Y.4
Y.5
Y.6
Y.7
Y.8
Z
280
360
350
450
350
550
400
700
80
140
2e+06
300
400
350
500
400
600

```
R version 3.4.4 (2018-03-15) -- "Someone to Lean On"   NC WEEK 2

> library(lme4)
Loading required package: Matrix
> longNC = read.table(file="http://rogosateaching.com/stat222/ncLong_data", h

>  longNC$timeInt = longNC$time -1   #define initial status
>  head(longNC)
      ID time   Y   Z timeInt
1 705810    1 380 120       0
2 705810    2 377 120       1
3 705810    3 460 120       2
4 705810    4 472 120       3
5 705810    5 495 120       4
6 705810    6 566 120       5
> attach(longNC)
> ncList = lmList(Y ~ timeInt |ID, data = longNC) #fit straight-line to each
> rate = coef(ncList)[2]
> stem(rate[,1]) #stem-and-leaf of rates-of-change

  The decimal point is 1 digit(s) to the right of the |

  0 |
  1 | 034
  1 | 89
  2 | 012224
  2 | 5566666677777788888899999999999999
  3 | 00000000111111111111112222222223333333333344444444444444
  3 | 55555555555555666666666666666666666666777777777777777777778888888888+8
  4 | 00000000000001111111111122222222222222333333333333444444
  4 | 555566667778889999999
  5 | 0000001334
  5 | 66
  6 | 4

> fivenum(rate[,1]) # rates-of-change
[1]  9.714286 31.559524 36.416667 41.023810 64.238095
> mean(rate[,1]) # rates-of-change
[1] 36.44808
> boxplot(rate[,1]) # rates-of-change
> Zt = tapply(longNC$Z, longNC$ID, mean) #get back the 277 Z-values
> cor(Zt,rate[,1])
[1] 0.6237752
> # strong relation, see it in the scatterplot, see the Conditional (using Z)
> plot(Zt,rate[,1], pch=20)
>
```

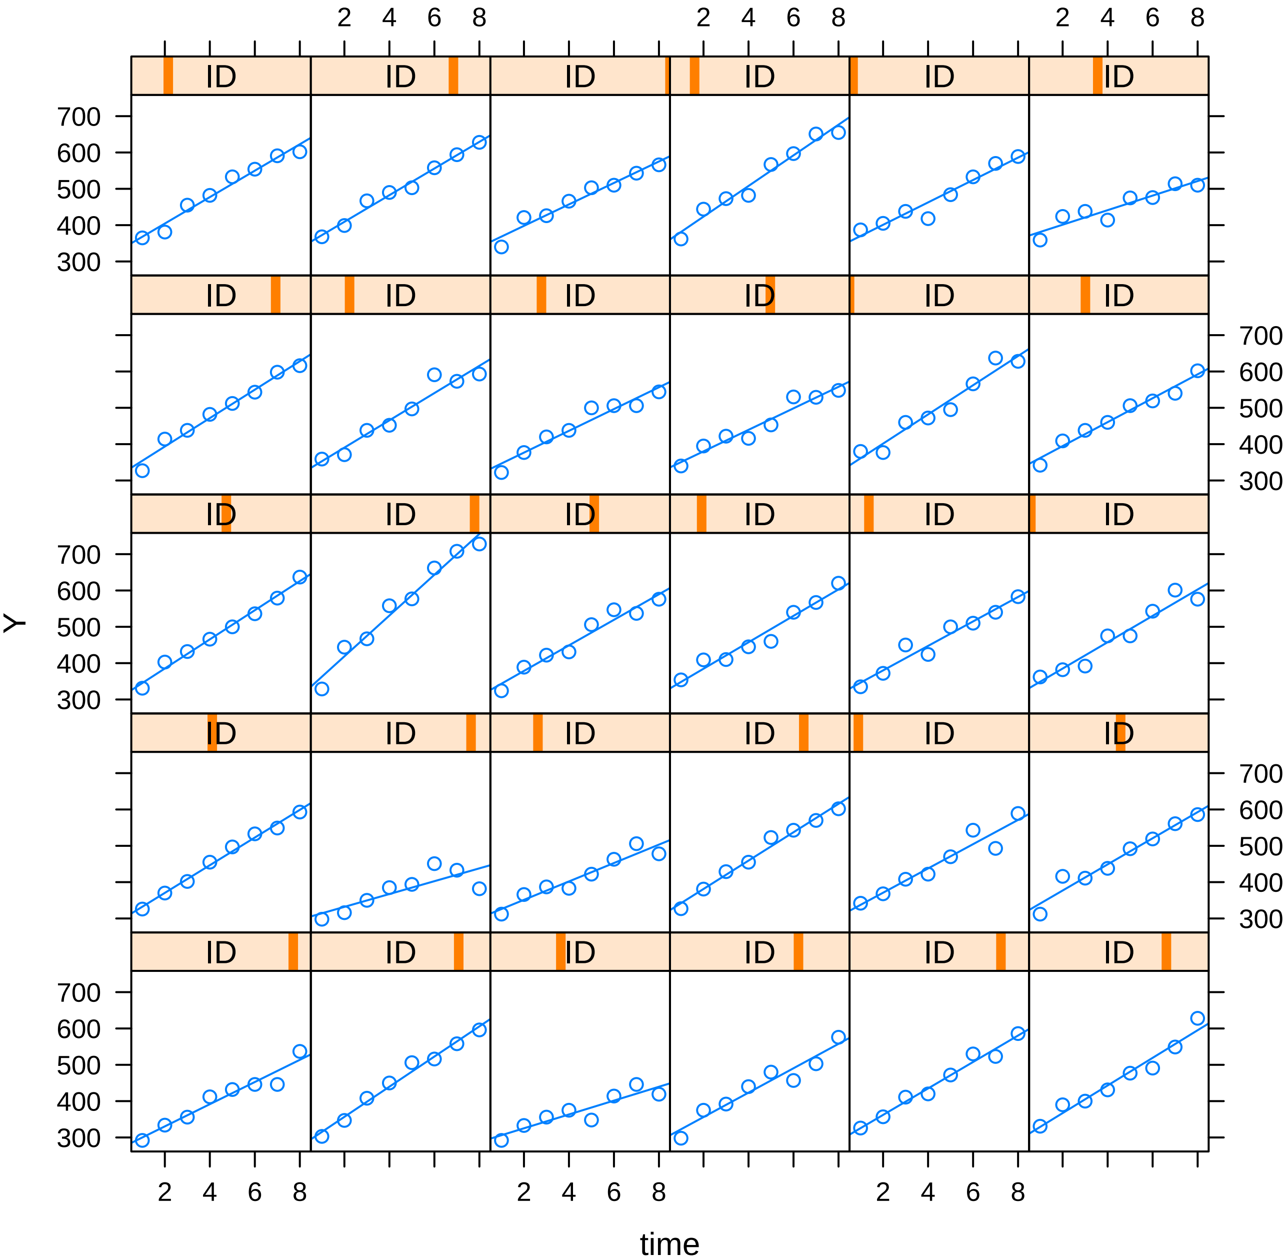

ID
Y
time
300
400
500
600
700
2
4
6
8

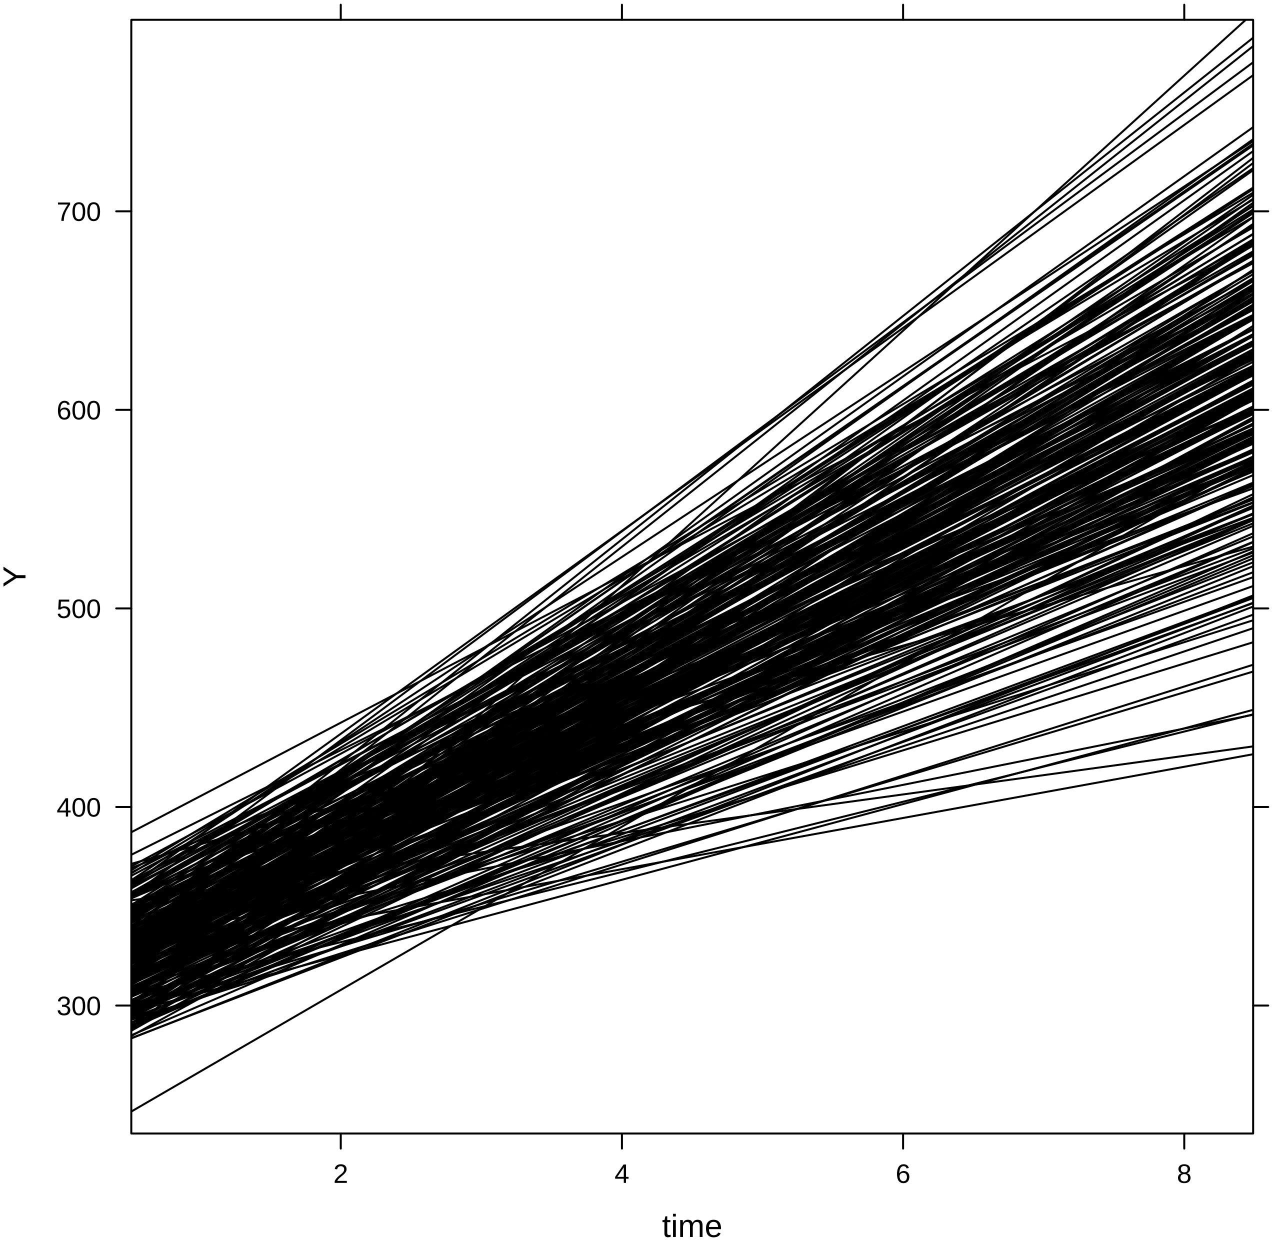
Y
700
600
500
400
300
2
4
6
8
time

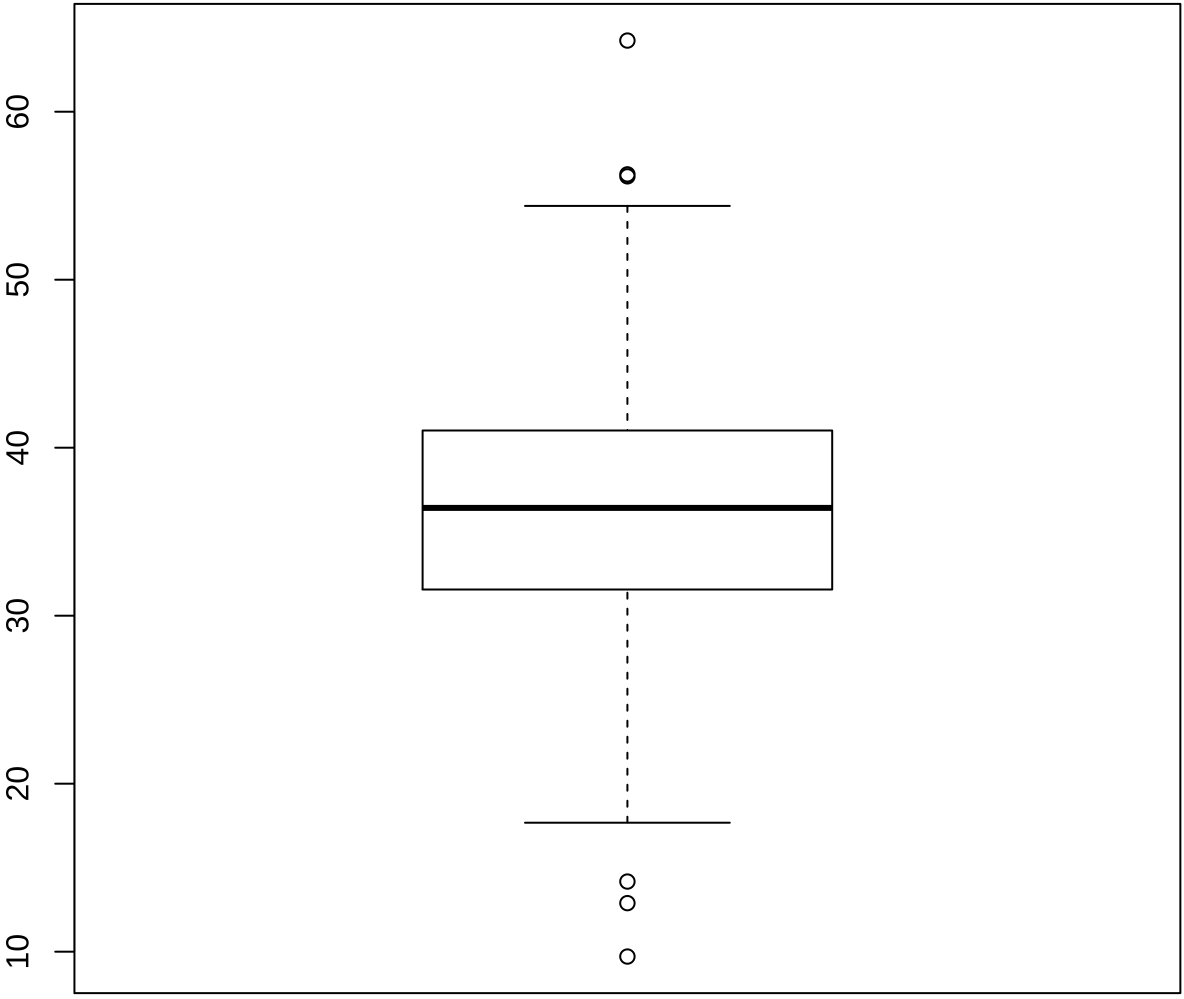
10
20
30
40
50
60

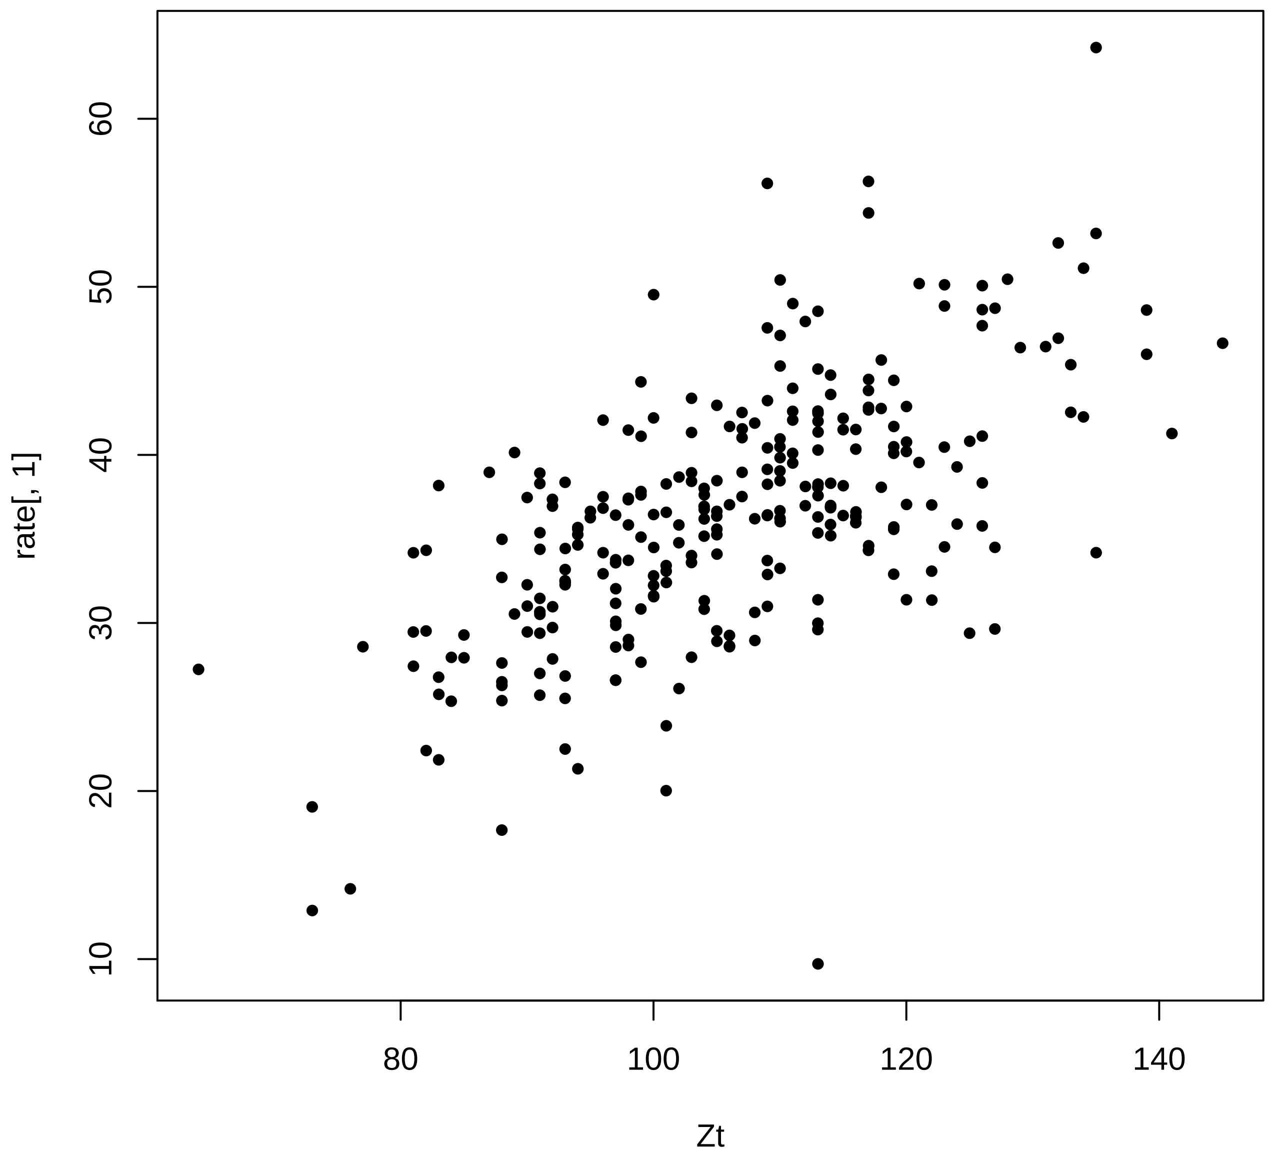

60
50
40
30
20
10
rate[, 1]
80
100
120
140
Zt

Analyses of Collections of Growth Curves

Stat 222 Week 2

Descriptive Analyses (smart first year student)

SFYS Fit Y on t regressions (lmList)

describe $\hat{\alpha}_0$ $\hat{\alpha}_1$ plots etc $\hat{\alpha}_1$ vs $\hat{\alpha}_0$ initial status, growth

systematic indiv differences (Z) in change

cor $(\hat{\alpha}_1, Z)$ plots $\hat{\alpha}_1$ vs Z etc

Mixed effects models (lmer)

A. Unconditional

sleep N/C

Level 1 $Y = \alpha_0 + \alpha_1 t + \varepsilon$

level at $t=0$ (define init status)

gradient

straight-line growth each unit

Level 2 $\alpha_0 = \gamma_{00} + u_0$

$\alpha_1 = \gamma_{10} + u_1$

"null"

Combined (estimation) model $Y = \gamma_{00} + \gamma_{10} t + [\varepsilon + u_0 + u_1]$

fixed random

B. Conditional Model

exog var Z, W

systematic indiv diffs in growth

Level 1 $Y = \alpha_0 + \alpha_1 t + \varepsilon$

NC ncCon2 (full)

$$\boxed{\begin{array}{l}\text{Level 2}\\ \alpha_0 = \gamma_{00} + \gamma_{01} Z + u_0\\ \alpha_1 = \gamma_{10} + \gamma_{11} Z + u_1\end{array}}$$

diff random terms

Combined model $Y = \gamma_{00} + \gamma_{01} Z + \gamma_{10} t + \gamma_{11} Z \cdot t + [\varepsilon + u_0 + u_1]$

model form $Y \sim Z * t$

see note on redundant time Int term in NC hand lmer

North Carolina Data   STAT 222 Week 2

```
> #NC comparing models
> week3NC = read.table(file="http://www-stat.stanford.edu/~rag/stat222/ncLong.dat", header = T)
```
long form

```
> week3NC$timeInt = week3NC$time -1
```
make initial time = 0 (implicit intercept)

```
> head(week3NC)
      ID time   Y   Z timeInt
1 705810    1 380 120       0
2 705810    2 377 120       1
3 705810    3 460 120       2
4 705810    4 472 120       3
5 705810    5 495 120       4
6 705810    6 566 120       5
> summary(week3NC)
       ID                time             Y               Z             timeInt
 Min.   :  705810   Min.   :1.00   Min.   :270.0   Min.   : 64.0   Min.   :0.00
 1st Qu.:  847813   1st Qu.:2.75   1st Qu.:395.0   1st Qu.: 97.0   1st Qu.:1.75
 Median : 1046817   Median :4.50   Median :464.0   Median :106.0   Median :3.50
 Mean   : 1461655   Mean   :4.50   Mean   :469.9   Mean   :106.1   Mean   :3.50
 3rd Qu.: 1290819   3rd Qu.:6.25   3rd Qu.:540.0   3rd Qu.:115.0   3rd Qu.:5.25
 Max.   :11090821   Max.   :8.00   Max.   :762.0   Max.   :145.0   Max.   :7.00
```
Descriptives in posted version

```
> attach(week3NC)
> xtabs(~ ID + timeInt, week3NC) # balanced, complete data, Bates does this way
```
see plots

```
# week 2 [redacted] descriptives, SFYS analyses; here revisit the mixed-models
> library(lme4)
> ncUnc = lmer(Y ~ timeInt + ( 1 + timeInt   | ID), data = week3NC)
> summary(ncUnc)
```
balanced data

no Level 2

```
Linear mixed model fit by REML
Formula: Y ~ timeInt + (1 + timeInt | ID)
   Data: week3NC
   AIC   BIC logLik deviance REMLdev
 20690 20724 -10339    20680   20678
Random effects:
 Groups   Name        Variance Std.Dev. Corr
 ID       (Intercept) 326.059  18.0571
          timeInt      46.229   6.7992  0.651
 Residual             403.487  20.0870
Number of obs: 2216, groups: ID, 277
> # best to transform time to be zero at the time point of interest
#####                  (cor(rate,initial status) correct, .651)
Fixed effects:
            Estimate Std. Error t value
(Intercept)  342.300      1.336  256.27
timeInt       36.448      0.449   81.18

Correlation of Fixed Effects:
        (Intr)
timeInt 0.279
```
$\hat{var}(\theta)$ 46.229 = 55.836 − msr/sst see below

gets this correct mle of cor(θ, π(initial time)) rate

match SFYS

$r_{Int, rate}$ from lmList

```
> ncCon = lmer(Y ~ timeInt + Z:timeInt + ( 1 + timeInt   | ID), data = week3NC) # incl Z in slope L2
>  summary(ncCon)
Linear mixed model fit by REML
Formula: Y ~ timeInt + Z:timeInt + (1 + timeInt | ID)
   Data: week3NC
   AIC   BIC logLik deviance REMLdev
 20582 20622 -10284    20565   20568
Random effects:
 Groups   Name        Variance Std.Dev. Corr
 ID       (Intercept) 326.062  18.0572
          timeInt      24.683   4.9681  0.262
 Residual             403.486  20.0870
Number of obs: 2216, groups: ID, 277

Fixed effects:
             Estimate Std. Error t value
(Intercept) 342.29994    1.33568  256.27
timeInt      -0.95502    2.70714   -0.35
timeInt:Z     0.35266    0.02531   13.93

Correlation of Fixed Effects:
          (Intr) timInt
timeInt   -0.010
timeInt:Z  0.000 -0.992
```

```
> ncList = lmList(Y ~ timeInt |ID, data = week3NC) #fit
>  rate = coef(ncList)[2]
> var(rate)
         timeInt
timeInt 55.83613
> ncList
Degrees of freedom: 2216 total; 1662 residual
Residual standard error: 20.08697
> sst = 2*(.5^2 + 1.5^2 + 2.5^2 +3.5^2)
> sst
[1] 42
> 55.836 - (20.087)^2/42
[1] 46.22915
```

Level 2 $\alpha_0 = \gamma_{00} + u_0$
$\alpha_1 = \gamma_{10} + \gamma_{11} Z + u_1$

see rate | Z plot

page 2 week 2 STAT222 North Carolina

```
> ncCon2 = lmer(Y ~ timeInt + Z*timeInt + ( 1 + timeInt  | ID), data = week3NC) # incl Z in level,slope L2
>  summary(ncCon2)
Linear mixed model fit by REML
Formula: Y ~ timeInt + Z * timeInt + (1 + timeInt | ID)
   Data: week3NC
   AIC   BIC logLik deviance REMLdev
 20501 20547 -10243    20478   20485
Random effects:
 Groups   Name        Variance Std.Dev. Corr
 ID       (Intercept) 194.731  13.9546
          timeInt      24.628   4.9626  0.379
 Residual             403.486  20.0870
Number of obs: 2216, groups: ID, 277

Fixed effects:
             Estimate Std. Error t value
(Intercept) 254.32454    8.83286  28.793
timeInt       0.84389    2.71301   0.311
Z             0.82948    0.08258  10.045
timeInt:Z     0.33569    0.02536  13.235

Correlation of Fixed Effects:
          (Intr) timInt Z
timeInt   -0.067
Z         -0.992  0.066
timeInt:Z  0.066 -0.992 -0.067

> confint(ncUnc)  #add-ons needed to do this
Error: $ operator not defined for this S4 class
```

new version works

```
> anova(ncUnc, ncCon, ncCon2) #formal model comparisons, nested trio
Data: week3NC
Models:
ncUnc: Y ~ timeInt + (1 + timeInt | ID)
ncCon: Y ~ timeInt + Z:timeInt + (1 + timeInt | ID)
ncCon2: Y ~ timeInt + Z * timeInt + (1 + timeInt | ID)
       Df   AIC   BIC logLik  Chisq Chi Df Pr(>Chisq)
ncUnc   6 20692 20727 -10340
ncCon   7 20579 20619 -10282 115.33      1  < 2.2e-16 ***
ncCon2  8 20494 20540 -10239  86.57      1  < 2.2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
> # ncCon2 seems a winner

> anova(ncCon, ncCon2) # just to show trio works as you would hope
Data: week3NC
Models:
ncCon: Y ~ timeInt + Z:timeInt + (1 + timeInt | ID)
ncCon2: Y ~ timeInt + Z * timeInt + (1 + timeInt | ID)
       Df   AIC   BIC logLik Chisq Chi Df Pr(>Chisq)
ncCon   7 20579 20619 -10282
ncCon2  8 20494 20540 -10239 86.57      1  < 2.2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
>
```

Level 1

$$Y = \alpha_0 + \alpha_1 t + \varepsilon$$

Level 2

$$\alpha_0 = \gamma_{00} + \gamma_{01} Z + u_0$$

$$\alpha_1 = \gamma_{10} + \gamma_{11} Z + u_1$$

Z w/ slope vs Z w/ int and slope

```
> #NC comparing models
> week3NC = read.table(file="http://www-stat.stanford.edu/~rag/stat222/ncLong.dat", header = T)

> week3NC$timeInt = week3NC$time -1
> head(week3NC)
      ID time   Y   Z timeInt
1 705810    1 380 120       0
2 705810    2 377 120       1
3 705810    3 460 120       2
4 705810    4 472 120       3
5 705810    5 495 120       4
6 705810    6 566 120       5
> summary(week3NC)
       ID                time             Y                 Z             timeInt
 Min.    :  705810   Min.    :1.00   Min.    :270.0   Min.    : 64.0   Min.    :0.00
 1st Qu.:  847813   1st Qu.:2.75   1st Qu.:395.0   1st Qu.: 97.0   1st Qu.:1.75
 Median : 1046817   Median :4.50   Median :464.0   Median :106.0   Median :3.50
 Mean   : 1461655   Mean    :4.50   Mean    :469.9   Mean    :106.1   Mean    :3.50
 3rd Qu.: 1290819   3rd Qu.:6.25   3rd Qu.:540.0   3rd Qu.:115.0   3rd Qu.:5.25
 Max.   :11090821   Max.    :8.00   Max.    :762.0   Max.    :145.0   Max.    :7.00

> attach(week3NC)
> xtabs(~ ID + timeInt, week3NC) # balanced, complete data, Bates does this way

# week 2 we did descripitives, SFYS analyses; here revisit the mixed-models
> library(lme4)
> ncUnc = lmer(Y ~ timeInt + ( 1 + timeInt  | ID), data = week3NC)
> summary(ncUnc)
Linear mixed model fit by REML
Formula: Y ~ timeInt + (1 + timeInt | ID)
   Data: week3NC
   AIC   BIC logLik deviance REMLdev
 20690 20724 -10339    20680   20678
Random effects:
 Groups   Name        Variance Std.Dev. Corr
 ID       (Intercept) 326.059  18.0571
          timeInt      46.229   6.7992  0.651
 Residual             403.487  20.0870
Number of obs: 2216, groups: ID, 277
> # best to transform time to be zero at the time point of interest
#####                (cor(rate,initial status) correct, .651)
Fixed effects:
            Estimate Std. Error t value
(Intercept)  342.300      1.336  256.27
timeInt       36.448      0.449   81.18

Correlation of Fixed Effects:
        (Intr)
timeInt 0.279

> ncCon = lmer(Y ~ timeInt + Z:timeInt + ( 1 + timeInt  | ID), data = week3NC) # incl Z in slope L2
>  summary(ncCon)
Linear mixed model fit by REML
Formula: Y ~ timeInt + Z:timeInt + (1 + timeInt | ID)
   Data: week3NC
   AIC   BIC logLik deviance REMLdev
 20582 20622 -10284    20565   20568
Random effects:
 Groups   Name        Variance Std.Dev. Corr
 ID       (Intercept) 326.062  18.0572
          timeInt      24.683   4.9681  0.262
 Residual             403.486  20.0870
Number of obs: 2216, groups: ID, 277

Fixed effects:
             Estimate Std. Error t value
(Intercept) 342.29994    1.33568  256.27
timeInt      -0.95502    2.70714   -0.35
timeInt:Z     0.35266    0.02531   13.93

Correlation of Fixed Effects:
          (Intr) timInt
timeInt   -0.010
timeInt:Z  0.000 -0.992
```

```
> ncCon2 = lmer(Y ~ timeInt + Z*timeInt + ( 1 + timeInt  | ID), data = week3NC) # incl Z in level,slope L2
>  summary(ncCon2)
Linear mixed model fit by REML
Formula: Y ~ timeInt + Z * timeInt + (1 + timeInt | ID)
   Data: week3NC
   AIC   BIC logLik deviance REMLdev
 20501 20547 -10243    20478   20485
Random effects:
 Groups   Name        Variance Std.Dev. Corr
 ID       (Intercept) 194.731  13.9546
          timeInt      24.628   4.9626  0.379
 Residual             403.486  20.0870
Number of obs: 2216, groups: ID, 277

Fixed effects:
             Estimate Std. Error t value
(Intercept) 254.32454    8.83286  28.793
timeInt       0.84389    2.71301   0.311
Z             0.82948    0.08258  10.045
timeInt:Z     0.33569    0.02536  13.235

Correlation of Fixed Effects:
          (Intr) timInt Z
timeInt   -0.067
Z         -0.992  0.066
timeInt:Z  0.066 -0.992 -0.067

> confint(ncUnc)  #add-ons needed to do this
Error: $ operator not defined for this S4 class

> anova(ncUnc, ncCon, ncCon2) #formal model comparisons, nested trio
Data: week3NC
Models:
ncUnc: Y ~ timeInt + (1 + timeInt | ID)
ncCon: Y ~ timeInt + Z:timeInt + (1 + timeInt | ID)
ncCon2: Y ~ timeInt + Z * timeInt + (1 + timeInt | ID)
       Df   AIC   BIC logLik  Chisq Chi Df Pr(>Chisq)
ncUnc   6 20692 20727 -10340
ncCon   7 20579 20619 -10282 115.33      1  < 2.2e-16 ***
ncCon2  8 20494 20540 -10239  86.57      1  < 2.2e-16 ***
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
> # ncCon2 seems a winner

> anova(ncCon, ncCon2) # just to show trio works as you would hope
Data: week3NC
Models:
ncCon: Y ~ timeInt + Z:timeInt + (1 + timeInt | ID)
ncCon2: Y ~ timeInt + Z * timeInt + (1 + timeInt | ID)
       Df   AIC   BIC logLik Chisq Chi Df Pr(>Chisq)
ncCon   7 20579 20619 -10282
ncCon2  8 20494 20540 -10239 86.57      1  < 2.2e-16 ***
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
>
===========================
page1 method of moments insert
> ncList = lmList(Y ~ timeInt |ID, data = week3NC) #fit straight-line to each subject
>  rate = coef(ncList)[2]
> var(rate)
         timeInt
timeInt 55.83613
> ncList
Degrees of freedom: 2216 total; 1662 residual
Residual standard error: 20.08697
> sst = 2*(.5^2 + 1.5^2 + 2.5^2 +3.5^2)
> sst
[1] 42
> 55.836 - (20.087)^2/42
[1] 46.22915
```

Do confint in a review Q

```
##ncCon2 without redundent model term

> week3NC = read.table(file="http://statweb.stanford.edu/~rag/stat2
> attach(week3NC)
> detach(week3NC)
> week3NC$timeInt = week3NC$time -1
> attach(week3NC)
> ncCon2 = lmer(Y ~  Z*timeInt + ( 1 + timeInt  | ID), data = week3
> summary(ncCon2)
Linear mixed model fit by REML ['lmerMod']
Formula: Y ~ Z * timeInt + (1 + timeInt | ID)
   Data: week3NC

REML criterion at convergence: 20485.2

Scaled residuals:
    Min      1Q  Median      3Q     Max
-3.5617 -0.6065 -0.0352  0.5934  3.1665

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 ID       (Intercept) 194.73   13.955
          timeInt      24.63    4.963   0.38
 Residual             403.49   20.087
Number of obs: 2216, groups:  ID, 277

Fixed effects:
             Estimate Std. Error t value
(Intercept) 254.32454    8.83285  28.793
Z             0.82948    0.08258  10.045
timeInt       0.84389    2.71313   0.311
Z:timeInt     0.33569    0.02537  13.235

Correlation of Fixed Effects:
          (Intr) Z      timInt
Z         -0.992
timeInt   -0.066  0.066
Z:timeInt  0.066 -0.066 -0.992
```

```
##Stat222, week2
>
> #NC residual plots
>  week3NC = read.table(file="http://rogosateaching.com/stat222/ncLong_data", header = T)
> week3NC$timeInt = week3NC$time -1
> ncUnc = lmer(Y ~ timeInt + ( 1 + timeInt  | ID), data = week3NC)
> summary(ncUnc)
Linear mixed model fit by REML ['lmerMod']
Formula: Y ~ timeInt + (1 + timeInt | ID)   Data: week3NC
REML criterion at convergence: 20677.8
Scaled residuals:
    Min      1Q  Median      3Q     Max
-3.6596 -0.6056 -0.0257  0.5941  3.2018

Random effects:
 Groups   Name        Variance Std.Dev. Corr  ## elements of Cov(alpha_0, alpha_1)
 ID       (Intercept) 326.06   18.057
          timeInt      46.23    6.799   0.65  ## .65 is mle estimate of Cor(alpha_0, alpha_1)
 Residual             403.49   20.087
Number of obs: 2216, groups:  ID, 277

Fixed effects:
            Estimate Std. Error t value
(Intercept)  342.300      1.336  256.27
timeInt       36.448      0.449   81.18

Correlation of Fixed Effects:
        (Intr)
timeInt 0.279          ## .28 is sample value of Cor(alpha_hat_0, alpha_hat_1)

> plot(ncUnc, id = .01)   #see plot

> ncCon2 = lmer(Y ~  Z*timeInt + ( 1 + timeInt  | ID), data = week3NC) # incl Z in level,slope L2
> summary(ncCon2)
Linear mixed model fit by REML ['lmerMod']
Formula: Y ~  Z * timeInt + (1 + timeInt | ID)   Data: week3NC
REML criterion at convergence: 20485.2
```

```
Scaled residuals:
    Min      1Q  Median      3Q     Max 
-3.5617 -0.6065 -0.0352  0.5934  3.1665 

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 ID       (Intercept) 194.73   13.955       
          timeInt      24.63    4.963   0.38
 Residual             403.49   20.087       
Number of obs: 2216, groups:  ID, 277

Fixed effects:
             Estimate Std. Error t value
(Intercept) 254.32454    8.83286  28.793
timeInt       0.84389    2.71313   0.311
Z             0.82948    0.08258  10.045
timeInt:Z     0.33569    0.02537  13.235

Correlation of Fixed Effects:
          (Intr) timInt Z     
timeInt   -0.066              
Z         -0.992  0.066       
timeInt:Z  0.066 -0.992 -0.066
> confint(ncCon2)
Computing profile confidence intervals ...
                  2.5 %      97.5 %
.sig01       11.6988560  16.1328515
.sig02        0.1691589   0.6087485
.sig03        4.3873745   5.5462211
.sigma       19.4229852  20.7896705
(Intercept) 237.0150973 271.6339817
timeInt      -4.4729277   6.1607130
Z             0.6676501   0.9913026
timeInt:Z     0.2859872   0.3854013
> plot(ncCon2, id = .01)   # see plot 2
```

```
##Stat222, week2
>
> #NC residual plots
>  week3NC = read.table(file="http://rogosateaching.com/stat222/ncLong_data", header = T)
> week3NC$timeInt = week3NC$time -1
> ncUnc = lmer(Y ~ timeInt + ( 1 + timeInt  | ID), data = week3NC)
> summary(ncUnc)
Linear mixed model fit by REML ['lmerMod']
Formula: Y ~ timeInt + (1 + timeInt | ID)   Data: week3NC
REML criterion at convergence: 20677.8
Scaled residuals:
    Min      1Q  Median      3Q     Max
-3.6596 -0.6056 -0.0257  0.5941  3.2018

Random effects:
 Groups   Name        Variance Std.Dev. Corr  ## elements of Cov(alpha_0, alpha_1)
 ID       (Intercept) 326.06   18.057
          timeInt      46.23    6.799   0.65  ## .65 is mle estimate of Cor(alpha_0, alpha_1)
 Residual             403.49   20.087
Number of obs: 2216, groups:  ID, 277

Fixed effects:
            Estimate Std. Error t value
(Intercept)  342.300      1.336  256.27
timeInt       36.448      0.449   81.18

Correlation of Fixed Effects:
        (Intr)
timeInt 0.279          ## .28 is sample value of Cor(alpha_hat_0, alpha_hat_1)

> plot(ncUnc, id = .01)   #see plot

> ncCon2 = lmer(Y ~  Z*timeInt + ( 1 + timeInt  | ID), data = week3NC) # incl Z in level,slope L2
> summary(ncCon2)
Linear mixed model fit by REML ['lmerMod']
Formula: Y ~  Z * timeInt + (1 + timeInt | ID)   Data: week3NC
REML criterion at convergence: 20485.2
```

```
Scaled residuals:
    Min      1Q  Median      3Q     Max
-3.5617 -0.6065 -0.0352  0.5934  3.1665

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 ID       (Intercept) 194.73   13.955
          timeInt      24.63    4.963   0.38
 Residual             403.49   20.087
Number of obs: 2216, groups:  ID, 277

Fixed effects:
             Estimate Std. Error t value
(Intercept) 254.32454    8.83286  28.793
timeInt       0.84389    2.71313   0.311
Z             0.82948    0.08258  10.045
timeInt:Z     0.33569    0.02537  13.235

Correlation of Fixed Effects:
          (Intr) timInt Z
timeInt   -0.066
Z         -0.992  0.066
timeInt:Z  0.066 -0.992 -0.066
> confint(ncCon2)
Computing profile confidence intervals ...
                  2.5 %      97.5 %
.sig01       11.6988560  16.1328515
.sig02        0.1691589   0.6087485
.sig03        4.3873745   5.5462211
.sigma       19.4229852  20.7896705
(Intercept) 237.0150973 271.6339817
timeInt      -4.4729277   6.1607130
Z             0.6676501   0.9913026
timeInt:Z     0.2859872   0.3854013
> plot(ncCon2, id = .01)   # see plot 2
```

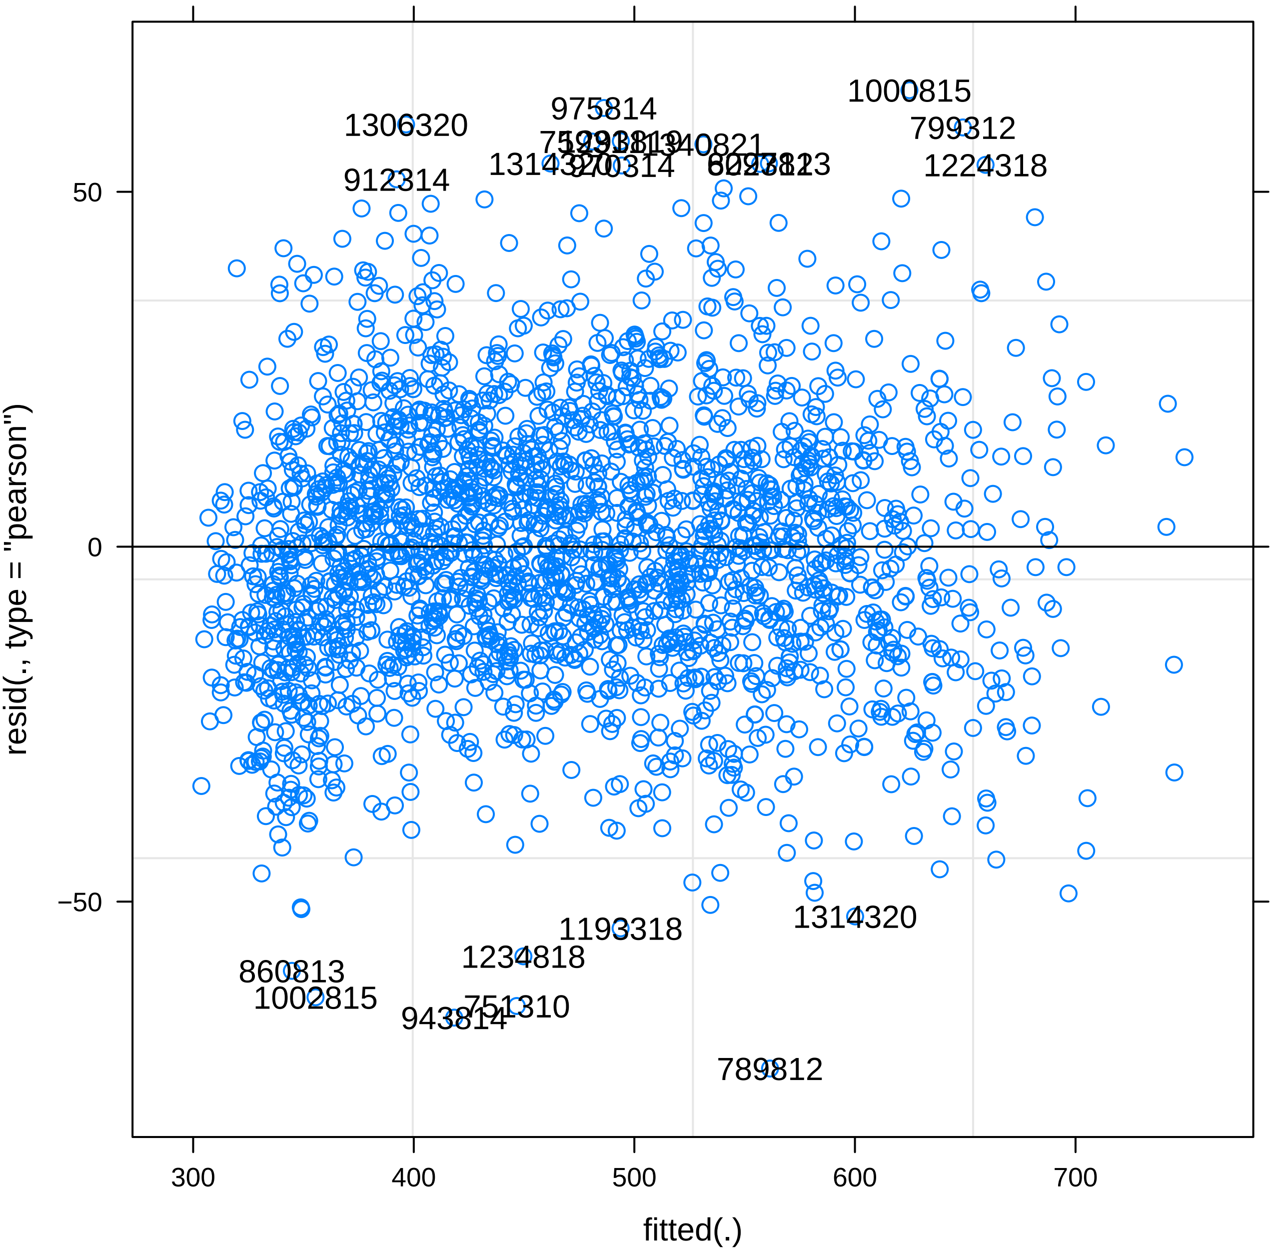

resid(., type = "pearson")
fitted(.)
300
400
500
600
700
50
0
−50
1000815
975814
1306320
799312
1224318
912314
1314320
1193318
1234818
860813
1002815
751310
943814
789812

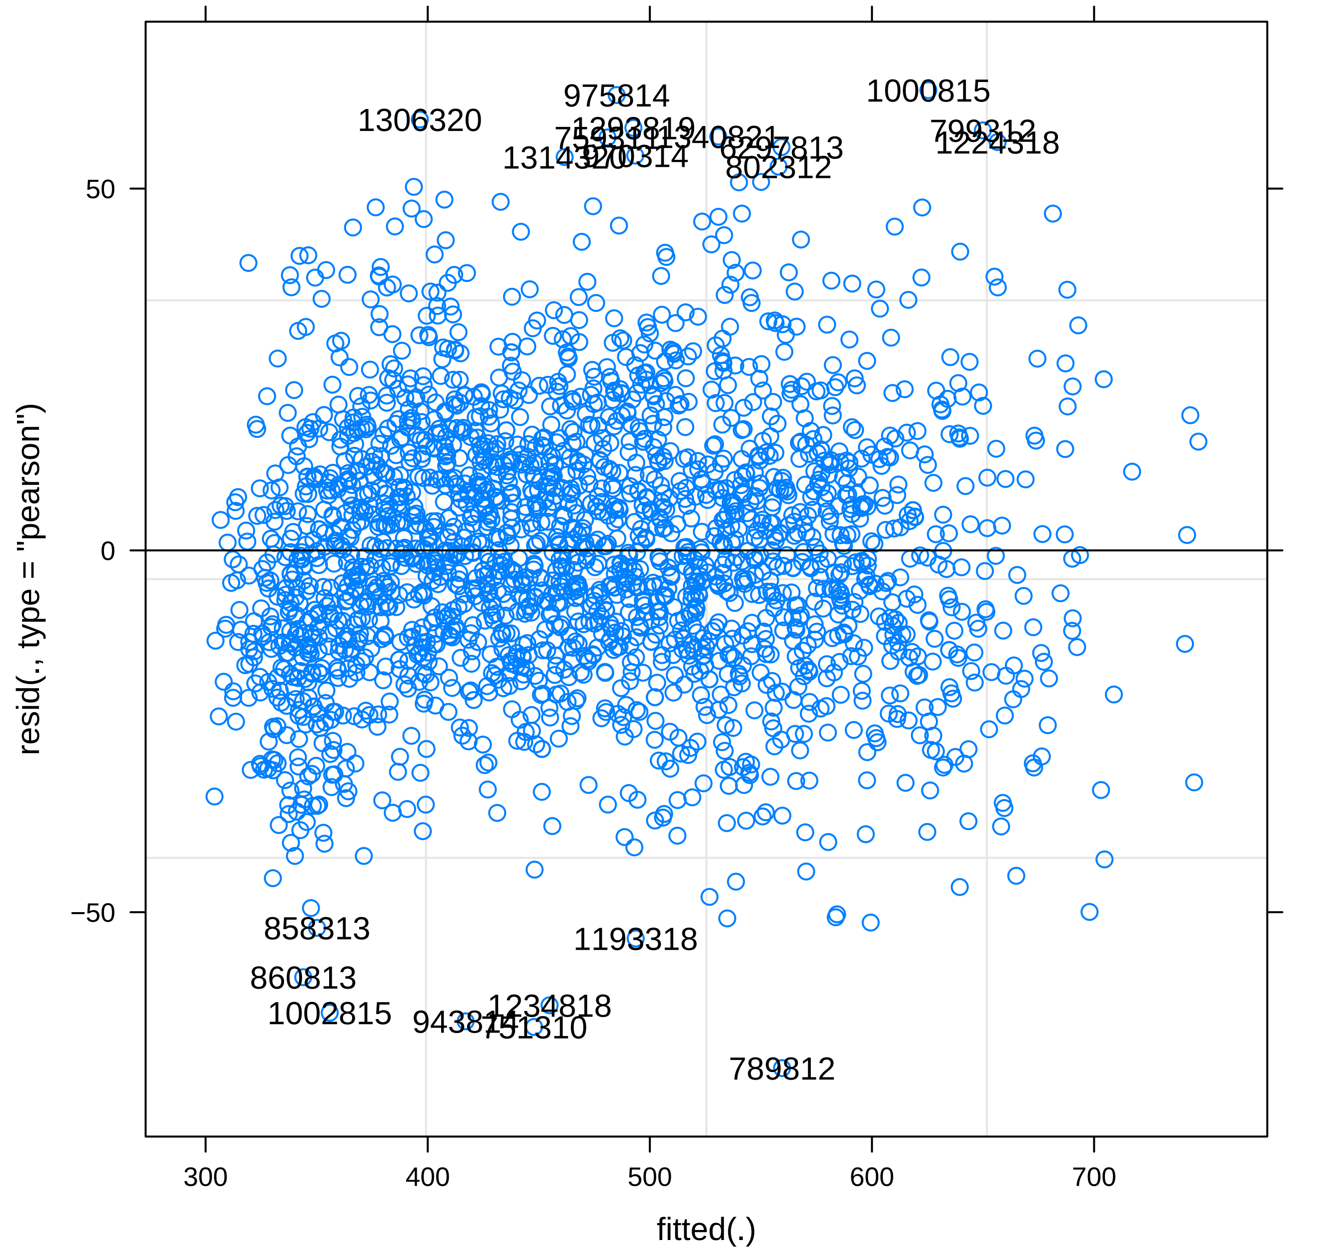
975814
1000815
1306320
1293819
799312
1224318
1314320
802312
50
0
−50
858313
1193318
860813
1002815
1234818
943814
751310
789812
300
400
500
600
700
fitted(.)
resid(., type = "pearson")

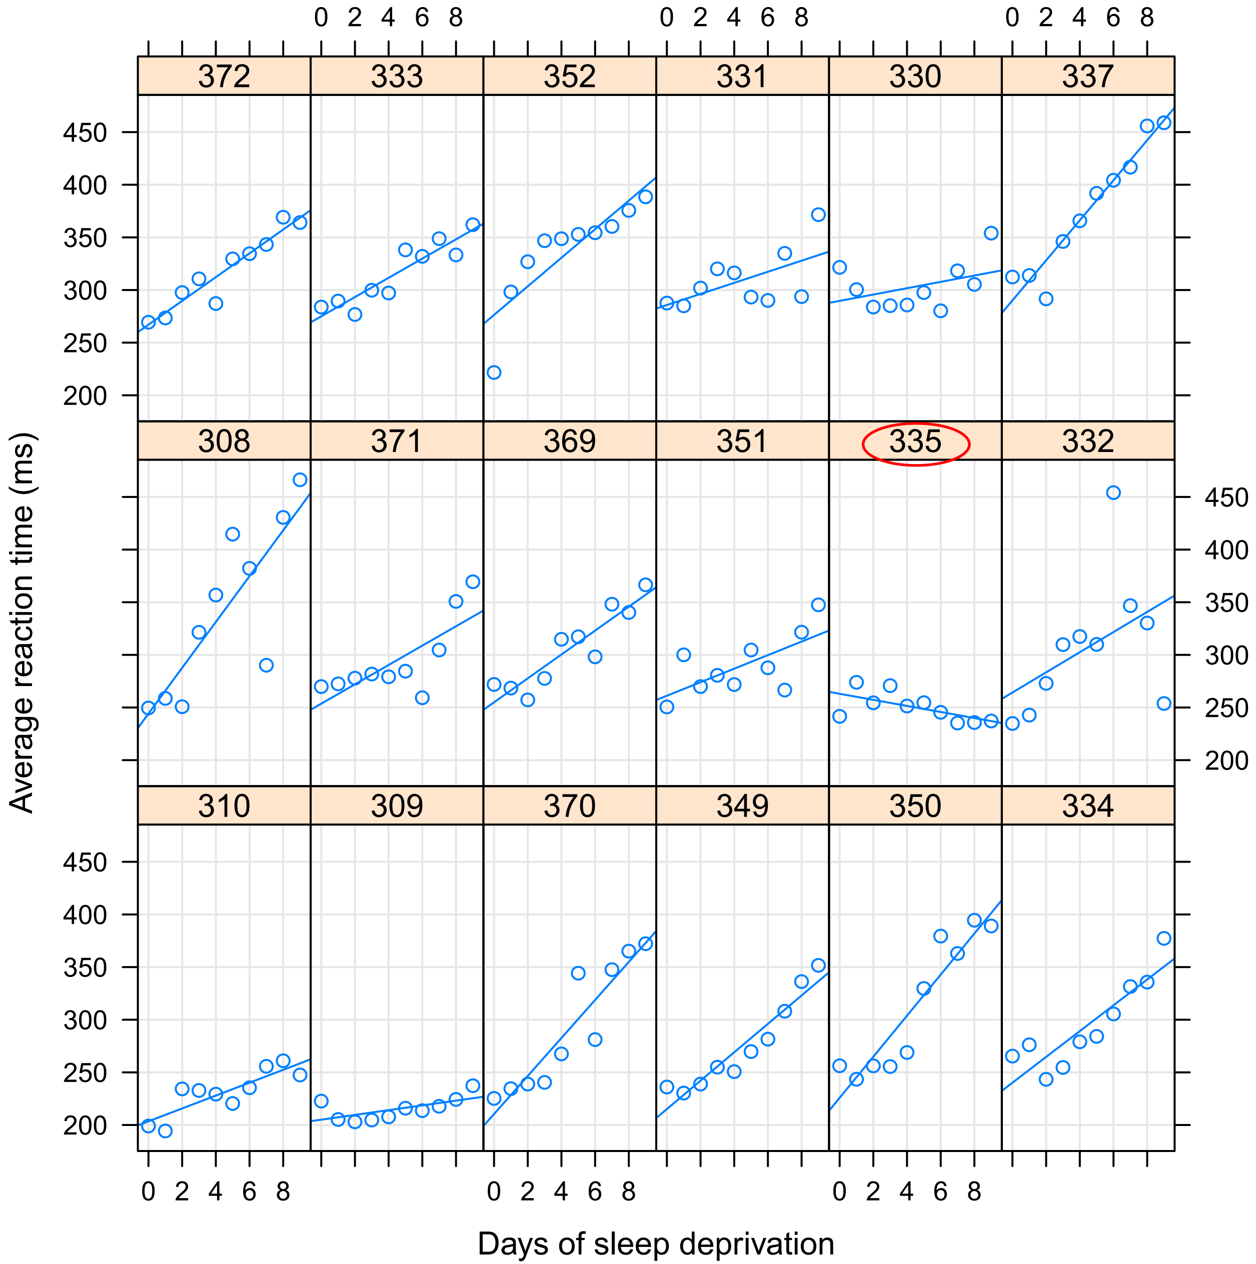

372
333
352
331
330
337
308
371
369
351
335
332
310
309
370
349
350
334
Average reaction time (ms)
Days of sleep deprivation

STAT 222 WEEK 3

```
R version 3.2.2 (2015-08-14) -- "Fire Safety"
Copyright (C) 2015 The R Foundation for Statistical Computing
Platform: x86_64-w64-mingw32/x64 (64-bit)

> library(lme4)
Loading required package: Matrix
> library(lattice)
> setwd("D:\\drr15\\ed401D\\week3\\")
> sleeplmList = lmList(Reaction ~ Days |Subject, data = sleepstudy)
> sleeplmList  # OLS each subject separately (individual 'fixed')
Call: lmList(formula = Reaction ~ Days | Subject, data = sleepstudy
Coefficients:
    (Intercept)      Days
308    244.1927 21.764702
309    205.0549  2.261785
310    203.4842  6.114899
330    289.6851  3.008073
331    285.7390  5.266019
332    264.2516  9.566768
333    275.0191  9.142045
334    240.1629 12.253141
335    263.0347 -2.881034
337    290.1041 19.025974
349    215.1118 13.493933
350    225.8346 19.504017
351    261.1470  6.433498
352    276.3721 13.566549
369    254.9681 11.348109
370    210.4491 18.056151
371    253.6360  9.188445
372    267.0448 11.298073

Degrees of freedom: 180 total; 144 residual
Residual standard error: 25.59182
> mean(coef(sleeplmList)[,1])   # sample means match lmer fixed eff
[1] 251.4051
>  mean(coef(sleeplmList)[,2])
[1] 10.46729
>  #mean int and slope match lmer Fixed effects results
> sleeplmer = lmer(Reaction ~ Days + (1 + Days|Subject), sleepstudy
> summary(sleeplmer)  # unconditional straight-line growth mixed mo
```

```
Linear mixed model fit by REML ['lmerMod']
Formula: Reaction ~ Days + (1 + Days | Subject)
   Data: sleepstudy

REML criterion at convergence: 1743.6

Scaled residuals:
    Min      1Q  Median      3Q     Max
-3.9536 -0.4634  0.0231  0.4634  5.1793

Random effects:
 Groups   Name        Variance Std.Dev. Corr
 Subject  (Intercept) 612.09   24.740
          Days         35.07    5.922   0.07
 Residual             654.94   25.592
Number of obs: 180, groups:  Subject, 18

Fixed effects:
            Estimate Std. Error t value
(Intercept)  251.405      6.825   36.84
Days          10.467      1.546    6.77

Correlation of Fixed Effects:
     (Intr)
Days -0.138

> plot(confint(sleeplmList)) # shows CI for slope and intercept lmL
> fixef(sleeplmer) # fixed: same for all units
(Intercept)        Days
  251.40510    10.46729
> ranef(sleeplmer) # individual; random effects (centered at 0) eac
$Subject
    (Intercept)        Days
308   2.2585654   9.1989719
309 -40.3985769  -8.6197032
310 -38.9602458  -5.4488799
330  23.6904985  -4.8143313
331  22.2602027  -3.0698946
332   9.0395259  -0.2721707
333  16.8404311  -0.2236244
334  -7.2325792   1.0745761
335  -0.3336958 -10.7521591
```

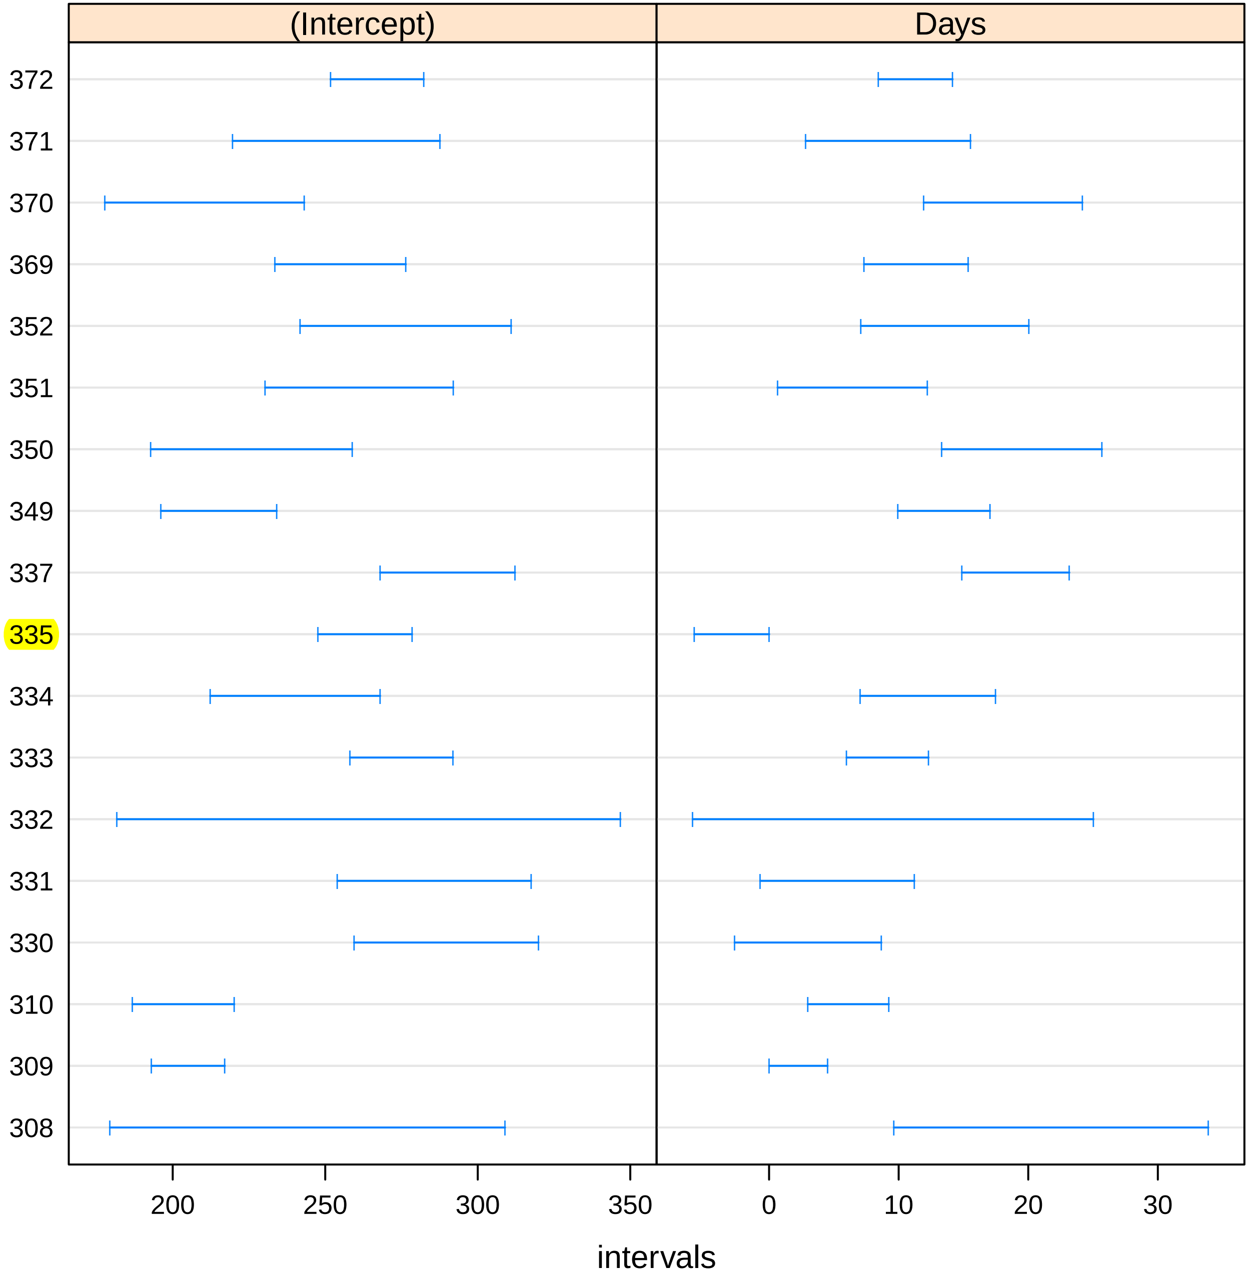
(Intercept)
Days
372
371
370
369
352
351
350
349
337
335
334
333
332
331
330
310
309
308
200
250
300
350
0
10
20
30
intervals

```
337  34.8903508   8.6282840
349 -25.2101104   1.1734142
350 -13.0699567   6.6142050
351   4.5778352  -3.0152572
352  20.8635924   3.5360133
369   3.2754530   0.8722166
370 -25.6128694   4.8224646
371   0.8070397  -0.9881551
372  12.3145393   1.2840297

> mean(unlist((ranef(sleeplmer))))
[1] 3.297408e-13


> dotplot(ranef(sleeplmer, postVar = TRUE))
$Subject
Warning message:
In ranef.merMod(sleeplmer, postVar = TRUE) :
  'postVar' is deprecated: please use 'condVar' instead

> dotplot(ranef(sleeplmer, condVar = TRUE)) # random effects and CI
$Subject

> qqmath(ranef(sleeplmer, condVar = TRUE)) # prefer when many rando
$Subject

###### make Bates BLUP plot to show shrinkage toward fixed effects
## adapt Bates code

> attach(sleepstudy)
> df <- coef(lmList(Reaction ~ Days | Subject, sleepstudy))
> fclow <- subset(df, `(Intercept)` < 251)
> fchigh <- subset(df, `(Intercept)` > 251)
> cc1 <- as.data.frame(coef(sleeplmer)$Subject)
> names(cc1) <- c("A", "B")
> df <- cbind(df, cc1)
> ff <- fixef(sleeplmer)
> with(df,
+      print(xyplot(`(Intercept)` ~ Days, aspect = 1,
+                   x1 = B, y1 = A,
+                   panel = function(x, y, x1, y1, subscripts, ...)
+                       panel.grid(h = -1, v = -1)
```

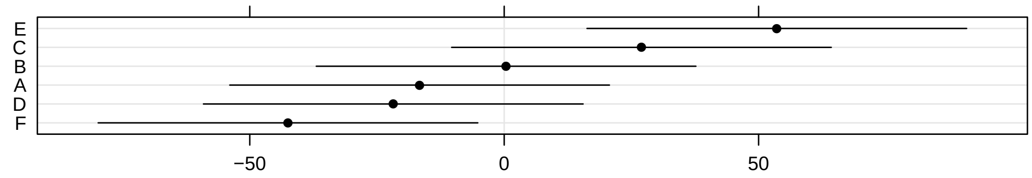

Fig. 1.10 95% prediction intervals on the random effects in `fm1ML`, shown as a dotplot.

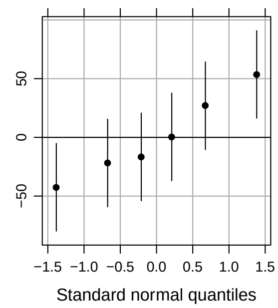

Fig. 1.11 95% prediction intervals on the random effects in `fm1ML` versus quantiles of the standard normal distribution.

interval. The `ranef` extractor takes an optional argument, `postVar = TRUE`, which adds these dispersion measures as an attribute of the result. (The name stands for "posterior variance", which is a misnomer that had become established as an argument name before I realized that it wasn't the correct term.)

We can plot these prediction intervals using

```
> dotplot(ranef(fm1ML, postVar = TRUE))
```

(Fig. 1.10), which provides linear spacing of the levels on the y axis, or using

```
> qqmath(ranef(fm1ML, postVar=TRUE))
```

(Fig. 1.11), where the intervals are plotted versus quantiles of the standard normal.

The dotplot is preferred when there are only a few levels of the grouping factor, as in this case. When there are hundreds or thousands of random effects the `qqmath` form is preferred because it focuses attention on the "important few" at the extremes and de-emphasizes the "trivial many" that are close to zero.

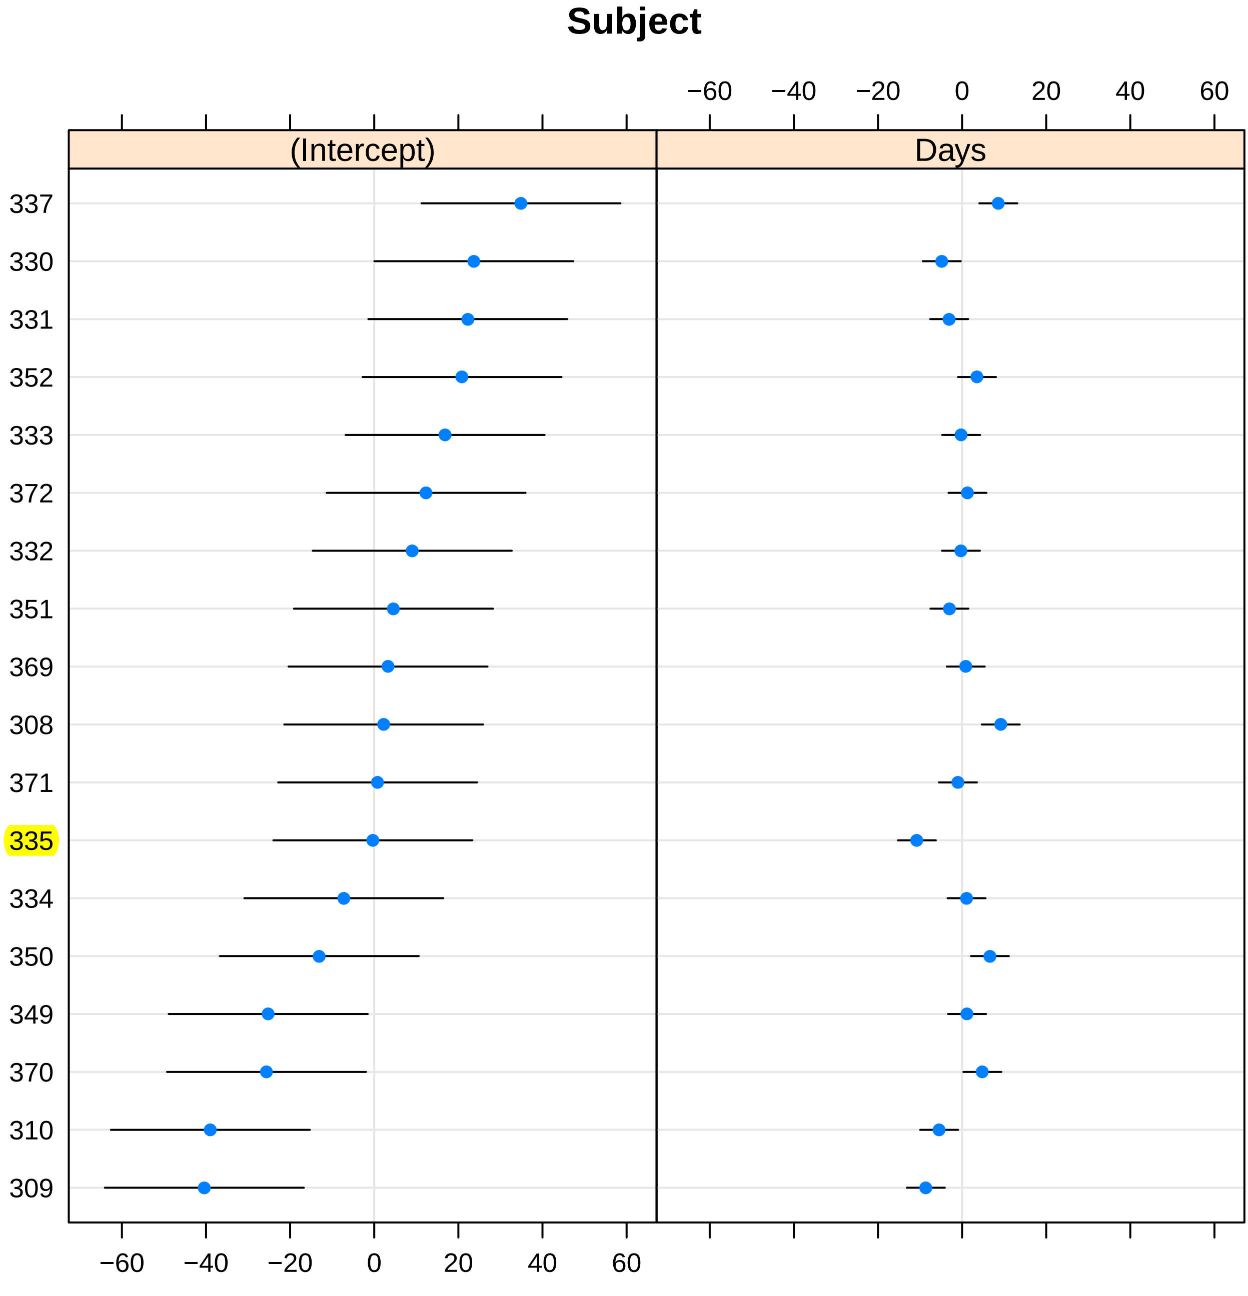

Subject
(Intercept)
Days
337
330
331
352
333
372
332
351
369
308
371
335
334
350
349
370
310
309
−60
−40
−20
0
20
40
60

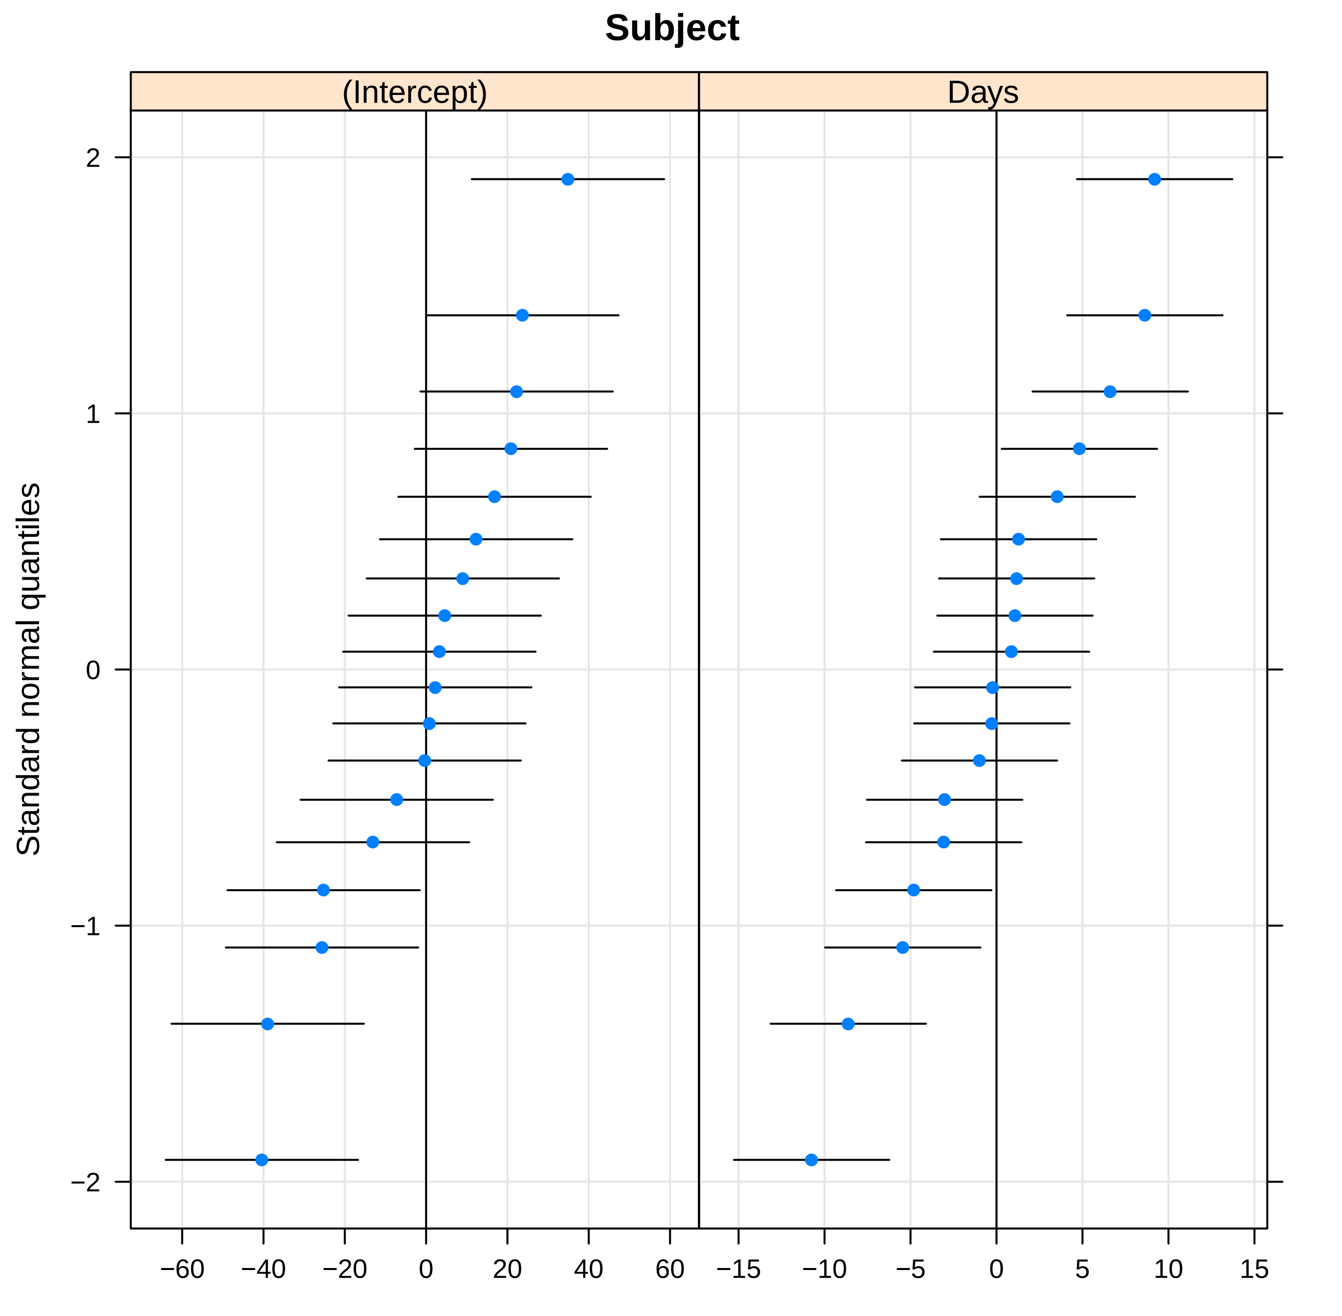
Subject
(Intercept)
Days
Standard normal quantiles
2
1
0
−1
−2
−60
−40
−20
0
20
40
60
−15
−10
−5
0
5
10
15

```
+                           x1 <- x1[subscripts]
+                           y1 <- y1[subscripts]
+                           larrows(x, y, x1, y1, type = "closed", leng
+                                   angle = 15, ...)
+                           lpoints(x, y,
+                                   pch = trellis.par.get("superpose.sy
+                                   col = trellis.par.get("superpose.sy
+                           lpoints(x1, y1,
+                                   pch = trellis.par.get("superpose.sy
+                                   col = trellis.par.get("superpose.sy
+                           lpoints(ff[2], ff[1],
+                                   pch = trellis.par.get("superpose.sy
+                                   col = trellis.par.get("superpose.sy
+                           ltext(fclow[,2], fclow[,1], row.names(fclow
+                                 adj = c(0.5, 1.7))
+                           ltext(fchigh[,2], fchigh[,1], row.names(fch
+                                 adj = c(0.5, -0.6))
+                       },
+                       key = list(space = "top", columns = 3,
+                       text = list(c("Mixed model", "Within-group", "P
+                       points = list(col = trellis.par.get("superpose.
+                       pch = trellis.par.get("superpose.symbol")$pch[1
+                       )))
>
## it actually worked
```

plot: blue = ranef + fixed, Stein estimation   borrowing strength

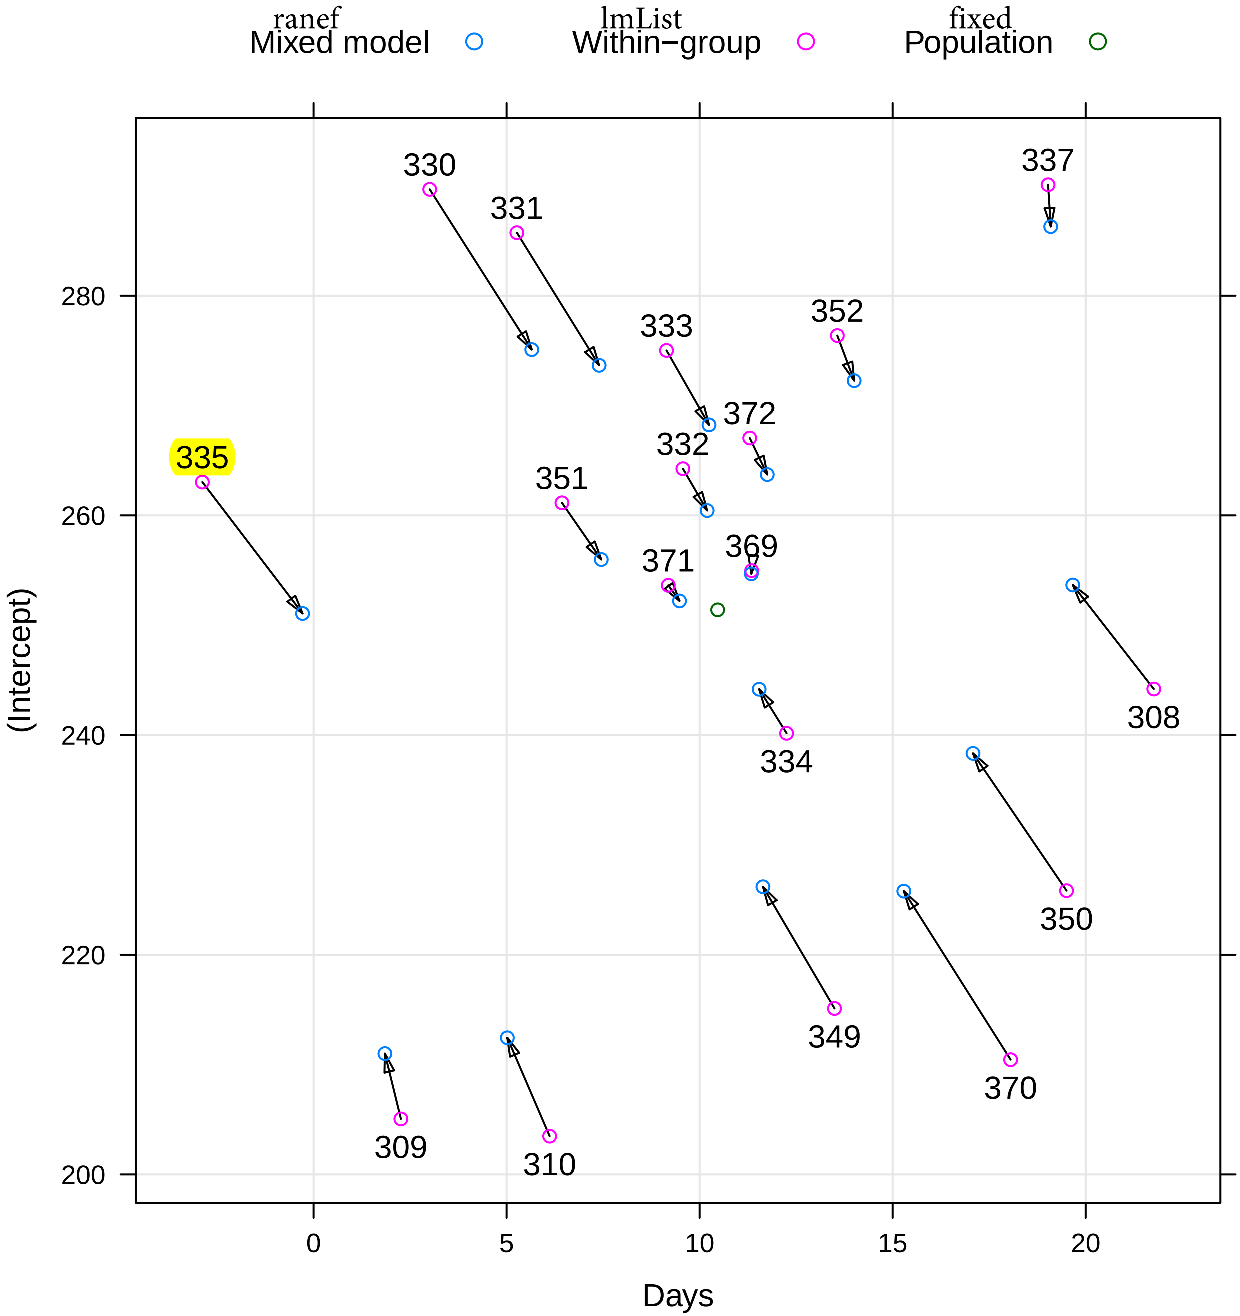
ranef
Mixed model
lmList
Within−group
fixed
Population
330
331
337
333
352
372
332
335
351
371
369
308
334
350
349
370
309
310
280
260
240
220
200
(Intercept)
0
5
10
15
20
Days

Brain Data

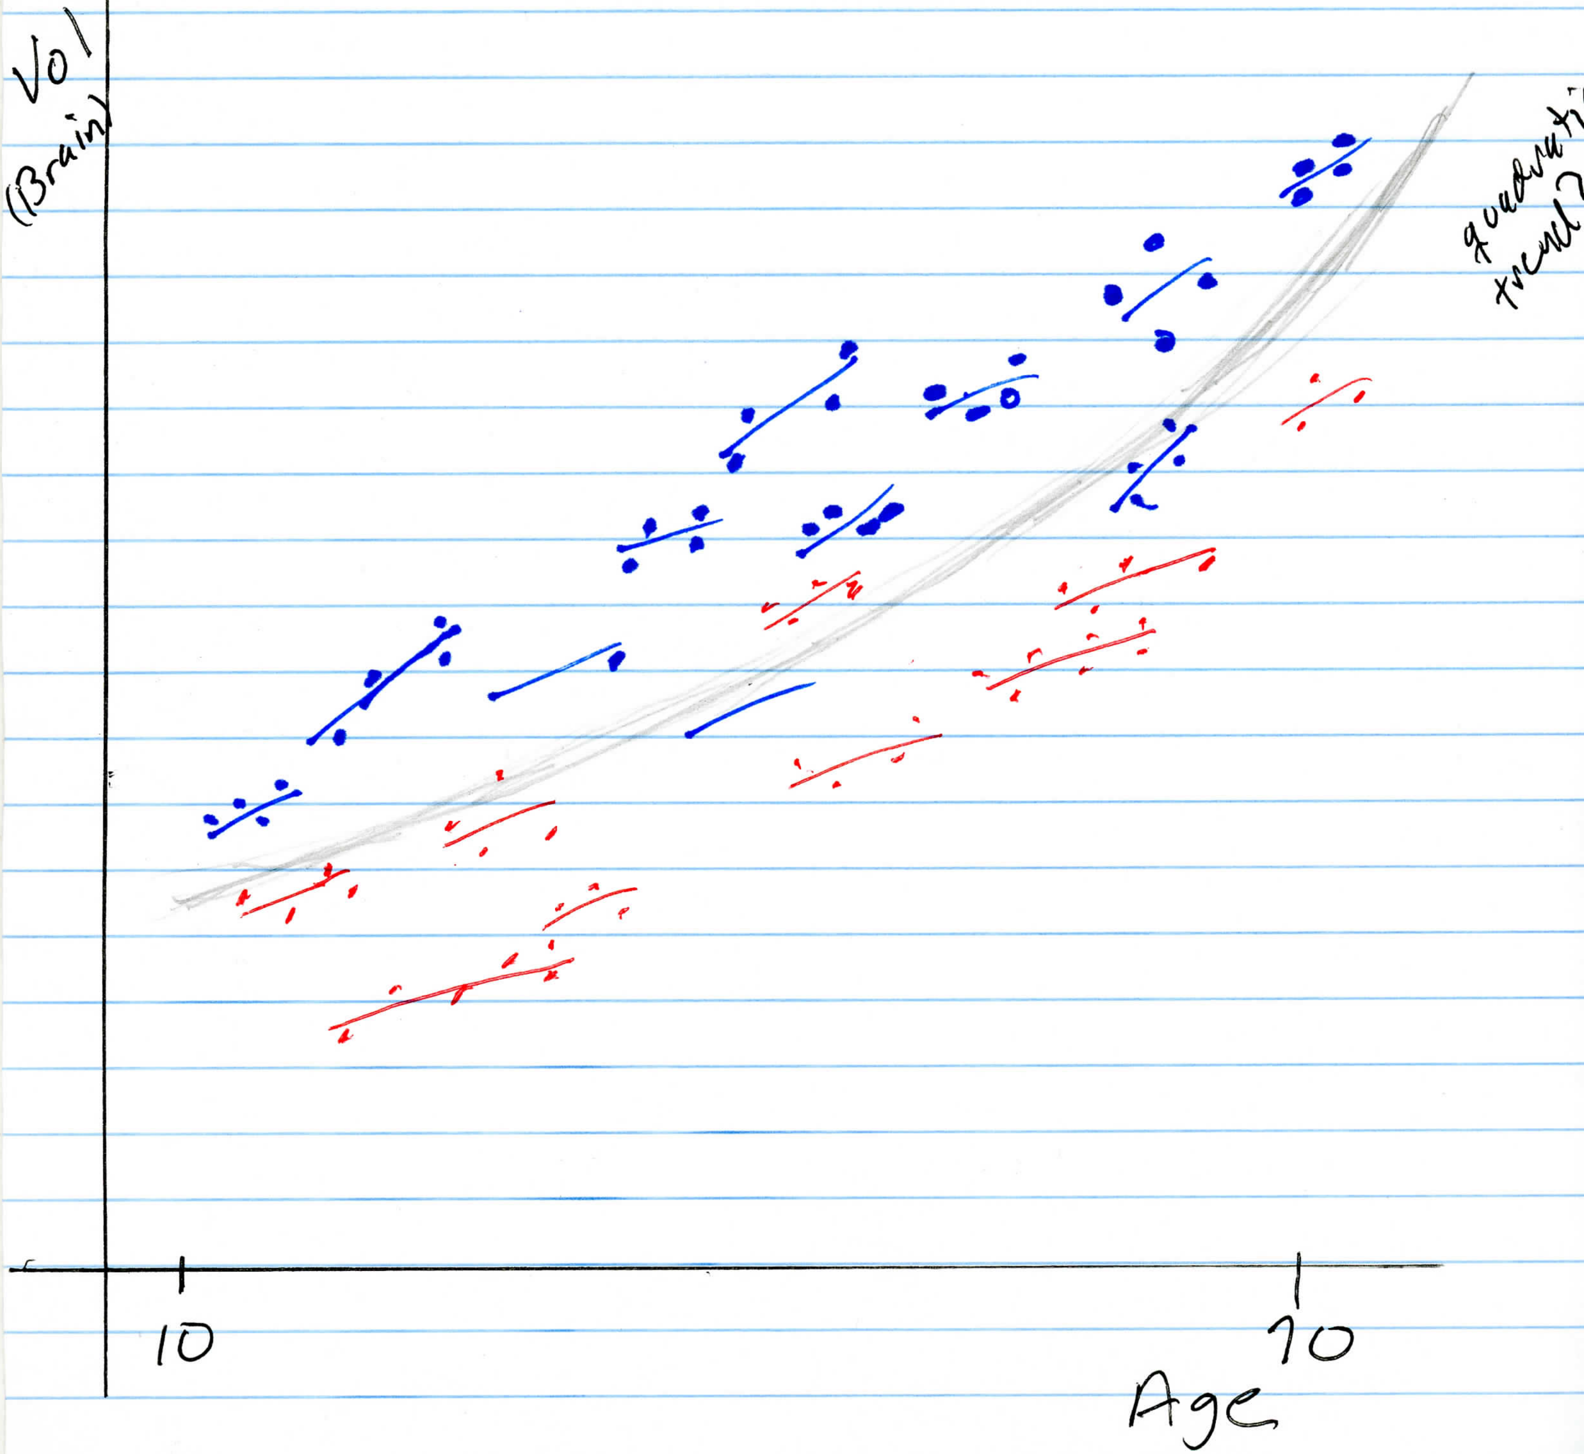

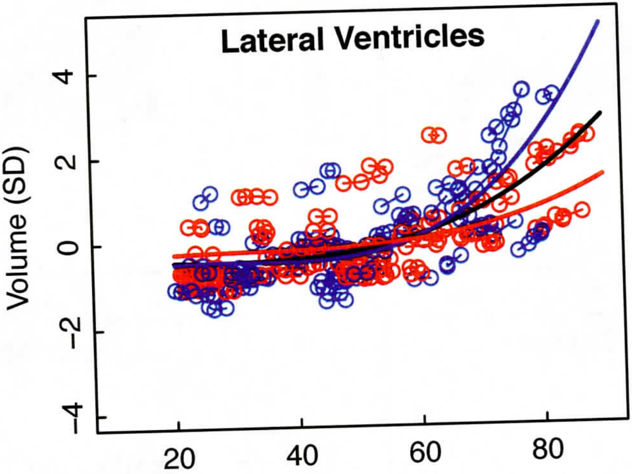
Lateral Ventricles
Volume (SD)
4
2
0
−2
−4
20
40
60
80

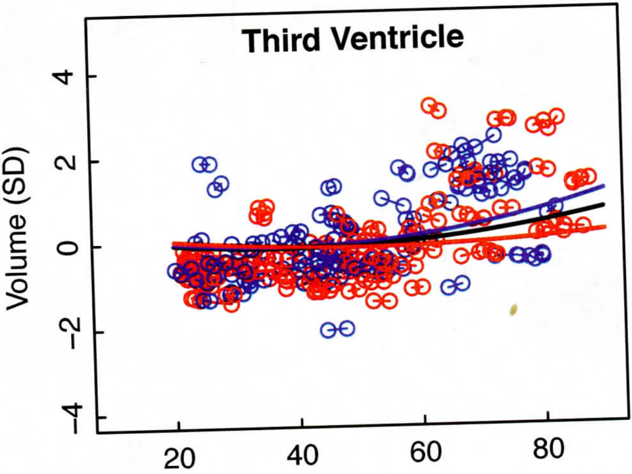
Third Ventricle
Volume (SD)
4
2
0
−2
−4
20
40
60
80

①

STAT 222
week 3
4/19/12

Models for LatVent trajectories

see linked plots, Males Females

Attributes for each individual

Gender : sexM = 1 if male

AgeMn : average age of observation
(i.e. if obs at age 40 42 46 48
AgeMn = 44

AgeC : observation times for an individual
centered for AgeMn
(i.e. for above. AgeC = {-4, -2, 2, 4}
for that individual)

Level I model (for individual i)
observations j

$$Vol_{ij} = \alpha_{0i} + \alpha_{1i}\, ageC_{ij} + \varepsilon_{ij}$$

α_{0i} mean Vol level for i

α_{1i} slope of Vol trajectory
on age for i

Level II model (basic)

level $\alpha_0 = \gamma_{00} + \gamma_{01}\, ageMn + \gamma_{02}\, sexM + u_0$

slope $\alpha_1 = \gamma_{10} + \gamma_{11}\, ageMn + \gamma_{12}\, sexM + u_1$

both meanVol and slope increase w/ AgeMn and
constant displacement for gender ($\gamma > 0$)

②

Combined Model (estimation model)

substitute
α_0, α_1 Level II into Level I

$$Vol \sim \gamma_{01}\, ageMn + \gamma_{02}\, sexM + \gamma_{10}\, ageC$$
$$+ \gamma_{11}\, ageMn \times ageC + \gamma_{12}\, sexM \times ageC$$
$$+ [\varepsilon + u_0 + u_1 \text{ terms}]$$

(x indicates multiplication)

---

* notation, A*B indicates A, B (main), A:B (interaction)

---

collect terms for lmer model

$$Vol \sim ageMn * ageC + sexM * ageC$$

summary for lmer object

parameter interpretation

param	lmer	interp
$\hat{\gamma}_{00}$	model intercept	
$\hat{\gamma}_{01}$	ageMn	increase in $\overline{Vol}$ by unit increase ageMn
$\hat{\gamma}_{02}$	sexM	increment in $\overline{Vol}$ for males
$\hat{\gamma}_{10}$	ageC	value of α_1 for age=0, Fem
$\hat{\gamma}_{11}$	ageMn:ageC	increase in α_{1i} for unit increase $ageMn_i$
$\hat{\gamma}_{12}$	sexM:ageC	increase in α_{1i} (slope) for Male

③

Even more, expand model to allow increase in α_{1i} (slope) for Male to depend on age (ageMn)

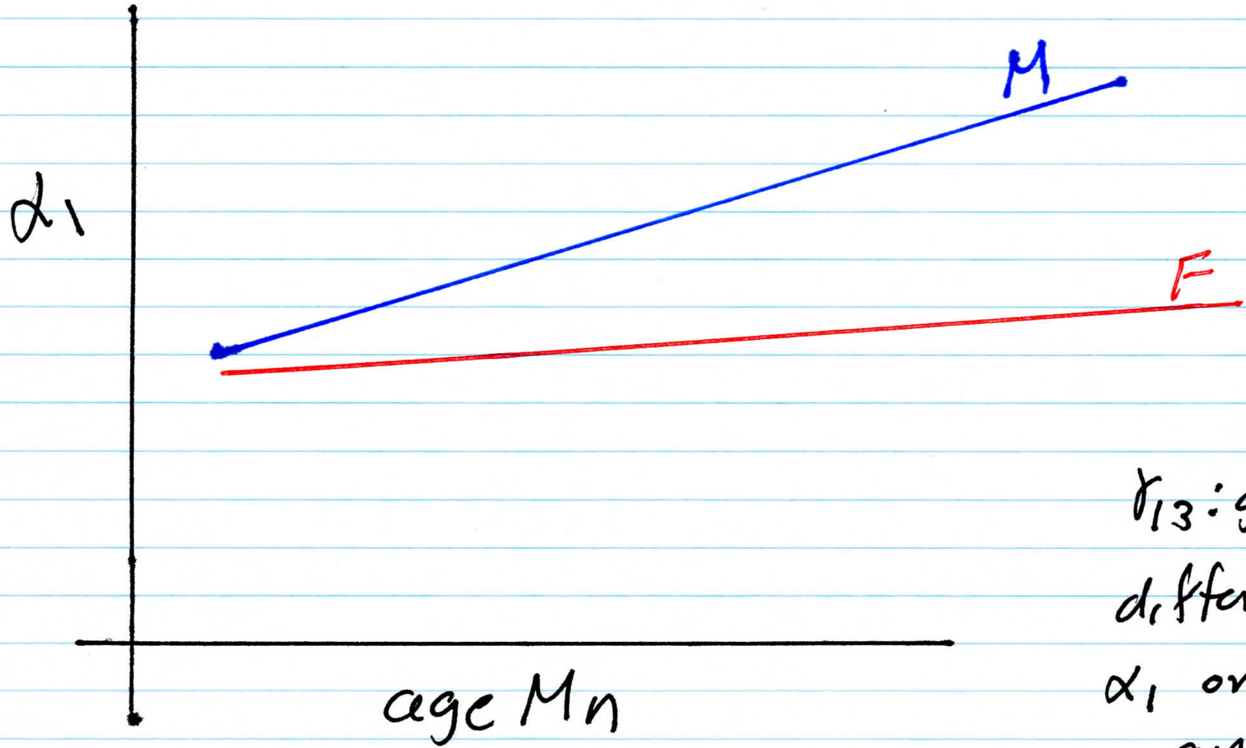

γ_{13}: gender difference in α_1 on ageMn gradient

Level II model extension

$$\alpha_1 = \gamma_{10} + \gamma_{11}\,\text{ageMn} + \gamma_{12}\,\text{sexM} + \gamma_{13}\,\text{sexM} \times \text{ageMn}$$

Combined (estimation) model

Vol ~ sexM * ageMn * ageC − sexM:ageMn

or equiv

Vol ~ ageMn*ageC + sexM*ageC + sexM:ageMn:ageC