

Mixed-Effects Model

Objectives

- Introduction
- Fixed-effect terms, Random-effect terms
- Fitting a model by **lmer** & **lme**
- Model Comparison & Diagnostics

Introduction

Ex 1. Shown below are weights of ten randomly selected people, “before” and 3-month “after” starting a diet program. These types of data are called **paired** or **matched** data. Let’s test if the diet is effective.

```
> data1 <- read.csv("U:\\STAT510\\dietDemo.csv")
> data1
  person time weight
1      1  bef   190
2      2  bef   208
3      3  bef   202
4      4  bef   189
5      5  bef   210
6      6  bef   240
7      7  bef   196
8      8  bef   233
9      9  bef   279
10     1  aft   178
11     2  aft   203
12     3  aft   194
13     4  aft   175
14     5  aft   202
15     6  aft   218
16     7  aft   185
17     8  aft   211
18     9  aft   229
> attach(data1)
> weight[time=="bef"]
[1] 190 208 202 189 210 240 196 233 279
> weight[time=="aft"]
[1] 178 203 194 175 202 218 185 211 229
> t.test(weight[time=="bef"],weight[time=="aft"],paired=T)
```

Paired t-test

```
data: weight[time == "bef"] and weight[time == "aft"]
t = 3.6819, df = 8, p-value = 0.006202
alternative hypothesis: true difference in means is not equal to 0
95 percent confidence interval:
 6.31136 27.46642
sample estimates:
mean of the differences
16.88889
```

Note that the t -test statistic= 3.6819 and the p -value= 0.006202. The diet program seems highly effective. Next, we show you how it can be analyzed by the “lmer” function of R.

```
> library(lme4)
> model1 <- lmer(weight~time+(1|person))
> summary(model1)
Linear mixed model fit by REML ['lmerMod']
Formula: weight ~ time + (1 | person)
```

REML criterion at convergence: 142.3

Random effects:

Groups	Name	Variance	Std.Dev.
person	(Intercept)	505.33	22.48
Residual		94.68	9.73

Number of obs: 18, groups: person, 9

Fixed effects:

	Estimate	Std. Error	t value
(Intercept)	199.444	8.165	24.427
timebef	16.889	4.587	3.682

Correlation of Fixed Effects:

	(Intr)
timebef	-0.281

```
> model2 <- lm(weight~time)
> anova(model2)
Analysis of Variance Table
```

Response: weight

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
time	1	1283.6	1283.56	2.1392	0.1629
Residuals	16	9600.2	600.01		

You will be learning how to write the “lmer” model as shown above. In this example, “time” is a fixed-effect term, and “person” is a random-effect term. We observe that the t -test statistic for “time” is the same as before, and the traditional MSE(= 600.01) was partitioned into two parts, i.e., variation due to the “person”(= 505.33) and random variation(= 94.68).

Fixed Effect & Random Effect

If the experiment were repeated	
Fixed effects:	Same levels would be used
Random effects:	Different levels would be used. “Nested” terms are often random effects.
Desired inference? The conclusions refer to	
Fixed effects:	The levels used.
Random effects:	A population from which the levels used are just a (random) sample.

Fixed effects influence only the **mean** of y , and random effects influence only the **variance** of y . Some more typical examples are:

Fixed effects	Random effects
Male or female	Individuals
Light vs shade	Families
One age vs another	Individuals
High altitude and low altitude	Trees
Drug administered or not	Genotypes
Insecticide sprayed or not	Brood
Nutrient added or not	Blocks within a field
One ethnic group vs another	Split plot within a plot

Replication & Pseudoreplication

To qualify as replicates, measurements must have the following properties:

- They must be independent.
- They must not form part of a time series (data collected from the same place on successive occasions are not independent).
- They must not be grouped together in one place (aggregating the replicates means that they are not spatially independent).
- They must be of an appropriate spatial scale.
- Ideally, one replicate from each treatment ought to be grouped together into a block, and each treatment repeated in many different blocks.
- Repeated measures (e.g., from the same individual or the same spatial location) are not replicates (this is probably the commonest cause of “pseudoreplication” in statistical work). Pseudoreplication occurs when you analyze the data as if you had more degrees of freedom than you really have. There are two kinds of pseudoreplication:
 - temporal pseudoreplication, involving repeated measurements from the same individual;
 - spatial pseudoreplication, involving several measurements taken from the same vicinity.

Pseudoreplication is a problem because one of the most important assumptions of standard statistical analysis is *independence of errors*. Repeated measures through time on the same individual will have non-independent errors because peculiarities of the individual will be reflected in all of the measurements made on it (i.e., the repeated measures will be temporally correlated with one another). Samples taken from the same vicinity will have non-independent errors because peculiarities of the location will be common to all the samples (e.g., yields will all be high in a good patch and all be low in a bad patch). There are various ways that you can do when your data are pseudoreplicated:

- Average away the pseudoreplication and carry out your statistical analysis on the means.
- Carry out separate analyses for each time period.
- Use proper time series analysis or mixed-effects models.

Ex 2. Spli-Plot Experiments: with categorical random-effect terms

The following example refers to an experiment on crop yield with three treatments: irrigation (with two levels, irrigated or not), sowing density (with three levels, low, medium and high), and fertilizer application (with three levels, low, medium and high). The largest plots were the four whole fields (**block**), each of which was split in half, and **irrigation** was allocated at random to one half of the field. Each irrigation plot was split into three, and one of three different seed-sowing **densities** (low, medium or high) was allocated at random (independently for each level of irrigation and each block). Finally, each density plot was divided into three, and one of three **fertilizer** nutrient treatments (N, P, or N & P together) was allocated at random.

```
> data1 <- read.table("U:\\STAT510\\splityield.txt",header=T)
> dim(data1)
[1] 72 5
> head(data1)
  yield block irrigation density fertilizer
1    90     A   control     low         N
2    95     A   control     low         P
3   107     A   control     low        NP
4    92     A   control    medium         N
5    89     A   control    medium         P
6    92     A   control    medium        NP
> attach(data1)
> par(mfrow=c(1,2))
> interaction.plot(fertilizer,irrigation,yield) # See Fig. 1
> interaction.plot(density,irrigation,yield)
```




Figure 1: There seem to be a significant interaction between **density** & **irrigation**, and between **fertilizer** & **irrigation**.

Here is how this scenario can be written with different error terms to test for various model terms.

```

> irr <- irrigation # To shorten variable names.
> fer <- fertilizer
> den <- density
> model1 <- aov(yield~irr*den*fer+Error(block/irr/den))
> summary(model1)

```

Error: block

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
Residuals	3	194.4	64.81		

Error: block:irr

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
irr	1	8278	8278	17.59	0.0247 *
Residuals	3	1412	471		

Error: block:irr:den

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
den	2	1758	879.2	3.784	0.0532 .
irr:den	2	2747	1373.5	5.912	0.0163 *
Residuals	12	2788	232.3		

Error: Within

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
fer	2	1977.4	988.7	11.449	0.000142 ***
irr:fer	2	953.4	476.7	5.520	0.008108 **
den:fer	4	304.9	76.2	0.883	0.484053
irr:den:fer	4	234.7	58.7	0.680	0.610667
Residuals	36	3108.8	86.4		

It is recommended that when there are missing values in a split-plot experiment you should use either `lmer` or `lme` instead of `aov` to fit the model. The way to write models in `lmer` and `lme` are slightly different. We will show first how to write in `lme` first, then `lmer`.

```

> library(nlme)
> model2 <- lme(yield~irr*den*fer, random = ~ 1|block/irr/den)
> summary(model2)

```

Linear mixed-effects model fit by REML

Data: NULL

	AIC	BIC	logLik
	481.6212	525.3789	-218.8106

Random effects:

Formula: ~1 | block
(Intercept)

StdDev: 0.000661967

Formula: ~1 | irr %in% block
(Intercept)

StdDev: 1.982465

Formula: ~1 | den %in% irr %in% block
(Intercept) Residual

StdDev: 6.975552 9.292805

Fixed effects: yield ~ irr * den * fer

	Value	Std.Error	DF	t-value	p-value
(Intercept)	80.50	5.893741	36	13.658558	0.0000
irrigigated	31.75	8.335008	3	3.809235	0.0318
denlow	5.50	8.216281	12	0.669403	0.5159
denmedium	14.75	8.216281	12	1.795216	0.0978
ferNP	5.50	6.571006	36	0.837010	0.4081
ferP	4.50	6.571006	36	0.684827	0.4978
irrigigated:denlow	-39.00	11.619576	12	-3.356405	0.0057
irrigigated:denmedium	-22.25	11.619576	12	-1.914872	0.0796
irrigigated:ferNP	13.00	9.292805	36	1.398932	0.1704
irrigigated:ferP	5.50	9.292805	36	0.591856	0.5576
denlow:ferNP	3.25	9.292805	36	0.349733	0.7286
denmedium:ferNP	-6.75	9.292805	36	-0.726368	0.4723
denlow:ferP	-5.25	9.292805	36	-0.564953	0.5756
denmedium:ferP	-5.50	9.292805	36	-0.591856	0.5576
irrigigated:denlow:ferNP	7.75	13.142011	36	0.589712	0.5591
irrigigated:denmedium:ferNP	3.75	13.142011	36	0.285344	0.7770
irrigigated:denlow:ferP	20.00	13.142011	36	1.521837	0.1368
irrigigated:denmedium:ferP	4.00	13.142011	36	0.304367	0.7626

```

Correlation:
<<< Omitted to save paper >>>
> library(lme4)
> model3 <- lmer(yield~irr*den*fer+(1|block)+(1|block:irr)+(1|block:irr:den))
> summary(model3)
<<< Omitted to save paper -- same printout as model2>>>

```

This output shows that the only significant effects are the main effect of irrigation (p -value= 0.0318) and the irrigation by density interaction (p -value= 0.0057). The threeway interaction is not significant so we remove it, fitting all terms up to two-way interactions:

```

> model2_1 <- lme(yield~(irr+den+fer)^2, random = ~ 1|block/irr/den)
> summary(model2_1)
Linear mixed-effects model fit by REML
Data: NULL
      AIC      BIC    logLik
503.1256 540.2136 -233.5628

```

Random effects:

```

Formula: ~1 | block
(Intercept)
StdDev: 0.0005567731

```

```

Formula: ~1 | irr %in% block
(Intercept)
StdDev: 1.982562

```

```

Formula: ~1 | den %in% irr %in% block
(Intercept) Residual
StdDev: 7.041303 9.142696

```

Fixed effects: yield ~ (irr + den + fer)^2

	Value	Std.Error	DF	t-value	p-value
(Intercept)	82.47222	5.443438	40	15.150760	0.0000
irrigigated	27.80556	7.069256	3	3.933307	0.0293
denlow	0.87500	7.256234	12	0.120586	0.9060
denmedium	13.45833	7.256234	12	1.854727	0.0884
ferNP	3.58333	5.278538	40	0.678850	0.5011

```

ferP                0.50000  5.278538 40  0.094723  0.9250
irrigated:denlow    -29.75000  8.800165 12 -3.380618  0.0055
irrigated:denmedium -19.66667  8.800165 12 -2.234807  0.0452
irrigated:ferNP      16.83333  5.278538 40  3.189014  0.0028
irrigated:ferP       13.50000  5.278538 40  2.557526  0.0144
denlow:ferNP         7.12500  6.464862 40  1.102112  0.2770
denmedium:ferNP      -4.87500  6.464862 40 -0.754076  0.4552
denlow:ferP          4.75000  6.464862 40  0.734741  0.4668
denmedium:ferP       -3.50000  6.464862 40 -0.541388  0.5912
<<< Omitted to save paper >>>

```

This output shows that the interaction between fertilizer and density is not significant, so we remove it:

```

> model2_2 <- lme(yield~irr*den+irr*fer, random = ~ 1|block/irr/den)
> summary(model2_2)
Linear mixed-effects model fit by REML
Data: NULL
      AIC      BIC    logLik
519.9035 549.6834 -245.9517

Random effects:
Formula: ~1 | block
      (Intercept)
StdDev: 0.0005551647

Formula: ~1 | irr %in% block
      (Intercept)
StdDev: 1.982612

Formula: ~1 | den %in% irr %in% block
      (Intercept) Residual
StdDev: 7.057132 9.105995

```

```

Fixed effects: yield ~ irr * den + irr * fer
              Value Std.Error DF   t-value p-value
(Intercept)  82.08333  4.994999 44  16.433103  0.0000
irrigated     27.80556  7.063995  3   3.936236  0.0292
denlow        4.83333  6.222653 12   0.776732  0.4524
denmedium    10.66667  6.222653 12   1.714167  0.1122
ferNP         4.33333  3.717507 44   1.165656  0.2500
ferP          0.91667  3.717507 44   0.246581  0.8064
irrigated:denlow -29.75000  8.800161 12 -3.380620  0.0055
irrigated:denmedium -19.66667  8.800161 12 -2.234808  0.0452
irrigated:ferNP   16.83333  5.257349 44   3.201867  0.0025
irrigated:ferP    13.50000  5.257349 44   2.567834  0.0137

```

We also use `anova` to compare mixed models with different fixed-effects structures. Then you must use maximum likelihood (method = "ML") rather than the default restricted maximum likelihood (REML). Here is the analysis again, but this time using `anova` to compare models with progressively simplified fixed effects.

```

> model2 <- lme(yield~irr*den*fer, random = ~ 1|block/irr/den,method="ML")
> model2_1 <- lme(yield~(irr+den+fer)^2, random = ~ 1|block/irr/den,method="ML")
> anova(model2,model2_1)
      Model df      AIC      BIC    logLik   Test  L.Ratio p-value

```

```

model2      1 22 573.5108 623.5974 -264.7554
model2_1    2 18 569.0046 609.9845 -266.5023 1 vs 2 3.493788 0.4788

```

This means that the “shorter” model (i.e., model2_1) is better.

```

> model2_2 <- lme(yield~irr*den+irr*fer, random = ~ 1|block/irr/den,method="ML")
> anova(model2_1,model2_2)
      Model df      AIC      BIC    logLik   Test  L.Ratio p-value
model2_1     1 18 569.0046 609.9845 -266.5023
model2_2     2 14 565.1933 597.0667 -268.5967 1 vs 2 4.188774 0.3811
> model2_3 <- lme(yield~irr*den+fer, random = ~ 1|block/irr/den,method="ML")
> anova(model2_2,model2_3)
      Model df      AIC      BIC    logLik   Test  L.Ratio p-value
model2_2     1 14 565.1933 597.0667 -268.5967
model2_3     2 12 572.3373 599.6573 -274.1687 1 vs 2 11.14397 0.0038

```

This means that “longer” model (i.e., model2_2) is better.

```

> model2 <- lme(yield~irr*den*fer, random = ~ 1|block/irr/den,method="ML")
> model2_1 <- lme(yield~(irr+den+fer)^2, random = ~ 1|block/irr/den,method="ML")
> model2_2 <- lme(yield~irr*den+irr*fer, random = ~ 1|block/irr/den,method="ML")
> model2_3 <- lme(yield~irr*den+fer, random = ~ 1|block/irr/den,method="ML")
> model2_4 <- lme(yield~den+irr*fer, random = ~ 1|block/irr/den,method="ML")
> anova(model2_2,model2_4)
      Model df      AIC      BIC    logLik   Test  L.Ratio p-value
model2_2     1 14 565.1933 597.0667 -268.5967
model2_4     2 12 572.9022 600.2221 -274.4511 1 vs 2 11.70883 0.0029

```

Again, this means that “longer” model (i.e., model2_2) is better.

```

> summary(model2_2)
Linear mixed-effects model fit by maximum likelihood
Data: NULL
      AIC      BIC    logLik
565.1933 597.0667 -268.5967

```

Random effects:

```

Formula: ~1 | block
(Intercept)
StdDev: 0.0005335774

```

```

Formula: ~1 | irr %in% block
(Intercept)
StdDev: 1.716893

```

```

Formula: ~1 | den %in% irr %in% block
(Intercept) Residual
StdDev: 5.722412 8.718327

```

```

Fixed effects: yield ~ irr * den + irr * fer
      Value Std.Error DF   t-value p-value
(Intercept)  82.08333  4.756285 44 17.257868 0.0000
irrirrigated  27.80556  6.726402  3  4.133793 0.0257
denlow       4.83333  5.807346 12  0.832279 0.4215
denmedium    10.66667  5.807346 12  1.836754 0.0911
ferNP         4.33333  3.835553 44  1.129781 0.2647
ferP          0.91667  3.835553 44  0.238992 0.8122
irrirrigated:denlow -29.75000  8.212827 12 -3.622382 0.0035

```



```

irr irrigated:denmedium -19.66667  8.212827 12 -2.394628  0.0338
irr irrigated:ferNP      16.83333  5.424290 44  3.103325  0.0033
irr irrigated:ferP       13.50000  5.424290 44  2.488805  0.0167

```

We note that most of the variation in the response is in **density** and the least is in **block**. Shown below is the same model done by `lmer`.

```

> model3_1 <- lmer(yield~irr*den+irr*fer+(1|block)+(1|block:irr)+(1|block:irr:den))
> summary(model3_1,cor=F)
Linear mixed model fit by REML ['lmerMod']
Formula: yield ~ irr * den + irr * fer + (1 | block) + (1 | block:irr) +      (1 | block:irr:den)

```

REML criterion at convergence: 491.9

Scaled residuals:

	Min	1Q	Median	3Q	Max
	-2.43777	-0.53126	0.06701	0.57279	2.00343

Random effects:

Groups	Name	Variance	Std.Dev.
block:irr:den	(Intercept)	49.803	7.057
block:irr	(Intercept)	3.931	1.983
block	(Intercept)	0.000	0.000
Residual		82.919	9.106

Number of obs: 72, groups: block:irr:den, 24; block:irr, 8; block, 4

Fixed effects:

	Estimate	Std. Error	t value
(Intercept)	82.0833	4.9950	16.433
irr irrigated	27.8056	7.0640	3.936
denlow	4.8333	6.2227	0.777
denmedium	10.6667	6.2227	1.714
ferNP	4.3333	3.7175	1.166
ferP	0.9167	3.7175	0.247
irr irrigated:denlow	-29.7500	8.8002	-3.381
irr irrigated:denmedium	-19.6667	8.8002	-2.235
irr irrigated:ferNP	16.8333	5.2574	3.202
irr irrigated:ferP	13.5000	5.2574	2.568

Finally, we look at some relevant diagnostic plots.

```

> plot(model2_2) # See Fig. 2
> plot(model2_2,yield~fitted(.),xlim=c(60,140),ylim=c(60,140),asp=1) # See Fig. 3
> s_residuals <- split(model2_2$residuals,block)
> par(mfrow=c(2,2))
> qqnorm(s_residuals[[1]],main="blk 1"); qqline(s_residuals[[1]]) # See Fig. 4
> qqnorm(s_residuals[[2]],main="blk 2"); qqline(s_residuals[[2]])
> qqnorm(s_residuals[[3]],main="blk 3"); qqline(s_residuals[[3]])
> qqnorm(s_residuals[[4]],main="blk 4"); qqline(s_residuals[[4]])

```

Ex 3. Repeated Measures Data: with continuous random-effect terms

A common cause of temporal pseudoreplication in growth experiments with fixed effects is when each individual is measured several times as it grows during the course of an experiment. In the

next example, we have a single fixed effect (a two-level categorical variable: with **fertilizer** added or not) and six replicate **plants** per each treatment, with each plant measured on five occasions (after 2, 4, 6, 8 or 10 **weeks** of growth). The response variable is **root** length. The random-effects formula needs to indicate that the week of measurement (a continuous random effect) represents pseudoreplication within each individual plant and will be written as **random = ~week|plant**. Because we have a continuous random effect (weeks) we write **~week** in the random-effects formula rather than the **~1** that we used with categorical random effects in Ex 2.

```
> data1 <- read.table("U:\\STAT510\\fertilizer.txt",header=T)
> dim(data1)
[1] 60 4
> head(data1)
  root week plant fertilizer
1  1.3   2  ID1      added
2  3.5   4  ID1      added
3  7.0   6  ID1      added
4  8.1   8  ID1      added
5 10.0  10  ID1      added
6  2.0   2  ID2      added
> attach(data1)
> summary(fertilizer)
added control
    30     30
> library(lme4)
> model1 <- lmer(root~fertilizer+(1|plant)+(week|plant))
> summary(model1)
Linear mixed model fit by REML ['lmerMod']
Formula: root ~ fertilizer + (1 | plant) + (week | plant)

REML criterion at convergence: 159

Scaled residuals:
    Min       1Q   Median       3Q      Max
-1.9928 -0.6587 -0.1004  0.6950  2.0225

Random effects:
Groups   Name              Variance Std.Dev. Corr
plant    (Intercept)  0.02162   0.1470
plant.1  (Intercept)  8.18077   2.8602
        week          0.87786   0.9369  -1.00
Residual                0.24664   0.4966
Number of obs: 60, groups: plant, 12

Fixed effects:
              Estimate Std. Error t value
(Intercept)      2.7997     0.1438   19.46
fertilizercontrol -1.0394     0.2034   -5.11

Correlation of Fixed Effects:
              (Intr)
frtlzrcntrl -0.707
```

The mean reduction in root size associated with the unfertilized controls is -1.0394 and this has a standard error of 0.2034 . **fertilizer** is a highly significant factor. Also, there is almost 9 times as

much variation in **plants** than in **weeks** (0.87786 vs 8.18077).

For the kind of data involved in mixed-effects models, there are some excellent built-in plotting functions (variously called panel plots, trellis plots, or lattice plots). To use trellis plotting, we begin by turning our dataframe into a **groupedData** object. To do this we specify the nesting structure of the random effects, and indicate the fixed effect by defining it as **outer** to this nesting.

```
> library(nlme)
> data2 <- groupedData(root~week|plant,outer=~fertilizer,data1)
> library(lattice)
> plot(data2) # See Fig. 5
> plot(data2,outer=T) #See Fig. 6
```

Ex 4. One more example: more practice writing a model in **lmer**

Three experimental treatments were administered to rats, and the glycogen content of the rats' livers was analyzed as the response variable. Two rats were allocated per treatment, and after each rat was killed, its liver was cut up into three pieces: a left-hand bit, a central bit and a right-hand bit. Finally, two separate preparations were made from each macerated bit of liver.

```
> data1 <- read.table("U:\\STAT510\\rats.txt",header=T)
> dim(data1)
[1] 36 4
> attach(data1)
> data1
```

	Glycogen	Treatment	Rat	Liver
1	131	1	1	1
2	130	1	1	1
3	131	1	1	2
4	125	1	1	2
5	136	1	1	3
6	142	1	1	3
7	150	1	2	1
8	148	1	2	1
9	140	1	2	2
10	143	1	2	2
11	160	1	2	3
12	150	1	2	3
13	157	2	1	1
14	145	2	1	1
15	154	2	1	2
16	142	2	1	2
17	147	2	1	3
18	153	2	1	3
19	151	2	2	1
20	155	2	2	1
21	147	2	2	2
22	147	2	2	2
23	162	2	2	3
24	152	2	2	3
25	134	3	1	1
26	125	3	1	1
27	138	3	1	2

```

28      138      3  1  2
29      135      3  1  3
30      136      3  1  3
31      138      3  2  1
32      140      3  2  1
33      139      3  2  2
34      138      3  2  2
35      134      3  2  3
36      127      3  2  3
> Treatment<-factor(Treatment)
> Liver<-factor(Liver)
> Rat<-factor(Rat)
> model1 <- lmer(Glycogen~Treatment+(1|Treatment:Rat)+(1|Treatment:Rat:Liver))
> summary(model1)
Linear mixed model fit by REML ['lmerMod']
Formula: Glycogen ~ Treatment + (1 | Treatment:Rat) + (1 | Treatment:Rat:Liver)

REML criterion at convergence: 219.6

Scaled residuals:
    Min       1Q   Median       3Q      Max
-1.48212 -0.47263  0.03062  0.42934  1.82935

Random effects:
 Groups                Name      Variance Std.Dev.
Treatment:Rat:Liver (Intercept) 14.17    3.764
Treatment:Rat      (Intercept) 36.06    6.005
Residual                                21.17    4.601
Number of obs: 36, groups: Treatment:Rat:Liver, 18; Treatment:Rat, 6

Fixed effects:
              Estimate Std. Error t value
(Intercept) 140.500      4.707  29.848
Treatment2   10.500      6.657   1.577
Treatment3   -5.333      6.657  -0.801

Correlation of Fixed Effects:
      (Intr) Trtmn2
Treatment2 -0.707
Treatment3 -0.707  0.500

```

We see that the treatment effect is not significant because $|t| < 2$.




Figure 2: Diagnostic plots: residuals vs fitted values




Figure 3: Diagnostic plots: observed values vs expected values

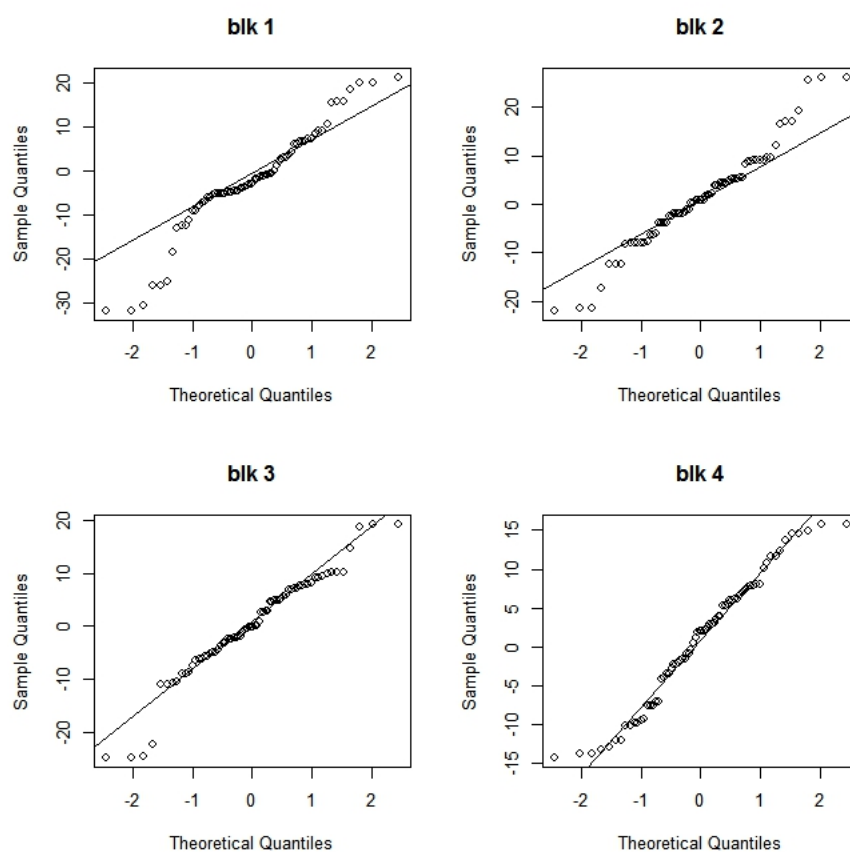


Figure 4: Diagnostic plots: normal QQ plots per block

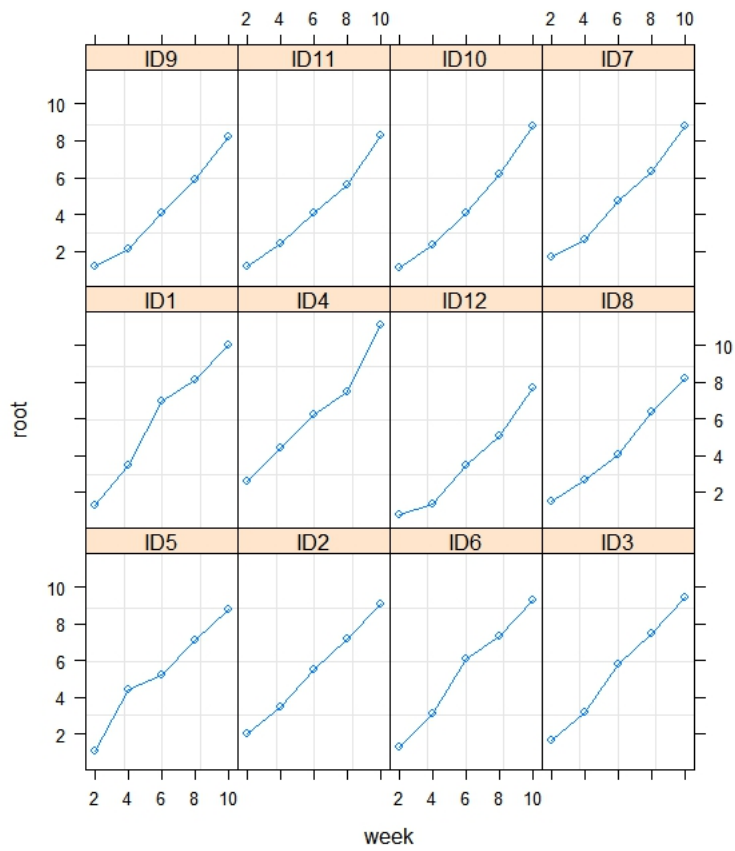


Figure 5: Plot for each individual




Figure 6: Another plot for each individual