Nested and nonnested grouping factors

560 Hierarchical modeling

Peter Hoff

Statistics, University of Washington

Nested groups

In some situations there are multiple grouping factors that are *nested*, having observations within groups within groups, etc:

- students within classrooms within schools within counties;
- cities within counties within states;
- medical measurements within patients within hospitals.

We will want to allow for across-group heterogeneity at each level of the hierarchy.

It is fairly straightforward to do this with random effects models, although some care has to be used in coding the group labels and fitting group-specific regression parameters.

Nested groups - ET example

A study examined the effects of two different instructional methods on three different exams.

- `itype` $\in \{1, 2\}$, instruction type, an unordered categorical factor.
- `etype` $\in \{1, 2, 3\}$, exam type, an unordered categorical factor.

Experimental design:

- $m_1 = 8$ different sessions (on 8 different days);
- $m_2 = 10$ subjects on each day (subjects were different across days);
- `itype=1` was given on odd days, `itype=2` was on even;
- Each subject given one of two instruction types; took all three exams.

Nested groups

```
etest[1:25,]

##    session itype subject sub.ses etype score
## 1        1     1       1       1     1    17
## 2        1     1       1       1     2    16
## 3        1     1       1       1     3    17
## 4        1     1       2       2     1    17
## 5        1     1       2       2     2    18
## 6        1     1       2       2     3    18
## 7        1     1       3       3     1    17
## 8        1     1       3       3     2    16
## 9        1     1       3       3     3    17
## 10       1     1       4       4     1    16
## 11       1     1       4       4     2    15
## 12       1     1       4       4     3    16
## 13       1     1       5       5     1    15
## 14       1     1       5       5     2    13
## 15       1     1       5       5     3    14
## 16       1     1       6       6     1    15
## 17       1     1       6       6     2    14
## 18       1     1       6       6     3    16
## 19       1     1       7       7     1    14
## 20       1     1       7       7     2    17
## 21       1     1       7       7     3    15
## 22       1     1       8       8     1    17
## 23       1     1       8       8     2    14
## 24       1     1       8       8     3    15
## 25       1     1       9       9     1    16
```

Nested groups

```
etest[20:45,]

##    session itype subject sub.ses etype score
## 20       1     1       7       7     2    17
## 21       1     1       7       7     3    15
## 22       1     1       8       8     1    17
## 23       1     1       8       8     2    14
## 24       1     1       8       8     3    15
## 25       1     1       9       9     1    16
## 26       1     1       9       9     2    16
## 27       1     1       9       9     3    15
## 28       1     1      10      10     1    16
## 29       1     1      10      10     2    13
## 30       1     1      10      10     3    16
## 31       2     2       1      11     1    15
## 32       2     2       1      11     2    14
## 33       2     2       1      11     3    15
## 34       2     2       2      12     1    13
## 35       2     2       2      12     2    11
## 36       2     2       2      12     3    12
## 37       2     2       3      13     1    13
## 38       2     2       3      13     2    15
## 39       2     2       3      13     3    16
## 40       2     2       4      14     1    15
## 41       2     2       4      14     2    13
## 42       2     2       4      14     3    15
## 43       2     2       5      15     1    16
## 44       2     2       5      15     2    13
## 45       2     2       5      15     3    15
```

Preliminary analysis

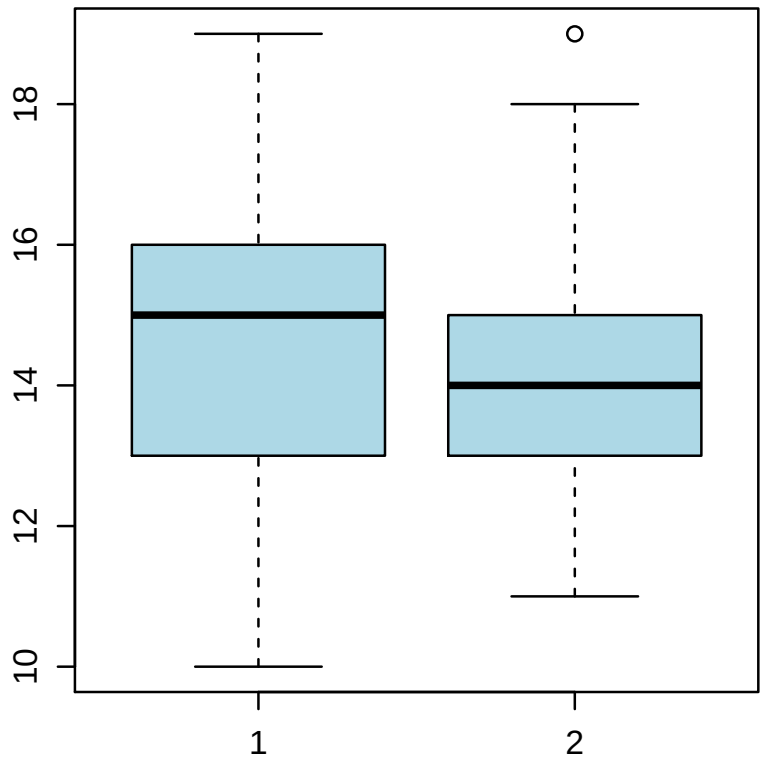

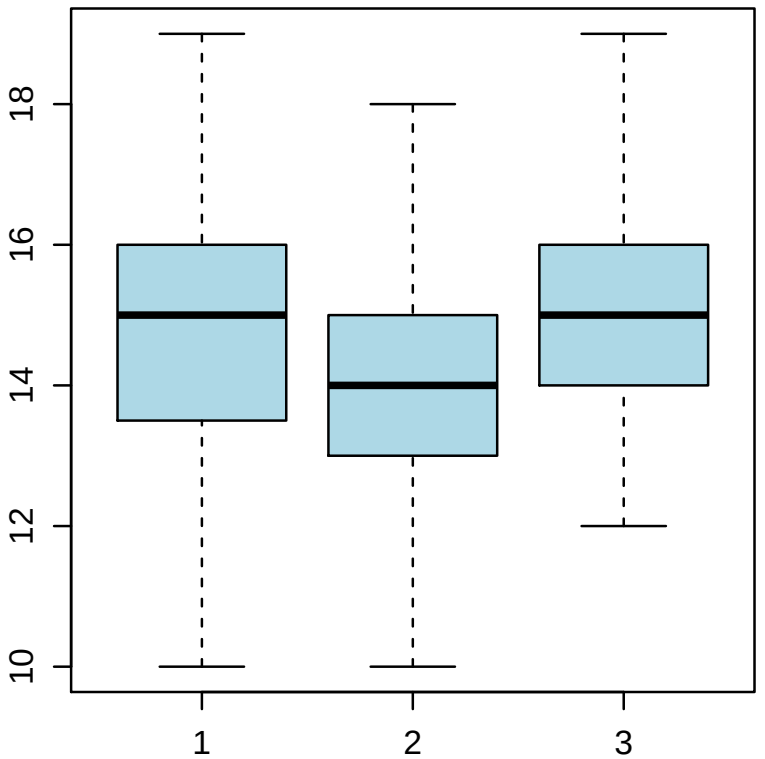

Preliminary analysis

```
anova(lm(score ~ as.factor(itype) + as.factor(etype) ,data=etest)  )

## Analysis of Variance Table
##
## Response: score
##                   Df Sum Sq Mean Sq F value   Pr(>F)
## as.factor(itype)   1  10.00 10.0042  3.2948 0.070770 .
## as.factor(etype)   2  37.07 18.5375  6.1052 0.002599 **
## Residuals        236 716.58  3.0364
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Controlling for heterogeneity

What if observations within a subject are correlated?

```
anova(lm(score ~ as.factor(sub.ses) + as.factor(itype) + as.factor(etype) ,data=etest))

## Analysis of Variance Table
##
## Response: score
##                      Df Sum Sq Mean Sq F value    Pr(>F)
## as.factor(sub.ses)  79 531.66  6.7299   5.455 < 2.2e-16 ***
## as.factor(etype)     2  37.07 18.5375  15.026 1.062e-06 ***
## Residuals          158 194.93  1.2337
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Problem:

- Subjects assigned to only one `itype`.
- Accounting for all subject variation leaves none for `itype` to explain.

Controlling for heterogeneity

What if observations within a session are correlated?

```
anova(lm(score ~ as.factor(session) + as.factor(itype) + as.factor(etype) ,data=etest))

## Analysis of Variance Table
##
## Response: score
##                      Df Sum Sq Mean Sq F value    Pr(>F)
## as.factor(session)    7 330.63  47.233  27.436 < 2.2e-16 ***
## as.factor(etype)      2  37.08  18.538  10.768 3.386e-05 ***
## Residuals           230 395.96   1.722
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Problem:

- Each day has only one `itype`.
- Accounting for all session variation leaves none for `itype` to explain .

A three-level model

Index the data as follows:

- $k = 1, \ldots, m_1 = 8$ indexes sessions;
- $j = 1, \ldots, m_2 = 10$ indexes subjects within a session;
- $i = 1, \ldots, n = 3$ indexes observations within a subject.

A simple multilevel model:

$$y_{i,j,k} = \mu + a_k + b_{j,k} + \texttt{itype}_k + \texttt{etype}_{i,j,k} + \epsilon_{i,j,k}$$
$$\{a_k\} \sim \text{i.i.d. normal}(0, \tau_1^2)$$
$$\{b_{j,k}\} \sim \text{i.i.d. normal}(0, \tau_2^2)$$
$$\{\epsilon_{i,j,k}\} \sim \text{i.i.d. normal}(0, \sigma^2)$$

- $\{a_k\}$ describes across-session heterogeneity;
- $\{b_{j,k}\}$ describes across-subject heterogeneity;
- $\{\epsilon_{i,j,k}\}$ describes within-subject heterogeneity.

As you might guess, τ_1^2 and τ_2^2 relate to within-session and within-subject correlation, repsectively.

Nested models in lme4

```
fit1<-lmer(score ~ as.factor(itype) + as.factor(etype) + (1|session) + (1|sub.ses) , data=e
```

```
summary(fit1)
```

```
## Linear mixed model fit by REML ['lmerMod']
## Formula: score ~ as.factor(itype) + as.factor(etype) + (1 | session) +
##     (1 | sub.ses)
##    Data: etest
##
## REML criterion at convergence: 818
##
## Scaled residuals:
##      Min       1Q   Median       3Q      Max
## -2.16205 -0.60066  0.04381  0.57415  2.73733
##
## Random effects:
##  Groups   Name        Variance Std.Dev.
##  sub.ses  (Intercept) 0.5195   0.7207
##  session  (Intercept) 1.6882   1.2993
##  Residual             1.2337   1.1107
## Number of obs: 240, groups:  sub.ses, 80; session, 8
##
## Fixed effects:
##                   Estimate Std. Error t value
## (Intercept)        14.7542     0.6750  21.859
## as.factor(itype)2  -0.4083     0.9437  -0.433
## as.factor(etype)2  -0.5000     0.1756  -2.847
## as.factor(etype)3   0.4625     0.1756   2.634
##
```

Alternative formulation

```
fit2<-lmer(score ~ as.factor(itype) + as.factor(etype) + (1|session/subject) , data=etest)
```

```
summary(fit2)
```

```
## Linear mixed model fit by REML ['lmerMod']
## Formula:
## score ~ as.factor(itype) + as.factor(etype) + (1 | session/subject)
##    Data: etest
##
## REML criterion at convergence: 818
##
## Scaled residuals:
##      Min       1Q   Median       3Q      Max
## -2.16205 -0.60066  0.04381  0.57415  2.73733
##
## Random effects:
##  Groups          Name        Variance Std.Dev.
##  subject:session (Intercept) 0.5195   0.7207
##  session         (Intercept) 1.6882   1.2993
##  Residual                    1.2337   1.1107
## Number of obs: 240, groups:  subject:session, 80; session, 8
##
## Fixed effects:
##                   Estimate Std. Error t value
## (Intercept)        14.7542     0.6750  21.859
## as.factor(itype)2  -0.4083     0.9437  -0.433
## as.factor(etype)2  -0.5000     0.1756  -2.847
## as.factor(etype)3   0.4625     0.1756   2.634
##
```

Nested index sets

```
BIC(fit1)
```

```
## [1] 856.328
```

```
BIC(fit2)
```

```
## [1] 856.328
```

The term `(1|session/subject)` here is a convenience feature.

Many datasets don't distinguish between (person 1,day 1) and (person 1,day 2).

If these people are different, you need to tell the software somehow:

- code them manually to be different;
- use the nesting feature in `lmer`

Beyond random intercepts

Do the effects of `itype`, `etype` vary across subjects or sessions?

Do we have enough data to detect such variance?

`itype`

- `itype` is a macro variable from the perspective of `session`
- `itype` is a macro variable from the perspective of `subject`

We do not have the data to detect variance in the effects of `itype` across either grouping factor.

`etype`

- `etype` is a micro variable from the perspective of `session`
- `etype` is a micro variable from the perspective of `subject`

We can estimate variance in the effects of `etype` across sessions.

We only have one rep per `etype` per `subject` - can't estimate variance in the effects of `etype` across subjects.

Some fits

```
fit<-lmer( score ~ as.factor(itype) + as.factor(etype) +
          ( as.factor(etype) | session )  +
          ( 1|sub.ses ) , data=etest )
```

```
drop1(fit,test="Chisq")

## Single term deletions
##
## Model:
## score ~ as.factor(itype) + as.factor(etype) + (as.factor(etype) |
##     session) + (1 | sub.ses)
##                  Df    AIC     LRT  Pr(Chi)
## <none>              837.35
## as.factor(itype)  1 835.61  0.2532 0.614817
## as.factor(etype)  2 846.71 13.3519 0.001261 **
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Do exam effects vary across sessions?

```
fit1<-lmer( score ~ as.factor(etype) +
          ( as.factor(etype) | session )  +
          ( 1|sub.ses ) , data=etest )

fit0<-lmer( score ~ as.factor(etype) +
          ( 1| session )  +
          ( 1|sub.ses ) , data=etest )
```

```
BIC(fit1)

## [1] 876.9898

BIC(fit0)

## [1] 852.7055
```

Is there excess group heterogeneity?

```
fit00<-lm( score ~ as.factor(etype) , data=etest)

fit10<-lmer( score ~ as.factor(etype) + ( 1| session ) , data=etest)

fit01<-lmer( score ~ as.factor(etype) + ( 1| sub.ses ) , data=etest)

fit11<-lmer( score ~ as.factor(etype) + ( 1| session ) +( 1| sub.ses ),data=etest)
```

```
BIC(fit00)

## [1] 968.8658

BIC(fit10)

## [1] 865.0546

BIC(fit01)

## [1] 896.9276

BIC(fit11)

## [1] 852.7055
```

Non nested grouping factors

Agricultural field trial:

Experimental material: $m_1 = 3$ plots of land for $m_2 = 6$ years.

Outcome: Crop yield

Treatments/explanatory variables:

- fertilizer type (`fert1`,`fert1`)
- seed variety (`seed1`,`seed2`)

Experimental design:

- `fert1` used in all plots in years 1-3, `fert2` used in all plots in years 4-6.
- each seed type assigned to two of four subplots, in each plot and year.

Exercise: Draw the design on the board.

Data

```
crops[1:25,]

##    yield year plot fert seed
## 1  12.16    1    1    0    0
## 2  12.64    1    1    0    0
## 3  12.89    1    1    0    1
## 4  11.57    1    1    0    1
## 5  11.15    1    2    0    0
## 6  10.35    1    2    0    0
## 7  10.53    1    2    0    1
## 8  10.90    1    2    0    1
## 9  12.43    1    3    0    0
## 10 10.27    1    3    0    0
## 11  9.18    1    3    0    1
## 12 10.92    1    3    0    1
## 13 12.04    2    1    0    0
## 14 10.88    2    1    0    0
## 15 11.02    2    1    0    1
## 16 10.22    2    1    0    1
## 17  8.64    2    2    0    0
## 18 10.35    2    2    0    0
## 19 10.45    2    2    0    1
## 20  9.43    2    2    0    1
## 21  9.88    2    3    0    0
## 22  7.02    2    3    0    0
## 23 11.58    2    3    0    1
## 24 10.17    2    3    0    1
## 25 11.59    3    1    0    0
```

Exploratory analysis

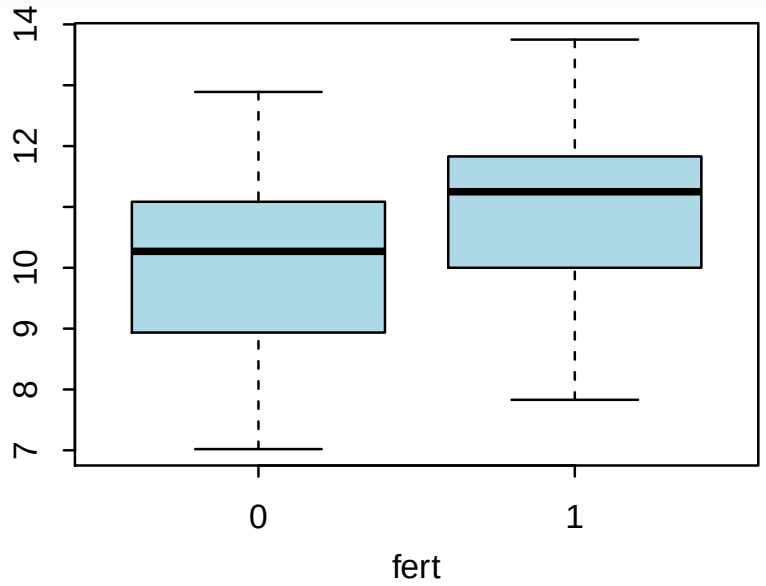

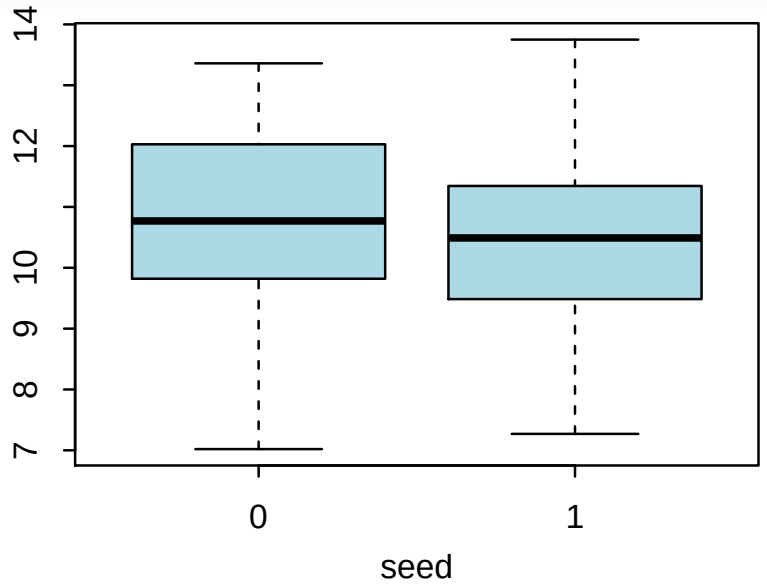


```
summary(lm(yield~fert+seed,data=crops))

##
## Call:
## lm(formula = yield ~ fert + seed, data = crops)
##
## Residuals:
##     Min      1Q  Median      3Q     Max
## -3.2589 -1.1164  0.1778  0.9261  2.9111
##
## Coefficients:
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept)  10.2789     0.2934  35.038  < 2e-16 ***
## fert          0.9067     0.3387   2.677  0.00929 **
## seed         -0.3000     0.3387  -0.886  0.37890
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
```

Exploratory analysis

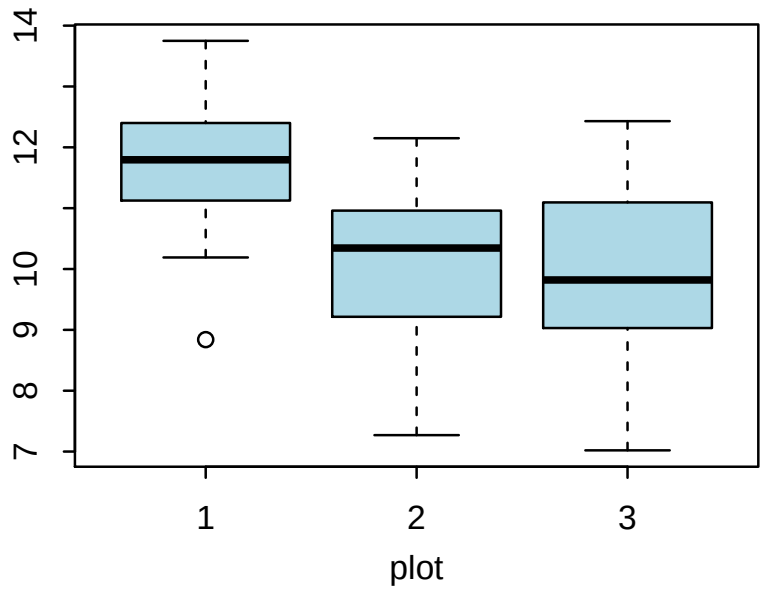


```
anova(lm(yield~as.factor(plot) ,data=crops))

## Analysis of Variance Table
##
## Response: yield
##                 Df Sum Sq Mean Sq F value    Pr(>F)
## as.factor(plot)  2  48.59 24.2951  15.192 3.411e-06 ***
## Residuals       69 110.34  1.5992
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Contolling for plot variaion

```
anova(lm(yield~ as.factor(plot) + fert + seed,data=crops))

## Analysis of Variance Table
##
## Response: yield
##                  Df Sum Sq Mean Sq F value   Pr(>F)
## as.factor(plot)   2 48.590 24.2951 17.3300 8.59e-07 ***
## fert              1 14.797 14.7968 10.5547 0.001813 **
## seed              1  1.620  1.6200  1.1556 0.286243
## Residuals        67 93.928  1.4019
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

The `fert` p-value assumes we have $3 \times 3 \times 4 = 36$ independent observations for both levels of `fert`.

Replication at different levels

- How many times was the fertilizer type obtained and applied?
- Ignoring `plot` and `seed`, how confident are we in the effects of `fert`?
- Could anything else cause the effects we are attributing to `fert`?

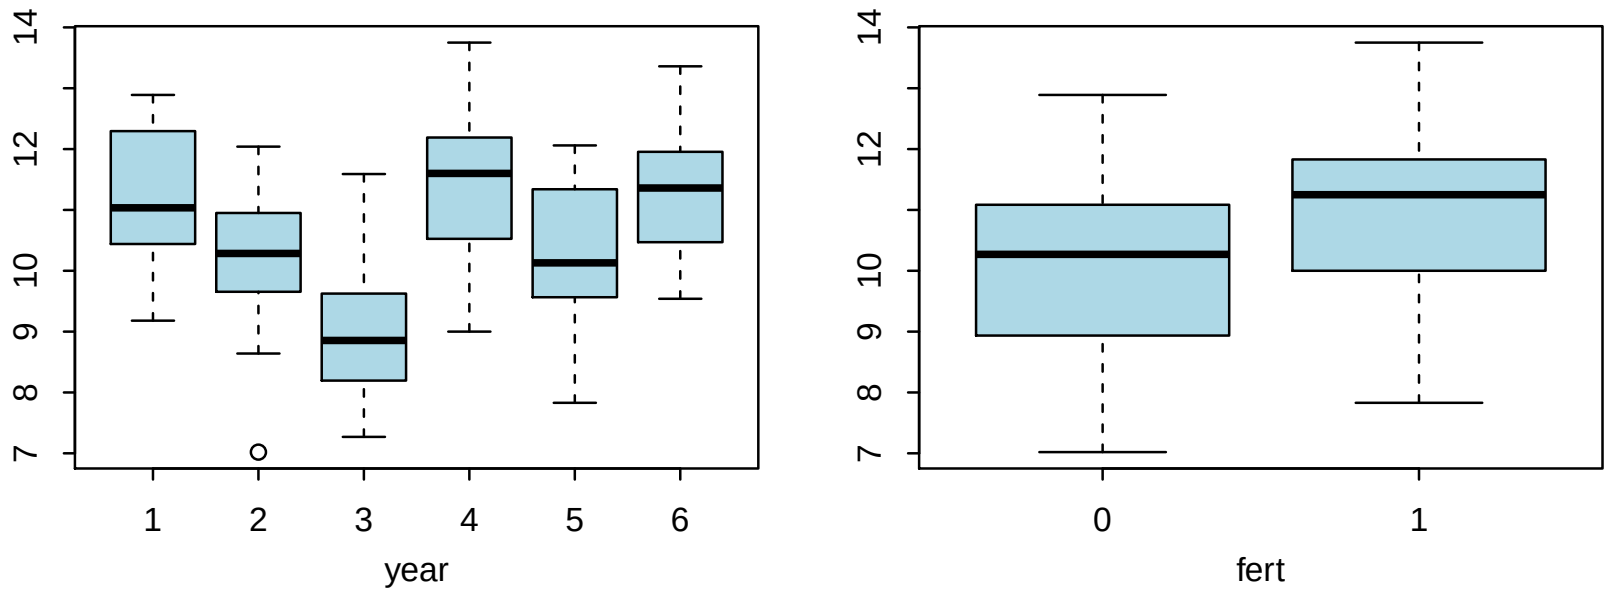

The "sample size" for `fert` is more like $m_1 = 6$, with 3 obs per level.

This issue is common in multilevel experiments (e.g. *split-plot* designs). See the notes for more details.

Accounting for year effects

```
anova(lm(yield~ as.factor(year) + as.factor(plot) + fert + seed,data=crops))

## Analysis of Variance Table
##
## Response: yield
##                 Df Sum Sq Mean Sq F value    Pr(>F)
## as.factor(year)  5 56.471 11.2942 13.6169 5.214e-09 ***
## as.factor(plot)  2 48.590 24.2951 29.2915 1.013e-09 ***
## seed             1  1.620  1.6200  1.9532    0.1671
## Residuals       63 52.254  0.8294
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Accounting for all year-to-year variability leaves none for `fert`.

Multilevel approach

$$\begin{aligned}
\texttt{yield}_{i,j,k} &= \mu + a_j + b_k + \beta_1 \times \texttt{fert}_k + \beta_2 \times \texttt{seed}_{i,j,k} + \epsilon_{i,j,k} \\
\{a_j\} &\sim \ \textit{iid}\ N(0, \tau_a^2) \\
\{b_k\} &\sim \ \textit{iid}\ N(0, \tau_b^2) \\
\{\epsilon_{i,j,k}\} &\sim \ \textit{iid}\ N(0, \sigma^2)
\end{aligned}$$

- $\{a_j\}$ represents heterogeneity across plots;
- $\{b_k\}$ represents heterogeneity across years;
- $\{\epsilon_{i,j,k}\}$ represents heterogeneity within years and plots.

Fitting with `lmer`

```
fit<-lmer( yield ~ fert + seed + (1|year) + (1|plot),  data=crops,REML=FALSE)
BIC(fit)

## [1] 236.3216

summary(fit)$coef

##               Estimate Std. Error   t value
## (Intercept) 10.2788889  0.6827211 15.055767
## fert         0.9066667  0.6600099  1.373717
## seed        -0.3000000  0.2130316 -1.408242
```

Other things to investigate:

- heterogeneity of `seed` effects across plots and years: `(seed|plots) + (seed|years)`
- heterogeneity of `fert` effects across plots, but not years. `(fert|plots)`