Outline

Two-factor design

Full factorial: all combinations of factor levels are observed.
Example: factor A (frog species) with $a = 5$ levels, factor B (carbaryl presence) with $b = 2$ levels. Full factorial balanced design with $r = 3$ replicates: total of $5 * 2 * 3 = 30$ observations.

Model

response = factorA + factorB + factorA:factorB + error

$$Y_i = \mu + \alpha_{j[i]} + \beta_{k[i]} + \gamma_{j[i],k[i]} + e_i \quad \text{with } e_i \sim \text{ i.i.d } \mathcal{N}(0, \sigma_e^2)$$

- μ is a population mean.
- α_j: main effect of A (frog species), compared to the mean across all species. Constrained to $\sum_{j=1}^a \alpha_j = 0$.
- β_k: main effect of B (carbaryl). Constraint: $\sum_{k=1}^b \beta_k = 0$.
- γ_{jk}: interactive effect of A and B. Constraints: $\sum_{j=1}^a \gamma_{jk} = 0$ for each carbaryl level k and $\sum_{k=1}^b \gamma_{jk} = 0$ for each frog species j.

Two factor design: ANOVA table

		$\mathbb{E}(\mathrm{MS})$	
Source	df	A, B fixed	A, B random
A	$a-1$	$\sigma_e^2 + br\,\frac{\sum_{j=1}^{a}\alpha_j^2}{a-1}$	$\sigma_e^2 + r\,\sigma_\gamma^2 + br\sigma_\alpha^2$
B	$b-1$	$\sigma_e^2 + ar\,\frac{\sum_{j=1}^{b}\beta_k^2}{b-1}$	$\sigma_e^2 + r\,\sigma_\gamma^2 + ar\sigma_\beta^2$
A:B	$(a-1)(b-1)$	$\sigma_e^2 + r\,\frac{\sum_{j,k}\gamma_{jk}^2}{(a-1)(b-1)}$	$\sigma_e^2 + r\sigma_\gamma^2$
Error	$ab(r-1)$	σ_e^2	σ_e^2
Total	$abr-1$		

- We may consider A and/or B as random. If so, their interaction must be considered random too, with $\gamma_{jk} \sim \mathcal{N}(0, \sigma_\gamma^2)$.
- The F test depends if A and B are considered fixed, or if one of them is considered random.

Plant alkaloid content

in wild and cultivated plants, different populations and seasons:

```
> dat = read.table("alkaloid.txt", header=T)
> with(dat, table(type, month, pop))
, , pop = A           , , pop = B           , , pop = C
             month                 month                 month
type         May Oct   type         May Oct   type         May Oct
  cultivated  9  10      cultivated 10  10      cultivated 10  10
  wild       10  10      wild       10  10      wild       10   8
```

```
> datB = subset(dat, pop=="B")
> str(datB)
’data.frame’:   40 obs. of  4 variables:
 $ pop     : Factor w/ 3 levels "A","B","C": 2 2 2 2 2 2 2 2 2 2 ...
 $ month   : Factor w/ 2 levels "May","October": 1 1 1 1 1 1 1 1 1 1 .
 $ type    : Factor w/ 2 levels "cultivated","wild": 2 2 2 2 2 2 2 2 2
 $ alkaloid: int  562 767 581 625 497 870 515 745 538 625 ...
```

We will focus on population (or ’block’) B, no missing values: full factorial design with $r = 10$ replicates.

Same type I and III F-tests

Because of the balance, the type I and type III SS are equal:

```
> fit1 = lm(sqrt(alkaloid) ~ type*month, data=datB)
> fit2 = lm(sqrt(alkaloid) ~ month*type, data=datB)

> anova(fit1)
           Df  Sum Sq Mean Sq F value    Pr(>F)
type        1   1.175   1.175  0.2447   0.62386
month       1 193.320 193.320 40.2661 2.406e-07 ***
type:month  1  25.064  25.064  5.2204   0.02832 *
Residuals  36 172.839   4.801

> anova(fit2)
           Df  Sum Sq Mean Sq F value    Pr(>F)
month       1 193.320 193.320 40.2661 2.406e-07 ***
type        1   1.175   1.175  0.2447   0.62386
month:type  1  25.064  25.064  5.2204   0.02832 *
Residuals  36 172.839   4.801

> drop1(fit2, test="F")
Single term deletions
           Df Sum of Sq    RSS    AIC
<none>                  172.84 66.539
month:type  1    25.064 197.90 69.956  5.2204 0.02832 *
```

Meaning of 'main effects'

F-test for the main effects of A (say, type):

H_0: all $\alpha_j = 0$ versus H_A: at least one $\alpha_j \neq 0$.

When there are interactive effects type:month, α_j is not the effect of "wild" type in October, nor in May. It is the average effect of "wild", taken across both months.

Main effect

$\alpha_j =$ predicted mean Y in treatment j across all values of the other factor(s) $-$ predicted mean $\hat{Y}$ across all treatment levels and all other factor levels.

This value might not be meaningful. It might, or might not. Accordingly, the associated f-test might not be meaningful.

Meaning of main effects

Why does the F-test for the main effect of ’type’ ($p = .62$) disagree with the t-test for the coefficient ’typeWild’ ($p = .057$)?

```
> anova(fit1)
           Df  Sum Sq Mean Sq F value    Pr(>F)    
type        1   1.175   1.175  0.2447   0.62386    
month       1 193.320 193.320 40.2661 2.406e-07 ***
type:month  1  25.064  25.064  5.2204   0.02832 *  
Residuals  36 172.839   4.801                      

> summary(fit1)
Coefficients:
                      Estimate Std. Error t value Pr(>|t|)    
(Intercept)            23.1219     0.6929  33.370  < 2e-16 ***
typewild                1.9259     0.9799   1.965  0.05712 .  
monthOctober           -2.8137     0.9799  -2.871  0.00681 ** 
typewild:monthOctober  -3.1663     1.3858  -2.285  0.02832 *  

with(datB, interaction.plot(month, type, sqrt(alkaloid)))
```

Meaning of 'main effects'

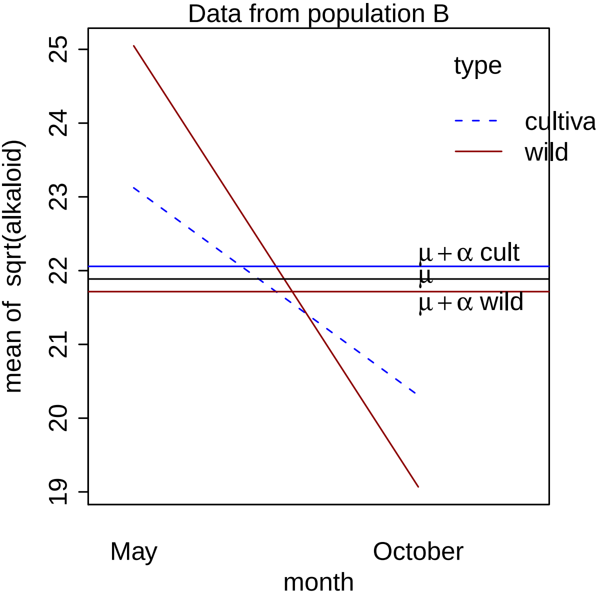

Meaning of ’main effects’

The F-test for the main effect of ’type’ corresponds to testing the coefficient ’type1’ when we use the sum contrast:

```
> options(contrasts=c("contr.sum", "contr.poly"))
> fit.sum.contrasts = lm(sqrt(alkaloid) ˜ type*month, data=datB)
> summary(fit.sum.contrasts)
Coefficients:
             Estimate Std. Error t value Pr(>|t|)
(Intercept)   21.8864     0.3464  63.174  < 2e-16 ***
type1         -0.1714     0.3464  -0.495   0.6239
month1         2.1984     0.3464   6.346 2.41e-07 ***
type1:month1  -0.7916     0.3464  -2.285   0.0283 *

mu        = 21.8864
alphaWild = -0.1714
alphaCult = +0.1714
with(datB, interaction.plot(month, type, sqrt(alkaloid)))
abline(h=mu);
abline(h=mu+alphaWild, col="darkred")
abline(h=mu+alphaCult, col="blue")
```

Effect of a few missing data

Analysis of the cultivated populations: factorial design but 1 missing obs.

```
> datCult = subset(dat, type=="cultivated")
> with(datCult, table(pop, month))
   month
pop May October
  A   9      10
  B  10      10
  C  10      10

> fit.Cult1 = lm(sqrt(alkaloid) ~ pop*month, data=datCult)
> fit.Cult2 = lm(sqrt(alkaloid) ~ month*pop, data=datCult)
```

Effect of a few missing data

Type I and type III sums of squares are not longer exactly equal:

```
> anova(fit.Cult1)
          Df  Sum Sq Mean Sq  F value    Pr(>F)    
pop        2 1368.52  684.26 173.3477 < 2.2e-16 ***
month      1  357.12  357.12  90.4703 4.704e-13 ***
pop:month  2   34.02   17.01   4.3097   0.01844 *  
Residuals 53  209.21    3.95                       

> anova(fit.Cult2)
          Df  Sum Sq Mean Sq  F value    Pr(>F)    
month      1  390.55  390.55  98.9409 1.013e-13 ***
pop        2 1335.09  667.54 169.1124 < 2.2e-16 ***
month:pop  2   34.02   17.01   4.3097   0.01844 *  
Residuals 53  209.21    3.95                       

> drop1(fit.Cult1, test="F")
Single term deletions
          Df Sum of Sq    RSS    AIC F value   Pr(F)  
<none>                 209.21 86.682                  
pop:month  2    34.024 243.23 91.572  4.3097 0.01844 *
```

Meaning of ’main effects’

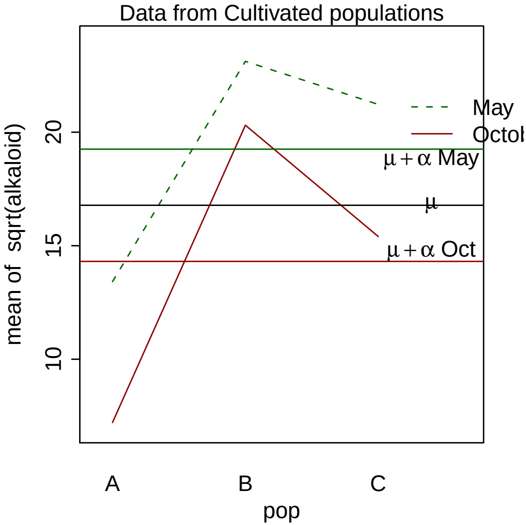

Main effects are meaningful here: the difference between May and October is consistent.

Meaning of ’main effects’

```
> summary(fit.Cult2)
Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept)  16.7818     0.2589  64.830  < 2e-16 ***
month1        2.4726     0.2589   9.552 4.08e-13 ***
pop1         -6.4660     0.3694 -17.504  < 2e-16 ***
pop2          4.9333     0.3644  13.538  < 2e-16 ***
month1:pop1   0.6295     0.3694   1.704  0.09421 .
month1:pop2  -1.0658     0.3644  -2.925  0.00507 **

> mu       = 16.7818
> alphaMay =  2.4726
> alphaOct = -2.4726
> with(datCult, interaction.plot(pop, month, sqrt(alkaloid)))
> abline(h=mu);
> abline(h=mu+alphaMay)
> abline(h=mu+alphaOct)
```

Outline

1. Two-factor design
    - Design and Model
    - ANOVA table and F test
    - Meaning of Main Effects
2. Split-plot design
    - Design and Model, CRD at whole-plot level
    - ANOVA table and F test
    - Split plot with RCBD at whole-plot level

Split-plot design

- Two of more factors, say A (irrigation) and C (barley variety)
- A blocking factor B, say large plots (’pan’)

- A single level of factor A (say) is applied to each block, making it difficult to compare the levels of A.
- All levels of C are observed in each block.

Split-plot design

Example: each irrigation treatment is applied to a large plot. The 4 varieties are planted in rows, randomized.

Plot 1: Irrigation 2	Plot 2: Irrigation 3	Plot 3: Irrigation 1
row 1: V4	row 1: V3	row 1: V4
row 2: V2	row 2: V1	row 2: V1
row 3: V1	row 3: V2	row 3: V3
row 4: V3	row 4: V4	row 4: V2

Here we cannot compare irrigation treatments, but can compare varieties. Why?

Split-plot design

6 large plots are used, randomly assigned to irrigation trts:

Plot 1: Irrigation 3	Plot 2: Irrigation 1	Plot 3: Irrigation 3
row 1: V3	row 1: V1	row 1: V4
row 2: V1	row 2: V4	row 2: V3
row 3: V4	row 3: V3	row 3: V1
row 4: V2	row 4: V2	row 4: V2

Plot 4: Irrigation 2	Plot 5: Irrigation 2	Plot 6: Irrigation 1
row 1: V2	row 1: V2	row 1: V4
row 2: V1	row 2: V4	row 2: V3
row 3: V3	row 3: V1	row 3: V1
row 4: V4	row 4: V3	row 4: V2

Now we can compare both the soil types and the varieties:

- 2 EUs (large plots) per irrigation treatment.
- 1 EU (row) per variety within each large plot.

Split-plot Model

Model for split plot design

response ∼ factor A * factor C + whole-plot (random) + error

$$Y_i = \mu + \alpha_j + v_k + (\alpha v)_{jk} + \gamma_l + e_i \quad \text{with } e_i \sim \mathcal{N}(0, \sigma_e^2)$$

- μ is a population mean across all treatments.
- α_j is a (fixed) main effect of A (irrigation) constrained to $\sum_{j=1}^{a} \alpha_j = 0$.
- v_k is a (fixed) main effect of C (variety) constrained to $\sum_{k=1}^{c} v_k = 0$.
- $(\alpha v)_{jk}$ is a (fixed) interaction effect of A and C constrained to $\sum_{j=1}^{a} (\alpha v)_{jk} = 0$ for each variety k and $\sum_{k=1}^{c} (\alpha v)_{jk} = 0$ for each irrigation treatment j.
- $\gamma_l \sim$ iid $\mathcal{N}(0, \sigma_\gamma^2)$ represents the large plot effect, nested in irrigation.

Split-plot design: ANOVA table

a: # levels for A
b: # whole-plots for each level of A
c: # levels for B, each observed once in each whole-plot.

Source	df	E(MS)
A	$a-1$	$\sigma_e^2 + c\sigma_\gamma^2 + bc\,\frac{\sum_j \alpha_j^2}{a-1}$
WPErr, or B	$a(b-1)$	$\sigma_e^2 + c\sigma_\gamma^2$
C	$c-1$	$\sigma_e^2 + ab\,\frac{\sum_k v_k^2}{c-1}$
AC	$(a-1)(c-1)$	$\sigma_e^2 + b\,\frac{\sum_{j,k}(\alpha v)_{jk}^2}{(a-1)(c-1)}$
SPErr	$a(b-1)(c-1)$	σ_e^2
Total	$abc-1$	

WPErr: Whole plot error –here at the large plot level
SPErr: Subplot error –here at the row level.

Split-plot design: f-tests

- Under H_0: $\alpha_j = 0$ for all j, i.e. no ‘main’ irrigation effect:

$$F = \frac{\text{MSA}}{\text{MSWPErr}} \sim F_{a-1,a(b-1)}$$

- Under H_0: $v_k = 0$ for all k, i.e. no ‘main’ variety effect:

$$F = \frac{\text{MSC}}{\text{MSSPErr}} \sim F_{c-1,a(b-1)(c-1)}$$

- Under H_0: $(\alpha v)_{jk} = 0$ for all j, k, i.e. additivity of irrigation and variety effects:

$$F = \frac{\text{MSAC}}{\text{MSSPErr}} \sim F_{(a-1)(c-1),a(b-1)(c-1)}$$

Oats example and R commands

`block`: fields grouped into 6 homogeneous blocks.
`field`: whole plots. 3 fields per block.
`variety`: of oats (Golden Rain, Marvellous and Victory). Assigned to fields randomly.
`nitro`: amount of nitrogen added to subplots. 4 subplots per field, 4 levels of nitrogen, randomized within each field.
`yield`: response.

We will first ignore the blocks, to illustrate the simplest split-splot design.

Oats example with `lmer`

```
> oats = read.table("oats.txt", header=T)
> oats$nitro    = ordered(oats$nitro) # factor with ordered levels
> oats$field123 = factor(oats$field123)
> oats$field    = factor(oats$field)
> oats
   block field field123    variety nitro yield
1      I     1        1 Marvellous     0   105
2      I     1        1 Marvellous   0.2   140
3      I     1        1 Marvellous   0.4   118
4      I     1        1 Marvellous   0.6   156
5      I     2        2    Victory     0   111
6      I     2        2    Victory   0.2   130
...
11     I     3        3 GoldenRain   0.4   161
12     I     3        3 GoldenRain   0.6   141
13    II     4        1 Marvellous     0    63
14    II     4        1 Marvellous   0.2    70
15    II     4        1 Marvellous   0.4   109
...
68    VI    17        2 GoldenRain   0.6   104
69    VI    18        3 Marvellous     0    97
70    VI    18        3 Marvellous   0.2    99
71    VI    18        3 Marvellous   0.4   119
72    VI    18        3 Marvellous   0.6   121
```

Oats example with `lmer`

```
> library(lme4)
> fit.lmer1 = lmer(yield ˜ variety * nitro + (1|field), data=oats)
> anova(fit.lmer1)
Analysis of Variance Table
              Df  Sum Sq Mean Sq F value
variety        2   216.6   108.3  0.6114
nitro          3 20020.5  6673.5 37.6856
variety:nitro  6   321.8    53.6  0.3028

> pf(0.6114, df1=2, df2=15, lower.tail=F)  # 15=3*(6-1)
[1] 0.5555                                 # no main variety effects

> pf(37.6856,df1=3, df2=45, lower.tail=F)  # 45=3*(6-1)*(4-1)
[1] 2.457759e-12                           # main nitro. effects

> pf(0.3028, df1=6, df2=45, lower.tail=F)
[1] 0.9322108                              # no interaction effects
```

Danger! coding for whole plot

```
> library(lme4)
> fit.wrong = lmer(yield ~ variety * nitro + (1|field123), data=oats)
> anova(fit.wrong)
Analysis of Variance Table
              Df  Sum Sq Mean Sq F value
variety        2  2004.3  1002.1  2.0475
nitro          3 20020.5  6673.5 13.6348
variety:nitro  6   321.7    53.6  0.1096
```

Ex: use `field` with ID from 1 through 18
or: use `field123` with ID from 1 through 3 AND indicate
nesting within ‘blocks’:

```
> library(lme4)
> fit.okay = lmer(yield ~ variety * nitro + (1|block:field123), data=o
> anova(fit.okay)
Analysis of Variance Table
              Df  Sum Sq Mean Sq F value
variety        2   216.6   108.3  0.6114
nitro          3 20020.5  6673.5 37.6856
variety:nitro  6   321.7    53.6  0.3028
```

Same analysis with aov

```
> fit.aov1 = aov(yield ~ variety*nitro + Error(field), data=oats)
> summary(fit.aov1)

Error: field
          Df  Sum Sq Mean Sq F value Pr(>F)
variety    2  1786.4   893.2  0.6121 0.5552
Residuals 15 21888.6  1459.2

Error: Within
              Df  Sum Sq Mean Sq F value    Pr(>F)
nitro          3 20020.5  6673.5 37.6856 2.458e-12 ***
variety:nitro  6   321.8    53.6  0.3028    0.9322
Residuals     45  7968.8   177.1
```

Same results from f-tests.

Again, make sure to use the `field` column with IDs from 1 to 18.

Danger! of not using random effects for whole plots

```
> fit.wrong = lm(yield ~ variety*nitro + field, data=oats)
> anova(fit.wrong)
              Df  Sum Sq Mean Sq F value    Pr(>F)
variety        2  1786.4   893.2  5.0438   0.01056 *
nitro          3 20020.5  6673.5 37.6856 2.458e-12 ***
field         15 21888.6  1459.2  8.2404 1.609e-08 ***
variety:nitro  6   321.8    53.6  0.3028   0.93220
Residuals     45  7968.7   177.1
```

SS are correct, but the f-value and f-test for variety are wrong.

Hint to catch the mistake: some coefficient were not estimated.

```
> summary(fit.wrong)
Coefficients: (2 not defined because of singularities)
                 Estimate Std. Error t value Pr(>|t|)
(Intercept)      103.9722     1.5683  66.297  < 2e-16 ***
...
field15           -2.2500     9.4097  -0.239 0.812101
field16                NA         NA      NA       NA
field17                NA         NA      NA       NA
variety1:nitro.L   0.7454     4.4358   0.168 0.867310
```

Split plot with RCBD at whole-plot level

Often, replication happens in blocks. Ex: in 2 different fields. Randomly assign factor A levels (irrigation) to whole plots within each block.

Plot 1: Irrigation 3	Plot 2: Irrigation 1	Plot 3: Irrigation 2
row 1: V3	row 1: V1	row 1: V4
row 2: V1	row 2: V4	row 2: V3
row 3: V4	row 3: V3	row 3: V1
row 4: V2	row 4: V2	row 4: V2

Plot 4: Irrigation 2	Plot 5: Irrigation 3	Plot 6: Irrigation 1
row 1: V2	row 1: V2	row 1: V4
row 2: V1	row 2: V4	row 2: V3
row 3: V3	row 3: V1	row 3: V1
row 4: V4	row 4: V3	row 4: V2

Split-plot Model with RCBD at whole-plot level

response $\sim$ factor A * factor C + block + whole-plot + error

$$Y_i = \mu + \alpha_j + v_k + (\alpha v)_{jk} + \beta_l + \gamma_{lm} + e_i \quad \text{with } e_i \sim \mathcal{N}(0, \sigma_e^2)$$

and random whole-plot effects.

μ, α_j, v_k and $(\alpha v)_{jk}$: as before. Overall population mean, fixed main effects of A (irrigation) and C (variety) and their interaction effect, properly constrained.

β_l: either fixed or random block effects. If fixed, constrained to sum to 0. If random, $\sim$ iid $\mathcal{N}(0, \sigma_\beta^2)$.

$\gamma_{lm} \sim$ iid $\mathcal{N}(0, \sigma_\gamma^2)$: whole-plot effect, nested in irrigation and in blocks. Must be random.

Split-plot Model with RCBD at whole-plot level

a: # levels for A
b: # blocks, one whole-plot for each level of A per block
c: # levels for B, each observed once in each whole-plot.

Source	df	E(MS)
A	$a-1$	$\sigma_e^2 + c\sigma_\gamma^2 + bc\,\frac{\sum_j \alpha_j^2}{a-1}$
B	$b-1$	$\sigma_e^2 + c\sigma_\gamma^2 + ac\,\frac{\sum_j \beta_j^2}{b-1}$
WPErr, or A:B	$(a-1)(b-1)$	$\sigma_e^2 + c\sigma_\gamma^2$
C	$c-1$	$\sigma_e^2 + ab\,\frac{\sum_k v_k^2}{c-1}$
AC	$(a-1)(c-1)$	$\sigma_e^2 + b\,\frac{\sum_{j,k}(\alpha v)_{jk}^2}{(a-1)(c-1)}$
SPErr	$a(b-1)(c-1)$	σ_e^2
Total	$abc-1$	

Oats example with blocks, `lmer`

Block effects as fixed:

```
> fit.lmer2 = lmer(yield ~ block + variety*nitro + (1|field), data=oat
> anova(fit.lmer2)
Analysis of Variance Table
              Df  Sum Sq Mean Sq F value
block          5  4673.2   934.6  5.2780
variety        2   525.9   262.9  1.4848
nitro          3 20020.5  6673.5 37.6856
variety:nitro  6   321.8    53.6  0.3028

> pf(1.4848, df1=2, df2=10, lower.tail=F)  # 10=(3-1)*(6-1)
[1] 0.2725004                              # no main variety effects

> pf(37.6856, df1=3, df2=45, lower.tail=F) # 45=3*(6-1)*(4-1)
[1] 2.457759e-12                           # main nitro. effects

> pf(0.3028, df1=6, df2=45, lower.tail=F)  # 45=3*(6-1)*(4-1)
[1] 0.9322108                              # no variety:nitro interact
```

Oats example with blocks, `lmer`

Block effects as random:

```
> fit.lmer3 = lmer(yield ~ variety*nitro + (1|block/field123), data=oa
> fit.lmer3 = lmer(yield ~ variety*nitro + (1|block/field), data=oats)
> fit.lmer3 = lmer(yield ~ variety*nitro + (1|field) + (1|block), data
> anova(fit.lmer3)
Analysis of Variance Table
              Df  Sum Sq Mean Sq F value
variety        2   526.1   263.0  1.4853
nitro          3 20020.5  6673.5 37.6856
variety:nitro  6   321.8    53.6  0.3028
```

We get the same f-values: doesn’t matter that block effects are random or fixed.

Oats example with `aov`, fixed blocks

```
> fit.aov2 = aov(yield ~ variety*nitro + block + Error(field), data=oa
> summary(fit.aov3)

Error: field
          Df  Sum Sq Mean Sq F value  Pr(>F)
variety    2  1786.4   893.2  1.4853 0.27239
block      5 15875.3  3175.1  5.2801 0.01244 *
Residuals 10  6013.3   601.3

Error: Within
              Df  Sum Sq Mean Sq F value    Pr(>F)
nitro          3 20020.5  6673.5 37.6856 2.458e-12 ***
variety:nitro  6   321.7    53.6  0.3028    0.9322
Residuals     45  7968.7   177.1
```

Make sure to use `field` and not `field123`!
Same f-values, same p-values as with `lmer`.

Oats example with `aov`

Random block effects. `aov` accepts only one `Error` term. Factors within this `Error()` must be nested.

```
> fit.aov3 = aov(yield ~ variety*nitro + Error(block/field), data=oats
Warning message:
In aov(yield ~ variety * nitro + Error(block/field), data = oats) :
  Error() model is singular
```

The ANOVA table would be correct.
To get rid of the warning: `aov` prefers field IDs that are repeated across blocks: 1,2,3,1,2,3 etc.

```
> fit.aov3 = aov(yield ~ variety*nitro + Error(block/field123), data=o
```

Oats example with aov

```
> summary(fit.aov3)

Error: block
          Df  Sum Sq Mean Sq F value Pr(>F)
Residuals  5 15875.3  3175.1

Error: block:field123
          Df Sum Sq Mean Sq F value Pr(>F)
variety    2 1786.4   893.2  1.4853 0.2724
Residuals 10 6013.3   601.3

Error: Within
              Df  Sum Sq Mean Sq F value    Pr(>F)
nitro          3 20020.5  6673.5 37.6856 2.458e-12 ***
variety:nitro  6   321.8    53.6  0.3028    0.9322
Residuals     45  7968.7   177.1
```