Split Plot Designs

Lukas Meier, Seminar für Statistik

What is a Split Plot Design? (Oehlert, 2000, Chapter 16.1)

- A split plot design is a **special case** of a **factorial treatment structure**.
- It is used when some factors are **harder** (or more expensive) to **vary** than others.
- Basically a split plot design consists of two experiments with **different experimental units** of different “size”.
- E.g., in agronomic field trials certain factors require “large” experimental units, whereas other factors can be easily applied to “smaller” plots of land.
- Let us have a look at an example…

Example I: Irrigation and Corn Variety (Oehlert, 2000)

- Consider the following factorial problem:
  - 3 different **irrigation levels**
  - 4 different **corn varieties**
  - Response: **biomass**
  - Available resources: **6 plots of land**
- By definition we **can not** vary the irrigation level on a too small scale.
- We are “forced” to use “large” experimental units for the irrigation level factor.
- Assume that we can use a specific irrigation level on each of the 6 plots.

Example I: Irrigation and Corn Variety

- Randomly assign each irrigation level to 2 of the plots (the so called **whole plots** or **main plots**).

- In every of the plots, randomly assign the 4 different corn varieties to the so called **split plots**.

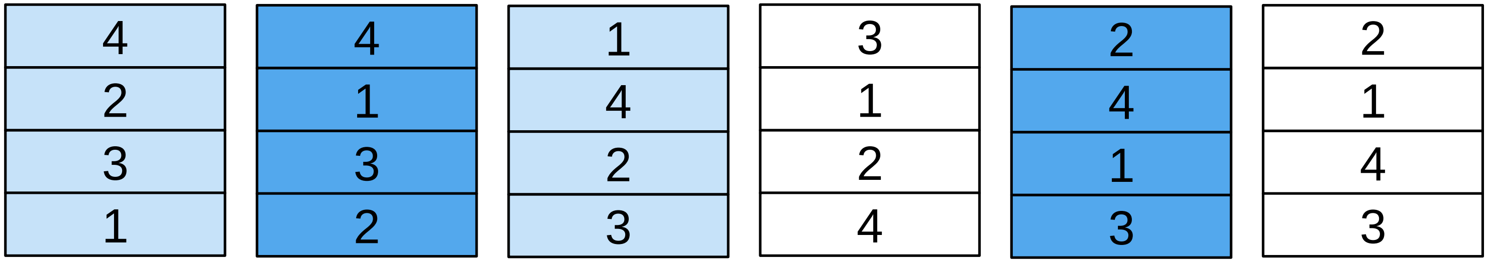

4	4	1	3	2	2
2	1	4	1	4	1
3	3	2	2	1	4
1	2	3	4	3	3

- **Two independent randomizations** are being performed!

- We also call irrigation level the **whole-plot factor** and corn variety the **split-plot factor**.

Example I: Irrigation and Corn Variety

- **Whole plots** (plots of land) are the experimental units for the **whole-plot factor** (irrigation level).
- **Split plots** (subplots of land) are the experimental units for the **split-plot factor**.
- In the split-plot “world”, whole plots act as **blocks**.
- Basically, we are performing **two different experiments in one**:
  - each experiment has its own **randomization**
  - each experiment has its own idea of **experimental unit**

Example I: Irrigation and Corn Variety

- How can we model such kind of data?
- We use a **mixed model** formulation with two **different errors**

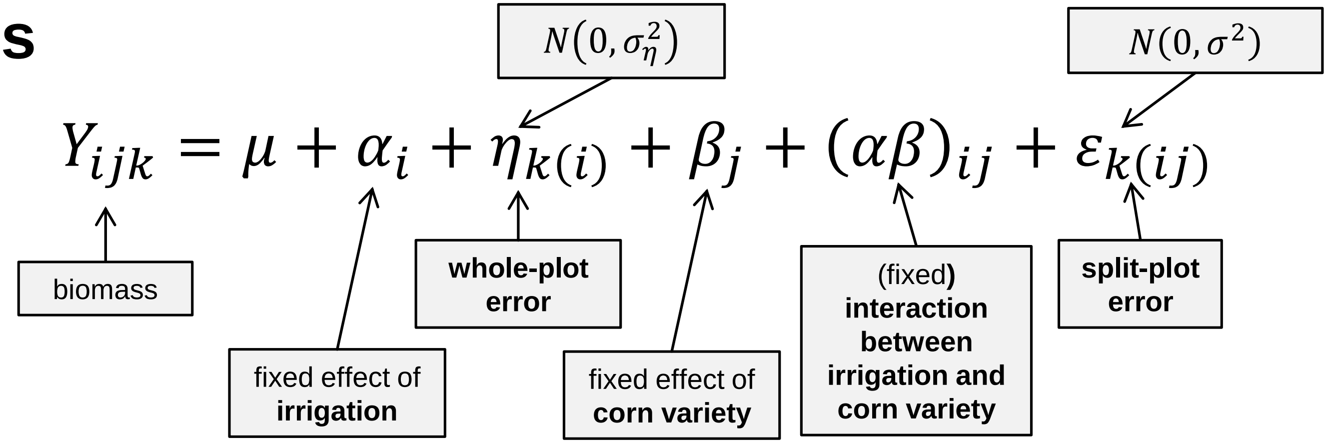

$N(0, \sigma_\eta^2)$

$N(0, \sigma^2)$

$$Y_{ijk} = \mu + \alpha_i + \eta_{k(i)} + \beta_j + (\alpha\beta)_{ij} + \varepsilon_{k(ij)}$$

biomass

fixed effect of **irrigation**

whole-plot error

fixed effect of **corn variety**

(fixed) **interaction between irrigation and corn variety**

split-plot error

- This means: Observations in the same whole plot share the same whole-plot error $\eta_{k(i)}$.
- In R, this model is easily fitted using `lmer` with a **random effect** (better terminology: **error**) of the form `(1|whole.plot)`

Example II: Pianos (Oehlert, 2000)

- **Two piano types** (baby grand / concert grand) from each of **4 manufacturers**.
- 40 music students are divided at random into **8 groups (“panels”)** of 5 students each.
- Two panels are assigned at random to each manufacturer (= 2 panels per manufacturer).
- Each panel goes to the concert hall and hears (blindfolded) the sound of **both pianos** (in random order).
- Response: Average rating of the 5 students in the panel (hence, student is “only” measurement unit here).

Example II: Pianos

- The **whole plots** are the 8 **panels**.
- The **whole-plot factor** is the **manufacturer**.
- The **split plots** are the **two sessions**.
- The **split-plot factor** is the **piano type** (baby vs. concert grand).

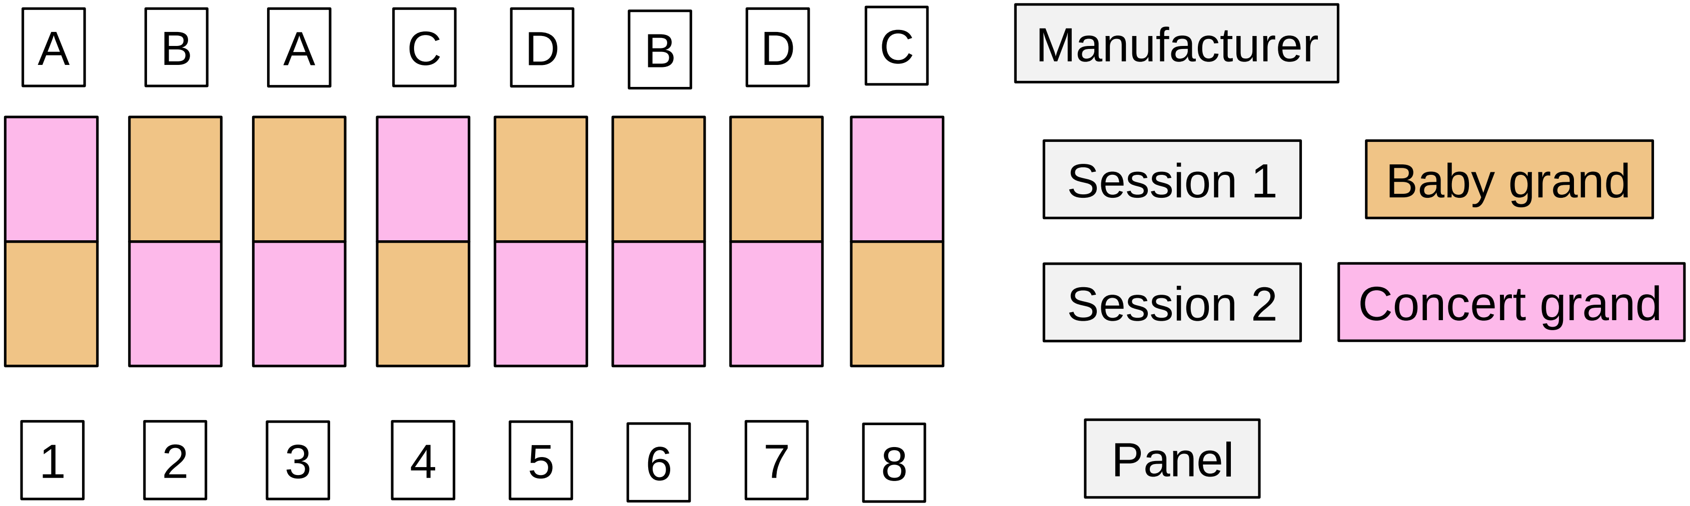

A B A C D B D C Manufacturer

Session 1 Baby grand

Session 2 Concert grand

1 2 3 4 5 6 7 8 Panel

Example II: Pianos

- The model is the same:

$N(0,\sigma_\eta^2)$

$N(0,\sigma^2)$

$$Y_{ijk} = \mu + \alpha_i + \eta_{k(i)} + \beta_j + (\alpha\beta)_{ij} + \varepsilon_{k(ij)}$$

average rating

fixed effect of **manufacturer**

whole-plot error

fixed effect of **piano type**

(fixed**) interaction between manufacturer and piano type**

split-plot error

- Again: This means that observations in the same whole-plot share the same whole-plot error $\eta_{k(i)}$ and are therefore **not independent**.

Example III: Oats

- Dataset `oats` from R-package `MASS`.

- As stated in the help file:

  *The yield of oats from a **split-plot field trial** using three varieties and four levels of manurial treatment. The experiment was laid out in 6 blocks of 3 main plots, each split into 4 sub-plots. The varieties were applied to the main plots and the manurial treatments to the sub-plots.*

- Overview of data:
  - 6 different **blocks (B)**
  - 3 different **varieties (V)**
  - 4 different **nitrogen treatments (N)**
  - **Response (Y):** Yields (in ¼ lbs per sub-plot, each of area $\frac{1}{80}$ acre).

- Let us first have a **graphical overview** of the experimental design.

Example III: Oats

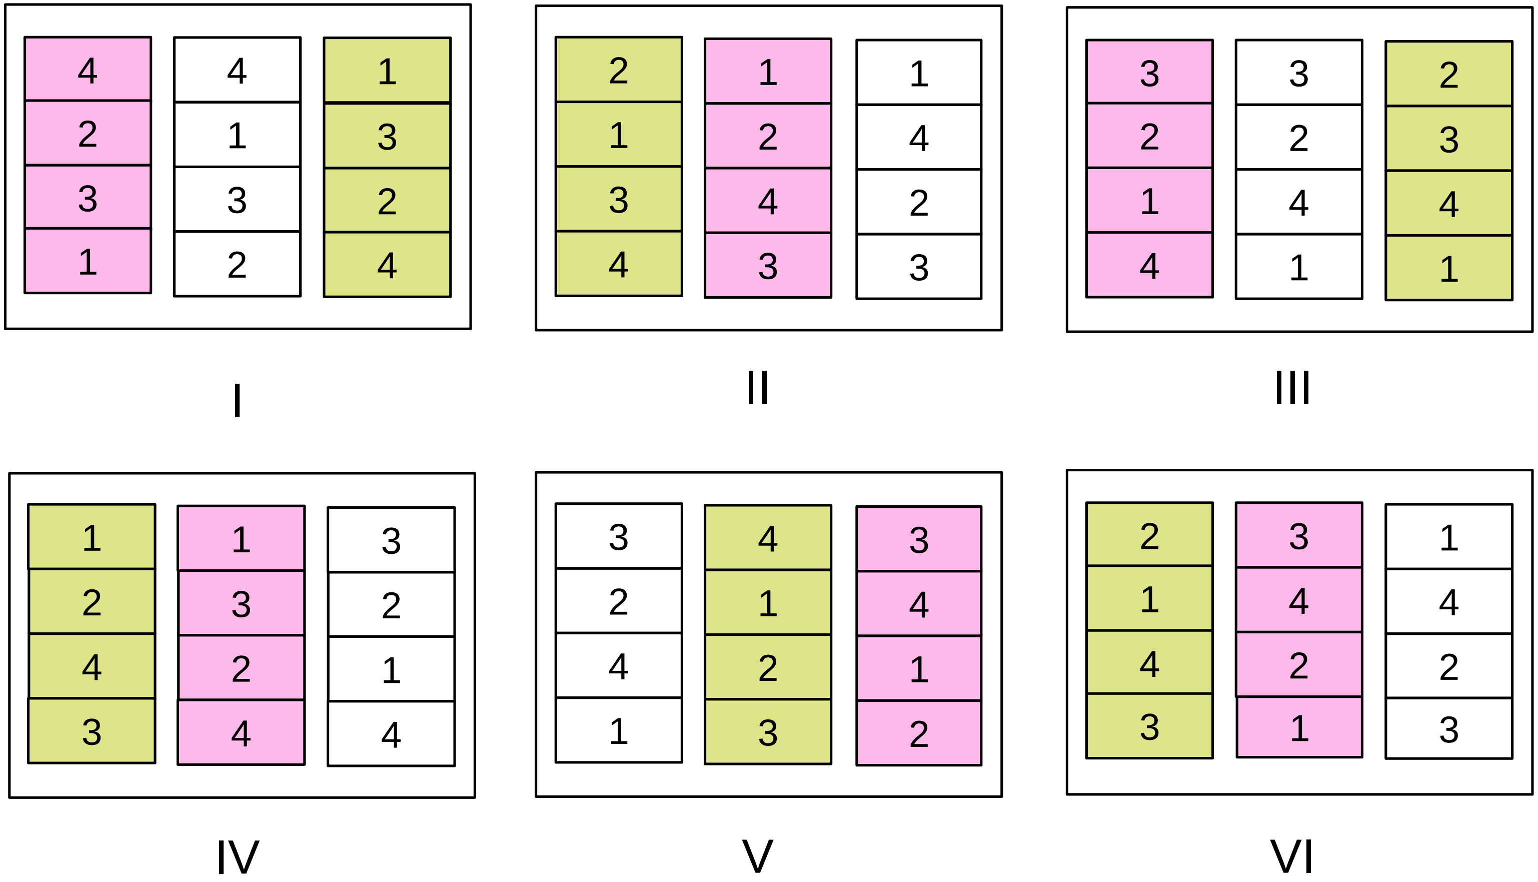

4	4	1
2	1	3
3	3	2
1	2	4

I

2	1	1
1	2	4
3	4	2
4	3	3

II

3	3	2
2	2	3
1	4	4
4	1	1

III

1	1	3
2	3	2
4	2	1
3	4	4

IV

3	4	3
2	1	4
4	2	1
1	3	2

V

2	3	1
1	4	4
4	2	2
3	1	3

VI

Example III: Oats

- This is a more complicated design as before as we have an additional **block factor**.
- A **whole-plot** is given by a plot of land in a block.
- The **whole-plot factor** is variety.
- A **block design (RCB)** was used at the whole-plot level.
- A **split plot** is given by a **subplot of land**.
- The **split-plot factor** is given by **nitrogen treatment**.

Example III: Oats

- We have an **RCB** for the whole-plot factor.
- The **experimental unit** on the whole-plot level is given by the **combination of block and variety**.
- We therefore use the model

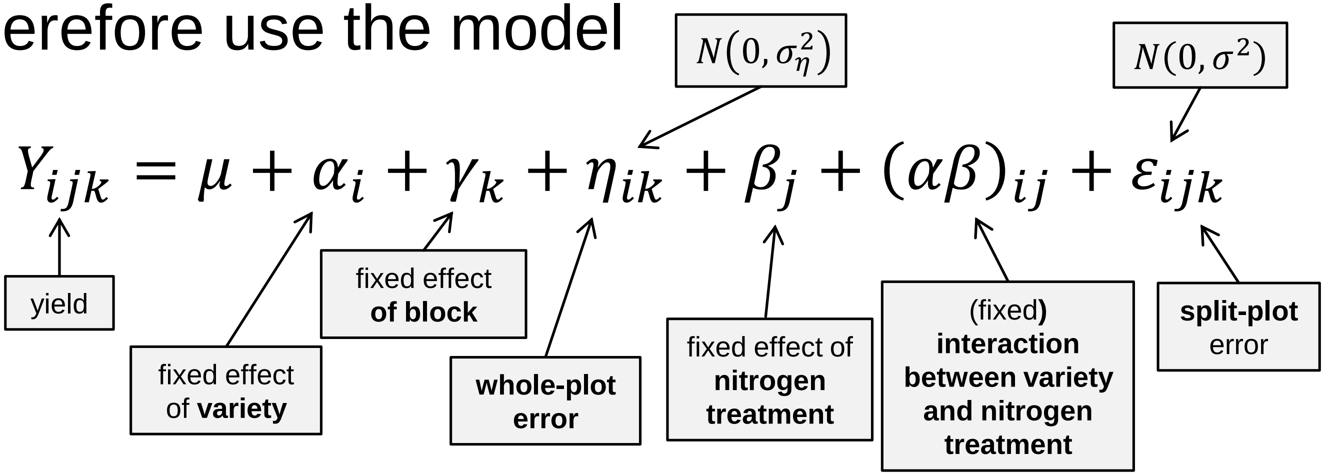

$$Y_{ijk} = \mu + \alpha_i + \gamma_k + \eta_{ik} + \beta_j + (\alpha\beta)_{ij} + \varepsilon_{ijk}$$

$N(0, \sigma_\eta^2)$

$N(0, \sigma^2)$

yield

fixed effect of **variety**

fixed effect **of block**

whole-plot error

fixed effect of **nitrogen treatment**

(fixed**)** **interaction between variety and nitrogen treatment**

split-plot error

Example III: Oats

- In R we use the `lmer` function with an extra random effect (error) per **combination** of **block** and **variety**.

- We get the following output

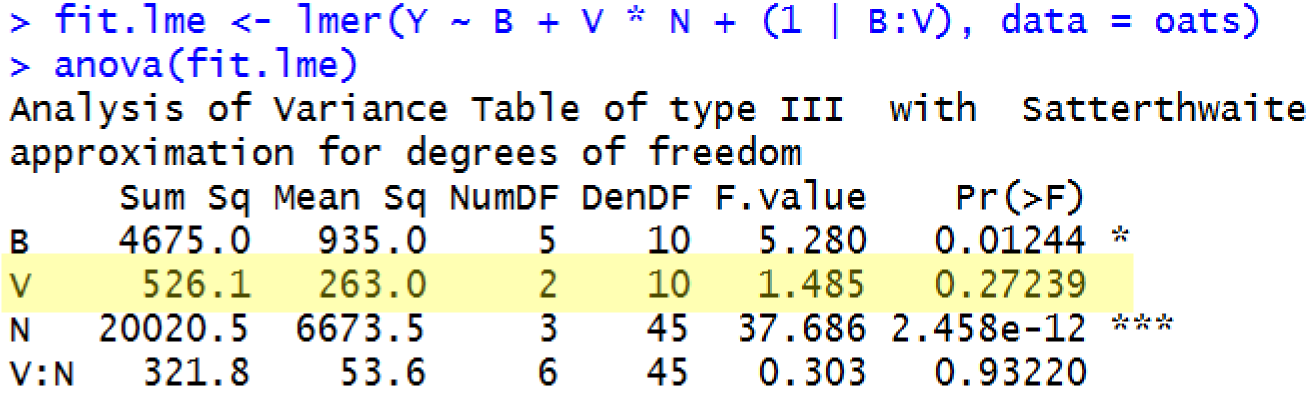

```
> fit.lme <- lmer(Y ~ B + V * N + (1 | B:V), data = oats)
> anova(fit.lme)
Analysis of Variance Table of type III  with  Satterthwaite
approximation for degrees of freedom
     Sum Sq Mean Sq NumDF DenDF F.value    Pr(>F)    
B    4675.0   935.0     5    10   5.280   0.01244 *  
V     526.1   263.0     2    10   1.485   0.27239    
N   20020.5  6673.5     3    45  37.686 2.458e-12 ***
V:N   321.8    53.6     6    45   0.303   0.93220    
```

- Observe that the test for variety uses 2 and 10 degrees of freedom.

- Why? Let us a have a closer look at the potential ANOVA table on the whole-plot level.

Example III: Oats

- On the whole-plot level we have the following ANOVA table:

Source	*df*
Block	5
Variety	2
Error (whole-plot)	$10\ (= 17 - 7)$
Total	$17\ (= 18 - 1)$

- Think of averaging “away” the nitrogen factor, hence we have **one observation** per **combination of block and variety**.
- Technically speaking, variety is tested against the **interaction** of block and variety.

Example III: Oats

- This also reveals a problem: We don’t have too many error df’s left to test the whole-plot factor (only 10).

- In contrast, we test everything involving the **split-plot** factor against the **residual error**, which has 45 df’s.

- Remember:

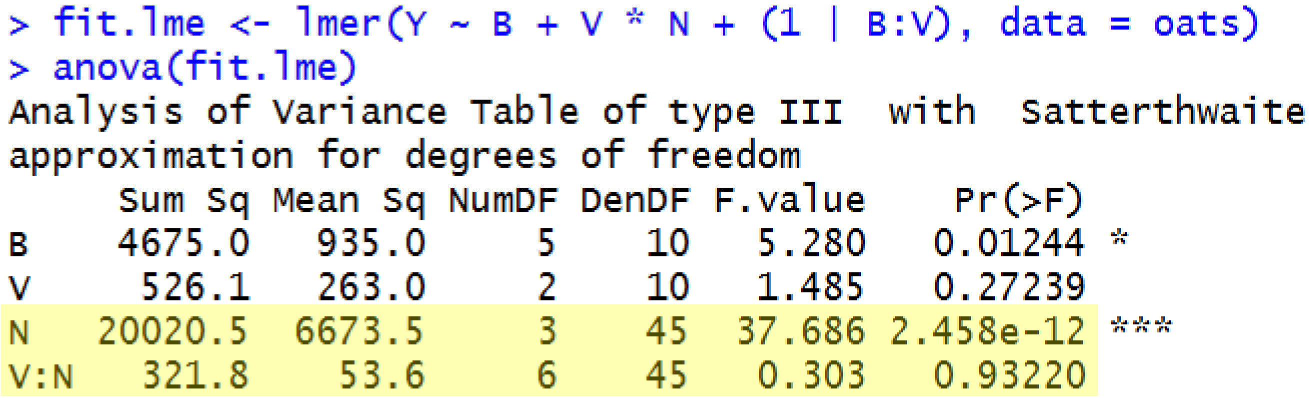

```
> fit.lme <- lmer(Y ~ B + V * N + (1 | B:V), data = oats)
> anova(fit.lme)
Analysis of Variance Table of type III  with  Satterthwaite
approximation for degrees of freedom
     Sum Sq Mean Sq NumDF DenDF F.value    Pr(>F)    
B    4675.0   935.0     5    10   5.280   0.01244 *  
V     526.1   263.0     2    10   1.485   0.27239    
N   20020.5  6673.5     3    45  37.686 2.458e-12 ***
V:N   321.8    53.6     6    45   0.303   0.93220    
```

- Hence, all effects involving the whole-plot factor are estimated **less precisely** and tests are **less powerful**.

General Situation

- Split-plot designs can also arise in (much) more complicated designs.
- There can be more than one whole-plot factor. E.g., think of a two-way factorial on the whole-plot level.
- In addition, there can be more than one factor on the split-plot level.
- To get the correct model we “only” have to **follow “the path of randomization”**.
- For every “level” (whole-plot / split-plot) of the experiment we have to introduce a corresponding random effect (better terminology: **error**) which acts as the experimental error on that level.

General Situation

- This means:
  - Start on the whole-plot level and forget about the split-plots.
  - Write down the corresponding model equation (incl. random effect / error).
  - Move on to the next level, expand equation with new terms (the upper level is now a block)
  - Etc.
- In R we just have to make sure that we tell `lmer` the correct random effects.
- In R it is sometimes useful to define new variables which identify the different experimental units on the different levels.

Example IV: Weed Biomass in Wetlands (Oehlert, 2000, Ex. 16.7)

- Experiment studies the effect of
  - **nitrogen** (4 levels of nitrogen)
  - **weed** (3 levels)
  - **clipping treatments** (2 levels: clipping / no clipping)

  on **plant growth** in wetlands.

- Experiment was performed as follows:
  - 8 **trays**, whereof each holds three artificial wetlands (rectangular wire baskets)
    - 4 of the trays were placed on **a table near the door** of the greenhouse
    - 4 of the trays on **a table in the center** of the greenhouse
  - On **each table**, we randomly assign one of the trays to each of the 4 **nitrogen treatments**.
  - Within **each tray**, we randomly assign the 3 **weed treatments**.
  - In addition, each wetland is split in half. One half is chosen at random and will be clipped, the other half is not clipped.
  - After 8 weeks: measure fraction of biomass that is nonweed.

Example IV: Weed Biomass in Wetlands

Experimental layout

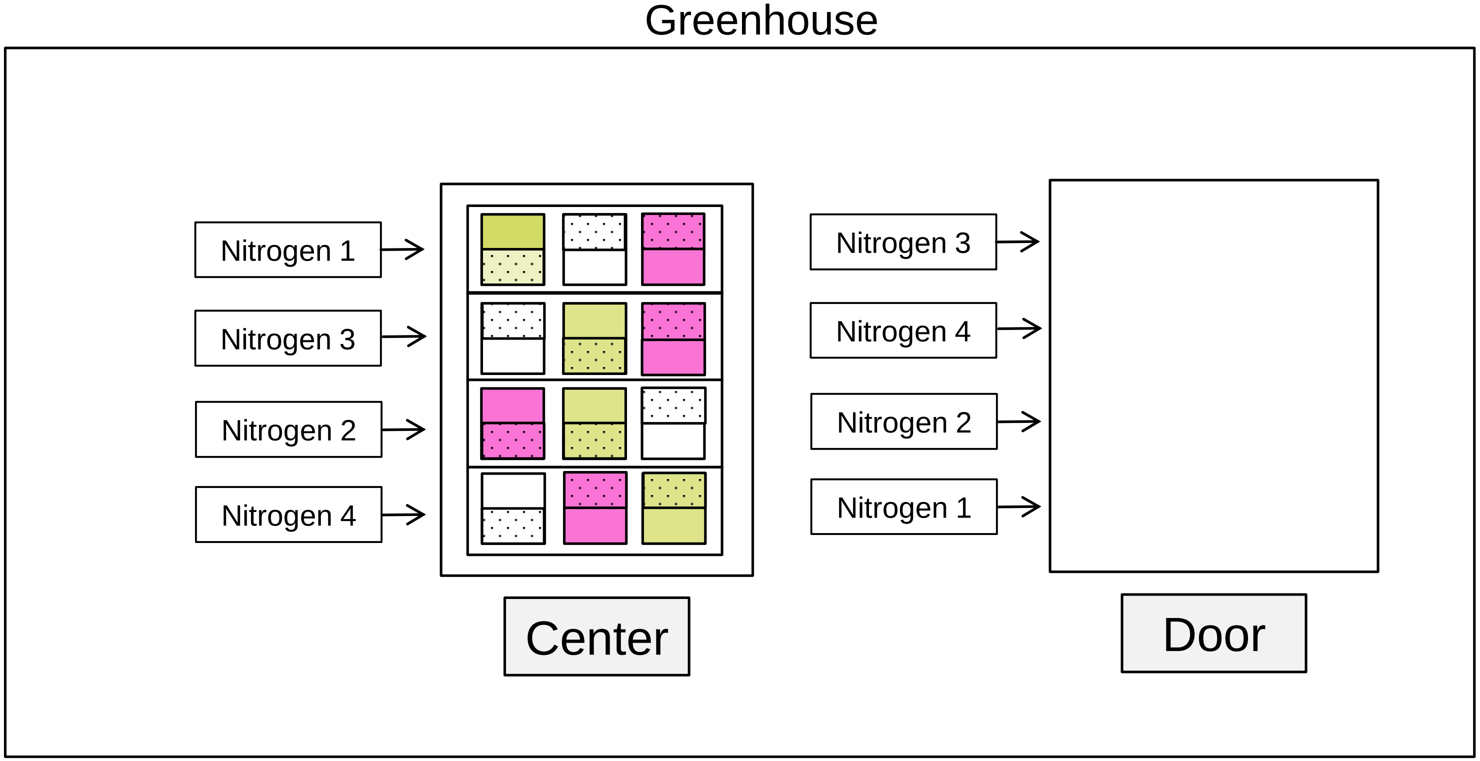

Example IV: Weed Biomass in Wetlands

- Let us follow the path of randomization:
  - Position in the greenhouse is a **block factor** (center / door)
  - **Trays** are **whole plots**, and **nitrogen level** is the **whole-plot factor**.
  - **Wetlands** are **split plots** and **weed treatment** is the **split-plot factor**.
  - **Wetland halves** are so called **split-split plots** and **clipping** is the **split-split-plot factor**.
- Hence, we have a so-called **split-split plot**.
- Let us now try to fit a model to this data-set in R.

Example IV: Weed Biomass in Wetlands

- We use the following model

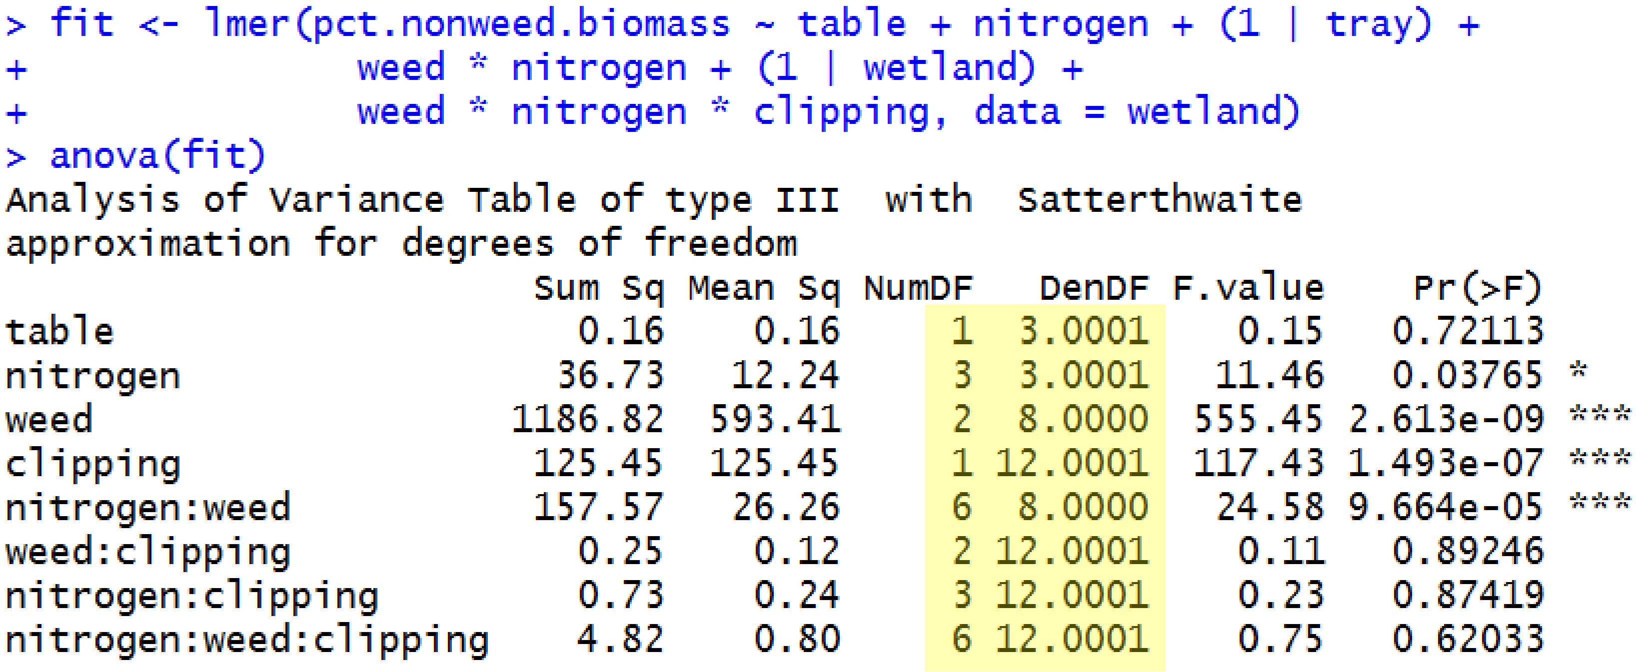

```
> fit <- lmer(pct.nonweed.biomass ~ table + nitrogen + (1 | tray) +
+              weed * nitrogen + (1 | wetland) +
+              weed * nitrogen * clipping, data = wetland)
> anova(fit)
Analysis of Variance Table of type III  with  Satterthwaite
approximation for degrees of freedom
                        Sum Sq Mean Sq NumDF   DenDF F.value    Pr(>F)
table                     0.16    0.16     1  3.0001    0.15   0.72113
nitrogen                 36.73   12.24     3  3.0001   11.46   0.03765 *
weed                   1186.82  593.41     2  8.0000  555.45 2.613e-09 ***
clipping                125.45  125.45     1 12.0001  117.43 1.493e-07 ***
nitrogen:weed           157.57   26.26     6  8.0000   24.58 9.664e-05 ***
weed:clipping             0.25    0.12     2 12.0001    0.11   0.89246
nitrogen:clipping         0.73    0.24     3 12.0001    0.23   0.87419
nitrogen:weed:clipping    4.82    0.80     6 12.0001    0.75   0.62033
```

- All **main-effects** and the **nitrogen** × **weed interaction** are significant.

Example IV: Weed Biomass in Wetlands

- We are here performing 3 experiments in 1.
- On the **whole-plot level** we have the "experiment"

Source	*df*
Table (block)	1
Nitrogen	3
Error (per tray)	$\mathbf{3}$ $(= 7 - 4)$
Total	$7\ (= 8 - 1)$

- On the **split-plot level** we have the "experiment"

Source	*df*
Block (=Tray)	7
Weed	2
Weed × Nitrogen	6
Error (per wetland)	$\mathbf{8}$ $(= 23 - 15)$
Total	$23\ (= 24 - 1)$

Example IV: Weed Biomass in Wetlands

- On the **split-split-plot level** we have the "experiment"

Source	*df*
Block (= wetland)	23
Clipping	1
Weed × Clipping	2
Nitrogen × Clipping	3
Nitrogen × Weed × Clipping	6
Error (per wetland half)	$\mathbf{12}\ \ (= 47 - 35)$
Total	$47\ (= 48 - 1)$

Summary

- Split plot designs and more complicated versions thereof are useful if some factors are harder (more expensive, …) to vary than others.
- To identify the correct design we have to know the randomization procedure.
- The general situation can be very complex, but by following the different randomization levels/steps, setting up a model is easy.
- Mixed effects software like `lmer` automatically identifies the correct denominator for tests if the random effects / errors are stated correctly.