Two-Factor Mixed Model ANOVA Example
Effectiveness of Sunscreens (§17.4)

• Evaluate effectiveness of 2 sunscreens. **Factor A: sunscreens (sun1, sun2), a fixed effect.**

• Experimental Units: A random sample of 40 people (20 randomly selected to receive sun1; the remainder getting sun2) . For each subject, a 1-inch square patch of skin was marked on back. A reading based on skin color was made prior to application of a fixed amount of sunscreen, and then again after a 2-hour exposure to sun. The difference in readings was recorded for each subject, with higher values indicating a greater degree of burning. **Response: burn.**

• Concerned that measurement of initial skin color is extremely variable. To assess variability due to the technicians taking the readings, 10 technicians were randomly selected and assigned 4 subjects each (2 receiving sun1, 2 receiving sun2). **Factor B: technicians (tech1,…,tech10), a random effect.**

• Result: CRD with factor A fixed (a=2), factor B random (b=10), and replication n=2 within each factor level combination. Total sample size is 2x10x2=40.

Trellis Panel Plot (from R)

8/1 = tech 8 and sun 1

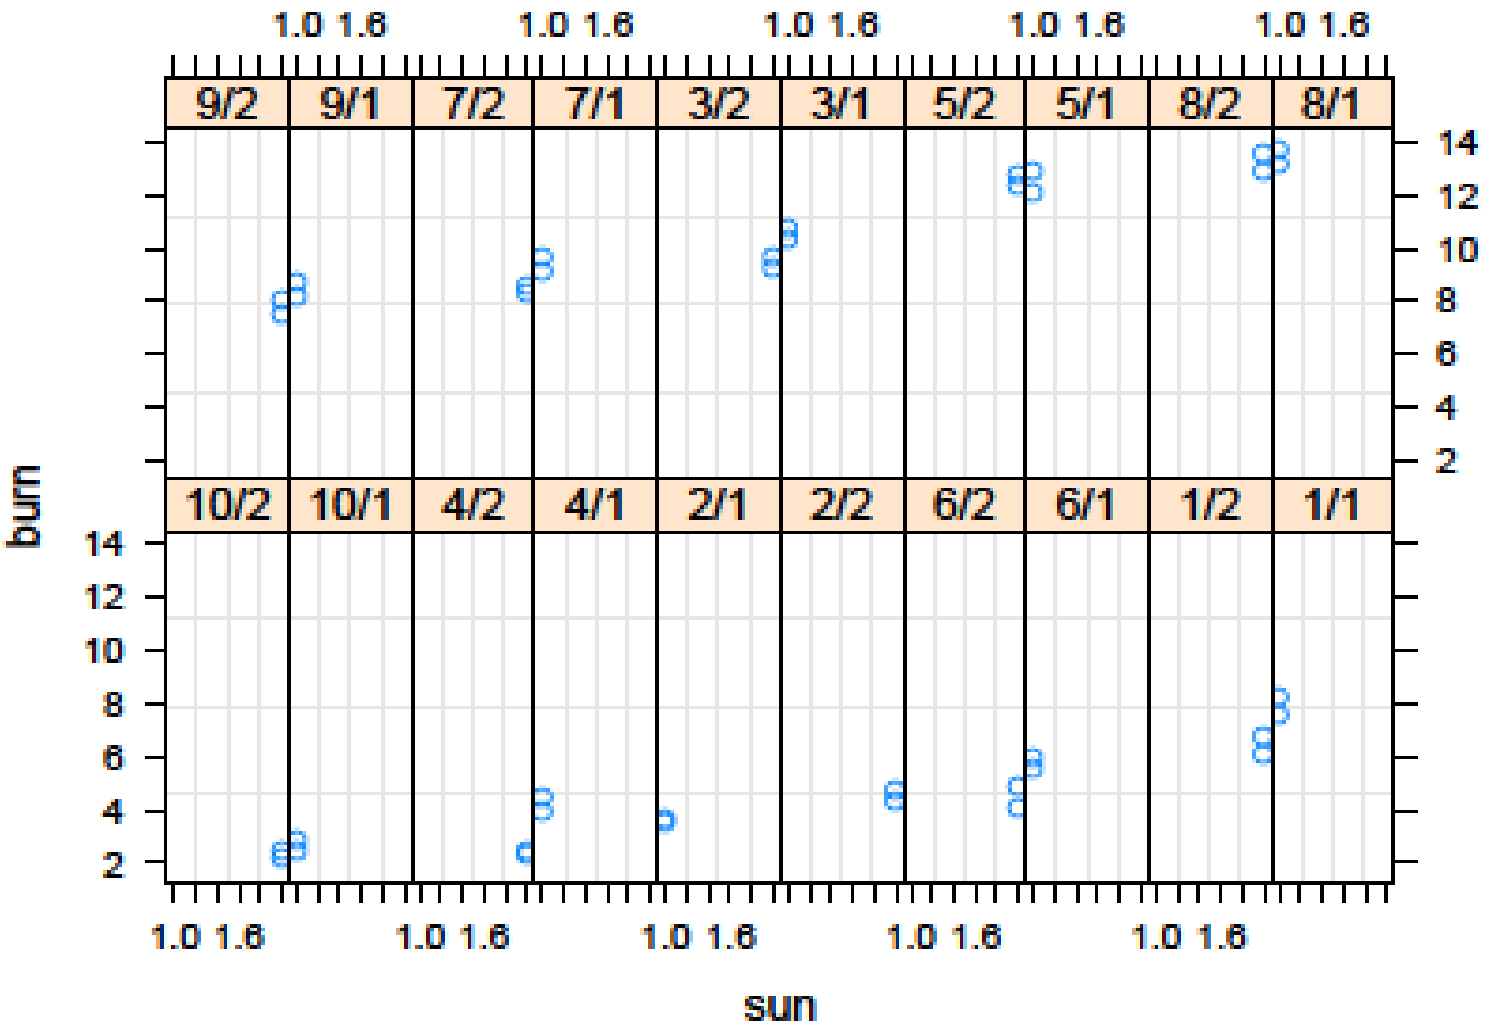

In MTB

Stat > ANOVA > Balanced ANOVA

- Response: “burn”
- Model: “sun tech sun*tech”
- Random Factors: “tech”
- Results: Display expected mean squares and variance components; Display means corresponding to the terms “sun tech”
- Options: Use restricted form of model

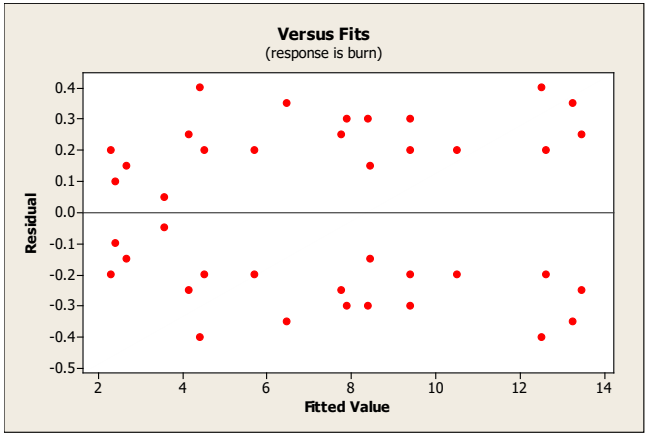

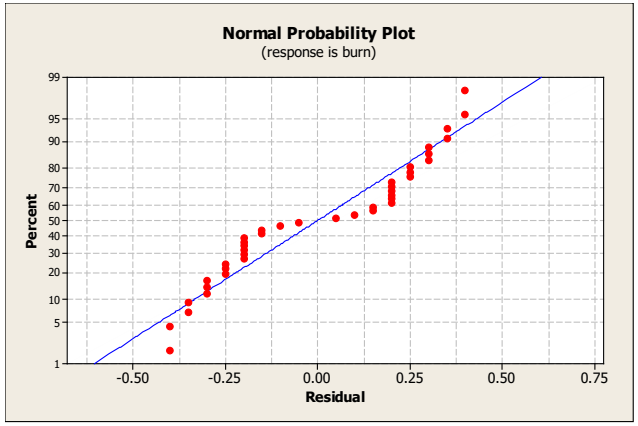

MTB Output: ANOVA table

ANOVA: burn versus sun, tech

```
Factor  Type    Levels  Values
sun     fixed        2  1, 2
tech    random      10  1,  2,  3,  4,  5,  6,  7,  8,  9, 10

Analysis of Variance for burn

Source     DF       SS       MS        F      P
sun         1    4.489    4.489     6.76  0.029
tech        9  517.486   57.498   435.59  0.000
sun*tech    9    5.976    0.664     5.03  0.001
Error      20    2.640    0.132
Total      39  530.591

S = 0.363318   R-Sq = 99.50%   R-Sq(adj) = 99.03%
```

sun differences

$\sigma_{\beta}^2 > 0$

$\sigma_{\alpha\beta}^2 > 0$

MTB Output: Variance components

	Source	Variance component	Error term	Expected Mean Square for Each Term (using restricted model)
1	sun		3	(4) + 2 (3) + 20 Q[1]
2	tech	14.3416	4	(4) + 4 (2)
3	sun*tech	0.2660	4	(4) + 2 (3)
4	Error	0.1320		(4)

$\hat{\sigma}^2_{\beta} = 14.3416$

Variability among technicians is substantial. (The variability is in determining initial skin color!)

$\hat{\sigma}^2_{\alpha\beta} = 0.2660$

$\hat{\sigma}^2_{\varepsilon} = 0.1320$

Variability among technicians is different for each of the two types of sunscreen. (This variability difference is significant, but not substantial.)

MTB Output: Means

```
Means

sun    N     burn
1     20   7.8200
2     20   7.1500

tech   N     burn
 1     4    7.175
 2     4    4.025
 3     4    9.950
 4     4    3.275
 5     4   12.550
 6     4    5.050
 7     4    8.925
 8     4   13.350
 9     4    8.075
10     4    2.475
```

Since there are sunscreen differences (ANOVA table), we **conclude sun 2 offers a greater amount of protection than sun 1**.

Large variation in technician means supports earlier finding, and testifies to the fact that measuring initial skin color is imprecise.

MTB Output: ANOVA table for model with both factors fixed

"Sun" p-value is now different

Two-way ANOVA: burn versus sun, tech

Source	DF	SS	MS	F	P
sun	1	4.489	4.4890	34.01	0.000
tech	9	517.486	57.4984	435.59	0.000
Interaction	9	5.976	0.6640	5.03	0.001
Error	20	2.640	0.1320		
Total	39	530.591			

S = 0.3633 R-Sq = 99.50% R-Sq(adj) = 99.03%

R Output: ANOVA

```
> library(nlme) # needed for lme function
> sunscreen <- read.csv("Data/Ott5thEdDataCh17/sunscreen.csv")
# first convert numbers to factor variables
> sunscreen$sun <- as.factor(sunscreen$sun)
> sunscreen$tech <- as.factor(sunscreen$tech)
> sun.lme <- lme(burn ~ sun, data=sunscreen, random=~1 | tech/sun,
   method="REML")
> anova(sun.lme)

Number of Observations: 40
Number of Groups:
         tech sun %in% tech
           10            20
> anova(sun.lme)
            numDF denDF  F-value p-value
(Intercept)     1    20 38.97512  <.0001
sun             1     9  6.76054  0.0287
```

sun differences

R Output: Variance components & fixed effects

```
> summary(sun.lme)
Linear mixed-effects model fit by REML
 Data: sunscreen
       AIC      BIC    logLik
  116.1123 124.3002 -53.05614

Random effects:
 Formula: ~1 | tech
        (Intercept)
StdDev:    3.769431

 Formula: ~1 | sun %in% tech
        (Intercept)  Residual
StdDev:   0.5157519 0.3633180

Fixed effects: burn ~ sun
            Value Std.Error DF   t-value p-value
(Intercept)  7.82  1.205845 20  6.485081  0.0000
sun2        -0.67  0.257682  9 -2.600104  0.0287
```

Note: standard deviations!

$$\hat{\sigma}_{\beta} = 3.7694$$

$$\hat{\sigma}_{\alpha\beta} = 0.5158$$

$$\hat{\sigma}_{\varepsilon} = 0.3633$$

95% confidence intervals for variance estimates

```
> intervals(sun.lme, which="var-cov")
Approximate 95% confidence intervals

 Random Effects:
  Level: tech
                      lower     est.    upper
sd((Intercept)) 2.362046 3.769431 6.015382
  Level: sun
                       lower      est.     upper
sd((Intercept)) 0.2882865 0.5157519 0.9226931

 Within-group standard error:
    lower      est.     upper
0.2665023 0.3633180 0.4953054
```

Note: standard deviations!

$$\hat{\sigma}_{\beta} = (2.36, 6.02)$$

$$\hat{\sigma}_{\alpha\beta} = (0.288, 0.923)$$

$$\hat{\sigma}_{\varepsilon} = (0.267, 0.495)$$

Diagnostic plots: qqnorm & resids vs. fitted

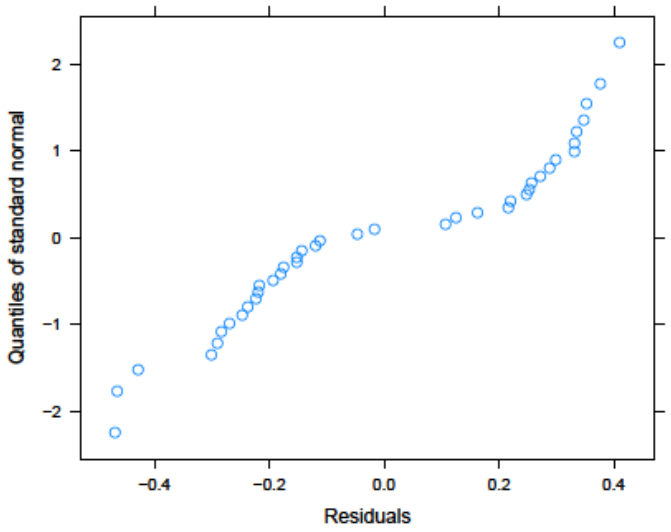

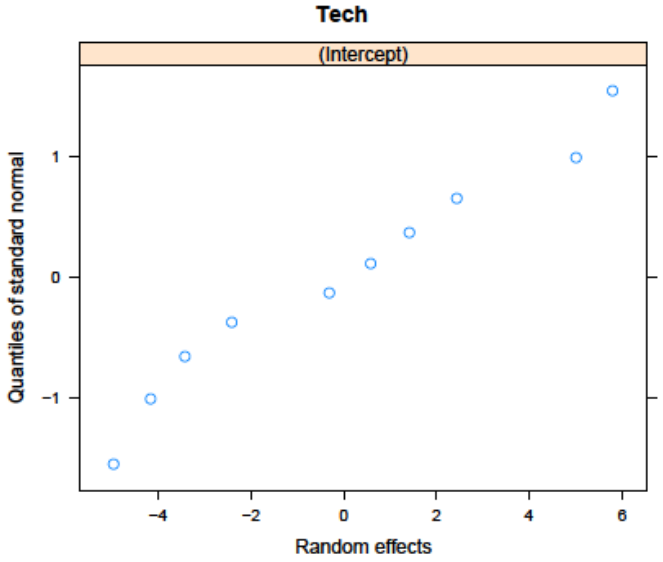

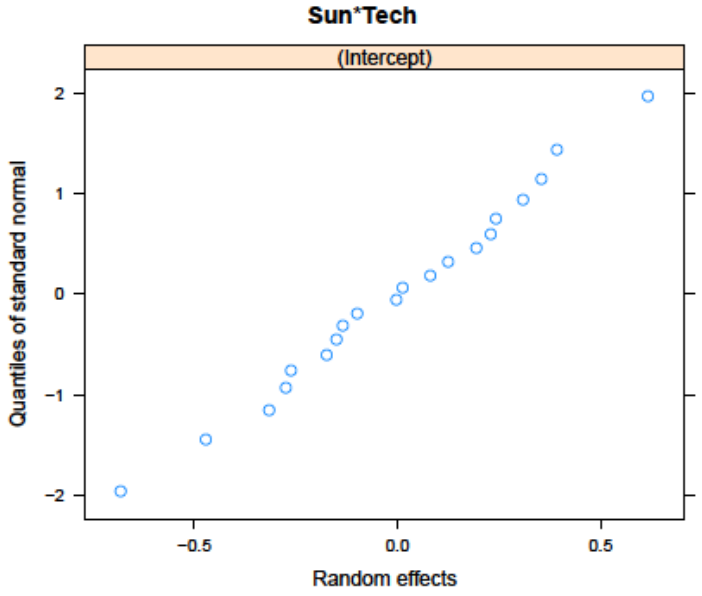

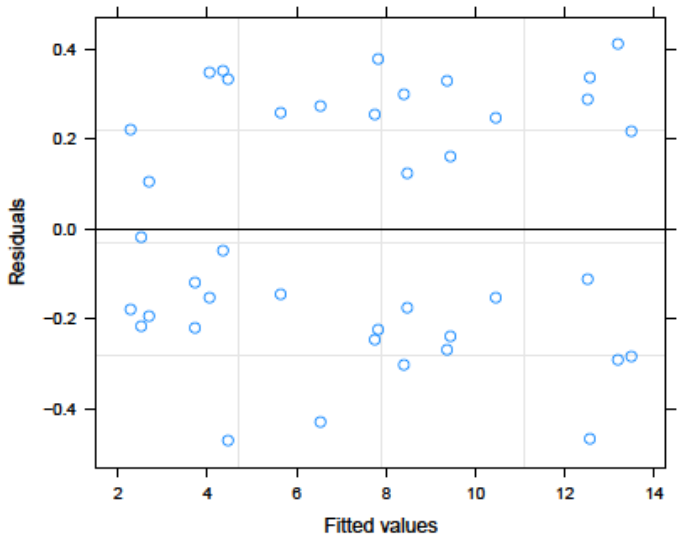

SAS

```
proc mixed;
class sun tech;
model burn = sun;
random tech sun*tech;
```

SPSS

```
proc mixed
Model fixed factors: sun
Model random factors: tech sun*tech
```

Random Effects ANOVA With Nesting Example
Content Uniformity of Drug Tablets (§17.6)

- **Response: Drug.** Content uniformity of drug tablets.
- **Factor A: Site (random).** Drug company manufactures at different sites; 2 are randomly chosen for analysis.
- **Factor B: Batch (random).** Three batches are randomly selected within each site (batch is nested within site).
- **Replicates:** 5 tablets are randomly selected from each batch for measurement.

In MTB

Stat > ANOVA > Balanced ANOVA

- Response: “Drug”
- Model: “Site Batch(Site)”
- Random Factors: “Site Batch”
- Results: Display expected mean squares and variance components
- Options: Use restricted form of model

MTB Output: ANOVA table

ANOVA: Drug versus Site, Batch

```
Factor        Type    Levels  Values
Site          random       2  1, 2
Batch(Site)   random       3  1, 2, 3

Analysis of Variance for Drug

Source        DF       SS        MS     F      P
Site           1  0.01825   0.01825  0.16  0.709
Batch(Site)    4  0.45401   0.11350  9.39  0.000
Error         24  0.29020   0.01209
Total         29  0.76247

S = 0.109962   R-Sq = 61.94%   R-Sq(adj) = 54.01%
```

$\sigma_{\alpha}^2 = 0$

$\sigma_{\beta(\alpha)}^2 > 0$

MTB Output: Variance components

	Source	Variance component	Error term	Expected Mean Square for Each Term (using restricted model)
1	Site	-0.00635	2	(3) + 5 (2) + 15 (1)
2	Batch(Site)	0.02028	3	(3) + 5 (2)
3	Error	0.01209		(3)

$\hat{\sigma}^2_{\alpha} = -0.00635$

Variability among sites is negligible. (Note negative estimate!)

$\hat{\sigma}^2_{\beta(\alpha)} = 0.02028$

$\hat{\sigma}^2_{\varepsilon} = 0.01209$

Considerable batch-to-batch variability in content uniformity of tablets.

R Output

```
> library(nlme) # needed for lme function
> content <- read.csv("Data/Ott5thEdDataCh17/ch17-Example17.10.csv")
# first convert numbers to factor variables
> content$Site <- as.factor(content$Site)
> content$Batch <- as.factor(content$Batch)
# fit random effects model with Batch nested in Site
> drug.lme <- lme(Drug~1, data=content, random=~1 | Site/Batch)
> summary(drug.lme)

Linear mixed-effects model fit by REML
 Data: content
        AIC       BIC   logLik
  -24.06435 -18.59516 16.03217

Number of Observations: 30
Number of Groups:
           Site Batch %in% Site
              2                6
```

R Output

```
Random effects:
 Formula: ~1 | Site
         (Intercept)
StdDev: 3.236734e-06

 Formula: ~1 | Batch %in% Site
        (Intercept)  Residual
StdDev:   0.1283446 0.1099621

Fixed effects: Drug ~ 1
               Value Std.Error DF  t-value p-value
(Intercept) 5.043333  0.056111 24 89.88136       0
```

$$\hat{\sigma}_{\alpha} = 0.000003236734 \Rightarrow \hat{\sigma}_{\alpha}^2 \approx 0$$

$$\hat{\sigma}_{\beta(\alpha)} = 0.1283446 \Rightarrow \hat{\sigma}_{\beta(\alpha)}^2 = 0.01647$$

$$\hat{\sigma}_{\varepsilon} = 0.1099621 \Rightarrow \hat{\sigma}_{\varepsilon}^2 = 0.01209$$

SAS

```
proc mixed;
class Site Batch;
model Drug = ;
random Site Batch(Site);
```

SPSS

proc mixed?

Split-Plot Example: *Soybean Yields* (§17.6, 5th Ed.)

- **Response: Yield.** Soybean yields in bushels per subplot unit.
- **Factor A: Fertilizer.** Two fertilizer types (1,2). Each fertilizer is randomly applied to 3 wholeplots (a=2).
- **Factor B (treatment): Variety.** Three varieties of soybean (1,2,3). Each wholeplot is divided into 3 subplots and each variety is randomly applied to each of the subplots. (t=3)
- **Wholeplots: WPlot**. Experiment is replicated 3 times (n=3). Each replicate consists of a pair of wholeplots (total of 6 wholeplots).
- **Note:** we are ignoring the Block (farm) factor in the original data. View as having 3 pairs of wholeplots (6 Wplots) in one farm.

The Data

EXAMPLE 17.11

Soybean yields (in bushels per subplot unit) are shown here for a two-factor split-plot design laid off in $b = 3$ blocks. Fertilizers (factor A) were applied at random to the wholeplot units within each farm. Soybean varieties (factor T) were then randomly allocated to the subplots within each wholeplot. Conduct an analysis of variance using these sample data. Give an approximate p-value for each test.

1			2			3		
	Fertilizers			Fertilizers			Fertilizers	
Varieties	1	2	Varieties	2	1	Varieties	1	2
1	10.6	10.9	2	11.9	11.5	3	9.5	9.8
2	11.4	11.7	3	12.6	12.1	1	8.1	8.2
3	11.8	12.4	1	11.6	10.8	2	8.7	9.3

Solution For these data with $a = 2$, $b = 3$, $t = 3$, and $n = 1$, the sum of squares are as shown (see the Type III SS column in the following computer output):

In MTB

Stat > ANOVA > General Linear Model

- Response: Yield
- Model: Fertilizer WPlot(Fertilizer) Variety Fertilizer*Variety
- Random Factors: WPlot
- Results: Display expected mean squares and variance components; Display means corresponding to the terms “Variety”.

MTB Output: ANOVA table

General Linear Model: Yield versus Fertilizer, Variety, WPlot

Factor	Type	Levels	Values
Fertilizer	fixed	2	1, 2
WPlot(Fertilizer)	random	6	1, 3, 5, 2, 4, 6
Variety	fixed	3	1, 2, 3

Analysis of Variance for Yield, using Adjusted SS for Tests

Source	DF	Seq SS	Adj SS	Adj MS	F	P
Fertilizer	1	0.8450	0.8450	0.8450	0.12	0.750
WPlot(Fertilizer)	4	28.9067	28.9067	7.2267	10.65	0.003
Variety	2	0.0233	0.0233	0.0117	0.02	0.983
Fertilizer*Variety	2	0.1233	0.1233	0.0617	0.09	0.914
Error	8	5.4267	5.4267	0.6783		
Total	17	35.3250				

S = 0.823610 R-Sq = 84.64% R-Sq(adj) = 67.36%

No Fertilizer differences

No Variety differences

MTB Output

Error Terms for Tests, using Adjusted SS

	Source	Error DF	Error MS	Synthesis of Error MS
1	Fertilizer	4.00	7.2267	(2)
2	WPlot(Fertilizer)	8.00	0.6783	(5)
3	Variety	8.00	0.6783	(5)
4	Fertilizer*Variety	8.00	0.6783	(5)

Variance Components, using Adjusted SS

Source	Estimated Value
WPlot(Fertilizer)	2.1828
Error	0.6783

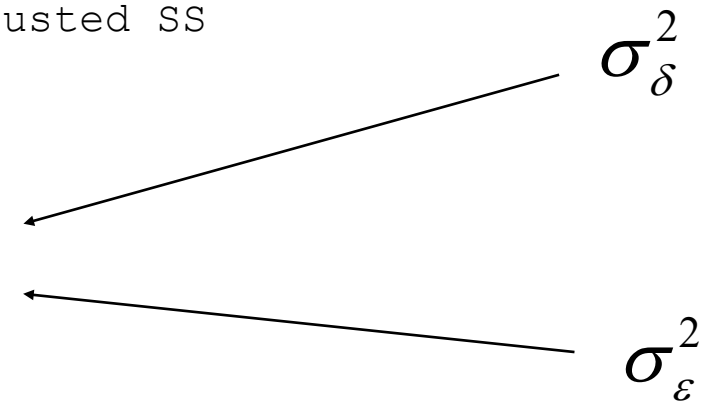

Least Squares Means for Yield

Variety	Mean
1	10.70
2	10.68
3	10.77

R code

```
> library(nlme) # needed for lme function
> soy <- read.csv("Data/Ott5thEdDataCh17/ch17-Example17.11.csv")
> # first convert numbers to factor variables
> soy$WPlot <- as.factor(soy$WPlot)
> soy$Fertilizer <- as.factor(soy$Fertilizer)
> soy$Variety <- as.factor(soy$Variety)
> # fit split-plot model with WPlot nested in Fertilizer (using lme
   to get random effects)
> soy.lme <- lme(Yield~Fertilizer*Variety, data=soy, random=~1 |
   WPlot)
> # fit split-plot model with WPlot nested in Fertilizer (using aov
   to get anova table)
> soy.lm <- aov(Yield~Fertilizer*Variety+Error(WPlot), data=soy)
```

Both soy.lm and soy.lme will give same fit, but latter will also estimate random effects

R Output: Variance components

```
> summary(soy.lme)

Random effects:
 Formula: ~1 | WPlot
        (Intercept)  Residual
StdDev:    1.477421 0.8236104

> intervals(soy.lme, which="var-cov")
Approximate 95% confidence intervals

 Random Effects:
  Level: WPlot
                       lower     est.    upper
sd((Intercept)) 0.6864762 1.477421 3.179676

 Within-group standard error:
    lower      est.     upper
0.5045427 0.8236104 1.3444535
```

Both random effects are significant (at the 5% level).

σ_δ

σ_ε

R Output: ANOVA

```
> anova(soy.lme)
                    numDF denDF   F-value p-value
(Intercept)             1     8 286.05857  <.0001
Fertilizer              1     4   0.11693  0.7496
Variety                 2     8   0.01720  0.9830
Fertilizer:Variety      2     8   0.09091  0.9140
```

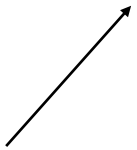

No evidence of Fertilizer or Variety differences…

R Output: LS means

```
> # table of estimated means
> model.tables(soy.lm, type="means")
Tables of means
Grand mean

10.71667

 Fertilizer
Fertilizer
     1      2
10.500 10.933

 Variety
Variety
     1      2      3
10.700 10.683 10.767

 Fertilizer:Variety
          Variety
Fertilizer 1      2      3
         1 10.533 10.533 10.433
         2 10.867 10.833 11.100
```

Fertilizer means

Variety means

All pairwise means

SAS

```
proc mixed;
class Fertilizer Variety WPlot;
model Yield = Fertilizer Variety Fertilizer*Variety / ddfm=satterth;
random WPlot(Fertilizer);
parms / nobound;
lsmeans Variety / pdiff cl;
```

SPSS

proc mixed?

Randomized Block Split-Plot Example: *Soybean Yields* (§17.6, 5th Ed.)

- **Response: Yield.** Soybean yields in bushels per subplot unit.
- **Factor A: Fertilizer.** Two fertilizer types (1,2). Each fertilizer is randomly applied to 3 wholeplots (a=2).
- **Factor B (treatment): Variety.** Three varieties of soybean (1,2,3). Each wholeplot is divided into 3 subplots and each variety is randomly applied to each of the subplots. (t=3)
- **Factor C: Blocks**. Experiment is replicated at each of 3 farms (b=3).

In MTB

Stat > ANOVA > General Linear Model

- Response: Yield
- Model: Fertilizer Block Fertilizer*Block Variety Fertilizer*Variety
- Random Factors: Block
- Results: Display expected mean squares and variance components; Display means corresponding to the terms “Variety”.

MTB Output: ANOVA table

General Linear Model: Yield versus Fertilizer, Block, Variety

Factor	Type	Levels	Values
Fertilizer	fixed	2	1, 2
Block	random	3	1, 2, 3
Variety	fixed	3	1, 2, 3

Analysis of Variance for Yield, using Adjusted SS for Tests

Source	DF	Seq SS	Adj SS	Adj MS	F	P
Fertilizer	1	0.8450	0.8450	0.8450	39.00	0.025
Block	2	28.8633	28.8633	14.4317	666.08	0.001
Fertilizer*Block	2	0.0433	0.0433	0.0217	0.03	0.969
Variety	2	0.0233	0.0233	0.0117	0.02	0.983
Fertilizer*Variety	2	0.1233	0.1233	0.0617	0.09	0.914
Error	8	5.4267	5.4267	0.6783		
Total	17	35.3250				

Fertilizer differences

No Variety differences

MTB Output: Variance Components

```
Variance Components, using Adjusted SS

                    Estimated
Source                  Value
Block                  2.4017
Fertilizer*Block      -0.2189
Error                  0.6783


Least Squares Means for Yield

Variety    Mean
1         10.70
2         10.68
3         10.77
```

Significant and substantial block to block variability

Confirms F-test of no Variety differences

R code

```
> soy.lme <- lme(Yield~Fertilizer*Variety, random=~1 |
   Block/Fertilizer, data=soy)
> anova(soy.lme)
                    numDF denDF   F-value p-value
(Intercept)             1     8 143.24368  <.0001
Fertilizer              1     2   1.54479  0.3399
Variety                 2     8   0.02133  0.9790
Fertilizer:Variety      2     8   0.11274  0.8948

> summary(soy.lme)
Random effects:
 Formula: ~1 | Block
        (Intercept)
StdDev:    1.521220

 Formula: ~1 | Fertilizer %in% Block
         (Intercept)  Residual
StdDev: 2.013288e-05 0.7395945
```

None of the fixed effects are significant under REML estimation!

But we do get positive random effects estimates!

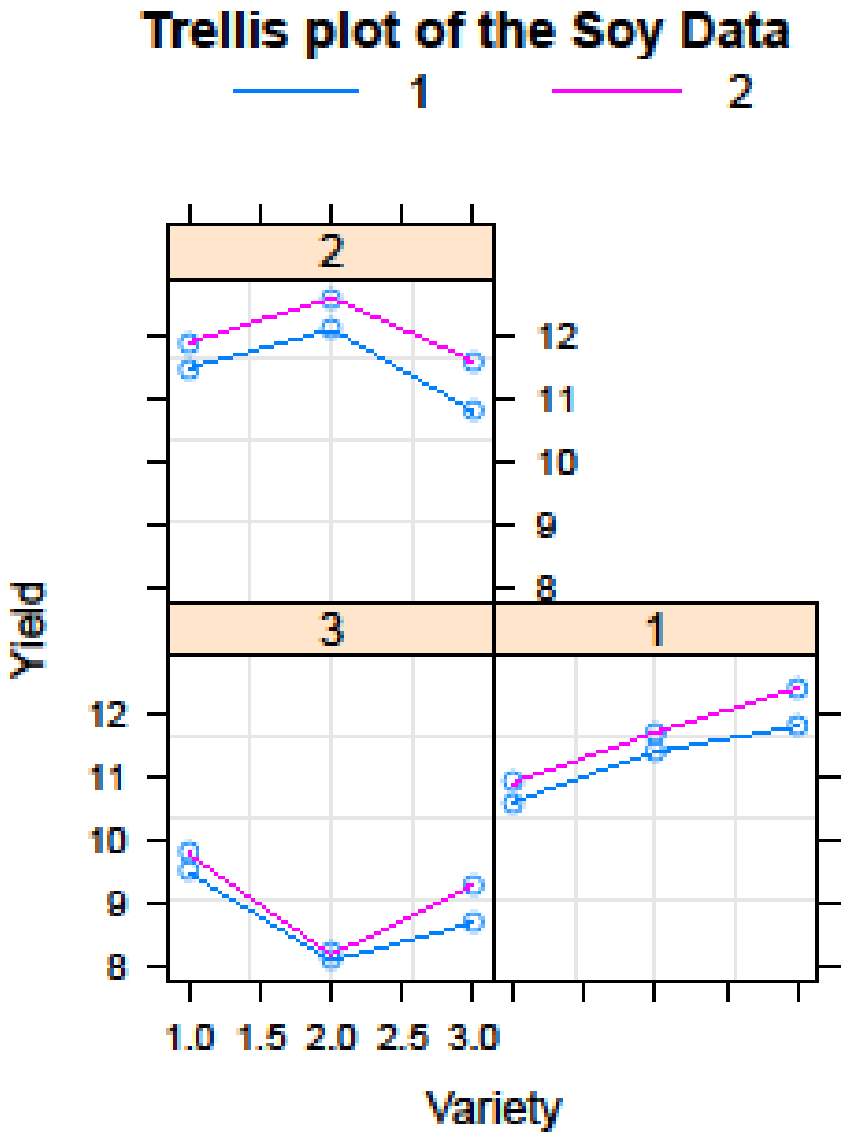

Blue (1) =Fertilizer 1.

Pink (2) =Fertilizer 2.

Repeated Measures Example: *Root Growth of Plants* (§18.3-4)

- **Response: root.** Root length.
- **Factor A: fertilizer.** Either “added” or not (“control”). Fixed.
- **Factor B: week.** Each of 6 plants was measured at weeks (2,4,6,8,10). Plants are nested in factor A. Random.
- **Factor C: plants**. 6 plants got fertilizer; 6 didn’t; acting as blocks. Random.

Panel plots of data

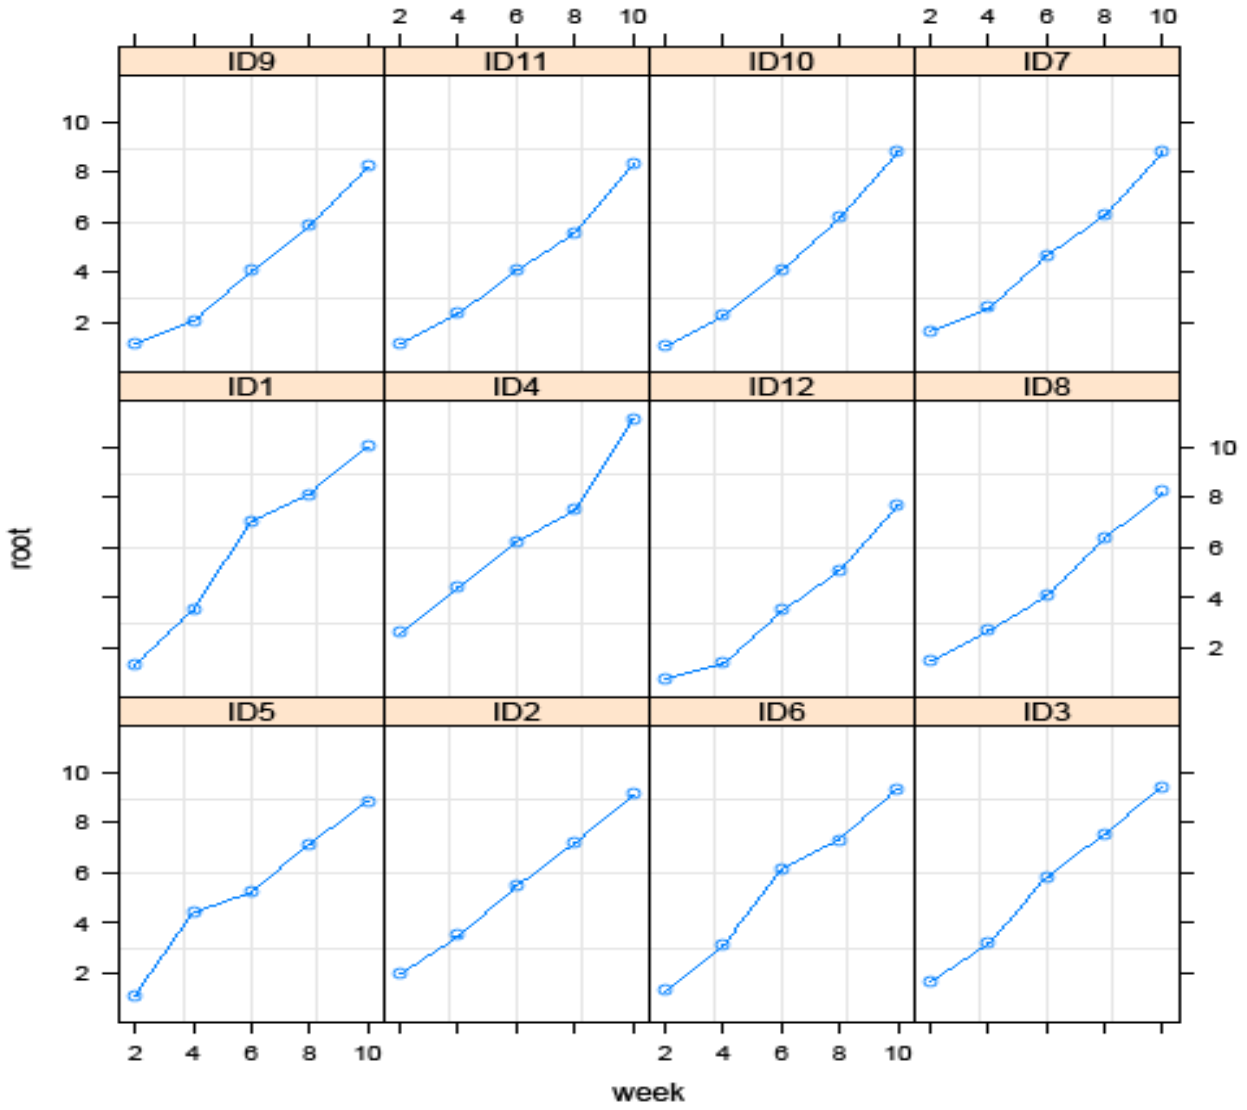

Panel plots: grouped by fertilizer treatment

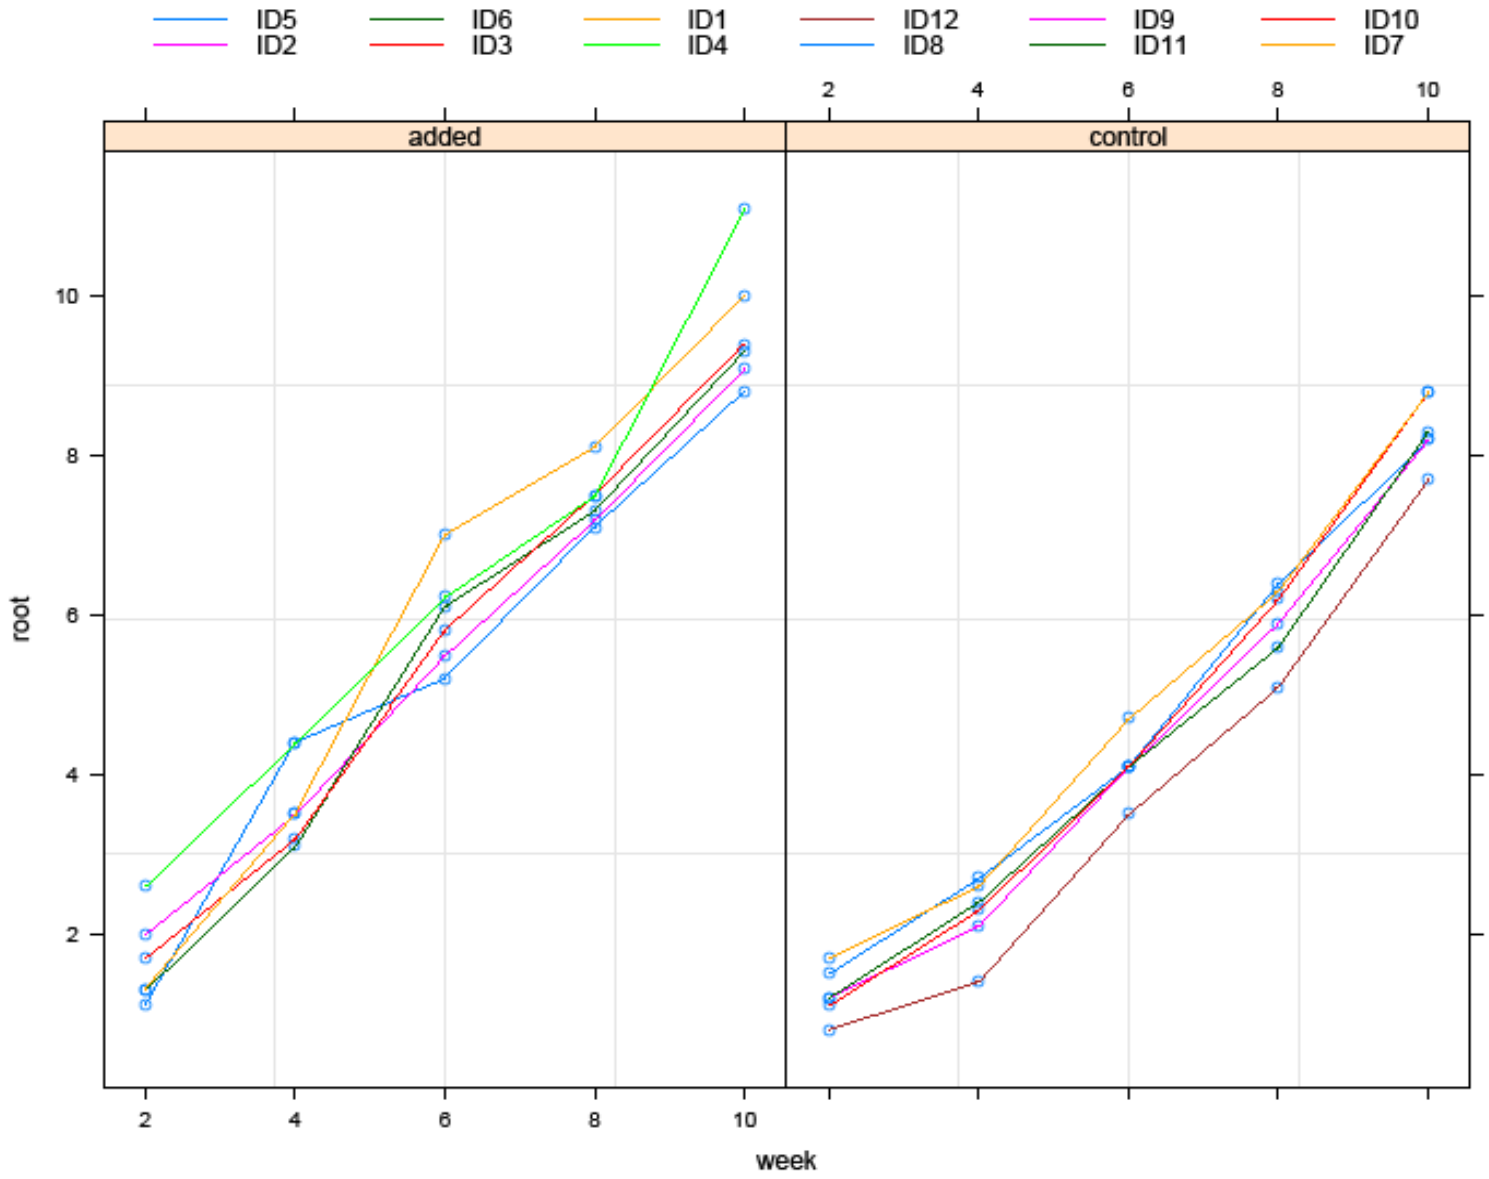

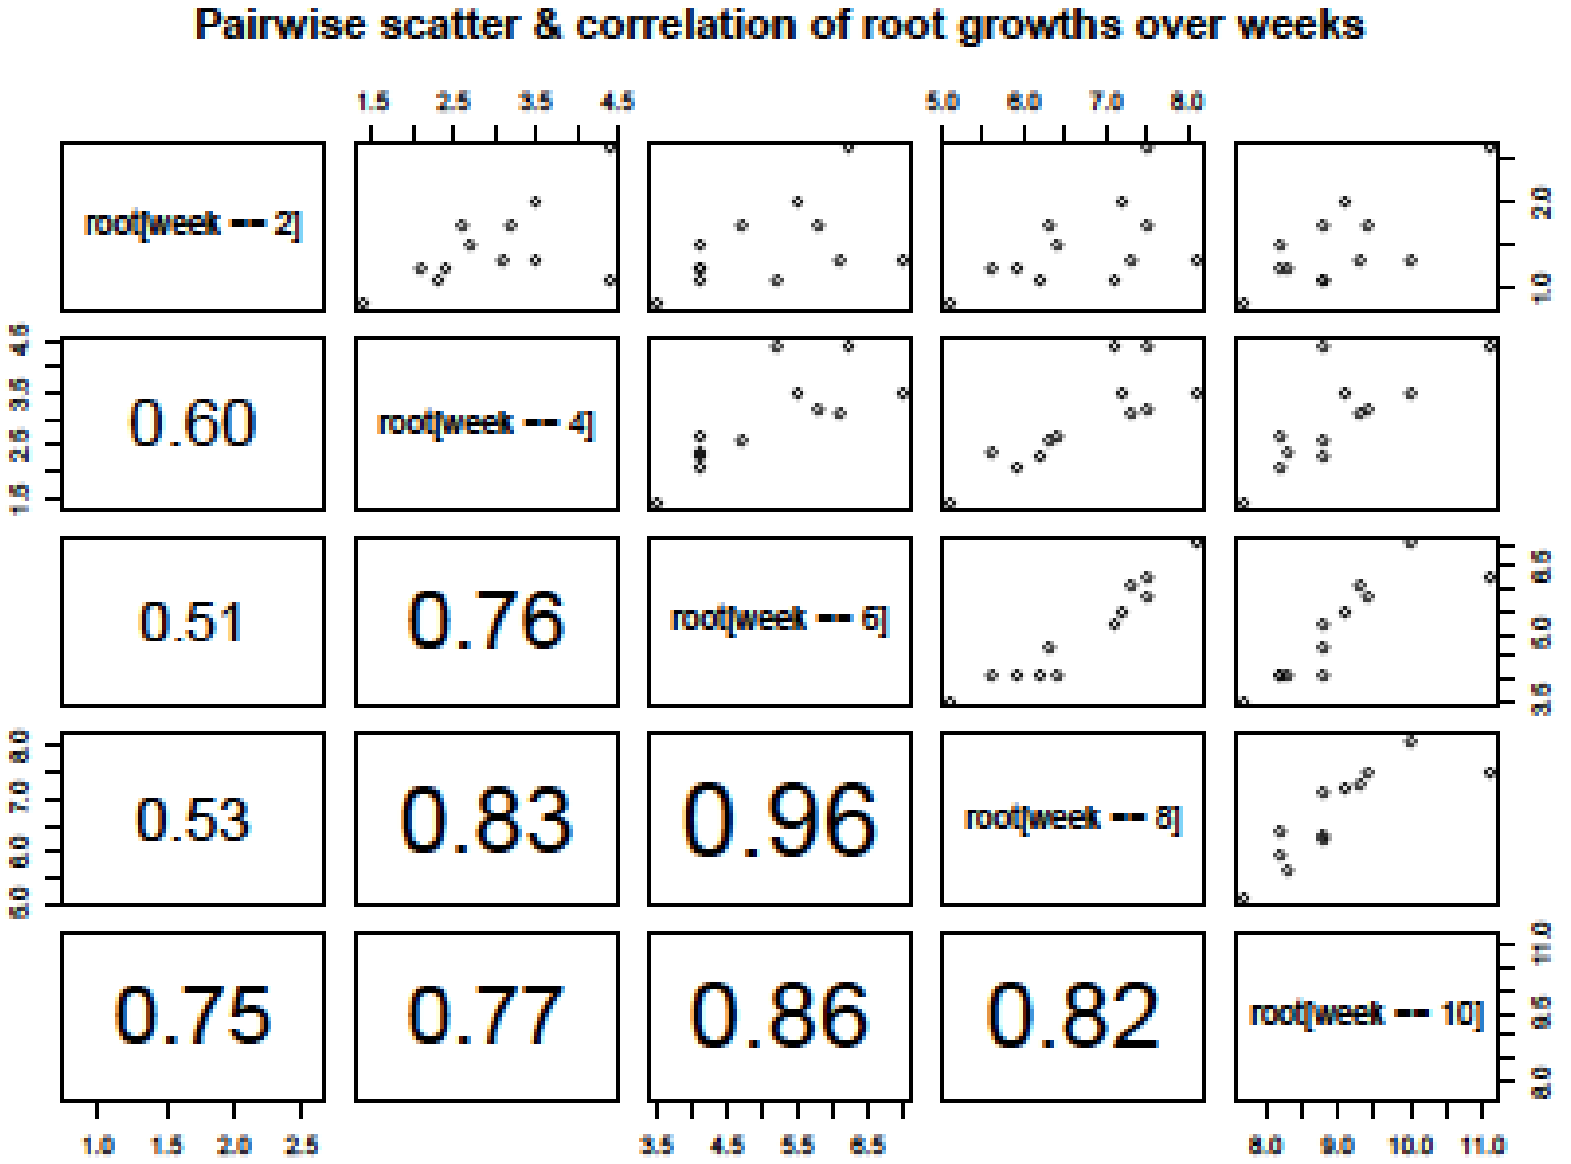
Pairwise scatter & correlation of root growths over weeks
root[week = 2]
root[week = 4]
root[week = 6]
root[week = 8]
root[week = 10]
0.60
0.51
0.76
0.53
0.83
0.96
0.75
0.77
0.86
0.82

R code: fit linear model in notes with plant nested in fertilizer, and default correlation structure for plants (compound symmetry)

```
> grow.lme <- lme(root~fertilizer*week, data=grow, random=~1 | plant)
> summary(grow.lme)
Linear mixed-effects model fit by REML
 Data: grow
       AIC      BIC    logLik
  105.0325 127.9767 -40.51623

Random effects:
 Formula: ~1 | plant
        (Intercept)  Residual
StdDev:   0.3541493 0.3855818
```

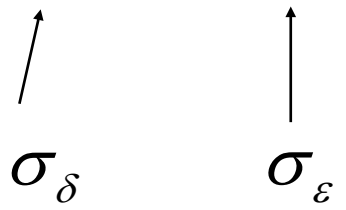

Model with AR(1) autocorrelation structure for plants

```
> grow.lme.3 <- lme(root~fertilizer*week,
   data=grow, random=~1 | plant,
                correlation=corAR1())
> summary(grow.lme.3)
Linear mixed-effects model fit by REML
 Data: grow
       AIC      BIC    logLik
  107.0169 131.8732 -40.50843

Random effects:
 Formula: ~1 | plant
        (Intercept)  Residual
StdDev:   0.3527663 0.3874222

Correlation Structure: AR(1)
 Formula: ~1 | plant
 Parameter estimate(s):
       Phi
0.02549701
```

AIC & BIC have increased a bit…

Little change in the variance components

Estimate of ϕ is small (maybe 2 weeks is long enough for carryover effects to wash out…)

Test if should go with lme (compound symmetry) or lme3 (AR1)

```
> grow.lme1 <- lme(root~fertilizer*week, data=grow, random=~1 |
   plant, method="ML")
> grow.lme2 <- lme(root~fertilizer*week, data=grow, random=~1 |
   plant, method="ML", correlation=corAR1())

> anova(grow.lme1,grow.lme2)
     Model df      AIC      BIC    logLik   Test    L.Ratio p-value
grow.lme1  1 12 88.79854 113.9307 -32.39927
grow.lme2  2 13 90.77983 118.0063 -32.38991 1 vs 2 0.01871329  0.8912
```

H_0: simpler model (lme) vs. **H_a:** more complex model (lme3)

P-value=0.8912 means that **lme (compound symmetry) suffices.**

Note: Must refit models via maximum likelihood (ML) so that the likelihood ratio test will be valid.

ANOVA table for fixed effects

```
> anova(grow.lme)
                 numDF denDF   F-value p-value
(Intercept)          1    40 1952.0103  <.0001
fertilizer           1    10   33.0633   2e-04
week                 4    40  712.5124  <.0001
fertilizer:week      4    40    5.9490   7e-04
```

Everything is significant!

The interaction will make interpretation more tricky…

Now fit this 2-way anova via AOV just to extract the LS means

```
> grow.lm <- aov(root~fertilizer*week+Error(plant), data=grow)
> model.tables(grow.lm, type="means")
```

Tables of means

```
Grand mean          5.023833

fertilizer
  added control
  5.678   4.370

week
    2     4     6     8    10
1.458 2.967 5.036 6.683 8.975

 fertilizer:week
           week
fertilizer 2     4     6     8     10
   added   1.667 3.683 5.972 7.450 9.617
   control 1.250 2.250 4.100 5.917 8.333
```

$\hat{\mu}$

Should not look at main effects (because of sig. interaction)

It seems more growth occurs when fertililizer is added (of course)

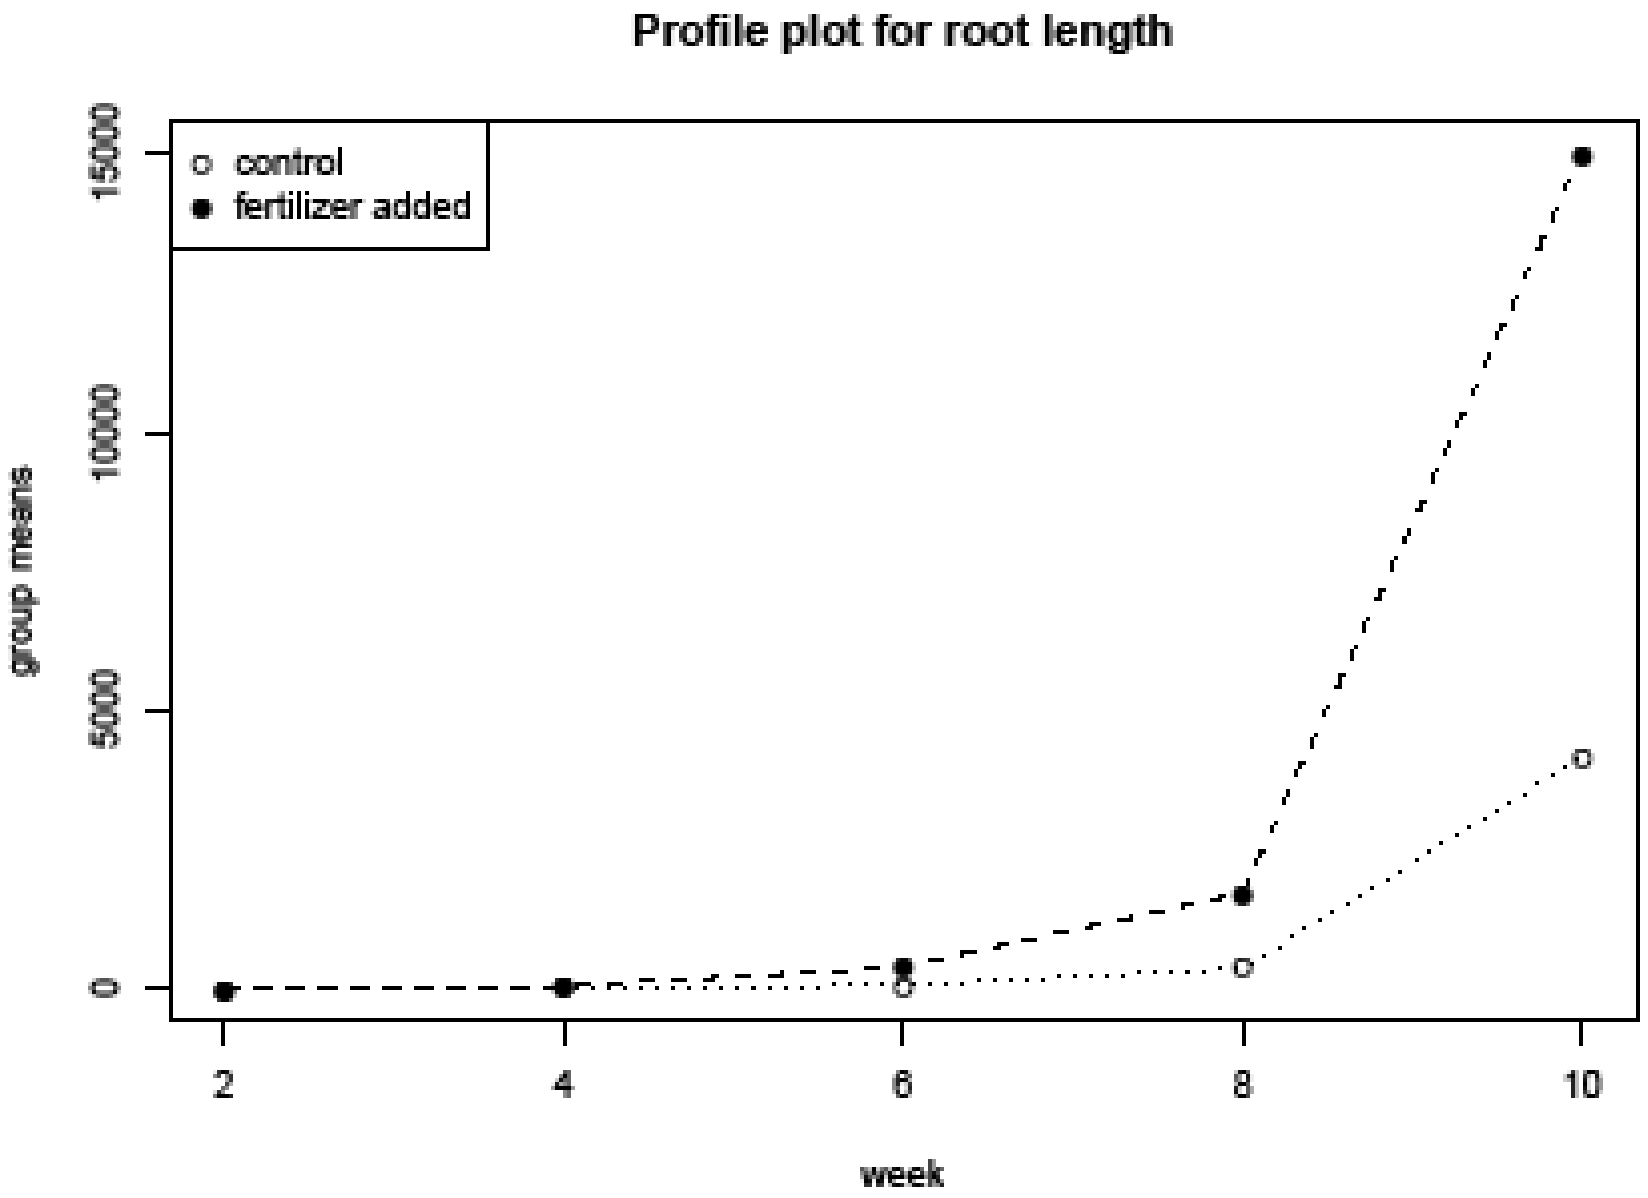
Profile plot for root length
control
fertilizer added
group means
0
5000
10000
15000
2
4
6
8
10
week

Diagnostics: Two sets, one for epsilon, the other for beta (plants)

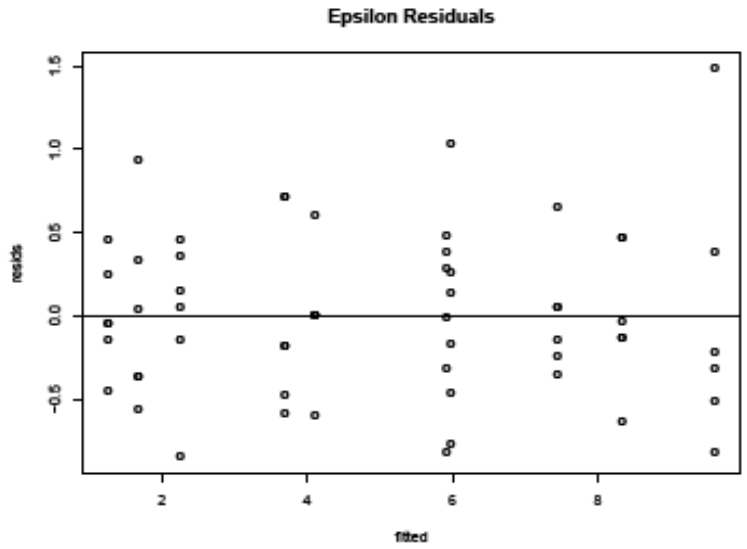

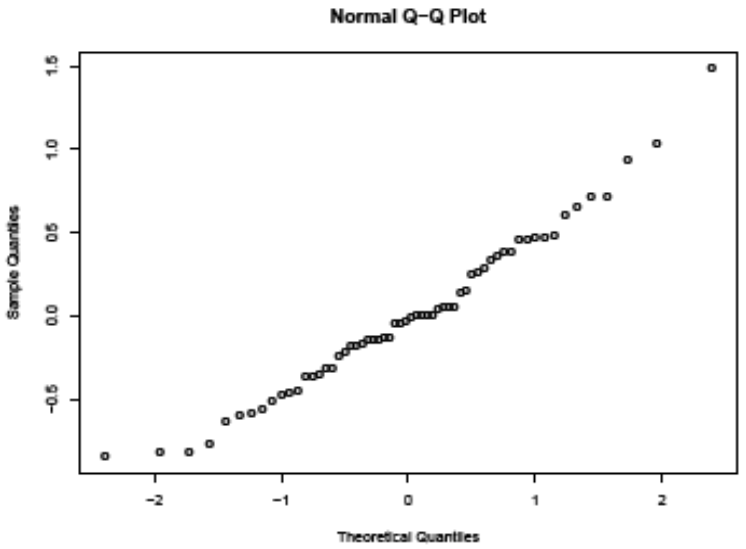

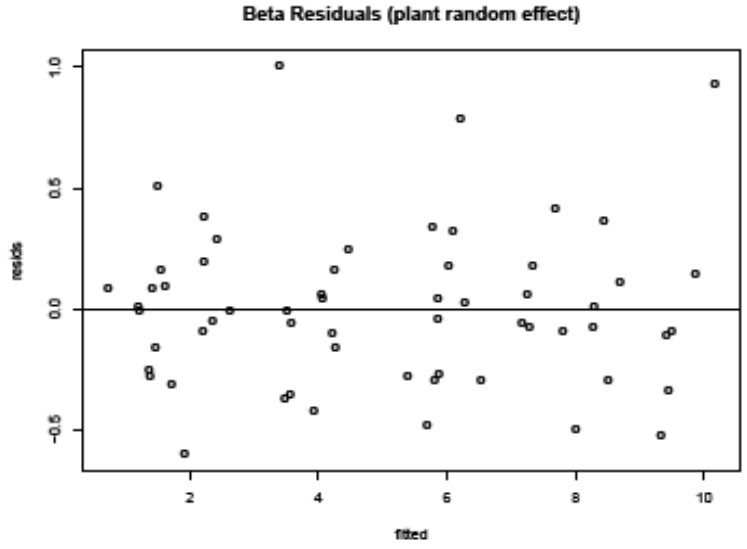

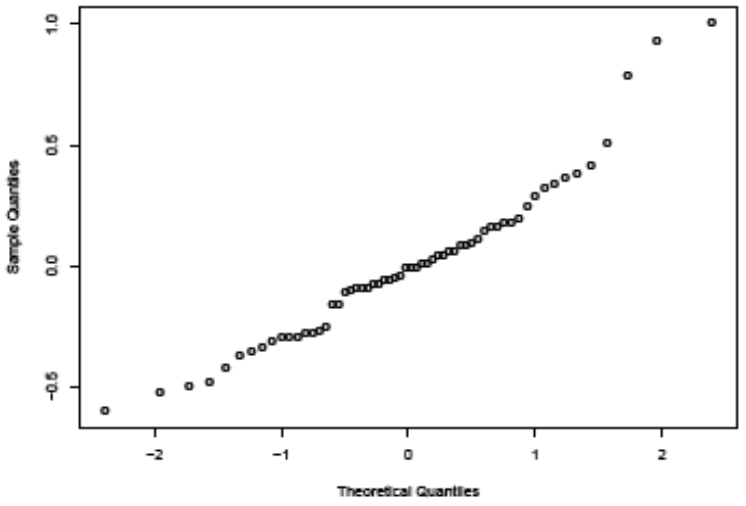

SAS

```
proc mixed;
class fertilizer week plant;
model root = fertilizer week fertilizer*week;
random plant(fertilizer);
repeated week / sub=plant(fertilizer) type=ar1 r rcorr;
lsmeans fertilizer*week / pdiff cl;
```

SPSS

proc mixed?