Chapter 19

More Complex ANOVA Designs

This chapter examines three designs that incorporate more factors and introduce some new elements of experimental design. They are three-way ANOVA, one-way nested ANOVA, and analysis of covariance (ANCOVA). These are common designs whose elements can be combined to generate even more elaborate ones. A useful guide to complex ANOVA designs is Winer et al. (1991), who provide a description and statistical model for each design. Once a particular design is identified, the statistical model can be used to program the analysis in SAS or other software.

19.1 Three-way ANOVA

We will first discuss three-way ANOVA, an analysis which examines how three different factors influence the means of the different groups. The three factors may be any combination of fixed or random effects and are typically referred to a Factors A, B, and C. In this design, there are one or more replicate observations for each combination of the three factors. The statistical analysis for three-way ANOVA designs may include F tests for the main effects of the factors as well as the interactions among them. For example, if the design has replication and all three factors are fixed, there are F tests for the main effects (Factor A, B, C), each pairwise interaction (A $\times$ B, A $\times$ C, B $\times$ C), and a three-way interaction (A $\times$ B $\times$ C). The additional complexity of this design with its many interactions can make interpretation

of the results quite challenging.

As an example of three-way ANOVA, we will analyze data from an experiment by Maestre & Reynolds (2007). This study examined how overall nutrient and water availability, and nutrient heterogeneity, affected grassland biomass production (Table 19.1). Nutrient heterogeneity was manipulated by placing the nitrogen at a particular location within the container vs. an even distribution. See Chapter 14 for further description of this experiment. We will use the notation Y_{ijkl} to reference the observations in three-way ANOVA designs. The i subscript refers to the group or treatment within Factor A (in this case nitrogen heterogeneity), j the treatment within Factor B (nitrogen levels), k the treatment within Factor C (water levels), while l refers to the observation within the treatment. For example, Y_{1134} refers to the fourth observation in the no nutrient heterogeneity, 40 mg N, 375 ml water treatment, which is 7.901.

Table 19.1: Example 1 - Effect of nitrogen heterogeneity, nitrogen availability, and water availability on the total biomass of grassland plants grown in microcosms (Maestre & Reynolds 2007). The table illustrates how the subscripts for Y_{ijkl} vary across treatments for a portion of the data set (see Chapter 21 for the full version).

N het. (Y/N)	N (mg)	Water (ml/week)	Y_{ijkl} = Biomass	i	j	k	l
N	40	125	4.372	1	1	1	1
N	40	125	4.482	1	1	1	2
N	40	125	4.221	1	1	1	3
N	40	125	3.977	1	1	1	4
N	40	250	7.400	1	1	2	1
N	40	250	8.027	1	1	2	2
N	40	250	7.883	1	1	2	3
N	40	250	7.769	1	1	2	4
N	40	375	7.226	1	1	3	1
N	40	375	8.126	1	1	3	2
N	40	375	6.840	1	1	3	3
N	40	375	7.901	1	1	3	4
etc.							
Y	120	250	10.731	2	3	2	1
Y	120	250	12.640	2	3	2	2
Y	120	250	10.350	2	3	2	3
Y	120	250	11.550	2	3	2	4
Y	120	375	14.697	2	3	3	1
Y	120	375	17.826	2	3	3	2
Y	120	375	14.711	2	3	3	3
Y	120	375	13.614	2	3	3	4

19.1.1 Three-way fixed effects model

Suppose that we want to model the observations in a study like Example 1, where there are Factors A, B, and C. Assume the design is factorial with every possible combination of the three factors, with $n > 1$ observations of each one. This design is often called three-way ANOVA with replication. A common model for the observations Y_{ijkl} in such designs (Winer et al. 1991) is

$$Y_{ijkl} = \mu + \alpha_i + \beta_j + + \gamma_k + (\alpha\beta)_{ij} + (\beta\gamma)_{jk} + (\alpha\gamma)_{ik} + (\alpha\beta\gamma)_{ijk} + \epsilon_{ijkl}. \quad (19.1)$$

Here μ is the grand mean of the observations, while α_i is the deviation from μ caused by the ith level or treatment of Factor A, β_j the deviation caused by the jth level of Factor B, and γ_k is the deviation caused by the kth level of Factor C. These terms are the **main effects** in the model. The terms $(\alpha\beta)_{ij}$, $(\beta\gamma)_{jk}$, and $(\alpha\gamma)_{ik}$ are pairwise or first-order interactions among Factors A and B, B and C, and A and C. These interactions are also symbolized as A $\times$ B, B $\times$ C, and A $\times$ C. They are similar to the interaction term in two-way ANOVA, but with three factors in the design there are more possibilities for interaction among them. The term $(\alpha\beta\gamma)_{ijk}$ models a second-order interaction (symbolized as A $\times$ B $\times$ C) among all three factors. It can be thought of as an interaction of interactions, i.e., the interaction between Factors A and B could change across levels of C. The ϵ_{ijkl} term represents the usual random departures from the mean value predicted by the main effects and interactions due to natural variability.

The objective in three-way ANOVA is to test whether Factor A, B, and C have an effect on the group means, and whether there are interactions among these factors. For Factor A this amounts to testing H_0 : all $\alpha_i = 0$, and similarly H_0 : all $\beta_j = 0$ for Factor B and H_0 : all $\gamma_k = 0$ for Factor C. For the A $\times$ B interaction, we would test $H_0 : (\alpha\beta)_{ij} = 0$, and similarly $H_0 : (\alpha\gamma)_{ik} = 0$ for the A $\times$ C and $H_0 : (\beta\gamma)_{jk} = 0$ for the B $\times$ C interactions. For the second-order interaction A $\times$ B $\times$ C, we are interested in testing H_0 : all $(\alpha\beta\gamma)_{ijk} = 0$. The F tests for these hypotheses can be constructed using various sums of squares and mean squares, similar to two-way ANOVA, and are also examples of likelihood ratio tests. We will not consider this process in detail but instead proceed directly to the analysis and interpretation of the Example 1 data set.

19.1.2 Three-way ANOVA for Example 1 - SAS demo

The first step in the program (see below) is to read in the observations using a `data` step, with the first variable (`nitrohet`) denoting the nitrogen heterogeneity treatment, while `nitrogen` and `water` represent the nitrogen and water levels. The variable `biomass` is then log-transformed before analysis, yielding the dependent variable `y = log10(biomass)`. Three separate plots then requested using `proc gplot` (SAS Institute Inc. 2014a), one for every pairwise combination of `nitrohet`, `nitrogen`, and `water`. These plots will allow us to examine the main effects and all first order (pairwise) interactions among the treatments. The choice as to whether a particular treatment is plotted on the x-axis or appears as separate groups (lines) on the graph is arbitrary. Like two-way ANOVA, if the lines are not parallel in a plot this suggests there is an interaction between the factors.

The second set of plots is intended to illustrate the second-order interaction among the three factors. Each plot illustrates the interaction between nitrogen and water at one level of nitrogen heterogeneity. If there is a second-order interaction, then the plots will appear different from one another.

The next section of the program conducts the three-way ANOVA using `proc glm` (SAS Institute Inc. 2014b). The `class` statement tells SAS that `nitrohet`, `nitrogen`, and `water` are used to classify the observations into the 18 different treatment groups. The `model` statement tells SAS the form of the ANOVA model. Recall that the model for fixed effects three-way ANOVA (Eq. 19.1). The statement `nitrohet|nitrogen|water` is SAS shorthand for this model, and will automatically generate all the possible main effects and interactions of the three factors.

The `lsmeans` statement causes `proc glm` to calculate quantities called least squares means for each level of `nitrohet`, `nitrogen`, and `water`. When the data are balanced these are equivalent to the means for each treatment group, but least squares means have some advantages for unbalanced data and other statistical models. The option `adjust=tukey` requests multiple comparisons among treatments using the Tukey method. This is useful for comparing the different levels of the main effects. However, tests for the main effects as well as multiple comparisons should be treated with caution in the presence of strong interaction (see Chapter 14 for discussion of this issue).

We now examine the results of the tests generated by SAS, examining the interactions first (see SAS output below). We are primarily interested in the results for Type III sums of squares. We see that the second-order

nitrogen heterogeneity × nitrogen × water interaction was nonsignificant ($F_{4,54} = 1.39, P = 0.2492$). The two graphs that illustrate this interaction appear similar, further indicating this interaction is weak or absent (Fig. 19.1, 19.2). Turning to the first order interactions, we see that the nitrogen heterogeneity × nitrogen interaction was nonsignificant ($F_{2,54} = 0.93, P = 0.4017$). In agreement with this result, the corresponding graph for this interaction (Fig. 19.3) suggests these two treatments are additive. The nitrogen × water interaction ($F_{4,54} = 12.90, P < 0.0001$) was highly significant. Examining Fig. 19.4, we see that the source of this interaction was a reduced effect of watering at lower nitrogen levels. The nitrogen heterogeneity × water interaction was also highly significant ($F_{2,54} = 13.10, P < 0.0001$). This interaction was apparently generated by a stronger effect of nitrogen heterogeneity at the lowest water level (Fig. 19.5). Overall, the significant interactions suggest that effects of these factors on biomass are not additive (Maestre & Reynolds 2007).

The SAS analysis also found highly significant main effects of nitrogen heterogeneity ($F_{1,54} = 144.14, P < 0.0001$), nitrogen ($F_{2,27} = 129.71, P < 0.0001$) and water ($F_{2,27} = 657.00, P < 0.0001$) on biomass. We can judge the strength of these effects through the interaction plots as well as the sum of squares values. Watering appears to have the largest effect on biomass, followed by nitrogen and nitrogen heterogeneity. The heterogeneity result is particularly intriguing, because more biomass was generated when this nutrient was heterogeneously distributed in space. Maestre & Reynolds (2007) suggest this occurred because nutrient patches encourage root proliferation, leading to increased nutrient uptake and overall growth. Even though there were significant interactions in this analysis, the main effects were larger and explained most of the variation in these data.

SAS program

```
* Maestre_biomass_3way.sas;
options pageno=1 linesize=80;
goptions reset=all;
title "Three-way ANOVA for biomass";
title2 "Data from Maestre and Reynolds (2007)";
data maestre;
        input nitrohet $ nitrogen water biomass;
        * Apply transformations here;
        y = log10(biomass);
        datalines;
N   40  125   4.372
N   40  125   4.482
N   40  125   4.221
N   40  125   3.977
N   40  250   7.400
N   40  250   8.027
N   40  250   7.883
N   40  250   7.769

etc.

Y  120  375  14.697
Y  120  375  17.826
Y  120  375  14.711
Y  120  375  13.614
;
run;
* Print data set;
proc print data=maestre;
run;
proc gplot data=maestre;
        plot y*nitrohet=nitrogen y*nitrogen=water y*nitrohet=water / vaxis=axis1
        haxis=axis1 legend=legend1;
        symbol1 i=std1mjt v=star height=2 width=3;
        axis1 label=(height=2) value=(height=2) width=3 major=(width=2) minor=none;
        legend1 label=(height=2) value=(height=2);
run;
* Sort data by nitrohet levels;
proc sort data=maestre;
        by nitrohet;
run;
* Plots to show three-way interaction;
proc gplot data=maestre;
```

```
        by nitrohet;
        plot y*nitrogen=water / vaxis=axis1 haxis=axis1 legend=legend1;
        symbol1 i=std1mjt v=star height=2 width=3;
        axis1 label=(height=2) value=(height=2) width=3 major=(width=2) minor=none;
        legend1 label=(height=2) value=(height=2);
run;
* Three-way ANOVA with all fixed effects;
proc glm data=maestre;
        class nitrohet nitrogen water;
        model y = nitrohet|nitrogen|water;
        lsmeans nitrohet nitrogen water / adjust=tukey cl lines;
        output out=resids p=pred r=resid;
run;
goptions reset=all;
title "Diagnostic plots to check anova assumptions";
* Plot residuals vs. predicted values;
proc gplot data=resids;
        plot resid*pred=1 / vaxis=axis1 haxis=axis1;
        symbol1 v=star height=2 width=3;
        axis1 label=(height=2) value=(height=2) width=3 major=(width=2) minor=none;run;
* Normal quantile plot of residuals;
proc univariate noprint data=resids;
        qqplot resid / normal waxis=3 height=4;
run;
quit;
```

SAS Output

```
                    Three-way ANOVA for biomass                        1
                 Data from Maestre and Reynolds (2007)
                                        09:49 Friday, June 7, 2013

        Obs    nitrohet    nitrogen    water    biomass       y

          1       N           40       125       4.372    0.64068
          2       N           40       125       4.482    0.65147
          3       N           40       125       4.221    0.62542
          4       N           40       125       3.977    0.59956
          5       N           40       250       7.400    0.86923
          6       N           40       250       8.027    0.90455
          7       N           40       250       7.883    0.89669
          8       N           40       250       7.769    0.89037
```

etc.

```
                    Three-way ANOVA for biomass                        3
                 Data from Maestre and Reynolds (2007)
                                        09:49 Friday, June 7, 2013

                          The GLM Procedure

                       Class Level Information

                 Class         Levels    Values

                 nitrohet           2    N Y

                 nitrogen           3    40 80 120

                 water              3    125 250 375


                Number of Observations Read          72
                Number of Observations Used          72
```

Three-way ANOVA for biomass 4
Data from Maestre and Reynolds (2007)
09:49 Friday, June 7, 2013

The GLM Procedure

Dependent Variable: y

Source	DF	Sum of Squares	Mean Square	F Value	Pr > F
Model	17	1.86010971	0.10941822	106.05	<.0001
Error	54	0.05571723	0.00103180		
Corrected Total	71	1.91582694			

R-Square	Coeff Var	Root MSE	y Mean
0.970917	3.492176	0.032122	0.919818

Source	DF	Type I SS	Mean Square	F Value	Pr > F
nitrohet	1	0.14872636	0.14872636	144.14	<.0001
nitrogen	2	0.26766625	0.13383312	129.71	<.0001
nitrohet*nitrogen	2	0.00191433	0.00095717	0.93	0.4017
water	2	1.35577897	0.67788949	657.00	<.0001
nitrohet*water	2	0.02702407	0.01351204	13.10	<.0001
nitrogen*water	4	0.05325694	0.01331423	12.90	<.0001
nitroh*nitroge*water	4	0.00574279	0.00143570	1.39	0.2492

Source	DF	Type III SS	Mean Square	F Value	Pr > F
nitrohet	1	0.14872636	0.14872636	144.14	<.0001
nitrogen	2	0.26766625	0.13383312	129.71	<.0001
nitrohet*nitrogen	2	0.00191433	0.00095717	0.93	0.4017
water	2	1.35577897	0.67788949	657.00	<.0001
nitrohet*water	2	0.02702407	0.01351204	13.10	<.0001
nitrogen*water	4	0.05325694	0.01331423	12.90	<.0001
nitroh*nitroge*water	4	0.00574279	0.00143570	1.39	0.2492

Three-way ANOVA for biomass 5
Data from Maestre and Reynolds (2007)
09:49 Friday, June 7, 2013

The GLM Procedure
Least Squares Means
Adjustment for Multiple Comparisons: Tukey

nitrohet	y LSMEAN	H0:LSMean1=LSMean2 Pr > |t|
N	0.87436837	<.0001
Y	0.96526708	

nitrohet	y LSMEAN	95% Confidence Limits	
N	0.874368	0.863635	0.885102
Y	0.965267	0.954534	0.976000

Least Squares Means for Effect nitrohet

i	j	Difference Between Means	Simultaneous 95% Confidence Limits for LSMean(i)-LSMean(j)	
1	2	-0.090899	-0.106077	-0.075720

Tukey Comparison Lines for Least Squares Means of nitrohet

LS-means with the same letter are not significantly different.

	y LSMEAN	nitrohet	LSMEAN Number
A	0.96526708	Y	2
B	0.87436837	N	1

etc.

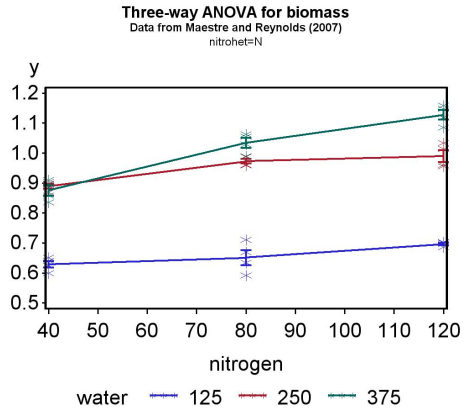

Figure 19.1: Means ± standard errors and data for the Example 1 experiment, where $Y = \log_{10}(\text{Biomass})$. This figure is the first of two figures illustrating any second order interaction, i.e., nitrogen heterogeneity × nitrogen × water.

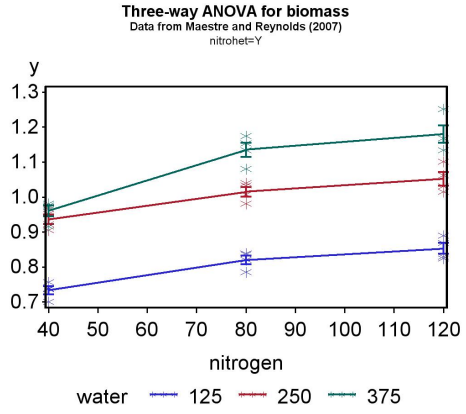

Figure 19.2: Means ± standard errors and data for the Example 1 experiment, where $Y = \log_{10}(\text{Biomass})$. This figure is the second of two figures illustrating any second-order interaction, i.e., nitrogen heterogeneity × nitrogen × water.

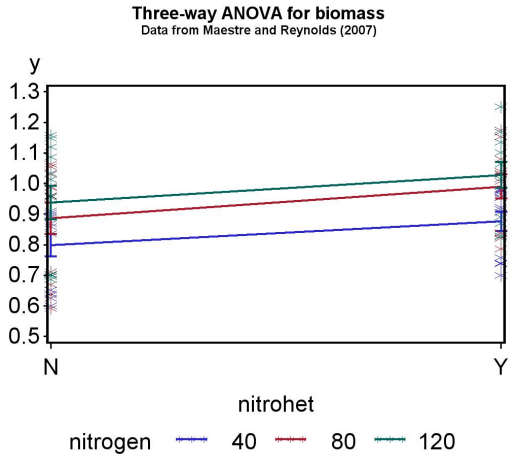

Figure 19.3: Means ± standard errors and data for the Example 1 experiment, where $Y = \log_{10}(\text{Biomass})$. This figure illustrates any nitrogen heterogeneity × nitrogen interaction, and the main effects of nitrogen heterogeneity and nitrogen.

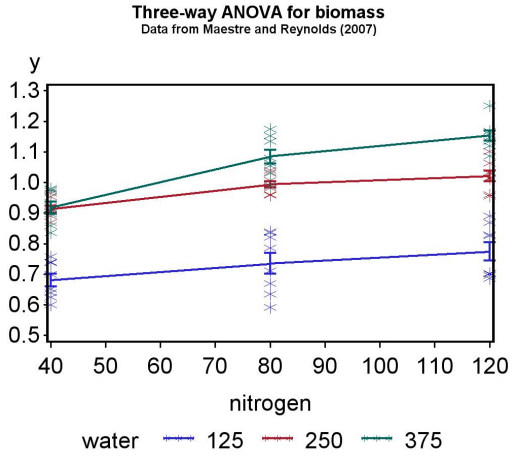

Figure 19.4: Means ± standard errors and data for the Example 1 experiment, where $Y = \log_{10}(\text{Biomass})$. This figure illustrates any nitrogen × water interaction, and the main effects of nitrogen and water.

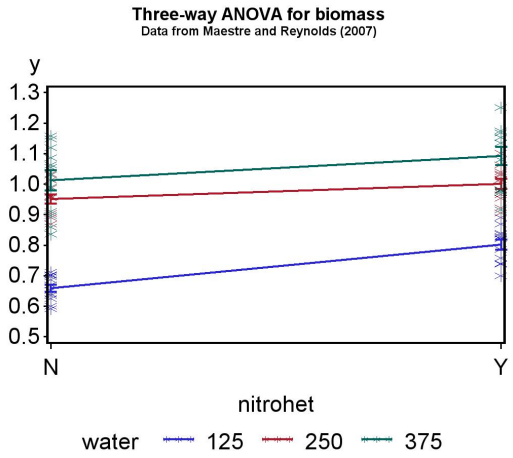

Figure 19.5: Means ± standard errors and data for the Example 1 experiment, where $Y = \log_{10}(\text{Biomass})$. This figure illustrates any nitrogen heterogeneity × water interaction, and the main effects of nitrogen heterogeneity and water.

19.1.3 Tests for main effects with interaction

As discussed in Chapter 14, there are questions as to whether tests of main effects are appropriate when interaction is significant, and these extend to three-way designs. As an alternative, we can use the `slice` option for `lsmeans` to avoid tests of the main effects. The modified SAS code is listed below along with the output. We first fit the full model including all the interactions, and observe that the nitrogen heterogeneity × nitrogen × water interaction is nonsignificant ($F_{4,54} = 1.39, P = 0.2492$), as is the nitrogen heterogeneity × nitrogen interaction ($F_{2,54} = 0.93, P = 0.4017$). We then drop these interactions and refit the model. The remaining two interactions are both highly significant in this reduced model (nitrogen heterogeneity × water, $F_{2,60} = 12.79, P < 0.0001$; nitrogen × water, $F_{4,60} = 12.61, P < 0.0001$). We skip the tests of the main effects because of these highly significant interactions, and instead use the `slice` option to test for a nitrogen heterogeneity effect at each water level, and vice versa. These tests were all highly significant, suggesting that nitrogen heterogeneity affects biomass at every water level, and water affects biomass at every nitrogen heterogeneity level. Similar tests could be conducted to examine the effects of nitrogen and water.

SAS Program

```
* Three-way ANOVA with interaction;
title3 "MODEL WITH ALL FOUR INTERACTIONS";
proc glm data=maestre;
        class nitrohet nitrogen water;
        model y = nitrohet|nitrogen|water / ss2;
        output out=resids p=pred r=resid;
run;
* Three-way ANOVA dropping ns interactions;
title3 "MODEL WITH ONLY SIGNIFICANT INTERACTIONS";
proc glm data=maestre;
        class nitrohet nitrogen water;
        model y = nitrohet nitrogen water nitrohet*water nitrogen*water / ss2;
        lsmeans nitrohet*water / slice=water slice=nitrohet;
run;
```

SAS Output

Three-way ANOVA for biomass 4
Data from Maestre and Reynolds (2007)
MODEL WITH ALL FOUR INTERACTIONS
15:27 Friday, November 8, 2013

The GLM Procedure

Dependent Variable: y

Source	DF	Sum of Squares	Mean Square	F Value	Pr > F
Model	17	1.86010971	0.10941822	106.05	<.0001
Error	54	0.05571723	0.00103180		
Corrected Total	71	1.91582694			

R-Square	Coeff Var	Root MSE	y Mean
0.970917	3.492176	0.032122	0.919818

Source	DF	Type II SS	Mean Square	F Value	Pr > F
nitrohet	1	0.14872636	0.14872636	144.14	<.0001
nitrogen	2	0.26766625	0.13383312	129.71	<.0001
nitrohet*nitrogen	2	0.00191433	0.00095717	0.93	0.4017
water	2	1.35577897	0.67788949	657.00	<.0001
nitrohet*water	2	0.02702407	0.01351204	13.10	<.0001
nitrogen*water	4	0.05325694	0.01331423	12.90	<.0001
nitroh*nitroge*water	4	0.00574279	0.00143570	1.39	0.2492

Three-way ANOVA for biomass 6
Data from Maestre and Reynolds (2007)
MODEL WITH ONLY SIGNIFICANT INTERACTIONS
15:27 Friday, November 8, 2013

The GLM Procedure

Dependent Variable: y

Source	DF	Sum of Squares	Mean Square	F Value	Pr > F
Model	11	1.85245259	0.16840478	159.44	<.0001
Error	60	0.06337435	0.00105624		
Corrected Total	71	1.91582694			

R-Square	Coeff Var	Root MSE	y Mean
0.966921	3.533291	0.032500	0.919818

Source	DF	Type II SS	Mean Square	F Value	Pr > F
nitrohet	1	0.14872636	0.14872636	140.81	<.0001
nitrogen	2	0.26766625	0.13383312	126.71	<.0001
water	2	1.35577897	0.67788949	641.80	<.0001
nitrohet*water	2	0.02702407	0.01351204	12.79	<.0001
nitrogen*water	4	0.05325694	0.01331423	12.61	<.0001

11:20 Monday, November 25, 2013

The GLM Procedure
Least Squares Means

nitrohet	water	y LSMEAN
N	125	0.65929804
N	250	0.95137559
N	375	1.01243148
Y	125	0.80223888
Y	250	1.00139663
Y	375	1.09216574

Three-way ANOVA for biomass 8
Data from Maestre and Reynolds (2007)
MODEL WITH ONLY SIGNIFICANT INTERACTIONS
11:20 Monday, November 25, 2013

The GLM Procedure
Least Squares Means

nitrohet*water Effect Sliced by water for y

water	DF	Sum of Squares	Mean Square	F Value	Pr > F
125	1	0.122592	0.122592	116.07	<.0001
250	1	0.015013	0.015013	14.21	0.0004
375	1	0.038145	0.038145	36.11	<.0001

Three-way ANOVA for biomass 9
Data from Maestre and Reynolds (2007)
MODEL WITH ONLY SIGNIFICANT INTERACTIONS
11:20 Monday, November 25, 2013

The GLM Procedure
Least Squares Means

nitrohet*water Effect Sliced by nitrohet for y

nitrohet	DF	Sum of Squares	Mean Square	F Value	Pr > F
N	2	0.854961	0.427481	404.72	<.0001
Y	2	0.527842	0.263921	249.87	<.0001

19.1.4 Other three-way designs

The Maestre & Reynolds (2007) experiment had four replicate containers for each treatment combination ($n = 4$), and so it was possible to fit a model with a second order interaction, namely nitrogen heterogeneity $\times$ nitrogen $\times$ water. Suppose now there was only observation for each treatment combination ($n = 1$). It is still possible to analyze these data using three-way ANOVA, but the data are not sufficient to fit a model with a second-order interaction. We would therefore use the model

$$Y_{ijk} = \mu + \alpha_i + \beta_j + +\gamma_k + (\alpha\beta)_{ij} + (\beta\gamma)_{jk} + (\alpha\gamma)_{ik} + \epsilon_{ijk}. \quad (19.2)$$

The equivalent model statement for `proc glm` would be

```
model y = nitrohet nitrogen water nitrohet*nitrogen nitrohet*water
nitrogen*water;
```

There is no shorthand method of specifying this model. The SAS output would be interpreted in the same way as the model with replication, except there would be no test for a second-order interaction.

Another common three-way design could have one or more factors that are random effects. For example, suppose that one manipulated nitrogen and water levels similar to Maestre & Reynolds (2007) but conducted the experiment in three different blocks, either different locations in the greenhouse or points in time. Block could be a random effect in this design, and the corresponding model would be

$$Y_{ijkl} = \mu + \alpha_i + \beta_j + C_k + (\alpha\beta)_{ij} + (\beta C)_{jk} + (\alpha C)_{ik} + (\alpha\beta C)_{ijk} + \epsilon_{ijkl}. \quad (19.3)$$

Here C stands for a random block effect, with $C \sim N(0, \sigma^2_C)$. Note that every interaction term involving C is also considered a random effect. This model could be analyzed with `proc mixed` (SAS Institute Inc. 2014b) using the following SAS statements:

```
proc mixed cl;
        class nitrogen water block;
        model y = nitrogen water nitrogen*water / ddfm=kr outp=resids;
        random block block*nitrogen block*water block*nitrogen*water;
run;
```

19.2 One-way nested ANOVA

The second design we will examine are called one-way nested designs. There are two factors in this design, a Factor A that may be a fixed or random effect, and a random nested Factor B. Nested means that for each level of Factor A, there are several levels of Factor B that are unique to that level of A. There are several replicate observations for each combination of Factor A and B.

As an example of this design, we will examine a genetic study of a minute parasitic wasp, *Anagrus delicatus* (Hymenoptera: Mymaridae). This wasp attacks eggs of the planthopper *Prokelisia marginata* (Homoptera: Delphacidae), a salt marsh insect that feeds on *Spartina* plants. Cronin & Strong (1996) were interested in the genetics of various wasp traits, including the number of eggs carried by the wasps themselves, ovipositor length, and various behavioral traits. They collected female wasps from three separate sites in San Franciso Bay and established genetically identical isolines from individual wasps collected from each site. They then measured the traits for a number of individuals from each isoline. Isolines are the nested factor in this design, because each isoline was established from a single site. Sites were classified as a fixed effect because there were essentially only three sites available for sampling, and so the sites were not randomly selected from a population of sites. Example 2 below shows a simulated data set based on this study, with three sites, 14 isolines per site, and eight individuals per isoline.

Table 19.2: Example 2 - Fecundity for *Anagrus delicatus* collected from three different sites, with 14 isolines per site and eight wasps per isoline. The data were simulated from results presented in Cronin and Strong (1996). Note that the values in the site, isoline, and wasp columns also correspond to the subscripts for Y_{ijk}. See Chapter 21 for the full version of this data set.

Site	Isoline	Wasp	Y_{ijk} = eggs
1	1	1	37
1	1	2	41
1	1	3	46
1	1	4	44
1	1	5	43
1	1	6	41
1	1	7	38
1	1	8	37
1	2	1	37
1	2	2	28
1	2	3	34
1	2	4	37
1	2	5	35
1	2	6	39
1	2	7	36
etc.			
3	13	1	36
3	13	2	39
3	13	3	36
3	13	4	30
3	13	5	37
3	13	6	32
3	13	7	38
3	13	8	39
3	14	1	32
3	14	2	34
3	14	3	41
3	14	4	33
3	14	5	35
3	14	6	35
3	14	7	34
3	14	8	31

19.2.1 Nested ANOVA models

Suppose that we want to model the observations in a study like Example 2, where there is a fixed Factor A and a nested Factor B. A common model for the observations Y_{ijk} in such designs (Winer et al. 1991) is

$$Y_{ijk} = \mu + \alpha_i + B_{j(i)} + \epsilon_{ijk}. \tag{19.4}$$

Here μ is the grand mean of the observations, α_i the deviation from μ caused by the ith level or treatment of Factor A, and $B_{j(i)}$ the random deviation caused by the jth level of Factor B within the ith level of Factor A. $B_{j(i)}$ is assumed to be normally distributed with mean zero and variance $\sigma^2_{B(A)}$, or $B_{j(i)} \sim N(0, \sigma^2_{B(A)})$, while $\epsilon_{ijk} \sim N(0, \sigma^2)$ as usual. $B_{j(i)}$ and ϵ_{ijk} are assumed to be independent. This model has two variance components, namely $\sigma^2_{B(A)}$ and σ^2.

The behavior of this model is illustrated in Fig. 19.6, for $a = 3$ levels of Factor A and $b = 4$ levels of Factor B nested within A. The figure illustrates how the value of α_i shifts the mean of the observations away from μ, similar to other ANOVA models. The $B_{j(i)}$ values, which are random variables, shift the observations for each nested level away from the values set by $\mu + \alpha_i$. Because they are random variables, the values of $B_{j(i)}$ are different for each level of Factor A.

The usual objectives for this nested ANOVA design are to test for Factor A effects, and estimate the variance components $\sigma^2_{B(A)}$ and σ^2. For Factor A this amounts to testing H_0 : all $\alpha_i = 0$. We will not consider this process in detail but proceed to the analysis and interpretation of the Example 2 data set. We will use `proc mixed` for the analysis because this design involves a mixed model.

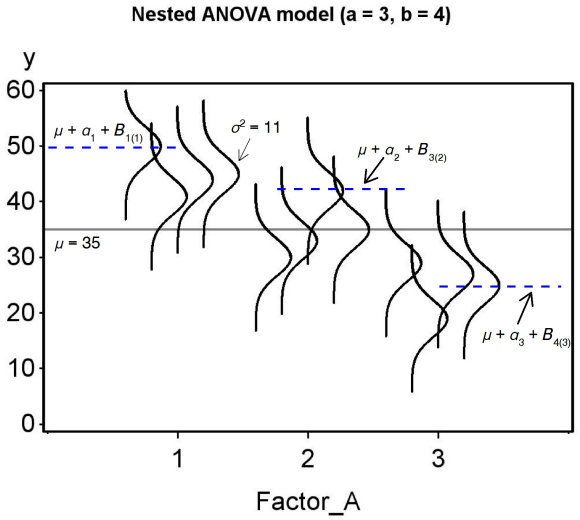

Figure 19.6: Mixed model for nested ANOVA showing the Factor A and B effects.

19.2.2 Nested ANOVA for Example 2 - SAS demo

The first step in analyzing the Example 2 data is to read the observations using a `data` step, with the variables `site` and `isoline` denoting the collection site and *Anagrus* isoline (see below). Although the isolines are numbered similarly across the three sites, note they are actually unique to each site and so are nested within sites. The variable `wasp` refers to a particular wasp within each isoline, but is not used in the analyses. Two plots are then requested using `proc gplot` (SAS Institute Inc. 2014a), one showing the mean for each site and so illustrating the site effect. The second plot shows the individual wasps color-coded by isoline, allowing for a visual comparison of variation among and within isolines. The x-axis position of each wasp is jittered to keep the points from overlapping. This involves adding a small random quantity to the `site` value, generating a new variable called `site_jit` that differs for each wasp.

The next section of the program conducts the nested ANOVA using `proc mixed` (SAS Institute Inc. 2014b). The `class` statement tells SAS that `site` and `isoline` are used to classify the observations. Next, the fixed effect site is listed in the `model` statement, while the random, nested effect of isoline

is incorporated in the `random` statement. SAS uses the syntax `isoline(site)` to indicate that isoline is nested within site. An `lsmeans` statement is used to compared the different sites using the Tukey method.

The analysis found no significant effect of site ($F_{2,39} = 2.3, P = 0.1323$) on the number of eggs per wasp (Fig. 19.7. The estimated variance among isolines within sites ($\hat{\sigma}^2_{B(A)} = \hat{\sigma}^2_{\text{site(isoline)}} = 10.17$) was substantial relative to the variance among wasps within isolines ($\hat{\sigma}^2 = 11.02$). This pattern can be observed in Fig. 19.8, with the observations for each isoline falling into discernable groups. This suggests that variation in egg number has a significant genetic component.

We can use the two variance components to estimate the heritability of egg number, which is the proportion of the variance due to genotypic vs. phenotypic differences among individuals (Falconer & Mackay 1996). The genotypic variance, V_G, is estimated by the variance among isolines within sites, because each isoline represents a different genetic group. For the wasp example, we have $V_G = \hat{\sigma}^2_{\text{site(isoline)}} = 10.17$. The environmental variance, V_E, is estimated by the variance among individuals within isolines, and represents variation among individuals not due to genotype. It is estimated by the variance among wasps within isolines, or $V_E = \hat{\sigma}^2 = 11.02$. The phenotypic variance is defined as the sum of the genotypic and environmental variance, or $V_P = V_G + V_E$. Heritability is then defined $h^2 = V_G/V_P = V_G/(V_G + V_E)$. It follows that $h^2 = 10.17/(10.17 + 11.02) = 0.48$ for the number of eggs in the wasps. This is relatively large value, suggesting that egg number could readily evolve in response to selection pressure.

SAS program

```
* Nested_ANOVA_Anagrus.sas;
options pageno=1 linesize=80;
goptions reset=all;
title "Nested ANOVA for fecundity";
title2 "Data simulated from Cronin and Strong (1996)";
data anagrus;
        input site isoline wasp eggs;
        * Apply transformations here;
        y = eggs;
        * Make jittered data for plots;
        site_jit = site + 0.1*rannor(0);
        datalines;
1   1  1  37
1   1  2  41
1   1  3  46
1   1  4  44
1   1  5  43
1   1  6  41
1   1  7  38
1   1  8  37

etc.

3  14  1  32
3  14  2  34
3  14  3  41
3  14  4  33
3  14  5  35
3  14  6  35
3  14  7  34
3  14  8  31
;
proc print data=anagrus;
run;
* Plot means and standard errors for each site;
proc gplot data=anagrus;
        plot y*site=1 / vaxis=axis1 haxis=axis1;
        symbol1 i=std1jmt v=none height=2 width=3;
        axis1 label=(height=2) value=(height=2) width=3 major=(width=2) minor=none;
run;
* Plot observations for each site and isoline;
proc gplot data=anagrus;
        plot y*site_jit=isoline / vaxis=axis1 haxis=axis1;
```

```
        symbol1 i=none v=dot height=0.5;
        axis1 label=(height=2) value=(height=2) width=3 major=(width=2) minor=none;
run;
* Nested ANOVA mixed model;
proc mixed cl data=anagrus;
        class site isoline;
        model y = site / ddfm=kr outp=resids;
        random isoline(site);
        * Compare levels of fixed effect using Tukey's HSD;
        lsmeans site / diff=all adjust=tukey cl adjdfe=row;
run;
goptions reset=all;
title "Diagnostic plots to check ANOVA assumptions";
* Plot residuals vs. predicted values;
proc gplot data=resids;
        plot resid*pred=1 / vaxis=axis1 haxis=axis1;
        symbol1 v=star height=2 width=3;
        axis1 label=(height=2) value=(height=2) width=3 major=(width=2) minor=none;
run;
* Normal quantile plot of residuals;
proc univariate noprint data=resids;
        qqplot resid / normal waxis=3 height=4;
run;
quit;
```

SAS output

Nested ANOVA for fecundity 1
Data simulated from Cronin and Strong (1996)
10:55 Tuesday, June 11, 2013

Obs	site	isoline	wasp	eggs	y	site_jit
1	1	1	1	37	37	1.10326
2	1	1	2	41	41	0.90939
3	1	1	3	46	46	1.18465
4	1	1	4	44	44	1.12283
5	1	1	5	43	43	1.09742
6	1	1	6	41	41	0.95798
7	1	1	7	38	38	1.11470
8	1	1	8	37	37	0.98907

etc.

329	3	14	1	32	32	2.89845
330	3	14	2	34	34	2.96535
331	3	14	3	41	41	3.02094
332	3	14	4	33	33	2.92618
333	3	14	5	35	35	2.93152
334	3	14	6	35	35	3.01175
335	3	14	7	34	34	3.06111
336	3	14	8	31	31	2.93977

Nested ANOVA for fecundity 8
Data simulated from Cronin and Strong (1996)
10:55 Tuesday, June 11, 2013

The Mixed Procedure

Model Information

Data Set	WORK.ANAGRUS
Dependent Variable	y
Covariance Structure	Variance Components
Estimation Method	REML
Residual Variance Method	Profile
Fixed Effects SE Method	Kenward-Roger
Degrees of Freedom Method	Kenward-Roger

Class Level Information

Class	Levels	Values
site	3	1 2 3
isoline	14	1 2 3 4 5 6 7 8 9 10 11 12 13 14

Dimensions

Covariance Parameters	2
Columns in X	4
Columns in Z	42
Subjects	1
Max Obs Per Subject	336

Number of Observations

Number of Observations Read	336
Number of Observations Used	336
Number of Observations Not Used	0

Iteration History

Iteration	Evaluations	-2 Res Log Like	Criterion
0	1	1965.68443676	
1	1	1841.14730382	0.00000000

Convergence criteria met.

Nested ANOVA for fecundity 9
Data simulated from Cronin and Strong (1996)
10:55 Tuesday, June 11, 2013

The Mixed Procedure

Covariance Parameter Estimates

Cov Parm	Estimate	Alpha	Lower	Upper
isoline(site)	10.1664	0.05	6.5003	18.1260
Residual	11.0187	0.05	9.4338	13.0417

Fit Statistics

-2 Res Log Likelihood	1841.1
AIC (smaller is better)	1845.1
AICC (smaller is better)	1845.2
BIC (smaller is better)	1848.6

Type 3 Tests of Fixed Effects

Effect	Num DF	Den DF	F Value	Pr > F
site	2	39	2.13	0.1323

Least Squares Means

Effect	site	Estimate	Standard Error	DF	t Value	Pr > |t|	Alpha
site	1	34.4821	0.9081	39	37.97	<.0001	0.05
site	2	34.2946	0.9081	39	37.77	<.0001	0.05
site	3	32.0982	0.9081	39	35.35	<.0001	0.05

Least Squares Means

Effect	site	Lower	Upper
site	1	32.6454	36.3188
site	2	32.4579	36.1313
site	3	30.2615	33.9349

Differences of Least Squares Means

Effect	site	_site	Estimate	Standard Error	DF	t Value	Pr > |t|	Adjustment

site	1	2	0.1875	1.2842	39	0.15	0.8847	Tukey
site	1	3	2.3839	1.2842	39	1.86	0.0710	Tukey

Nested ANOVA for fecundity 10
Data simulated from Cronin and Strong (1996)
08:29 Monday, November 11, 2013

The Mixed Procedure

Differences of Least Squares Means

Effect	site	_site	Estimate	Standard Error	DF	t Value	Pr > |t|	Adjustment
site	2	3	2.1964	1.2842	39	1.71	0.0951	Tukey

Differences of Least Squares Means

Effect	site	_site	Adj P	Alpha	Lower	Upper	Adj Lower	Adj Upper
site	1	2	0.9883	0.05	-2.4100	2.7850	-2.9411	3.3161
site	1	3	0.1651	0.05	-0.2136	4.9814	-0.7447	5.5126
site	2	3	0.2142	0.05	-0.4011	4.7939	-0.9322	5.3251

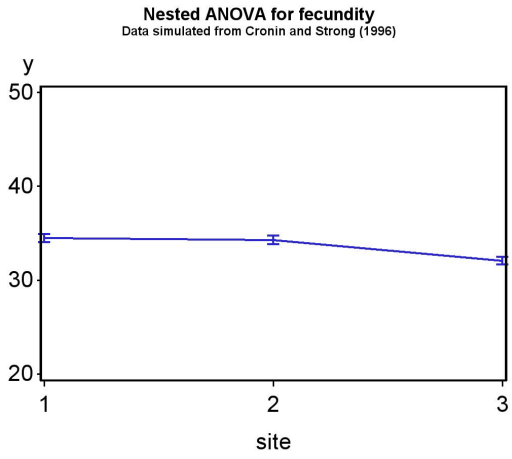

Figure 19.7: Means ± standard errors for each site in the Example 2 study, where Y = eggs.

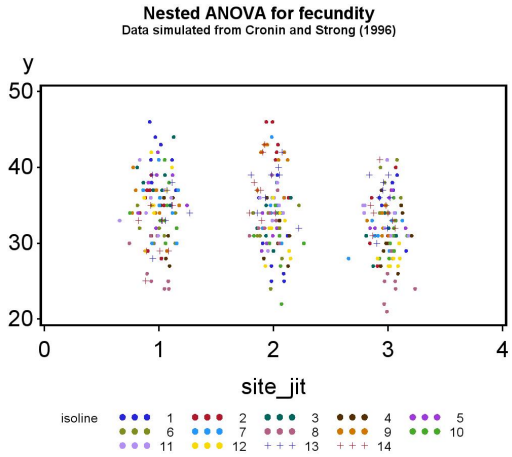

Figure 19.8: Observations for each site and isoline in the Example 2 study, where Y = eggs.

19.3 Analysis of covariance

Analysis of covariance, or ANCOVA, is a design that combines elements of ANOVA and regression. The simplest ANCOVA design is a combination of one-way ANOVA and linear regression. Factor A in the design is typically a fixed effect. For each observation in the design, a covariate X is measured and along with the dependent variable Y. The covariate X is thought to explain some level of variation in Y, and by including it in the design this may increase the power to detect treatment effects. Y is often assumed to be linearly related to X, although nonlinear relationships can be accomodated. More generally, a study might involve a mixture of factors and covariates, and the covariate effects may be of equal or greater interest than the factors.

As an example of ANCOVA, we will analyze a study of the fitness of adult *Thanasimus dubius*, a bark beetle predator, reared on an artificial diet vs. individuals collected from the wild (Reeve et al. 2003). The fitness variables measured were the total number of eggs laid (fecundity) and elytral length (Table 19.3). Body size and fecundity are often related in insects, so elytral length was used as a covariate in the analysis. This helps control for natural variation in body size to better see the treatment effect. The three treatments in the study were (1) artificial diet as larvae and *Ips grandicollis* as adults (`DietIG`), (2) artificial diet and cowpea weevils (`DietCPW`), and (3) wild adults fed cowpea weevils (`WildCPW`). The wild adults were collected from the field and so reared on natural prey as larvae. We will use the notation Y_{ij} to reference the observations in ANCOVA designs, with the i subscript refering to the Factor A or treatment group, while j is the observation within the treatment.

Table 19.3: Example 3 - Fitness of the predator *T. dubius*, reared on an artificial diet as larvae vs. wild individuals collected from the field (Reeve et al. 2003). See Chapter 21 for the full data set.

Y_{ij} = Eggs	X_{ij} = Length (mm)	Treatment	i	j
290	5.7	DietIG	1	1
99	5.2	DietIG	1	2
340	5.5	DietIG	1	3
271	4.8	DietIG	1	4
200	5.2	DietIG	1	5
etc.				
66	4.6	DietCPW	2	1
93	5.0	DietCPW	2	2
9	5.4	DietCPW	2	3
404	5.4	DietCPW	2	4
244	5.1	DietCPW	2	5
etc.				
62	4.7	WildCPW	3	1
290	5.0	WildCPW	3	2
488	5.8	WildCPW	3	3
336	5.2	WildCPW	3	4
337	5.8	WildCPW	3	5
etc.				

19.3.1 ANCOVA model

The following model is commonly used for simple ANCOVA designs (Winer et al. 1991). We have

$$Y_{ij} = \mu + \alpha_i + \beta(X_{ij} - \bar{X}) + \epsilon_{ij}, \tag{19.5}$$

where μ is the grand mean and α_i is the deviation from μ caused by the ith level of Factor A. The term X_{ij} is the value of the covariate for observation Y_{ij}, while $\bar{X}$ is the average of all the covariate values. The parameter β is the slope of the relationship between Y_{ij} and X_{ij}. This slope is assumed to be the same across all levels of Factor A. We will later see how to test this assumption. As usual, the model assumes $\epsilon_{ij} \sim N(0, \sigma^2)$.

The model can also be written in the form

$$Y'_{ij} = Y_{ij} - \beta(X_{ij} - \bar{X}) = \mu + \alpha_i + \epsilon_{ij}. \tag{19.6}$$

Displayed this way, we can see that ANCOVA is equivalent to carrying out a one-way ANOVA on values of Y_{ij} that have been adjusted for the covariate X, namely the values of Y'_{ij}.

Another adjustment of the model is needed by SAS and other statistical software. Combining some elements, the model can be written as

$$Y_{ij} = \mu' + \alpha_i + \beta X_{ij} + \epsilon_{ij}, \tag{19.7}$$

where $\mu' = \mu - \beta\bar{X}$. The quantity μ' represents a grand mean adjusted for the effect of the covariate. The objective in ANCOVA is to test whether Factor A and the covariate have an effect, and so test H_0 : all $\alpha_i = 0$ and $H_0 : \beta = 0$. However, before conducting these F tests we will first test whether the slopes across Factor A groups are identical by including an interaction term in the SAS model. If the slopes are significantly different, we have a scenario similar to ANOVA when interaction is present (see Chapter 14). Like ANOVA, when the interaction is significant tests of the main effects in ANCOVA, namely Factor A and the covariate X, may not make sense.

19.3.2 ANCOVA for Example 3 - SAS demo

The first step in the analysis (see program below) is to plot the number of eggs (`y`) for each treatment (`treat`) against elytral length, the covariate (`x`), using `proc gplot` (SAS Institute Inc. 2014a). This gives some idea whether

each treatment group has the same slope, a key assumption of ANCOVA. The slopes do appear to be similar (Fig. 19.9).

We then fit the ANCOVA model using `proc glm`, because all the effects in the model are fixed effects (SAS Institute Inc. 2014b). The first step is to fit a model with an interaction between the treatment and covariate, and examine the test for the interaction (see first SAS output below). We see that it is non-significant ($F_{2,35} = 0.02, P = 0.9781$), and so can assume the slopes are the same across treatments. We then rerun the program using the model without interaction. We see a highly significant effect of the covariate ($F_{1,37} = 9.99, P = 0.0031$), illustrating the typical strong relationship between body size and fecundity in insects. The treatment effect was nonsigificant ($F_{2,37} = 0.52, P = 0.5976$), implying the treatments themselves had no effect on egg numbers. Predators reared on the artificial diet are apparently similar to wild predators on this measure of fitness, controlling for elytral length and so body size.

The program also includes an `lsmeans` statement to calculate the least squares means for each treatment group, and test for differences among them using the Tukey method. Least squares means are means adjusted for the effect of other variables in the model, and in the case of ANCOVA are the treatment means adjusted for the covariate. In particular, they have the form

$$\bar{Y}_i(adj) = \bar{Y}_i - \hat{\beta}(\bar{X}_i - \bar{\bar{X}}). \tag{19.8}$$

We can see they are composed of two terms, the treatment means and the adjustment for the covariate. Treatment groups that have covariate means ($\bar{X}_i$ values) far from the overall covariate mean ($\bar{\bar{X}}$) receive a larger adjustment. No significant differences were found among the treatment groups, which is not surprising given the overall treatment effect was nonsignificant.

SAS Program

```
* ANCOVA_fitness.sas;
options pageno=1 linesize=80;
goptions reset=all;
title 'ANCOVA for T. dubius fitness';
data fitness;
        input eggs length treat $;
        * Choose y and x variables;
        y = eggs;
        x = length;
        datalines;
290    5.7   DietIG
 99    5.2   DietIG
340    5.5   DietIG
271    4.8   DietIG
200    5.2   DietIG

etc.

;
run;
* Print data set;
proc print data=fitness;
run;
* Plot data and regression line;
proc gplot data=fitness;
        plot y*x=treat / vaxis=axis1 haxis=axis1 legend=legend1;
        symbol1 i=rl v=star height=2 width=3;
        axis1 label=(height=2) value=(height=2) width=3 major=(width=2) minor=none;
        legend1 label=(height=2) value=(height=2);
run;
* ANCOVA;
proc glm data=fitness;
        class treat;
        * Model with interaction;
        *model y = treat x treat*x;
        * Model without interaction;
        model y = treat x;
        lsmeans treat / pdiff=all adjust=tukey cl lines;
        output out=resids p=pred r=resid;
run;
goptions reset=all;
title "Diagnostic plots to check ANCOVA assumptions";
* Plot residuals vs. predicted values;
```

```
proc gplot data=resids;
        plot resid*pred=1 / vaxis=axis1 haxis=axis1;
        symbol1 v=star height=2 width=3;
        axis1 label=(height=2) value=(height=2) width=3 major=(width=2) minor=none;
run;
* Normal quantile plot of residuals;
proc univariate noprint data=resids;
        qqplot resid / normal waxis=3 height=4;
run;
quit;
```

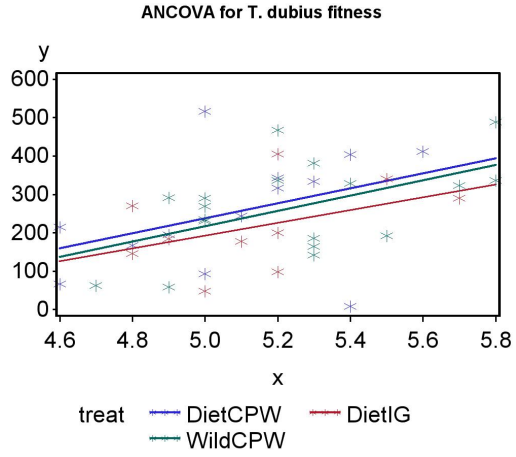

Figure 19.9: Eggs laid by adult *T. dubius* in three treatments vs. elytra length.

SAS Output - Model with Interaction

ANCOVA for T. dubius fitness 3
13:26 Thursday, September 26, 2013

The GLM Procedure

Dependent Variable: y

Source	DF	Sum of Squares	Mean Square	F Value	Pr > F
Model	5	149241.7740	29848.3548	2.13	0.0845
Error	35	489963.3479	13998.9528		
Corrected Total	40	639205.1220			

R-Square	Coeff Var	Root MSE	y Mean
0.233480	47.29918	118.3172	250.1463

Source	DF	Type I SS	Mean Square	F Value	Pr > F
treat	2	16193.0211	8096.5105	0.58	0.5661
x	1	132427.1693	132427.1693	9.46	0.0041
x*treat	2	621.5837	310.7918	0.02	0.9781

Source	DF	Type III SS	Mean Square	F Value	Pr > F
treat	2	396.6464	198.3232	0.01	0.9859
x	1	114086.8726	114086.8726	8.15	0.0072
x*treat	2	621.5837	310.7918	0.02	0.9781

SAS Output - Model without Interaction

ANCOVA for T. dubius fitness 1
08:29 Monday, November 11, 2013

Obs	eggs	length	treat	y	x
1	290	5.7	DietIG	290	5.7
2	99	5.2	DietIG	99	5.2
3	340	5.5	DietIG	340	5.5
4	271	4.8	DietIG	271	4.8
5	200	5.2	DietIG	200	5.2

etc.

ANCOVA for T. dubius fitness 2
08:29 Monday, November 11, 2013

The GLM Procedure

Class Level Information

Class	Levels	Values
treat	3	DietCPW DietIG WildCPW

Number of Observations Read	41
Number of Observations Used	41

ANCOVA for T. dubius fitness 3
08:29 Monday, November 11, 2013

The GLM Procedure

Dependent Variable: y

Source	DF	Sum of Squares	Mean Square	F Value	Pr > F
Model	3	148620.1904	49540.0635	3.74	0.0193
Error	37	490584.9316	13259.0522		

Corrected Total	40	639205.1220

R-Square	Coeff Var	Root MSE	y Mean
0.232508	46.03224	115.1480	250.1463

Source	DF	Type I SS	Mean Square	F Value	Pr > F
treat	2	16193.0211	8096.5105	0.61	0.5484
x	1	132427.1693	132427.1693	9.99	0.0031

Source	DF	Type III SS	Mean Square	F Value	Pr > F
treat	2	13846.2749	6923.1375	0.52	0.5976
x	1	132427.1693	132427.1693	9.99	0.0031

ANCOVA for T. dubius fitness 4
08:29 Monday, November 11, 2013

The GLM Procedure
Least Squares Means
Adjustment for Multiple Comparisons: Tukey-Kramer

treat	y LSMEAN	LSMEAN Number
DietCPW	270.496170	1
DietIG	221.056056	2
WildCPW	251.610513	3

Least Squares Means for effect treat
Pr > |t| for H0: LSMean(i)=LSMean(j)

Dependent Variable: y

i/j	1	2	3
1		0.5702	0.8992
2	0.5702		0.7839

```
         3        0.8992        0.7839


treat           y LSMEAN      95% Confidence Limits

DietCPW       270.496170      205.021331   335.971009
DietIG        221.056056      147.207956   294.904156
WildCPW       251.610513      195.890594   307.330433


          Least Squares Means for Effect treat

                Difference          Simultaneous 95%
                   Between       Confidence Limits for
i    j               Means        LSMean(i)-LSMean(j)

1    2          49.440114       -69.094011   167.974239
1    3          18.885656       -85.958935   123.730248
2    3         -30.554457      -142.397015    81.288100


Tukey-Kramer Comparison Lines for Least Squares Means of treat

LS-means with the same letter are not significantly different.

                                          LSMEAN
                     y LSMEAN    treat    Number

                A     270.496    DietCPW       1
                A
                A     251.611    WildCPW       3
                A
                A     221.056    DietIG        2
```

19.4 References

Cronin, J. T. & Strong, D. R. (1996) Genetics of oviposition success of a thelytokous fairyfly parasitoid, *Anagrus delicatus*. *Heredity* 76: 43-54.

Falconer, D. S. & MacKay, T. F. C. (1996) *Introduction to Quantitative Genetics*, 4th edition. Longman Group Ltd., Essex, England.

Maestre, F. T. & Reynolds, J. F. (2007) Amount or pattern? Grassland responses to the heterogeneity and availability of two key resources. *Ecology* 88: 501-511.

Reeve, J. D., Rojas, M. G. & Morales-Ramos, J. A. (2003) Artificial diet and rearing methods for *Thanasimus dubius* (Coleoptera: Cleridae), a predator of bark beetles (Coleoptera: Scolytidae). *Biological Control* 27: 315-322.

SAS Institute Inc. (2014a) *SAS/GRAPH 9.4: Reference, Third Edition.* SAS Institute Inc., Cary, NC.

SAS Institute Inc. (2014b) *SAS/STAT 13.2 Users Guide.* SAS Institute Inc., Cary, NC.

Winer, B. J., Brown, D. R. & Michels, K. M. (1991) *Statistical Principles in Experimental Design*, 3rd edition. McGraw-Hill, Inc., Boston, MA.

19.5 Problems

1. A limnologist wants to examine the length of a zooplankton species reared using four different algal growth media (1, 2, 3, and 4). She is also interested in whether there is variation among the containers used to rear the organisms. An experiment is conducted where three containers are used for each rearing medium, for a total of 12 different containers. The containers were randomly selected from a box of containers. The length of four animals was determined for each container, yielding the following data:

Medium	Container	Lengths 1-4 (mm)
1	1	3.1, 3.0, 3.2, 3.0
1	2	3.3, 3.6, 2.8, 2.5
1	3	3.7, 3.4, 3.4, 3.6
2	1	2.7, 2.9, 3.2, 3.0
2	2	2.9, 3.4, 3.5, 2.9
2	3	3.5, 3.5, 3.7, 4.0
3	1	2.8, 2.7, 1.8, 2.5
3	2	2.6, 2.5, 3.2, 2.4
3	3	2.6, 2.9, 1.8, 2.4
4	1	4.1, 4.6, 3.3, 4.5
4	2	3.7, 3.9, 4.0, 3.9
4	3	4.4, 4.4, 3.9, 4.6

   (a) Write an appropriate ANOVA model for this design, stating which factors are fixed, random, and possibly nested.

   (b) Use SAS to analyze these data using your ANOVA model, transforming the observations only if necessary. Is there a significant difference among the four media in zooplankton length?

   (c) Use the Tukey method to compare the media treatments. Interpret your results.

   (d) Compare the magnitude of your variance components. Does there appear to be much variation among containers?

2. An ecologist is interested in the effect of three management treatments (labeled 1, 2, and 3) on the abundance of an endangered snail. Treatment 2 is a control treatment. Twenty-four plots are established and

the three treatments assigned at random to the plots. The density of snails is then measured at a later time, as well as a covariate in the form of a habitat index. Larger values of the habitat index are thought to indicate better snail habitat. See data set below.

Treatment	Index	Snails
1	9.3	23.0
1	9.8	24.9
1	9.9	24.7
1	10.1	24.6
1	8.9	23.4
1	10.8	27.1
1	9.6	25.4
1	10.7	25.4
2	11.9	21.8
2	9.6	18.8
2	10.3	21.0
2	10.8	21.5
2	9.9	20.9
2	10.9	22.6
2	8.9	19.8
2	10.2	22.4
3	11.2	23.4
3	10.3	18.5
3	11.1	22.3
3	9.8	20.5
3	11.2	20.5
3	8.7	18.4
3	8.4	18.7
3	10.5	19.2

(a) Test for equality of slopes among the different treatment groups using SAS. Is this key assumption of ANCOVA satisfied?

(b) Use ANCOVA and SAS to test for overall treatment and covariate effects in this experiment, and the Tukey method to compare the different treatments. Interpret and discuss your results. Is there a significant treatment and covariate effect? How do the different treatments compare?

3. A scientist interested in aquaculture raises fish using three kinds of treatments in a factorial design. There were two fish diets (A and B), two strains of fish (1 and 2), and three temperatures (22^o, 24^o, and 26^oC). Two fish were reared for each combination of the treatments. The following data were obtained:

Diet	Strain	Temp	Weight (lb)
A	1	22	5.5
A	1	22	5.8
A	1	24	5.9
A	1	24	5.7
A	1	26	6.2
A	1	26	5.9
A	2	22	5.2
A	2	22	5.0
A	2	24	5.4
A	2	24	5.6
A	2	26	5.0
A	2	26	4.9
B	1	22	5.4
B	1	22	4.8
B	1	24	5.4
B	1	24	5.4
B	1	26	5.7
B	1	26	5.5
B	2	22	5.2
B	2	22	4.8
B	2	24	5.1
B	2	24	5.1
B	2	26	4.8
B	2	26	4.5

   (a) Write an appropriate ANOVA model for this design, stating which factors are fixed or random.

   (b) Use SAS to analyze these data using your ANOVA model, transforming the observations only if necessary. Interpret the results of your analysis.