

Mixed Models

Jing Cheng Gayla Olbricht Nilupa Gunaratna Rebecca Kendall
Alex Lipka Sudeshna Paul Benjamin Tyner

May 19, 2005

Contents

1	Introduction	3
2	Two-Way Mixed Effects Models	3
2.1	Pearl Data Example	3
2.1.1	PROC MIXED	7
2.2	SPSS	8
2.2.1	GLM Procedure: Comparable to PROC GLM in SAS	9
2.2.2	Mixed Procedure: Comparable to PROC MIXED in SAS	10
3	Comparison of PROC GLM and PROC MIXED in SAS	13
4	Crossover designs	14
4.1	Introduction	14
4.2	Why Crossover and its disadvantage	15
4.3	Analysis of structure and confounding	16
4.3.1	Example 1	16
4.3.2	Data summary	17
5	Repeated Measures	20
5.1	Covariance Structure	20
5.1.1	Unstructured	20
5.1.2	ARMA(1,1)	21
5.1.3	Equicorrelation	21
5.1.4	Uncorrelated	21
5.2	Software Packages	22
5.2.1	Dental Data Set Analysis in SAS	22
5.2.2	SAS PROC GLM	23
5.2.3	SAS PROC MIXED	24
5.2.4	Example: Three-Factor Factorial Design with One Repeated Measure	26
5.2.5	Repeated Measures in SPSS	30
6	Appendix	31
6.1	Pearl data as entered into SAS	31

1 Introduction

In many experimental design situations, one or more of the factors in the study may be random factors. That is, the levels of those factors are actually a sample from a larger population of levels and inferences are desired about the population of factor levels (e.g., the variance of the population of factor levels). When one or more of these random factors are examined along with one or more fixed factors, a mixed model approach is needed to analyze such data. In this paper, we give a basic introduction of a two-way mixed effects model. Our main focus is to demonstrate how to use different procedures in SPSS and SAS to analyze such data.

2 Two-Way Mixed Effects Models

2.1 Pearl Data Example

The data set we consider here is obtained from Neter et. al. (1996). The data pertain to the production of imitation pearls. Researchers were interested in studying the effects of two factors on the market value of the pearl. The main factor of interest (Factor A) is the number of coats of a special lacquer applied to the opalescent plastic bead used as the base of the pearl. Factor A has three levels (6, 8, and 10 coats) that were fixed in advance. The other factor (Factor B) was the batch of the beads used in the study. Each batch consisted of 12 beads equally divided among the 3 levels of factor A. There were a total of 4 batches which can be regarded as a random sample of batches from the bead production process. Information about the variance between batches could give researchers an indication of the quality of their production process since each batch was made at the same time. The data as entered into SAS is given in the appendix.

The model for these data is given below:

$$\begin{aligned} Y_{ijk} &= \mu + A_i + B_j + (AB)_{ij} + \varepsilon_{ijk} \\ \sum A_i &= 0 \\ B_j &\sim \mathcal{N}(0, \sigma_B^2) \\ (AB)_{ij} &\sim \mathcal{N}(0, \sigma_{AB}^2) \\ \varepsilon_{ijk} &\sim \mathcal{N}(0, \sigma^2) \end{aligned}$$

where Y_{ijk} is the market value of the k^{th} pearl receiving the i^{th} level of factor A and the j^{th} level of factor B; $i = 1, 2, 3$; $j = 1, 2, 3, 4$; $k = 1, 2, 3, 4$.

It is important to note that two types of models are commonly used for mixed models. They are called the restricted and unrestricted models. The unrestricted model assumptions are limited to those listed above, while the restricted model imposes the additional assumption that $\sum_{i=1}^3 (AB)_{ij} = 0$ for all j .

Data analysis for the pearl data is given below for both SAS and SPSS. Both of these software packages base calculations on the unrestricted mixed effects model. Based on expected mean squares (EMS), the $MS(\text{coat} \times \text{batch})$ is used as the denominator in the test for the fixed coat effect and the random batch effect. The available options for running mixed models are very comparable between SAS and SPSS. These similarities are noted in the descriptions below. Finally, since the output from the two programs is also similar, output from only one of the programs is given per procedure.

SAS: There are two procedures that can be used to obtain results for mixed models. These are: PROC GLM and PROC MIXED. Examples of how to use these procedures are given below. It should be noted that PROC MIXED is recommended for this type of analysis to avoid the pitfalls of PROC GLM.

PROC GLM: Uses the method of least squares to fit general linear models. No other parameter estimation method can be specified. This procedure is comparable to analyzing mixed models in SPSS by clicking: Analyze >> General Linear Models >> Univariate

```
proc glm data=pearl;
class coat batch;
model market_value=coat batch coat*batch;
random batch coat*batch/test;
means coat/tukey lines e=coat*batch;
run;
```

Figure 1: Code

Explanation:

Class Specify all factors in the model.

Model Specify the model by equating the response variable to all fixed and random effects.

Random Specify which effects are random (including all random interactions).

Test Performs hypothesis tests for each effect specified in the model, using appropriate error terms as determined by the expected mean squares.

Means Specify the fixed factors of interest. The arithmetic means and standard deviations for all factors that are specified will be calculated for each factor level. Note that the arithmetic means are not adjusted for other effects in the model. For adjusted means, the `lsmeans` statement should be used. This is especially important for unbalanced data.

Tukey Specify the type of multiple comparison adjustment desired for testing all pairwise comparison. The `lines` option specifies how the output should be displayed. There are several options to adjust for multiple comparisons. These are: `BON`, `DUNCAN`, `DUNNETT`, `DUNNETTL`, `DUNNETTU`, `GABRIEL`, `GT2`, `LSD`, `REGWQ`,

E= Specify the error mean square to use in the multiple comparisons. By default, the overall residual or MSE is used.

The GLM Procedure

Dependent Variable: market_value

Source	DF	Sum of Squares	Mean Square	F Value	Pr > F
Model	11	305.0916667	27.7356061	5.75	<.0001
Error	36	173.6250000	4.8229167		
Corrected Total	47	478.7166667			

R-Square	0.637312	Coeff Var	2.904593	Root MSE	2.196114	market_value Mean	75.60833
----------	----------	-----------	----------	----------	----------	-------------------	----------

Source	DF	Type III SS	Mean Square	F Value	Pr > F
coat	2	150.3879167	75.1939583	15.59	<.0001
batch	3	152.8516667	50.9505556	10.56	<.0001
coat*batch	6	1.8520833	0.3086806	0.06	0.9988

Figure 2: Output (1)

Output (1) shows tests of all model effects. However, this output does not use the correct error term to test the coat and batch effects. The only test that should be interpreted is the test for the interaction of coat and batch. This output reveals that the variance term for the coat*batch interaction is not significantly greater than zero (p-value 0.9988)

The GLM Procedure

Source	Type III Expected Mean Square
coat	$\text{Var}(\text{Error}) + 4 \text{Var}(\text{coat*batch}) + Q(\text{coat})$
batch	$\text{Var}(\text{Error}) + 4 \text{Var}(\text{coat*batch}) + 12 \text{Var}(\text{batch})$
coat*batch	$\text{Var}(\text{Error}) + 4 \text{Var}(\text{coat*batch})$

Figure 3: Output (2)

Output (2) shows the expected mean squares for each of the model terms.

```

The GLM Procedure

Tests of Hypotheses for Mixed Model Analysis of Variance

Dependent Variable: market_value

Source                DF      Type III SS      Mean Square      F Value      Pr > F
coat                  2        150.387917        75.193958        243.60       <.0001
batch                 3        152.851667        50.950556        165.06       <.0001
Error                 6          1.852083          0.308681
Error: MS(coat*batch)

Source                DF      Type III SS      Mean Square      F Value      Pr > F
coat*batch            6          1.852083          0.308681          0.06        0.9988
Error: MS(Error)     36        173.625000        4.822917

```

Figure 4: Output (3)

Output (3) shows tests of all model effects with correct error terms. It is apparent there is a significant coat effect (p-value < 0.0001) and the variance term for batch is significantly greater than zero (p-value < 0.0001). The test for coat*batch is the same as in output (1).

```

Tukey's Studentized Range (HSD) Test for market_value

NOTE: This test controls the Type I experimentwise error rate,
but it generally has a higher Type II error rate than REGWQ.

Alpha                0.05
Error Degrees of Freedom      6
Error Mean Square          0.308681
Critical Value of Studentized Range  4.33902
Minimum Significant Difference    0.6027

Means with the same letter are not significantly different.

Tukey Grouping      Mean      N      coat
A                   76.9250   16     3
A
A                   76.7938   16     2
B                   73.1063   16     1

```

Figure 5: Output (4)

Output (4) gives the Tukey pairwise comparison results. The means for each of the coat levels are given. It appears the means for pearls with eight coat applications (coat = 2) and ten coat applications (coat=3) are not significantly different. However, the mean for pearls given six coats (coat=1) is significantly different from the other two coat applications.

```

proc mixed data=pearl method=type3;
class coat batch;
model market_value=coat;
random batch coat*batch;
lsmeans coat/alpha=0.05 cl diff adjust=tukey;
run;

```

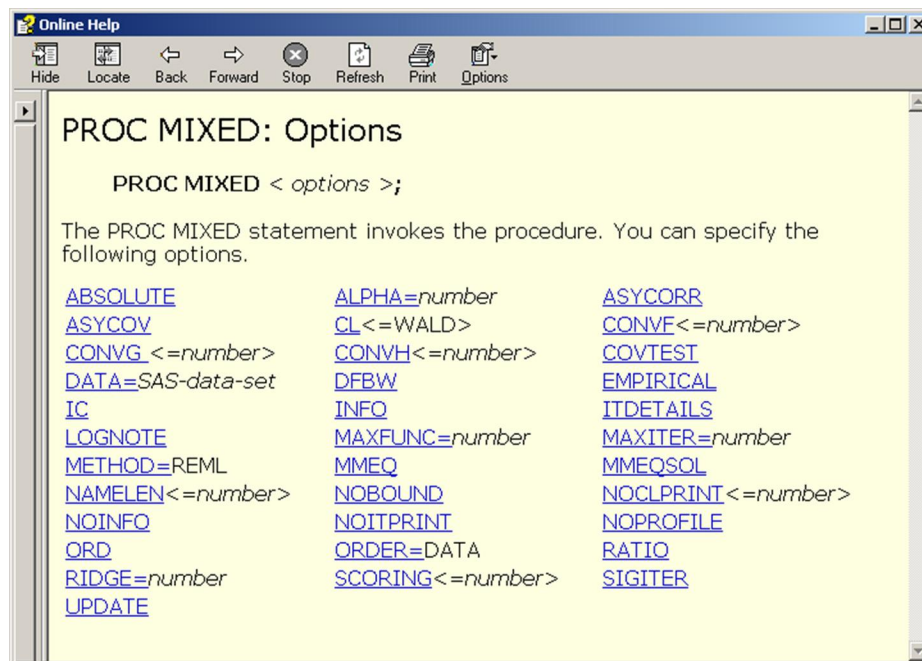
Figure 6: Code

2.1.1 PROC MIXED

Fits a variety of mixed linear models to data and allows specification of the parameter estimation method to be used. This procedure is comparable to analyzing mixed models in SPSS by clicking: Analyze >> Mixed Models >> Linear

Explanation:

The following window from the SAS help menu shows the options available within the PROC MIXED statement. To gain a better understanding of these options, exploring the SAS help menu is recommended.



An important option to note from this menu is the **METHOD=** option.

This allows the user to specify the estimation method for the covariance parameters. Note in this example **TYPE3** was specified to conform with GLM for this balanced data, but this type of estimation is not necessarily recommended and results from REML are shown in SPSS. The following estimation procedures are available:

METHOD= REML (Residual (Restricted) Maximum Likelihood) – Default Method

METHOD= ML (Maximum Likelihood)

METHOD= MIVQUE0 (Minimum Variance Quadratic Unbiased Estimation) – Same as MINQUE in SPSS

METHOD= TYPE1 (Method of Moments – ANOVA Table, Type I SS)

METHOD= TYPE2 (Method of Moments – ANOVA Table, Type II SS)

METHOD= TYPE3 (Method of Moments – ANOVA Table, Type III SS)

Explanation (Cont):

Class Specify all factors in the model.

Model Specify the model by equating the response variable to all fixed effects. Note how this differs from PROC GLM. There are also several options for this statement that can be further explored via the SAS help menu.

Random Specify which effects are random (including all random interactions). No Test statement is available as in PROC GLM. However, there are also several options for this statement that can be further explored via the SAS help menu.

Lsmeans Specify the fixed factors of interest. This statement will generate the least-squares means of fixed effects. The Means statement is not an option in PROC MIXED.

Cl Requests that t-type confidence limits be constructed for each of the LS-means. The confidence level is 0.95 by default, however this can be changed with the Alpha= option.

Diff Requests that differences of the LS-means be displayed.

Adjust= Specify the type of multiple comparison adjustment that is desired for testing all pairwise comparisons. There are several options to adjust for multiple comparisons. These are shown in the SAS help menu window below.

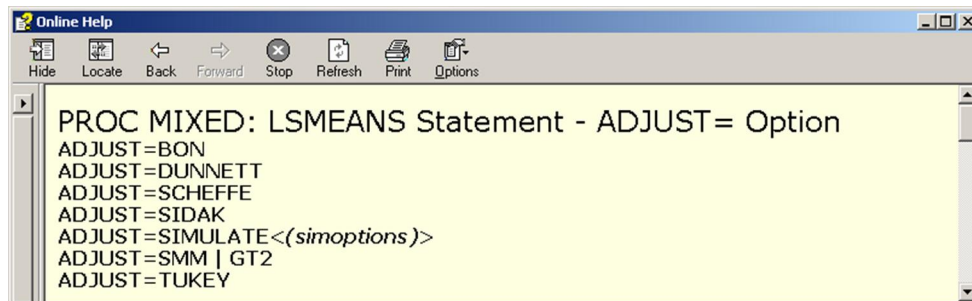


Figure 7: Output: Shown for SPSS (using REML)

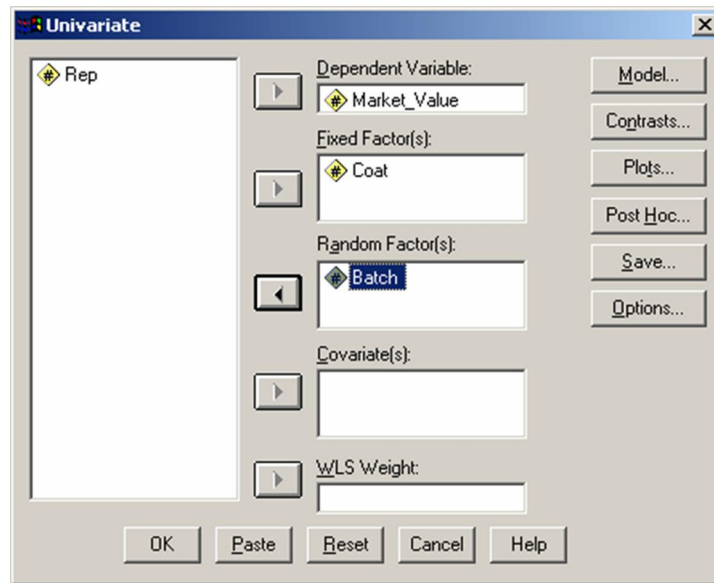
2.2 SPSS

There are also two procedures in SPSS that can be used to obtain results for mixed models. As noted in the SAS example, these procedures are similar to PROC GLM and PROC MIXED in SAS. Recall that SPSS also uses the unrestricted model. Default settings can be used for those features not discussed.

2.2.1 GLM Procedure: Comparable to PROC GLM in SAS

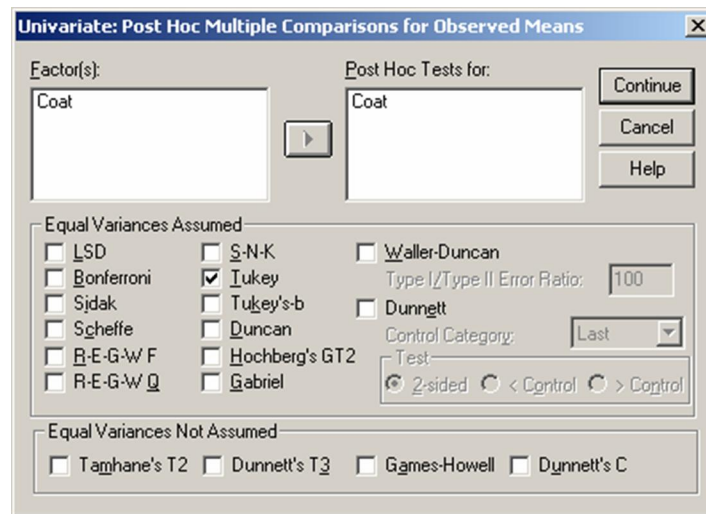
Click on: Analyze >> General Linear Models >> Univariate

Step 1. Define Dependent Variable, Fixed Factors, and Random Factors.



Step 2. Specify additional features. By clicking the buttons on the right in the above menu, several other options are available besides those listed below.

Post Hoc – Gives multiple comparison adjustments to test for mean differences in fixed effects. The following window shows the possible options.



Possible Options – Descriptive Statistics, Parameter Estimates, Means with Confidence Interval Adjustments, Estimates of Effect Size, Diagnostics

NOTE: Like SAS PROC GLM, there is no option to specify the type of estimation to use for variance components in this procedure.

Step 3. Click OK.

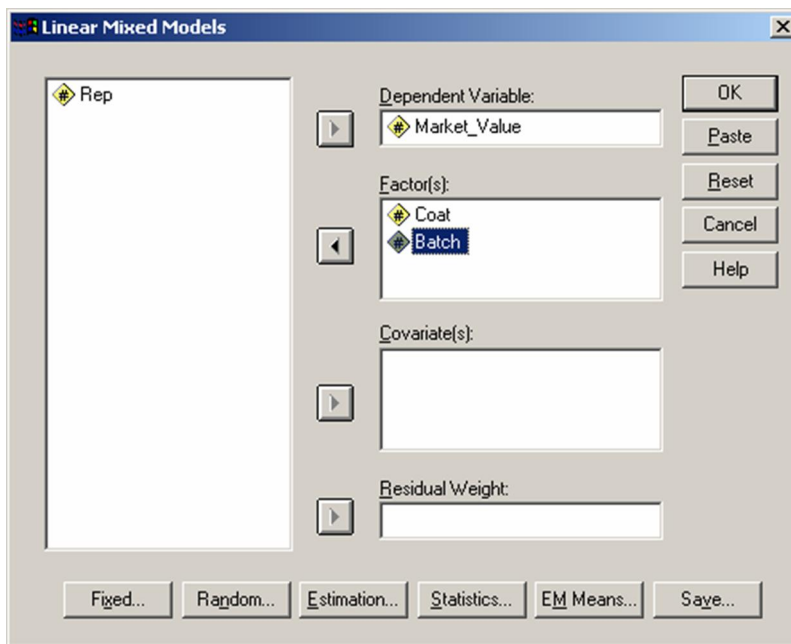
Output: Shown for SAS. Note that even though in SPSS there is not an equivalent to the Test option under the Random statement in SAS, the tests that are shown in the SPSS output are calculated using the correct error terms. The output for the factor effect tests is like that of the SAS output (3).

2.2.2 Mixed Procedure: Comparable to PROC MIXED in SAS

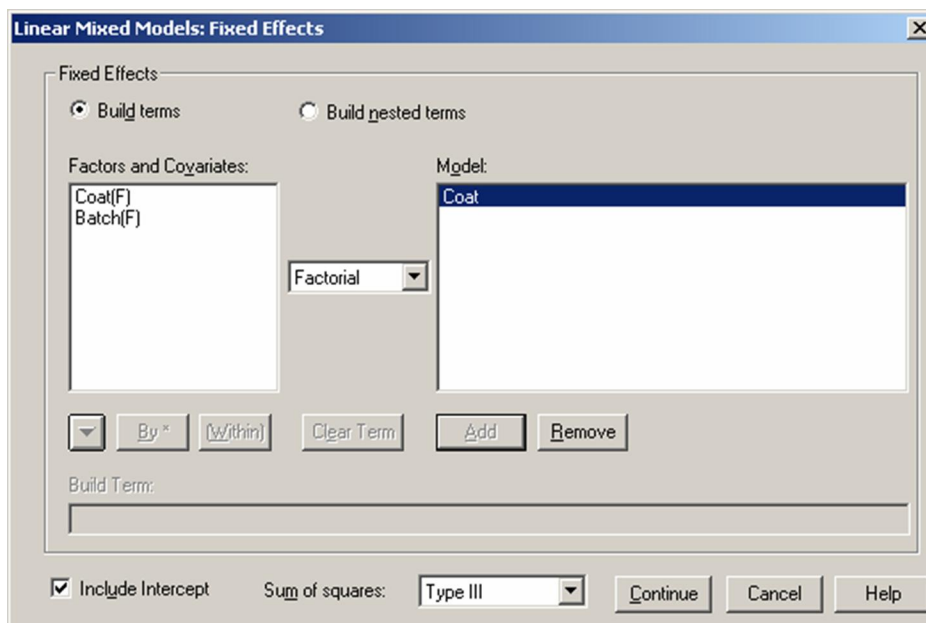
Click on: Analyze >> Mixed Models >> Linear

Step 1. Click the Continue button on the first screen that appears. Since this model does not require any special correlation structure, nothing needs to be specified.

Step 2. Define dependent variable and factors in model. These factors will be defined as fixed or random by clicking on the buttons at the bottom of the menu.



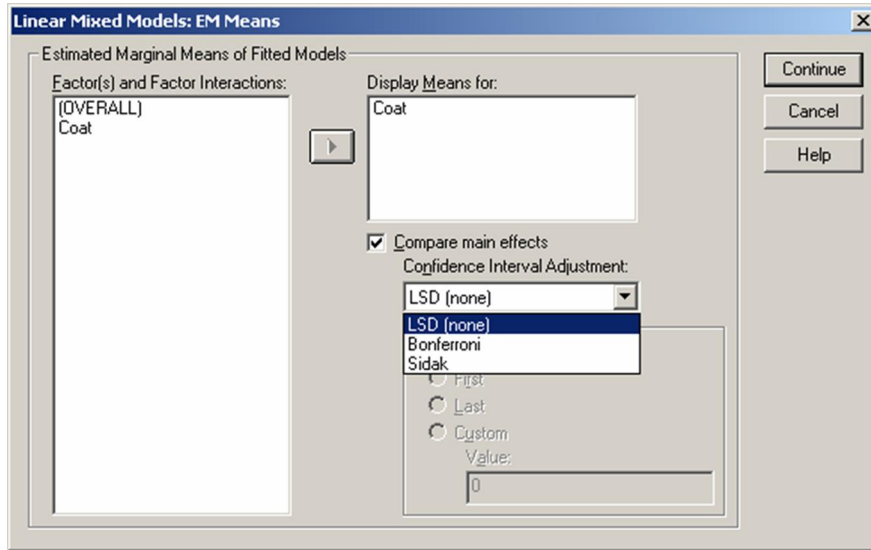
Step 3. Define fixed factor(s). Click on: Fixed in the menu from Step 1. The following menu will then appear. In the box on the left, click on the fixed factor (Coat) and then click the Add button, so that Coat appears in the box on the right (as shown). Then, click Continue.



Step 4. Define random factor(s). Click on: Random in the menu from Step 1. A menu that is similar to the fixed factor menu will appear. Simply define your random factor as you did your fixed factor in Step 3.

Step 5. Determine the estimation procedure to be used for the variance components. Click on: Estimation in the menu from Step 1. Choose between the options: REML or ML. (Note: Not as many options as SAS PROC MIXED).

Step 6. Specify pairwise comparisons and multiple comparison procedure for fixed factors. Click on: EM Means in the menu from Step 1. The menu below will appear.



Step 6. (Cont). Click on the fixed effect (Coat), then Compare main effects. Notice the only two options for multiple comparison are Bonferroni and Sidak. (LSD doesn't adjust for multiple comparisons).

Step 7. Click OK.

Information Criteria^a

-2 Restricted Log Likelihood	207.868
Akaike's Information Criterion (AIC)	211.868
Hurvich and Tsai's Criterion (AICC)	212.154
Bozdogan's Criterion (CAIC)	217.481
Schwarz's Bayesian Criterion (BIC)	215.481

The information criteria are displayed in smaller-is-better forms.

a. Dependent Variable: Market_Value.

Figure 8: Output (REML estimation)

Output (5) shows five different information criteria.

Type III Tests of Fixed Effects^a

Source	Numerator df	Denominator df	F	Sig.
Intercept	1	3.000	5385.570	.000
Coat	2	42.000	17.997	.000

a. Dependent Variable: Market_Value.

Figure 9: Output (6)

Output (6) gives the test for the fixed effect (Coat). It is clear that the number of Coats applied to the pearl is a significant factor in the model (p-value < 0.0001).

Estimates of Covariance Parameters^a

Parameter	Estimate	Std. Error
Residual Variance	4.1780258	.9117200
Batch Variance	3.8977108	3.4675786

a. Dependent Variable: Market_Value.

Figure 10: Output (7)

Output (7) shows the variance component estimates (calculated by REML) and their standard errors.

Estimates^a

Coat	Mean	Std. Error	df	95% Confidence Interval	
				Lower Bound	Upper Bound
1	73.106	1.112	4.057	70.037	76.175
2	76.794	1.112	4.057	73.725	79.863
3	76.925	1.112	4.057	73.856	79.994

a. Dependent Variable: Market_Value.

Figure 11: Output (8)

Output (8) gives the means for each of the coat levels with 95% confidence intervals. The pearls for which 6 coats (Coat=1) has the lowest mean.

Pairwise Comparisons^b

(I) Coat	(J) Coat	Mean Difference (I-J)	Std. Error	df	Sig. ^a	95% Confidence Interval for Difference ^a	
						Lower Bound	Upper Bound
1	2	-3.688*	.723	42.000	.000	-5.146	-2.229
	3	-3.819*	.723	42.000	.000	-5.277	-2.360
2	1	3.688*	.723	42.000	.000	2.229	5.146
	3	-.131	.723	42.000	.857	-1.590	1.327
3	1	3.819*	.723	42.000	.000	2.360	5.277
	2	.131	.723	42.000	.857	-1.327	1.590

Based on estimated marginal means

*. The mean difference is significant at the .05 level.

a. Adjustment for multiple comparisons: Least Significant Difference (equivalent to no adjustments).

b. Dependent Variable: Market_Value.

Figure 12: Output (9)

Output (9) gives the pairwise comparisons for the three coat levels. In this case, no corrections were made for multiple comparisons (the LSD method was employed). As in the Tukey output (4) from PROC GLM, the correction for multiple comparisons yielded the same overall conclusion as when no correction was made. Pearls with 6 coat applications had a significantly different mean market value than pearls with 8 or 10 coat applications (p-values < 0.0001 for both comparisons). However, there was no significant difference in the mean market value for pearls with 8 or 10 applications (p-value=0.857).

3 Comparison of PROC GLM and PROC MIXED in SAS

PROC GLM fits standard linear models and PROC MIXED fits the wider class of random and mixed-effect linear models in SAS.

The default fitting method used in PROC MIXED maximizes the restricted likelihood of the data under the assumption that the data are normally distributed and any missing data are missing at random. This general framework accommodates many common correlated-data methods, including variance component models and repeated measures analyses.

PROC GLM fits some random-effects and repeated-measures models, although its methods are based on method-of-moments estimation and a portion of the output applies only to the fixed-effects model. The effects specified in the RANDOM statement are still treated as fixed as far as the model fit is concerned, and they serve only to produce corresponding expected mean squares.

Given below is a brief comparison between PROC GLM and PROC MIXED in SAS.

PROC GLM	PROC MIXED
(1) Designed for “fixed effects” models with allowance for some adjustments when random effects are present.	(1) Designed for “random” and “mixed effects” models
(2) Estimation of fixed effects: Estimates are based on Ordinary Least Squares.	(2) Estimation of fixed effects: Based on Generalized Least Squares in a normal error model.(estimates are maximum likelihood estimates under normality).
(3) Estimation of variance components via the RANDOM statement: Uses method of moments estimation of solving expected mean squares for variance components.	(3) Estimation of variance components: Uses maximum likelihood or restricted maximum likelihood estimation methods for variance components (unless type 3 is requested).
(4) Both fixed and random effects are listed in the model statement.	(4) Only the fixed effects are listed in the model statement. The random effects are listed in the random statement.
(5) The LSMEANS statement interprets all effects as fixed, even the random effects. While the least square means are correct, their standard errors are not necessarily correct. This is true even if a RANDOM statement is used.	(5) Effects in the MODEL statement are assumed fixed. The standard error estimates for least square means account for the random effects.
(6) The RANDOM statement under PROC GLM invokes the calculation of expected mean squares for the listed effects and the appropriate test using the “test” option. The randomness of the effect is not incorporated into the tests of main effects, lsmeans, contrasts etc.	(6) Signals incorporation of the listed effects in all aspects of inference. Various correlation structures, describing the dependencies of multiple random effects can be selected using different options.
(7) The TEST statement is an important statement to use correct error terms in testing model effects (e.g. Subsampling designs, split-plot design).	(7) No TEST statement under PROC MIXED
(8) The REPEATED statement in PROC GLM is used to specify various transformations with which to conduct the traditional univariate or multivariate tests.	(8) The REPEATED statement in PROC MIXED is used to specify covariance structures for repeated measurements on subjects. The approach used in PROC MIXED is more flexible and more widely applicable than either the univariate or multivariate approaches.
(9) PROC GLM is not efficient in handling missing values.	(9) PROC MIXED has a better mechanism for handling missing values.

4 Crossover designs

4.1 Introduction

Crossover designs are a special type of repeated measures experiment where the experimental units are given different treatments in different sequences over time.

	Sequence	Sequence
Treatment	Group 1	Group 2
Period 1	A	B
Period 2	B	A

Figure 13: Basic Pattern 1: 2-treatment 2-group 2-period

	Sequence	Sequence
Treatment	Group 1	Group 2
Period 1	A	B
Period 2	B	A
Period 3	B	A

Figure 14: Basic Pattern 2: 2-treatment 2-group 3-period

	Sequence	Sequence
Treatment	Group 1	Group 2
Period 1	A	B
Period 2	B	A
Period 3	B	A
Period 4	A	B

Figure 15: Basic Pattern 3: 2-treatment 2-group 4-period

4.2 Why Crossover and its disadvantage

Crossover design is frequently used in animal behavior experiments because it offers some real advantages over traditional parallel subject designs when comparing different treatments. For two treatments (A and B) studies, traditional parallel studies test the treatments in parallel. One group of subjects receives only treatment A; the other group receives only the other treatment B. At the end of the study, the two groups are compared on some quantitative outcome measure (dependent variable of primary interest), most often by using a t test for independent samples.

If all subjects received the two treatments in the same order, observed differences between treatments would be confounded with any other changes that occur over time. In a study of the effect of treatments on cholesterol levels, for example, subjects might change their diet and exercise behavior for the better as a result of heightened awareness of health issues. This would likely manifest itself as a decrease in cholesterol levels over the later portion of the study and might end up being attributed to the second treatment.

The two treatment, two-period crossover study seeks to overcome this difficulty by having half of the subjects receive treatment A followed by treatment B while the other half receive B followed by A. The order of administration is incorporated into the formal analysis. In essence, any temporal change that might favor B over A in one group will favor A over B in the other group and cancel out of the treatment comparison.

A crossover design enables each subject to receive both treatments and thus requires fewer subjects for the same number of observations than a parallel study. Also the repeated structure increases precision of the treatment comparisons by having extra information on within subject variation. Since the comparison of treatments is no longer be contaminated by the variability between subjects when the comparison is carried

out within each individual.[g1] (In crossover, we can look at the treatment difference on each subject. But before the observed difference which could be a result from the subject or the treatment or both.

There are a few major disadvantages of the crossover design. Generally, a long wash out period is needed to wash out the carryover effect between treatments, which will force the experiment to last for a long time. During the long duration, missing data could result. When the carryover effect can not be completely avoid by the wash out, more complicated design with more time periods should be considered to balance the carryover. For example, 3 periods: ABB and BAA or 4 periods ABBA and BAAB.

4.3 Analysis of structure and confounding

Although the structure of the crossover design is conceptually quite simple, it is still hard to read the literature due to varied terminology and notation. The confusion is coming from for the treatment combination. For that treatment combination, some people call it sequence and some people call it group.

Recall, in the design, a group of people take sequence 1 and the other group of people takes sequence 2. The sequence means a combination of treatment in specified order. In this discussion, it was denoted as Group. All the design factors are denoted as the following:

- treatment (TREATMENT), with levels “A” and “B”,
- time period (PERIOD), with levels “1” and “2”,
- Sequence group (GROUP), with levels “A then B” and “B then A”.

If crossover studies were full-factorial designs (with factors GROUP, TREATMENT, and PERIOD), it could be possible to evaluate not only the main effects, but also the GROUP*TREATMENT, PERIOD*TREATMENT, GROUP*PERIOD, and PERIOD*TREATMENT*GROUP interactions. However, crossover studies are not full factorial designs. Not all combinations of factors appear in the study (there is no GROUP=’A then B’, PERIOD=’1’, TREATMENT=’B’ combination, for example). Also, each estimate of a main effect is confounded with a two-factor interaction. For example, when the TREATMENT and PERIOD are fixed, then GROUP is fixed also, which means we cannot separate the main effect of group and the two-factor interaction of TREATMENT*PERIOD. That is, if one of the main effects is significant, it is impossible to tell whether the effect, its alias, or both are generating the significant result. One could argue that there is no reason to expect a significant affect involving GROUP because subjects are assigned to GROUPS at random. Therefore, a significant GROUP effect should be interpreted as resulting from a PERIOD*TREATMENT interaction and not from a difference between GROUPS. For similar reasons, a significant PERIOD effect is not considered to be the result of a GROUP*TREATMENT interaction, nor is a significant TREATMENT effect considered to be the result of a GROUP*PERIOD interaction.

4.3.1 Example 1

In summary, the experimental design to test 2 treatment effects is a crossover experiment, with two treatments and two blocking factors – period and subject. The layout of this design is presented in Figure 1, where the row blocking factor is “Period”, the column blocking factor is “Subject”, and the two treatments are denoted by “A” and “B”.

Latin Squares	1	2	...	n
Subject	1 2	3 4	...	$2n - 1$ $2n$
Period 1	A B	B A	...	A B
Period 2	B A	A B	...	B A

Figure 16: Crossover design

The statistical model of this crossover experiment is:

$$y_{ijk} = \mu + P_i + \tau_j + S_k + \varepsilon_{ijk}$$

where μ is the overall mean, P_i is the i^{th} row (period) effect, $i = 1, 2$, τ_j is the j^{th} treatment effect, $j = 1, 2$, and S_k is the k^{th} column (subject) effect, $k = 1, 2, \dots, 2n$, ε_{ijk} is the random error with zero mean and constant variance σ^2 . y_{ijk} is the observation in the i^{th} row and k^{th} column for the j^{th} treatment, i.e., the time for the k^{th} subject to finish the j^{th} treatment in the i^{th} period.

4.3.2 Data summary

```

title " Data Statistics Summary";
proc sort data=a3;
by sequence period treatment;run;
proc means data=a3 mean std cv;
var Y;
by sequence period treatment;
output out=mean1 mean=avY;
run;

/*plot 1*/
proc sort data=mean1;
by period;
run;
title "Mean effect plot";
symbol1 i=j v=dot c=red pointlabel=(h=1.5) ;
symbol2 i=j v=star c=blue pointlabel=(h=1.5);
proc gplot data=mean1;
plot avy*period=treatment;
format avy 6.2;
run;

```

Analysis Variable : Y		
Mean	Std Dev	Coeff of Variation
14.2891630	2.7232130	19.0578903

Figure 17: sequence=1 period=1 Treatment=A

Analysis Variable : Y		
Mean	Std Dev	Coeff of Variation
15.2608359	2.9628641	19.4148220

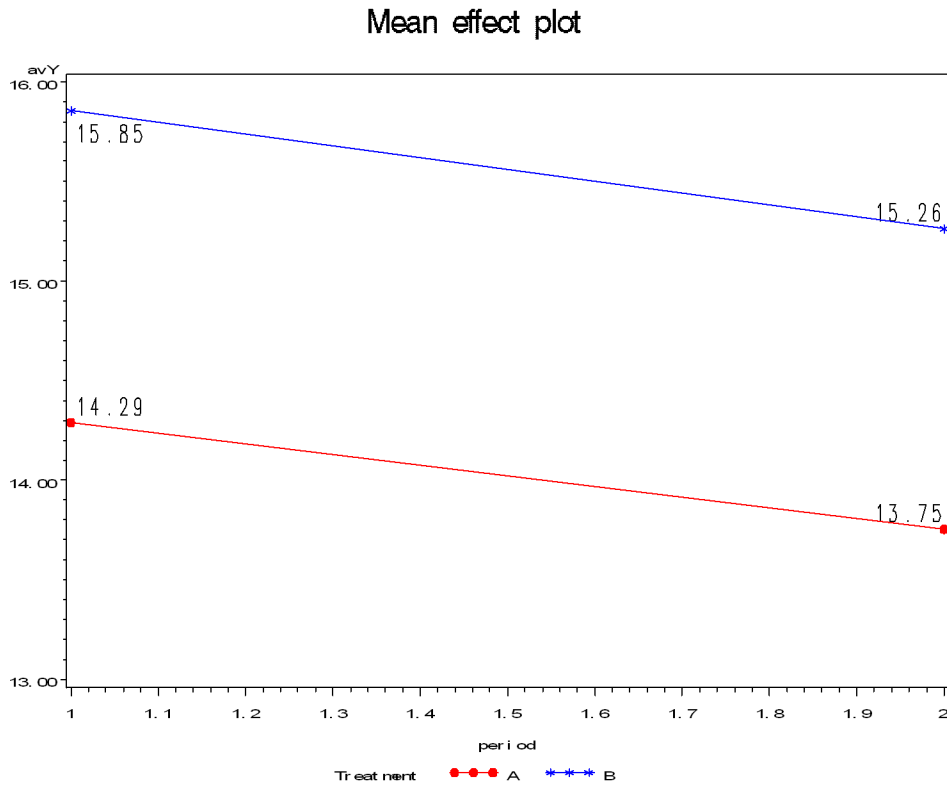
Figure 18: sequence=1 period=2 Treatment=B

Analysis Variable : Y		
Mean	Std Dev	Coeff of Variation
15.8548191	3.0215940	19.0578903

Figure 19: sequence=2 period=1 Treatment=B

Analysis Variable : Y		
Mean	Std Dev	Coeff of Variation
13.7538354	2.6702827	19.4148220

Figure 20: sequence=2 period=2 Treatment=A



```

PROC MIXED data=a3;
CLASSES sequence cow period treatment;
MODEL Y = sequence period treatment / solution ;
RANDOM cow(sequence);
*lsmeans period/alpha=0.05 cl diff adjust=tukey;
run;

```

Figure 21: Code

In the example, sequence and treatment*period are equivalent (sequence= treatment*period) in the model because when both treatment and period fixed, we know exactly which sequence it is. lsmeans statement can give values of the period average and we can do similar things with treatment and sequence.

Cow, which is the subject, was considered as random effect. Since half of the cow are assigned to one sequence and the rest are assigned to the other half, so they are nested within sequence. Accordingly, we would like to know the variation of the cow effect.

Covariance Parameter Estimates	
Cov Parm	Estimate
cow(sequence)	7.5934
Residual	0.5203

Figure 22: Output

Fit Statistics	
-2 Res Log Likelihood	230.0
AIC (smaller is better)	234.0
AICC (smaller is better)	234.2
BIC (smaller is better)	236.8

Solution for Fixed Effects								
Effect	Treatment	sequence	period	Estimate	Standard Error	DF	t Value	Pr > t
Intercept				15.2902	0.7355	28	20.79	< .0001
sequence		1		-0.02933	1.0233	28	-0.03	0.9773
sequence		2		0
period			1	0.5647	0.1862	28	3.03	0.0052
period			2	0
Treatment	A			-1.5363	0.1862	28	-8.25	< .0001
Treatment	B			0

Type 3 Tests of Fixed Effects				
Effect	Num DF	Den DF	F Value	Pr > F
sequence	1	28	0.00	0.9773
period	1	28	9.19	0.0052
Treatment	1	28	68.04	< .0001

Overall, the two sequences are not significantly different with each other. It means that it is not statistical important which treatment the cow receive in the first period and which treatment to receive in the second period. Also, period 1 is significantly different with period 2 with p-value of 0.0052. The treatment effect is also significantly different with each other.

5 Repeated Measures

Any measurement on an experimental unit that can be repeated (either across time or across space) can be analyzed under this broad heading. Crowder and Hand[3] describe *repeated measures* as the situation in which measurements “are made of the same characteristic on the same observational unit but on more than one occasion.” This is what is meant by the term *longitudinal data*. The scope of repeated measures can be expanded to include *clustered data* as well; that is, measurements on members of a cluster that are related in some way. For longitudinal data, the observational units are subjects, and a common occurrence is for measurements on the subject to be correlated. For example, multiple measurements of milk production are made on each cow in a study over time. For cluster data, there may be two layers of correlation: there is correlation between different subjects within the same cluster, and if measurements are repeated on subjects, then there is correlation between measurements on the same subject. For example, subjects (which are clusters) may have both of their eyes (which are observational units, being members of one cluster) are measured; there is correlation between the eyes of the same subject in addition to the correlation across measurements of the same eye. In both of these cases, the usual model assumption of independent errors may be violated, so a model that can incorporate this lack of independence is needed.

5.1 Covariance Structure

We assume N observational units (individuals or clusters) and n_i observations of the response for the i^{th} unit, $i = 1, \dots, N$ (if there are no missing values, then $n_i = n$). The observations for the i^{th} unit are coded in the vector \mathbf{y}_i , which has length n_i . The design matrices \mathbf{X}_i and \mathbf{Z}_i consist of n_i rows, \mathbf{X}_i having q columns and \mathbf{Z}_i having s columns. The entries of \mathbf{X}_i are the values taken by the q continuous explanatory variables and/or design indicators of fixed effects across the n_i measurements of the i^{th} unit; if any between-subjects random effects exist, their design indicators appear in \mathbf{Z}_i , and their parameter estimates appear in γ_i . For the response $\mathbf{y}_i = (y_{i1}, \dots, y_{in_i})$, the model

$$\mathbf{y}_i = \mathbf{X}_i\boldsymbol{\beta} + \mathbf{Z}_i\boldsymbol{\gamma}_i + \boldsymbol{\epsilon}_i$$

$$\begin{bmatrix} \boldsymbol{\gamma}_i \\ \boldsymbol{\epsilon}_i \end{bmatrix} \sim \mathcal{N} \left(\mathbf{0}, \begin{bmatrix} \mathbf{G}_i & \mathbf{0} \\ \mathbf{0} & \mathbf{R}_i \end{bmatrix} \right)$$

the term Covariance Structure is used to describe how the matrices \mathbf{G}_i and \mathbf{R}_i are constrained in the (Normal case of the) general linear mixed model:

$$\mathbf{y}_1, \dots, \mathbf{y}_n \stackrel{\text{ind}}{\sim} \mathcal{N}_{n_i}(\mathbf{X}_i\boldsymbol{\beta}, \mathbf{V}_i)$$

where $\mathbf{V}_i = \mathbf{Z}_i\mathbf{G}_i\mathbf{Z}_i^{\text{T}} + \mathbf{R}_i$ is called the *Variance-Covariance Matrix* of the i^{th} unit. Typically, it is assumed that all of the \mathbf{V}_i s take the same form[4]. The examples that follow are forms of \mathbf{R}_i , but can be used for \mathbf{G}_i as well:

5.1.1 Unstructured

This is the most general form:

$$\mathbf{R}_i = \sigma^2 \begin{bmatrix} 1 & \rho_{12} & \cdots & \rho_{1n_i} \\ \rho_{21} & 1 & & \rho_{2n_i} \\ \vdots & & \ddots & \vdots \\ \rho_{n_i1} & \rho_{n_i2} & \cdots & 1 \end{bmatrix}$$

In this case the variance-covariance matrix contains $n_i(n_i - 1)/2 + 1$ nuisance parameters to be estimated, so in practice, estimation of this structure may only converge for $N \gg n_i$. Likewise the statistical power under this structure is reduced since the only “constraint” on \mathbf{R}_i is that it be symmetric.

5.1.2 ARMA(1,1)

The more structured forms of \mathbf{R}_i involve assuming that some or all of the ρ_{jks} are a function of the “distance” between observations j and k . If measurements are repeated across time, then the time variable will be prominent in this distance function. If measurements are repeated across location, then the distance function will involve some spatial metric reflecting the experimental design’s geometry. One of the more common structures is the *first order autoregressive, first order moving-average model*, abbreviated ARMA(1,1):

$$\mathbf{R}_i = \sigma^2 \begin{bmatrix} 1 & \gamma\rho^0 & \cdots & \gamma\rho^{n_i-2} \\ \gamma\rho^0 & 1 & & \gamma\rho^{n_i-3} \\ \vdots & & \ddots & \vdots \\ \gamma\rho^{n_i-2} & \gamma\rho^{n_i-3} & \cdots & 1 \end{bmatrix}$$

Notice that now there are only three parameters to be estimated, σ^2 , γ , and ρ . A special case of ARMA(1,1) occurs when $\gamma = \rho$; this is called ARMA(1,0) or just simply AR(1), and only two covariance parameters σ^2 and ρ are estimated.

5.1.3 Equicorrelation

Another model assumes all repeated measurements are equally correlated:

$$\mathbf{R}_i = \sigma^2 \begin{bmatrix} 1 & \rho & \cdots & \rho \\ \rho & 1 & & \rho \\ \vdots & & \ddots & \vdots \\ \rho & \rho & \cdots & 1 \end{bmatrix}$$

This structure (also called *spherical* or *exchangeable*) might be applicable to cluster data, where ρ is called the *intra-class correlation coefficient* between two members of the same cluster and is “a relative measure of the within-cluster similarity.” [4] A special case of equicorrelation, called *compound symmetry*, arises by enforcing $\rho = \frac{\sigma_a^2}{\sigma_a^2 + \sigma_e^2}$ for some σ_a^2 and σ_e^2 . In that case, if $\sigma^2 = \sigma_a^2 + \sigma_e^2$, then

$$\mathbf{R}_i = \begin{bmatrix} \sigma_a^2 + \sigma_e^2 & \sigma_a^2 & \cdots & \sigma_a^2 \\ \sigma_a^2 & \sigma_a^2 + \sigma_e^2 & & \sigma_a^2 \\ \vdots & & \ddots & \vdots \\ \sigma_a^2 & \sigma_a^2 & \cdots & \sigma_a^2 + \sigma_e^2 \end{bmatrix}$$

5.1.4 Uncorrelated

Finally, we compare to the fixed effects case where independence is assumed over repeated measurements:

$$\mathbf{R}_i = \sigma^2 \begin{bmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{bmatrix}$$

This helps to show (when $\mathbf{G}_i = \mathbf{0}$) that the standard linear model can be thought of as a special case of repeated measures.

5.2 Software Packages

5.2.1 Dental Data Set Analysis in SAS

For all of our examples, we will use the dental measurements of Pothaff and Roy (1964). The subjects are children who each had measurements taken at 8, 10, 12, and 14 years of age. Thus, the $N = 27$ children are the observational units, there are $n_i = n = 4$ measures on each child. Since there are no random effects, $\mathbf{Z}_i\boldsymbol{\gamma}_i$ does not appear in the model:

$$Y_{ijk} = \mu + \text{gender}_i + \text{age}_j + \text{gender*age}_{ij} + \epsilon_{ijk}$$

$$\epsilon_{ijk} \sim \mathcal{N}(0, \mathbf{R}_j)$$

where $\mathbf{R}_j = \sigma^2 \begin{pmatrix} 1 & \cdots & \rho_{14} \\ \vdots & \ddots & \vdots \\ \rho_{41} & \cdots & 1 \end{pmatrix}$

and the ρ_{lm} are constrained according to the covariance structure we specify with the TYPE= option of the REPEATED statement. The following SAS code to enter the data is from [9]:

<pre>data forglm(keep=person gender y1-y4) formixed(keep=person gender age y); input person gender\$ y1-y4; output forglm; y=y1; age=8; output formixed; y=y2; age=10; output formixed; y=y3; age=12; output formixed; y=y4; age=14; output formixed; datalines; 1 F 21.0 20.0 21.5 23.0 2 F 21.0 21.5 24.0 25.5 3 F 20.5 24.0 24.5 26.0 4 F 23.5 24.5 25.0 26.5 5 F 21.5 23.0 22.5 23.5 6 F 20.0 21.0 21.0 22.5 7 F 21.5 22.5 23.0 25.0 8 F 23.0 23.0 23.5 24.0 9 F 20.0 21.0 22.0 21.5 10 F 16.5 19.0 19.0 19.5 11 F 24.5 25.0 28.0 28.0</pre>	<pre>12 M 26.0 25.0 29.0 31.0 13 M 21.5 22.5 23.0 26.5 14 M 23.0 22.5 24.0 27.5 15 M 25.5 27.5 26.5 27.0 16 M 20.0 23.5 22.5 26.0 17 M 24.5 25.5 27.0 28.5 18 M 22.0 22.0 24.5 26.5 19 M 24.0 21.5 24.5 25.5 20 M 23.0 20.5 31.0 26.0 21 M 27.5 28.0 31.0 31.5 22 M 23.0 23.0 23.5 25.0 23 M 21.5 23.5 24.0 28.0 24 M 17.0 24.5 26.0 29.5 25 M 22.5 25.5 25.5 26.0 26 M 23.0 24.5 26.0 30.0 27 M 22.0 21.5 23.5 25.0 ;</pre>
--	--

Both PROC GLM and PROC MIXED provide a REPEATED statement by which one or more repeated measures can be specified. However, there are many nontrivial differences between these two procedures. Some of the more important differences are:

- Notice that the `forglm` data set has a multivariate response, while the `formixed` data set has a univariate response. This is because PROC GLM requires there to be a separate response for each repeated measure, while PROC MIXED allows a separate variable (in our case, `age`) to identify to which a particular measure an observation corresponds.
- PROC GLM has as a model assumption that \mathbf{R}_i is either uncorrelated or of “Type H” structure[7], whereas PROC MIXED allows a choice between over 30 different structures for \mathbf{R}_i through the TYPE= option of the REPEATED statement.
- PROC GLM ignores observational units with missing observations, whereas PROC MIXED includes them. (Though note that they must be missing at random for the estimators to remain unbiased.)
- PROC GLM assumes covariates are constant within observational units, whereas in PROC MIXED they are allowed to vary. Since there are no time-varying covariates in our example data set we can use either PROC GLM or PROC MIXED.

See[9] for a longer treatment of the differences.

5.2.2 SAS PROC GLM

Here is the SAS code to perform a repeated measures analysis on the data using PROC GLM. Notice the multivariate response, and the specification of the four values taken by age. The `/ nouni` option suppresses the traditional univariate GLM output. This is included since it would contain incorrect standard error estimates.

```
proc glm data=forglm;
class gender;
model y1-y4=gender / nouni;
repeated age 4 (8 10 12 14) / printe;
run;
```

It is important to note that this is only applicable when \mathbf{R}_i follows either an uncorrelated or type H structure. In order to test this model assumption, we requested a test of sphericity through the `/ printe` option of the `repeated` statement:

Sphericity Tests				
Variables	DF	Mauchly's Criterion	Chi-Square	Pr > ChiSq
Transformed Variates	5	0.4998695	16.449181	0.0057
Orthogonal Components	5	0.7353334	7.2929515	0.1997

The p-value of 0.1997 indicates that we do not have evidence to suspect that the type H structure is inappropriate. If the test were significant, then it is likely that PROC GLM would not be appropriate for the analysis of this data.[2]

5.2.3 SAS PROC MIXED

Here is the code for the same analysis using PROC MIXED, but instead of being uncorrelated, the residuals are fit to a specific covariance structure instead of being uncorrelated.

Compound Symmetry Note that `repeated age/type=cs sub=person;` has been shortened to `repeated /type=cs sub=person;` this is because for most data sets and one-level covariance structures, SAS seems to be able to automatically determine which effect is the repeated one. `subject=person` identifies `person` as the variable corresponding to observational units, `type=cs` requests that the covariance structure be compound symmetry. The `r` option tells SAS to report the estimated R_i matrix.

```
proc mixed data=formixed;
class gender age person;
model y = gender|age;
repeated / type=cs subject=person r;
run;
```

Partial SAS output:

Estimated R Matrix for person 1				
Row	Col1	Col2	Col3	Col4
1	5.2604	3.2854	3.2854	3.2854
2	3.2854	5.2604	3.2854	3.2854
3	3.2854	3.2854	5.2604	3.2854
4	3.2854	3.2854	3.2854	5.2604
Covariance Parameter Estimates				
Cov Parm	Subject	Estimate		
CS	person	3.2854		
Residual		1.9750		

The above output shows that for a given subject, the variance at a particular age is $\sigma_a^2 + \sigma_e^2 = 5.2604$, and the covariance between any two nonequal ages is constrained to be the same, $\sigma_a^2 = 3.2854$; the residual variance is $\sigma_e^2 = 1.9750$.

First Order Autoregressive If we change the covariance structure to AR(1) by setting `/ type=ar(1)`, the output is shown in Figure 5.2.3.

Estimated R Matrix for person 1				
Row	Col1	Col2	Col3	Col4
1	5.2467	3.2283	1.9863	1.2222
2	3.2283	5.2467	3.2283	1.9863
3	1.9863	3.2283	5.2467	3.2283
4	1.2222	1.9863	3.2283	5.2467
Covariance Parameter Estimates				
Cov Parm	Subject	Estimate		
AR(1)	person	0.6153		
Residual		5.2467		

Figure 23:

Now, using our earlier ARMA(1,1) notation, $\gamma = \rho$ here, and we see $\sigma^2 = 5.2467$, and R is constrained by ρ , ρ^2 , and ρ^3 where $\rho = \frac{3.2283}{5.2467} = 0.6153$.

Unstructured If we change the covariance structure to be unstructured by setting / `type=un`, the output contains the $5 \cdot 4/2 = 10$ unique elements of \mathbf{R}_i , and from the output below, we can calculate $R_{21} = \sigma^2 \rho_{21}$, so $\rho_{21} = \frac{2.7168}{5.4155} = 0.5016711$ See Figure 5.2.3

Estimated R Matrix for person 1				
Row	Col1	Col2	Col3	Col4
1	5.4155	2.7168	3.9102	2.7102
2	2.7168	4.1848	2.9272	3.3172
3	3.9102	2.9272	6.4557	4.1307
4	2.7102	3.3172	4.1307	4.9857
Covariance Parameter Estimates				
Cov Parm	Subject	Estimate		
UN(1,1)	person	5.4155		
UN(2,1)	person	2.7168		
UN(2,2)	person	4.1848		
UN(3,1)	person	3.9102		
UN(3,2)	person	2.9272		
UN(3,3)	person	6.4557		
UN(4,1)	person	2.7102		
UN(4,2)	person	3.3172		
UN(4,3)	person	4.1307		
UN(4,4)	person	4.9857		

Figure 24:

Notice that \mathbf{R}_i contains all of the covariance parameter estimates reported by SAS by default, so it is actually not necessary to ask for the \mathbf{R}_i matrix in the unstructured case.

5.2.4 Example: Three-Factor Factorial Design with One Repeated Measure

This example will use data from an experiment with three between-subject factors: A, B, and C. Factors A and B each have two levels, and factor C has four levels, resulting in 16 treatment combinations. Four subjects (“subj”) are nested within each treatment combination. A response, Y, is measured at each of four sites (“site”) on each subject’s body. The primary interest of this study is to compare measurements of the response across sites within the body.

SAS’s PROC MIXED has many covariance structures available to model the dependence of observations among the sites. Here is a procedure for selecting a suitable covariance structure:

Begin with an unstructured covariance structure (`type = un` in the code below). This requires estimation of the greatest number of parameters, and the code may not always run. Examination of the covariance and correlation matrices for measurements among sites may suggest simpler structures.

Next, try a compound symmetry (`type = cs`) or heterogeneous compound symmetry (`type = csh`) structure. Compound symmetry assumes the same covariance between observations at any two pairs of sites, while heterogeneous compound symmetry assumes the same correlation between observations at any two pairs of sites. The heterogeneous compound symmetry structure also allows unequal variances at the four sites.

Because of their relative simplicity, compound symmetry and heterogeneous compound symmetry structures are recommended if a better fitting covariance structure cannot be found. In this example, heterogeneous

compound symmetry may be preferred over compound symmetry because variances appear to differ across sites.

Other covariance structures can be tried, based on observed patterns in the covariance and correlation matrices and on knowledge of the experimental system under study. It is not advisable to try every available covariance structure. A final covariance structure can be selected with a preference for simple structures (which require fewer degrees of freedom for the estimation of fewer parameters) and by comparing various fit statistics given in the output. Recall that for information criteria such as AIC or BIC, a smaller value is better.

We begin by fitting an unstructured covariance structure:

```
proc mixed data=example covtest cl;
class subj A B C site;
model Y = A|B|C|site / ddfm=kr;
repeated site / subject=subj type=un r rcorr;
run;
```

A few comments on the code:

CL (proc statement): This generates confidence limits for the variance-covariance estimates.

DDFM=KR (model statement): The Kenward-Roger procedure is recommended for calculating the denominator degrees of freedom in repeated measures analyses. This procedure may result in denominator degrees of freedom that are not integers.

R (repeated statement): This prints the covariance matrix.

RCORR (repeated statement): This prints the correlation matrix.

Estimated R Matrix for subj 11				
Row	Col1	Col2	Col3	Col4
1	0.04102	-0.00930	0.004626	0.009375
2	-0.00930	0.1029	0.03861	0.02113
3	0.004626	0.03861	0.1326	0.04486
4	0.009375	0.02113	0.04486	0.09459

Estimated R Correlation Matrix for subj 11				
Row	Col1	Col2	Col3	Col4
1	1.0000	-0.1431	0.06272	0.1505
2	-0.1431	1.0000	0.3306	0.2142
3	0.06272	0.3306	1.0000	0.4005
4	0.1505	0.2142	0.4005	1.0000

Looking at the main diagonal of the estimated R matrix (the covariance matrix for the four sites), we see that the variances do indeed look unequal between the sites. However, the covariance and correlation matrices do not otherwise suggest suitable alternative covariance structures. Note that for the unstructured model, we must estimate ten covariance parameters.

From the Type 3 tests of fixed effects, we find that **A*C*site** is significant, while no terms involving factor B are significant. The non-integer denominator degrees of freedom are the result of the Kenward-Roger procedure.

Type 3 Tests of Fixed Effects				
Effect	Num DF	Den DF	F Value	Pr > F
A	1	48	240.14	<.0001
B	1	48	2.06	0.1577
A*B	1	48	2.14	0.1503
C	3	48	72.56	<.0001
A*C	3	48	0.36	0.7817
B*C	3	48	1.45	0.2392
A*B*C	3	48	0.33	0.8012
site	3	46	300.66	<.0001
A*site	3	46	22.76	<.0001
B*site	3	46	0.06	0.9801
A*B*site	3	46	1.78	0.1640
C*site	9	69.2	14.71	<.0001
A*C*site	9	69.2	5.36	<.0001
B*C*site	9	69.2	1.16	0.3339
A*B*C*site	9	69.2	1.47	0.1768

Next, we try a heterogeneous compound symmetry structure, which requires the estimation of only five parameters:

```
proc mixed data=example covtest cl;
class subj A B C site;
model Y = A|B|C|site / ddfm=kr;
repeated site / subject=subj type=csh r rcorr;
run;
```

We note the change in the covariance and correlation matrices:

Estimated R Matrix for subj 11				
Row	Col1	Col2	Col3	Col4
1	0.04352	0.01165	0.01290	0.01092
2	0.01165	0.1043	0.01997	0.01690
3	0.01290	0.01997	0.1278	0.01871
4	0.01092	0.01690	0.01871	0.09152

Estimated R Correlation Matrix for subj 11				
Row	Col1	Col2	Col3	Col4
1	1.0000	0.1730	0.1730	0.1730
2	0.1730	1.0000	0.1730	0.1730
3	0.1730	0.1730	1.0000	0.1730
4	0.1730	0.1730	0.1730	1.0000

We can compare this model to the unstructured case by comparing the fit statistics:

type = un:	
-2 Res Log Likelihood	143.7
AIC (smaller is better)	163.7
AICC (smaller is better)	165.0
BIC (smaller is better)	185.3
type = csh:	
-2 Res Log Likelihood	153.8
AIC (smaller is better)	163.8
AICC (smaller is better)	164.2
BIC (smaller is better)	174.6

We see that the simpler csh structure barely changes the AIC, while decreasing BIC. We conclude that a heterogeneous compound symmetry structure is preferable in this case.

A compound symmetry structure requires the estimation of only two parameters. Below are the associated covariance and correlation matrices:

Estimated R Matrix for subj 11				
Row	Col1	Col2	Col3	Col4
1	0.09278	0.01822	0.01822	0.01822
2	0.01822	0.09278	0.01822	0.01822
3	0.01822	0.01822	0.09278	0.01822
4	0.01822	0.01822	0.01822	0.09278

Estimated R Correlation Matrix for subj 11				
Row	Col1	Col2	Col3	Col4
1	1.0000	0.1964	0.1964	0.1964
2	0.1964	1.0000	0.1964	0.1964
3	0.1964	0.1964	1.0000	0.1964
4	0.1964	0.1964	0.1964	1.0000

We compare the fit statistics for the heterogeneous compound symmetry and compound symmetry structures:

type = csh:	
-2 Res Log Likelihood	153.8
AIC (smaller is better)	163.8
AICC (smaller is better)	164.2
BIC (smaller is better)	174.6
type = cs:	
-2 Res Log Likelihood	167.9
AIC (smaller is better)	171.9
AICC (smaller is better)	171.9
BIC (smaller is better)	176.2

The higher AIC and BIC for type = cs suggests that the heterogeneous compound symmetry structure is preferable. Other covariance structures could be investigated in similar fashion.

5.2.5 Repeated Measures in SPSS

SPSS is very similar to SAS with respect to the fact that SPSS has equivalent procedures to PROC GLM and PROC MIXED. The differences between these two procedures are equivalent to the differences between PROC GLM and PROC MIXED in SAS. The SPSS version of PROC GLM can be obtained either through using the syntax editor, or by going to Analyze > General Linear Model > Univariate, and then selecting the desired model. Only the newest versions of SPSS have the equivalent of PROC MIXED. This new SPSS procedure can be obtained either through the syntax editor, or by going to Analyze > Mixed Models > Linear.

6 Appendix

6.1 Pearl data as entered into SAS

data pearl;				77.2	2	2	4
input market_value coat batch rep;				80.2	2	3	1
datalines;				76.6	2	3	2
72.0	1	1	1	77.3	2	3	3
74.6	1	1	2	79.9	2	3	4
67.4	1	1	3	74.3	2	4	1
72.8	1	1	4	77.6	2	4	2
72.1	1	2	1	74.4	2	4	3
76.9	1	2	2	72.9	2	4	4
74.8	1	2	3	76.3	3	1	1
73.3	1	2	4	74.1	3	1	2
75.2	1	3	1	77.1	3	1	3
73.8	1	3	2	75.0	3	1	4
75.7	1	3	3	80.9	3	2	1
77.8	1	3	4	73.7	3	2	2
70.4	1	4	1	78.6	3	2	3
68.1	1	4	2	80.2	3	2	4
72.4	1	4	3	79.2	3	3	1
72.4	1	4	4	78.0	3	3	2
76.9	2	1	1	77.6	3	3	3
78.1	2	1	2	81.2	3	3	4
72.9	2	1	3	71.6	3	4	1
74.2	2	1	4	77.7	3	4	2
80.3	2	2	1	75.2	3	4	3
79.3	2	2	2	74.4	3	4	4
76.6	2	2	3	;			

References

- [1] Sas library/ proc glm and proc mixed of sas. <http://www.ats.ucla.edu/stat/sas/library/>.
- [2] B.A. Craig. Stat 514: Experimental design class notes, topic 22, 2004.
- [3] M.J. Crowder and D.J. Hand. *Analysis of Repeated Measures*. Chapman & Hall, 1990.
- [4] Annette J. Dobson. *An Introduction to Generalized Linear Models*. Chapman & Hall/CRC, 2002.
- [5] B.S. Everitt. *Making Sense of Statistics in Psychology: A Second-Level Course*. Oxford University Press, Oxford, 1996.
- [6] J. Neter, M.H. Kutner, C.J. Nachtsheim, and W. Wasserman. *Applied Linear Statistical Models*. The McGraw Hill Companies, Inc., 4th edition, 1996.

- [7] SAS Institute Inc., Cary, NC. *SAS OnlineDoc®*, Version 8, 1999.
- [8] SPSS. *SPSS 12.0 for Windows, System Help*.
- [9] R.D. Wolfinger and M. Chang. Comparing the SAS GLM and MIXED procedures for repeated measures. Technical report, SAS Institute Inc., Cary, NC, 1998.
- [10] Y. Zhu. Stat 514: Experimental design class notes, latin square and related designs, 2004.