

Linear Mixed Models with Random Effects

Introduction and Analysis of a Split – Plot Experiment with SAS/STAT® Software

Jerry W. Davis, University of Georgia, Griffin Campus

Introduction

Linear mixed models allow for modeling fixed, random and repeated effects in analysis of variance models. “Factor effects are either fixed or random depending on how levels of factors that appear in the study are selected. An effect is called fixed if the levels in the study represent all possible levels of the factor, or at least all levels about which inference is to be made.” “Factor effects are random if the levels of the factor that are used in the study represent only a random sample of a larger set of potential levels.” (Littell et al, 1996).

Those are the “official” definitions for fixed and random effects, and as you may imagine, there is a lot of overlap between the two. Therefore, for this tutorial and seminar, I am going to discuss fixed and random effects within the context of agricultural experiments I have worked with over the past 30 years.

Repeated effects come from measuring the same experimental unit over time. Subjects, often the same as experimental units, are identified and the correlation between them modeled with appropriate covariance structures. Analysis of variance with repeated effects will be covered in a separate document.

Maximum likelihood estimation (MLE) is used to estimate parameters in the mixed model procedures. Unlike least squares methods that solve the normal equations, maximum likelihood methods seek parameter estimates that maximize a likelihood function. However, it is easier to estimate the parameters by working with the negative of the natural logarithm of the likelihood function. It now becomes a minimization problem. The process converges when parameter estimates are found that minimize the $-\text{LogLikelihood}$. (JMP 13, “Likelihood, AICc, and BIC”).

Analysis of Variance with Random Effects

Random effects in agricultural experiments are often factors like replications, blocks, years, locations or subjects. They are random when one would expect to get similar results regardless of which block, year or location an experimental unit appeared. Random effects are usually not tested in simple analysis of variance models. The variability associated with random effects adjusts the standard errors for tests on the fixed effects. Adjusting the standard errors make the tests more general (broad inference), implying that the results apply to the larger population from which the random effects were drawn. The interaction between a random and fixed effect is typically considered random. This is important when fitting hierarchical models such as split-plots.

If a researcher wants to test for year, location or any other factor that could be random or fixed, then it can be modeled as a fixed effect and tested accordingly. This reinforces the statement made above that there is often no absolute dichotomy between fixed and random factors.

The Random Statement

Random effects are listed in the random statement. There are situations where it is desirable to also list random and interaction effects in the model statement and test related hypotheses. However, those models are more complex and won't be discussed here.

While there are no default tests for random effects, if they don't affect the dependent variable, their covariance parameter estimate may be zero. If the estimate is zero, it can be removed from the random statement. Removing it won't change the results, but the model diagnostics may change slightly.

When a model has two random effects, it is usually not necessary to include the interaction between the two in the random statement. For example, given a model with rep and year as random effects, they may account for all random variation so rep*year may not be needed.

random rep year rep*year;

random rep year; *rep*year*; <- this term may not improve the model.

I hesitate to say that this will always be the case, so I usually include all the terms and remove the interaction if one (or more) of the factors has an estimate of zero.

The random statement has many options to change how random factors are modeled. The simplest way to identify random factors is shown in the first statement. It works well for simple models and small data sets. A common alternative is to use the subject option when specifying a random effect. These two statements are equivalent:

random rep;

random intercept / subject=rep;

The second statement is more complicated but mathematically more efficient because the mixed model procedures are optimized to process random effects by subjects. Use this syntax for complicated models, large data sets or when there are convergence problems.

Multiple random statements are allowed, so each random factor can appear in a separate statement.

Again, there are many more options available for modeling random effects but these provide a good starting point.

Least Squares Means (LS-means)

Least squares means are the only option for calculating treatment level means within the mixed model procedures. LS-means are estimated from the model while regular means are an average of the data

values. LS-means are equivalent to arithmetic means when the design is balanced, there are no missing values and the means are not adjusted for other factors in the model such as covariates. If there are missing values, LS-means are theoretically better estimates of population means because they are less sensitive to missing values. LS-means do not have an accompanying value for **N**. See Appendix A for an example comparing regular and LS-mean calculations.

If the design is unbalanced such that a whole treatment combination is missing, the LS-mean for the treatment level with the empty cell are non-estimable. This is easier to demonstrate than explain so an example data set and analysis is included.

Tests for differences between LS-means are done using pair wise t-tests. There is no single value analogous to Fisher's LSD that accompanies the tests. Standard errors are based on sample size not **N**. There are options, i.e. **adjust=Tukey**, to make the tests more conservative when the number of comparisons is large.

Contrast, estimate and lsmestimate statements are available for testing comparisons that are more complex.

When a model has random effects, the LS-means are called conditional means because they are *conditioned* by the random effects. Otherwise, they are called marginal means.

Least squares means are usually referred to as LS-means now because the mixed model procedures do not use least squares for analysis of variance calculations. LS differentiates LS-means from regular means, i.e. the **lsmeans** statement vs the **means** statement.

A Split-plot Example

The following program analyzes data from a split-plot experiment. Nitrogen is the main plot, green manure is the sub plot, there are three replications and the response variable is yield. (Little and Hill, 1978).

```
/*
/ A SAS Software program to analyze data from a split-plot experiment
/ using PROC MIXED. The data comes from Little and Hill page 91.
/ Nitrogen is the main plot; gmanure is the subplot.
/
/ SplitPlotFull13.sas
/
/=====*/
```

```
ods html close;
ods html;

data one;
  input nitrogen gmanure $ rep yield;
datalines;
0 f 1 13.8
0 f 2 13.5
0 f 3 13.2
0 b 1 15.5
0 b 2 15.0
0 b 3 15.2
0 v 1 21.0
0 v 2 22.7
0 v 3 22.3
0 bv 1 18.9
0 bv 2 18.3
0 bv 3 19.6
1 f 1 19.3
1 f 2 18.0
1 f 3 20.5
1 b 1 22.2
1 b 2 24.2
1 b 3 25.4
1 v 1 25.3
1 v 2 24.8
1 v 3 28.4
1 bv 1 25.9
1 bv 2 26.7
1 bv 3 27.6
;
```

In SAS versions 9.3 and later, the default output is in html format. These two commands close the current output file and open a new one. Otherwise output from subsequent runs is appended to the previous output file.

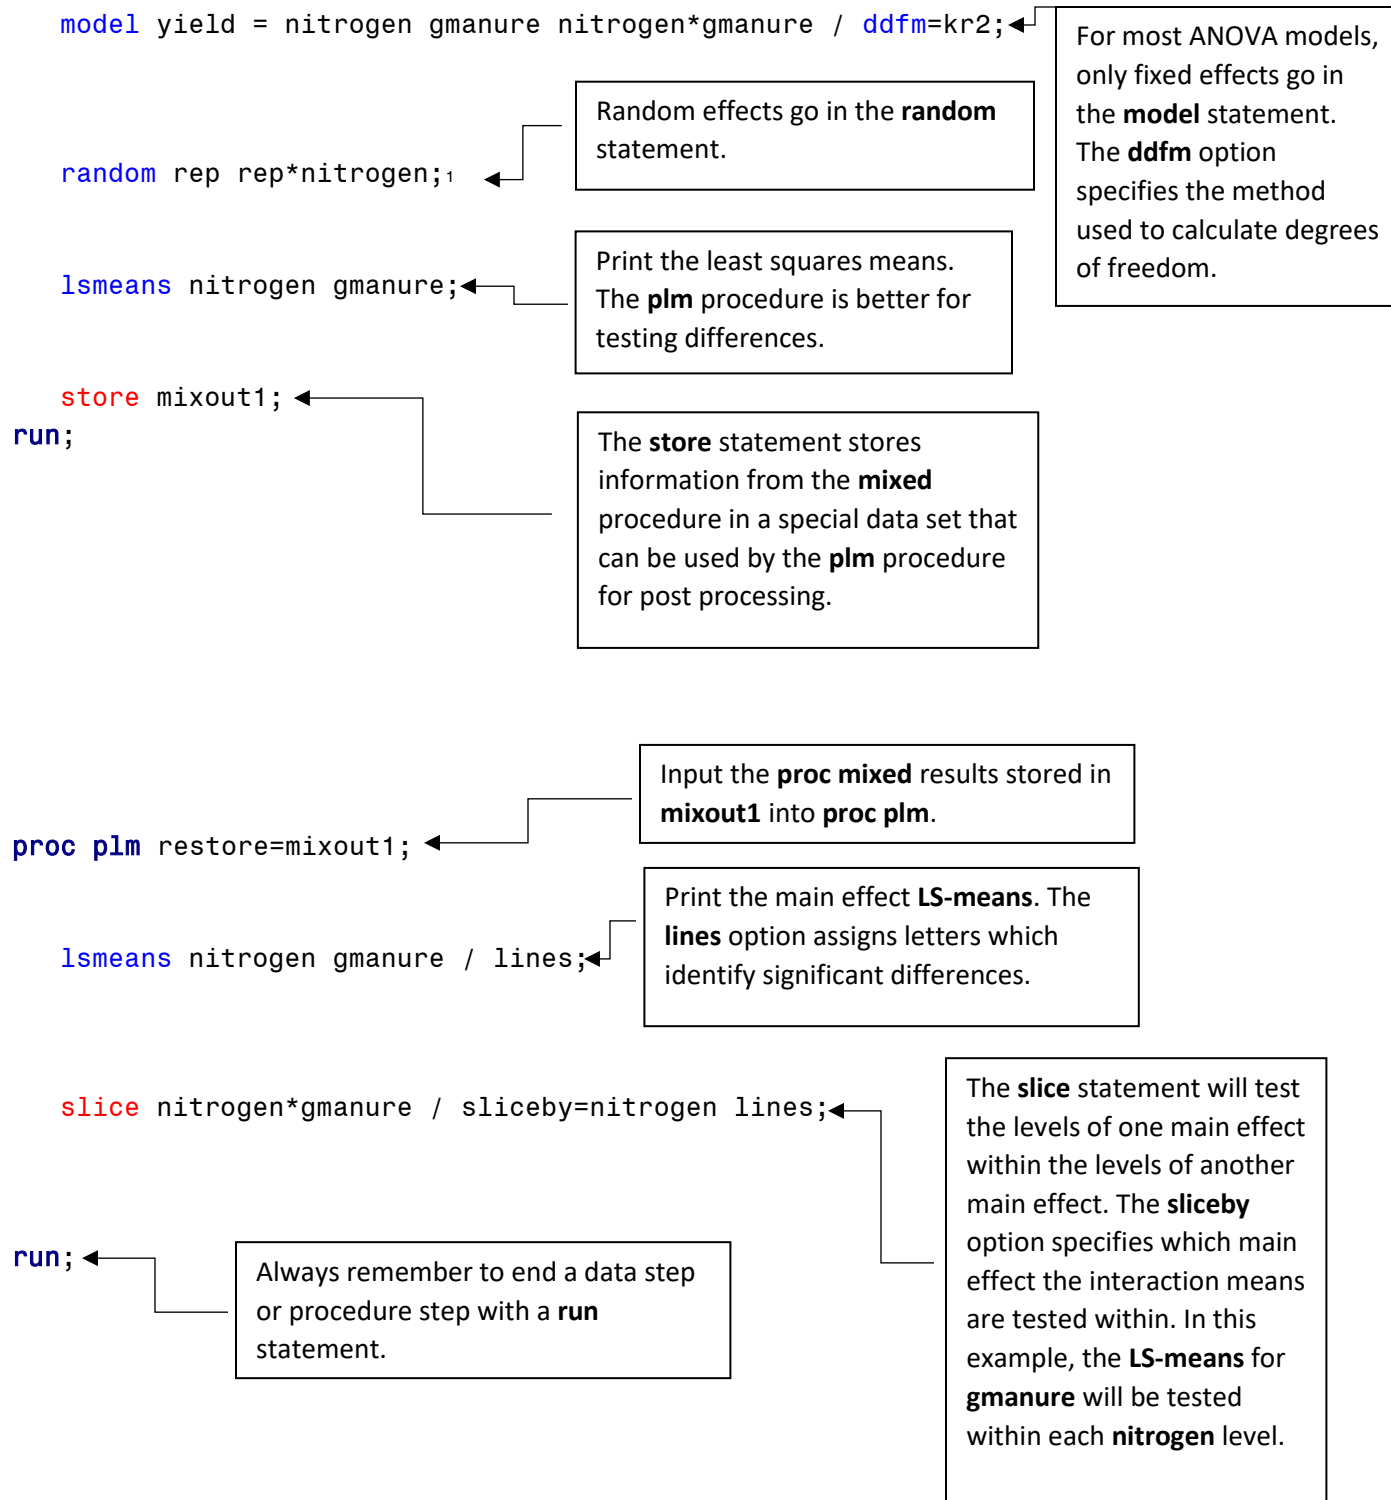
```
proc print;
run;

proc mixed plots=residualpanel;

  class rep nitrogen gmanure;
```

The **plots** option produces a nice set of diagnostic plots.

The **class** statement is the same as for **proc glm**.



¹ I chose to use the simple random statement syntax since this is an introductory seminar. The number of subjects in the **Dimensions** table changes if the complex syntax is used but results the will be the same.

ODS statements. **ODS** is short for Output Delivery System. This was added to SAS to provide an easy way to customize output files. HTML, PDF, RTF and other document types can be written with simple commands. The amount of customization afforded by **ODS** is nearly limitless.

The plots option. The plots option can generate many diagnostic and descriptive graphs. It's part of the **ODS** Statistical Graphics subsystem that provides an easy and flexible way to visualize data.

The model statement. Unlike GLM, only one dependent variable is allowed in the model statement. Usually for ANOVA models, only fixed effects are listed in the model statement because these are the factors to be tested. There are models where it's appropriate to also list random effects, but after working with mixed models for twenty years, I've encountered few of these situations. So, let's focus on the more common ones.

The ddfm option. This option defines the method used to calculate the denominator degrees of freedom used in F tests for the fixed effects. The default is residual error DF which is the traditional ANOVA method and what **GLM** uses. Satterthwaite is a good method because it doesn't assume equal variances among treatment levels. The example program uses KenwardRoger2 (**kr2**) which takes the Satterthwaite method further by adjusting the degrees of freedom for the number of treatment levels. The statisticians who developed the mixed model procedures recommend using Kenward-Roger when possible. Degrees of freedom calculated by these newer methods may not be integers. Satterthwaite and KenwardRoger2 are not compatible with some maximum likelihood estimation methods. When this happens, omit the **ddfm** option to use the default residual error DF method.

The random statement. This is a split-plot experiment so the appropriate random effects are **rep** and **rep*nitrogen** (nitrogen is the main plot). In classic agricultural statistics books the replication by main plot effect is referred to as **error(A)**. This is the appropriate error term for testing the main plot effect. **GLM** has a **test** statement that can be used to specify that the main plot is tested with **error(A)**. In **MIXED**, **error(A)** is listed in the **random** statement.

PROC PLM. The **PLM** procedure performs post fitting statistical analyses for the contents of a SAS item store that was previously created in some other SAS/STAT procedure (Tobias, R., Cai, W., 2010). The **MIXED** procedure doesn't have the ability to assign letters to identify significant LS-mean differences. This is easily accomplished with **PROC PLM** using the **lsmeans** statement with the **lines** option.

The **slice** statement will print and test interaction means. The **sliceby** option will test the interaction means of one factor within the levels of another factor and assign letter groups. This is much better than doing the analysis **by** one of the factors because it uses all the error degrees of freedom in the experiment for testing LS-means instead of sub-setting the data for each factor level. Curiously, when an analysis is done for **by** groups, **PLM** inherits the by group information and does the post fitting correctly but doesn't identify the by groups in the output tables.

PROC MIXED Output Tables

The SAS System

The Mixed Procedure

Model Information	
Data Set	WORK.ONE
Dependent Variable	yield
Covariance Structure	Variance Components
Estimation Method	REML
Residual Variance Method	Profile
Fixed Effects SE Method	Kenward-Roger2
Degrees of Freedom Method	Kenward-Roger2

Model Information. Lists the dependent variable, then the Covariance Structure, Estimation Method and Residual Variance Method. Program statements and options set these values. The notes about Kenward-Roger2 result from specifying `ddfm=kr2` in the model statement.

Class Level Information		
Class	Levels	Values
rep	3	1 2 3
nitrogen	2	0 1
gmanure	4	b bv f v

Class Level Information. Like the table for PROC GLM, it lists the class variables, the number of levels for each variable and the values of each level in sorted order.

Dimensions	
Covariance Parameters	3
Columns in X	15
Columns in Z	9
Subjects	1
Max Obs per Subject	24

Dimensions. Specifies the size of the elements in the model's design matrix. See the footnote below for how the numbers are determined.¹

Number of Observations	
Number of Observations Read	24
Number of Observations Used	24
Number of Observations Not Used	0

Number of Observations. A similar table appears in PROC GLM output.

2

Covariance Parameters = 3 (1 (rep) + 1 (rep*nitrogen) + 1 (residual error)).

Columns in X (fixed elements) = 15 (1 (grand mean) + 2 (nitrogen levels) + 4 (gmanure levels) + 8 (gmanure*nitrogen levels)).

Columns in Z (random elements) = 9 (3 (rep levels) + 6 (rep*nitrogen levels)).

Subjects = 1 (default) (no subjects are identified in this model).

Max Obs per Subject = 24 (number of observations since there was only one subject).

Iteration History			
Iteration	Evaluations	-2 Res Log Like	Criterion
0	1	57.87683801	
1	1	52.71436959	0.00000000

Convergence criteria met.

Iteration History. The model converges when there is no more improvement in reducing (-2 * the Residual Log Likelihood) function.

If the note below the **Iteration History** table is different from this, **beware!** When running a program with more than one **mixed** procedure, carefully check the results to ensure that each procedure converged.

Covariance Parameter Estimates	
Cov Parm	Estimate
rep	0.1769
rep*nitrogen	0.4787
Residual	0.6032

The **Covariance Parameter Estimates** are estimates of the random effects and residual error. If a random effect does not affect the dependent variable its estimate may be zero.

Fit Statistics	
-2 Res Log Likelihood	52.7
AIC (Smaller is Better)	58.7
AICC (Smaller is Better)	60.7
BIC (Smaller is Better)	56.0

The **Fit Statistics** are relative measures of how well the current model fits the data compared to other model parameterizations. The concept of a good or bad AIC, for example, does not exist. If all the model parameterizations fit the data poorly, the fit statistics will not reflect this. They only indicate which of the models provide the better fit. (Wikipedia, "Akaike Information Criterion").

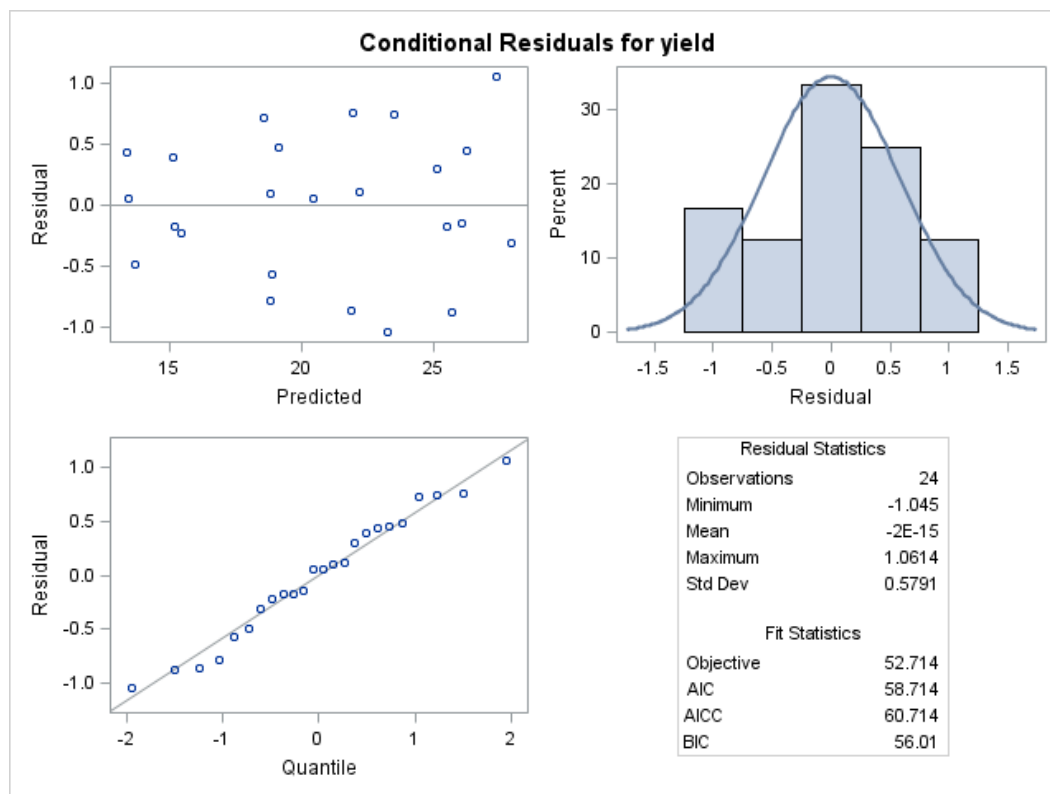
These values can be negative and a negative sign indicates a value smaller than zero. I once thought, "smaller is better" referred to the absolute value but that is not correct.

Type 3 Tests of Fixed Effects				
Effect	Num DF	Den DF	F Value	Pr > F
nitrogen	1	2	104.06	0.0095
gmanure	3	12	118.96	<.0001
nitrogen*gmanure	3	12	10.33	0.0012

Type 3 Tests of Fixed Effects. This is the **ANOVA** table. It shows the tests for the fixed effects. The money table!

Least Squares Means							
Effect	gmanure	nitrogen	Estimate	Standard Error	DF	t Value	Pr > t
nitrogen		0	17.4167	0.5184	3.82	33.59	<.0001
nitrogen		1	24.0250	0.5184	3.82	46.34	<.0001
gmanure	b		19.5833	0.4892	4.12	40.04	<.0001
gmanure	bv		22.8333	0.4892	4.12	46.68	<.0001
gmanure	f		16.3833	0.4892	4.12	33.49	<.0001
gmanure	v		24.0833	0.4892	4.12	49.23	<.0001

Least Squares Means. These are the estimates for the treatment level means. The t-tests test whether the estimates are different from 0. In this example, all of the means are significantly different from 0. Note that the degrees of freedom (**DF**) are not integers. This often happens when using option **DDFM=kr2** in the model statement.



Conditional Residuals for yield. These graphs are the result of using the **plots=residualpanel** option in the **PROC MIXED** statement. These are conditional residuals because they have been conditioned by the random effects. The residual vs predicted graph looks good because there should not be patterns above or below the 0.0 reference line. The histogram in the upper right corner is somewhat blocky because **N** is relatively small. The points in the quantile plot should follow the 45 degree diagonal reference line, which these do acceptably. The residual statistics are a numerical summary of the information in the residual vs predicted graph.

This residual panel indicates that the model fits the data well.

The PLM Procedure

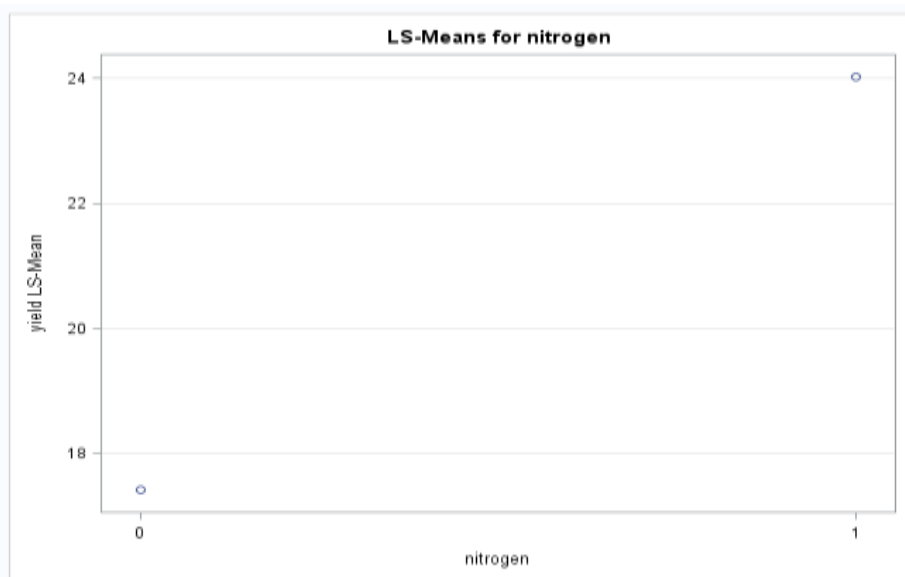
Store Information	
Item Store	WORK.MIXOUT1
Data Set Created From	WORK.ONE
Created By	PROC MIXED
Date Created	18JAN17:15:09:04
Response Variable	yield
Distribution	Normal
Class Variables	rep nitrogen gmanure
Model Effects	Intercept nitrogen gmanure nitrogen*gmanure
Degrees of freedom Method	Kenward-Roger

Class Level Information		
Class	Levels	Values
rep	3	1 2 3
nitrogen	2	0 1
gmanure	4	b bv fv

Store Information. The PLM procedure prints the model information it inherited from the **mixout1** data set (item store).

nitrogen Least Squares Means					
nitrogen	Estimate	Standard Error	DF	t Value	Pr > t
0	17.4167	0.5184	3.816	33.59	<.0001
1	24.0250	0.5184	3.816	46.34	<.0001

Nitrogen Least Squares Means. Yield LS-means for each nitrogen level.



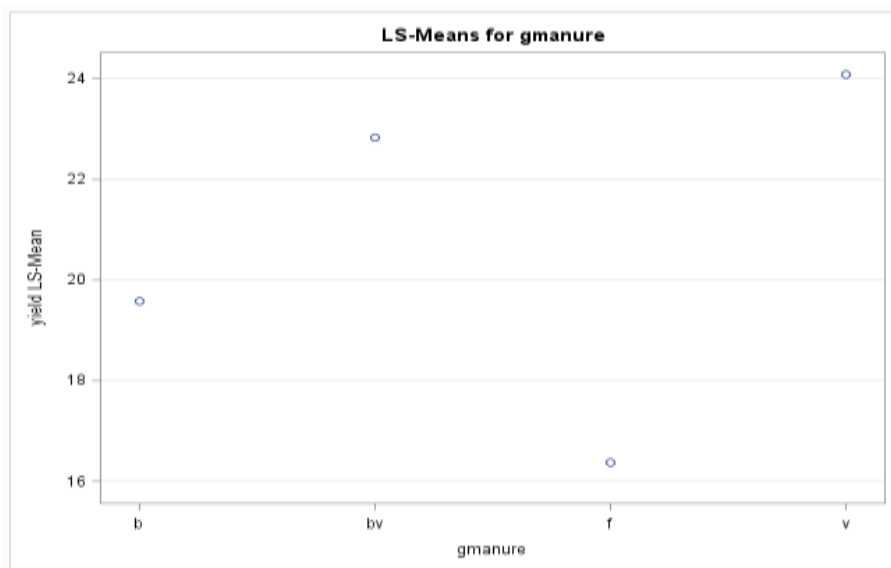
LS-Means for nitrogen. A graph showing yield LS-means for each nitrogen level. Incidentally, these default graphs can be turned off by using the **ods graphics off** statement.

T Grouping for nitrogen Least Squares Means (Alpha=0.05)		
LS-means with the same letter are not significantly different.		
nitrogen	Estimate	
1	24.0250	A
0	17.4167	B

T Grouping for nitrogen. Differences for the LS-means are determined by pairwise t-tests. Letters identify different LS-mean groups.

gmanure Least Squares Means					
gmanure	Estimate	Standard Error	DF	t Value	Pr > t
b	19.5833	0.4892	4.119	40.04	<.0001
bv	22.8333	0.4892	4.119	46.68	<.0001
f	16.3833	0.4892	4.119	33.49	<.0001
v	24.0833	0.4892	4.119	49.23	<.0001

Gmanure Least Squares Means. Yield LS-means for each gmanure level.



LS-Means for gmanure. A graph showing yield LS-means for each gmanure level.

T Grouping for gmanure Least Squares Means (Alpha=0.05)		
LS-means with the same letter are not significantly different.		
gmanure	Estimate	
v	24.0833	A
bv	22.8333	B
b	19.5833	C
f	16.3833	D

T Grouping for gmanure. Differences for the LS-means are determined by pairwise t-tests. Letters identify different LS-mean groups.

F Test for nitrogen*gmanure Least Squares Means Slice				
Slice	Num DF	Den DF	F Value	Pr > F
nitrogen 0	3	12	71.97	<.0001

F test for nitrogen*gmanure. A test for whether there are differences between gmanure levels at nitrogen level 0. This and the following table are the result of using the **slice** statement and **sliceby** option in **PLM** procedure.

T Grouping for nitrogen*gmanure Least Squares Means Slice (Alpha=0.05)			
LS-means with the same letter are not significantly different.			
Slice	gmanure	Estimate	
nitrogen 0	v	22.0000	A
nitrogen 0			
nitrogen 0	bv	18.9333	B
nitrogen 0			
nitrogen 0	b	15.2333	C
nitrogen 0			
nitrogen 0	f	13.5000	D

T Grouping for nitrogen*gmanure. LS-mean pairwise t-tests for gmanure at nitrogen level 0.

F Test for nitrogen*gmanure Least Squares Means Slice				
Slice	Num DF	Den DF	F Value	Pr > F
nitrogen 1	3	12	57.31	<.0001

F test for nitrogen*gmanure. A test for whether there are differences between gmanure levels at nitrogen level 1.

T Grouping for nitrogen*gmanure Least Squares Means Slice (Alpha=0.05)			
LS-means with the same letter are not significantly different.			
Slice	gmanure	Estimate	
nitrogen 1	bv	26.7333	A
nitrogen 1			A
nitrogen 1	v	26.1667	A
nitrogen 1			
nitrogen 1	b	23.9333	B
nitrogen 1			
nitrogen 1	f	19.2667	C

T Grouping for nitrogen*gmanure. LS-mean
pairwise t-tests for gmanure at nitrogen level 1.

Problems and Some Solutions

A Zero Covariance Parameter Estimate

Covariance Parameter Estimates	
Cov Parm	Estimate
rep	0
rep*nitrogen	0.1993
Residual	0.6032

A zero estimate for a random factor usually indicates that it does not change how the fixed effects influence the response variable. **Proc mixed** prints this message in the **SASLOG** when an estimate is zero:

NOTE: Convergence criteria met.

NOTE: Estimated G matrix is not positive definite.

NOTE: Asymptotic variance matrix of covariance parameter estimates has been found to be singular and a generalized inverse was used.
Covariance parameters with zero variance do not contribute to degrees of freedom computed by DDFM=KENWARDROGER2.

In other words. The covariance parameter estimate for **rep** was a negative value. By definition, variances cannot be negative, so the procedure set the estimate to zero. The procedure used an alternate algorithm to minimize the $-\text{LogLikelihood}$ function due to the zero estimate. It converged and the zero estimate did not affect the results of the analysis of variance.

Does this situation require any action from the data analyst? Not really. When using complicated or multiple random statements with subjects and other options, I will often remove terms with zero estimates to simplify the procedure. For simple models, like this example, I usually leave them in because they do not affect the results and the client can see what the random effects are.

Non-estimable LS-means

This is the same data set used in the previous example except the yield values for **nitrogen=0** and **gmanure=f** were changed to missing values.

Obs	nitrogen	gmanure	rep	yield
1	0	f	1	.
2	0	f	2	.
3	0	f	3	.
4	0	b	1	15.5
5	0	b	2	15.0
6	0	b	3	15.2
7	0	v	1	21.0
8	0	v	2	22.7
9	0	v	3	22.3
10	0	bv	1	18.9
11	0	bv	2	18.3
12	0	bv	3	19.6
13	1	f	1	19.3
14	1	f	2	18.0
15	1	f	3	20.5
16	1	b	1	22.2
17	1	b	2	24.2
18	1	b	3	25.4
19	1	v	1	25.3
20	1	v	2	24.8
21	1	v	3	28.4
22	1	bv	1	25.9
23	1	bv	2	26.7
24	1	bv	3	27.6

The missing cell. **Yield** values for all three **reps** are missing for **nitrogen=0** and **gmanure=f**.

Least Squares Means							
Effect	gmanure	nitrogen	Estimate	Standard Error	DF	t Value	Pr > t
nitrogen		0	Non-est
nitrogen		1	24.0250	0.5380	3.34	44.65	<.0001
gmanure	b		19.5833	0.5388	3.61	36.35	<.0001
gmanure	bv		22.8333	0.5388	3.61	42.38	<.0001
gmanure	f		Non-est
gmanure	v		24.0833	0.5388	3.61	44.70	<.0001

Consequently, the **LS-means** for **nitrogen=0** and **gmanure=f** are **non-estimable**.

Differences of nitrogen Least Squares Means						
nitrogen	_nitrogen	Estimate	Standard Error	DF	t Value	Pr > t
0	1	Non-est

Differences of gmanure Least Squares Means						
gmanure	_gmanure	Estimate	Standard Error	DF	t Value	Pr > t
b	bv	-3.2500	0.4704	10.25	-6.91	<.0001
b	f	Non-est
b	v	-4.5000	0.4704	10.25	-9.57	<.0001
bv	f	Non-est
bv	v	-1.2500	0.4704	10.25	-2.66	0.0235
f	v	Non-est

The **lines** option does not work when one or more treatment level means are **non-estimable**. However, the **pdiff** option will test the **estimable** means.

Notice that even though **yield** values are missing for only one level of gmanure, comparisons between the two nitrogen levels are invalid because the nitrogen values cannot be averaged over the gmanure levels. Likewise, **yield** cannot be averaged over the two nitrogen levels when **gmanure=f**, so any test involving **f** is invalid.

Unfortunately, this problem does not have a good solution. Sometimes, when a design is more complex, it may be that a factor or factor level can be deleted to make the experiment balanced or the analysis is done within the factor with missing values. Some information is lost, but it may be worthwhile to get better tests on the other factors. "This is the price one pays for having missing data." (Goodnight et al. 1978).

If there is a significant interaction between two factors or it is otherwise desirable to test the interaction means within the levels of a factor with missing values, **PROC PLM** will ignore the missing cell and test the estimable LS-means.

T Grouping for nitrogen*gmanure Least Squares Means Slice (Alpha=0.05)			
LS-means with the same letter are not significantly different.			
Slice	gmanure	Estimate	
nitrogen 0	v	22.0000	A
nitrogen 0			
nitrogen 0	bv	18.9333	B
nitrogen 0			
nitrogen 0	b	15.2333	C

When **nitrogen=0** the LS-mean value for **f** is non-estimable, so it is left out of the comparison table.

T Grouping for nitrogen*gmanure Least Squares Means Slice (Alpha=0.05)			
LS-means with the same letter are not significantly different.			
Slice	gmanure	Estimate	
nitrogen 1	bv	26.7333	A
nitrogen 1			A
nitrogen 1	v	26.1667	A
nitrogen 1			
nitrogen 1	b	23.9333	B
nitrogen 1			
nitrogen 1	f	19.2667	C

When **nitrogen=1** the LS-mean for **f** is estimable and is in the comparison table.

If this was a two-factor Randomized Complete Block experiment, one could use the **sliceby** option for either factor but not both. That would be making too many tests. However, this is a split-plot experiment and nitrogen is the main plot, so only test gmanure within nitrogen.

The procedure fails to converge.

Non-convergence is somewhat like a goodness of fit test. It indicates that the model does not fit the data. Convergence should not be a problem for simple experiments like this one. The data are Gaussian (normally distributed or close to it) and random effects are modeled correctly. For experiments with repeated measures and models with multiple subjects and complex covariance structures, convergence problems may arise.

References

Cai, Weiji. (2014), "Making Comparisons Fair: How LS-Means Unify the Analysis of Linear Models," Proceedings of the SAS Global Forum 2014 Conference, Cary NC: SAS Institute Inc.

Gomez, Kwanchai, A., A. A. Gomez. 1984. Statistical Procedures for Agricultural Research. John Wiley and Sons, Inc.

Goodnight, J. H., Walter R. Harvey. (1978), Least Squares Means in the Fixed Effects General Linear Model, SAS Technical Report R-103, SAS Institute Inc., Cary, NC.

JMP 13 Online Documentation (English). "Likelihood, AICc, and Bic." JMP Statistical Discovery™. From SAS. Accessed Jan 27, 2017. http://www.jmp.com/support/help/Likelihood_AICc_and_BIC.shtml

Littell, Ramon C., George A. Millikin, Walter W. Stroup, Russell D. Wolfinger, and Oliver Schabenbenger. 2006. *SAS® for Mixed Models, Second Edition*. Cary, NC: SAS Institute Inc.

Little, Thomas M., F. Jackson Hill. 1978. *Agricultural Experimentation: Design and Analysis*. John Wiley and Sons, Inc.

Saxton, Arnold, Bob Auge, Mike Reed. "Least Squares Means," Glossary, Design and Analysis Web Guide (DAWG), The University of Tennessee. Accessed Feb 28, 2017.
http://dawg.utk.edu/glossary/g_least_squares_means.htm

Tobias, Randy, Weiji Cai. (2010), "Introducing PROC PLM and Postfitting Analysis for Very General Linear Models in SAS/STAT® 9.22," Proceedings of the SAS Global Forum 2010 Conference, Cary NC: SAS Institute Inc.

Wikipedia contributors. "Akaike information criterion." *Wikipedia, The Free Encyclopedia*. Wikipedia, The Free Encyclopedia, 4 Jan. 2017. Web. 27 Jan. 2017.

Appendix A

LS-means versus regular means: The case for LS-means

This example appears in:

Cai, Weiji. (2010). "Making Comparisons Fair: How LS-Means Unify the Analysis of Linear Models," Proceedings of the SAS Global Forum 2010 Conference, Cary, NC: SAS Institute Inc.

The following simple example illustrates the concept of LS-means. In a study of salaries of faculty members selected from different departments at a university, two factors are considered: gender (male and female) and rank (associate and full). Table 1 provides salary means in thousands and sample sizes (shown in parentheses) for each combination of gender and rank.

Table 1. Salary Information for Tenured Professors at a University

Gender	Associate	Full
Male	130 (12)	136 (2)
Female	118 (4)	120 (5)

The mean salary for associate professors is $((130 \times 12) + (118 \times 4)) / 16 = 127$, and the mean salary for full professors is $((136 \times 2) + (120 \times 5)) / 7 = 124.6$. The overall mean salary for associate professors is higher than for full professors even though associate professors earn less than full professors in each gender category. The seeming contradiction is caused by the imbalance in the data. The associate professors are predominantly male, and all male professors earn more than their female colleagues in this particular sample. To correct the imbalance, you can compute LS-means for associate and full professors. The LS-means are simply arithmetic means over genders. For associate professors, the LS-mean is $(130 + 118) / 2 = 124$. For full professors, the LS-mean is $(136 + 120) / 2 = 128$. Thus, the least squares mean of salary for associate professors is lower than for full professors. The LS-means ignore the sample size information in each group and assume a balanced gender distribution in the underlying population. If a balanced design had been available (that is, the sample sizes in all the groups had been the same), then the LS-means would be equivalent to the means. Thus, an LS-mean can be thought of as the mean that would be calculated if a balanced design had been obtainable.

Appendix B

Random statements for experimental designs commonly used for agricultural experiments

CRD – Completely Randomized Design

```
proc mixed;
class A B;
model response = A B A*B;
```

A CRD has neither blocks nor replications and consequently, no random effects.

RCB – Randomized Complete Block

```
proc mixed;
class block A B;
model response = A B A*B;
random block;
```

Split-plot – A is the main plot, B is the sub-plot

```
proc mixed;
class block A B;
model response = A B A*B;
random block block*A;
```

Split-block or Strip-plot

```
proc mixed;
class block A B;
model response = A B A*B;
random block block*A block*B;
```

Split-split-plot - A is the main plot, B is the sub-plot and C is the sub-sub-plot

```
proc mixed;
class block A B C;
model response = A B C A*B A*C B*C;
random block block*A block*A*B;
```

RCB at multiple locations and location is a random effect.

If the experiment is for multiple years instead of locations and year is random, substitute year for location.

```
proc mixed;
class location block A B;
model response = A B A*B;
random location block;
```

Split-plot at multiple locations. Location is a random effect and A is the main plot.

```
proc mixed;
class location block A B;
model response = A B A*B;
random location block block*A;
```

Notice that except for treatment C in the split-split-plot model, the model statements are the same for the different experimental designs. The terms in the random statement differentiate the models and determine how the standard errors are adjusted to test the fixed effects listed in the model statement. Random factors such as block, location or year broaden the inference gained from statistical tests by implying that one would get the same results regardless of the block, location or year in which the experiment occurred. Factors such as block*A, are error terms derived from hierarchical or nested designs like split-plots. They are the same terms that statistical texts for agriculture refer to as error(A), error(B), etc.

Appendix C

GLM statements for experimental designs commonly used for agricultural experiments

CRD – Completely Randomized Design

```
proc glm;
class A B;
model response = A B A*B;
means A B / lsd lines;
```

RCB – Randomized Complete Block

```
proc glm;
class block A B;
model response = block A B A*B;
means A B / lsd lines;
```

Split-plot – A is the main plot, B is the sub-plot

```
proc glm;
class block A B;
model response = block A block*A B A*B;
test h=A e=block*A;
means A / lsd lines e=block*A
```

In the test statement, h is short for hypothesis and e is short for error. The statement tells SAS to test the effect of **A** on the response variable and use **block*A** as the error term for the test.

Use the main effect error term to test the main effect means.

Split-block or Strip-plot

```
proc glm;
class block A B;
model response = block A block*A B block*B A*B;
test h=A e=block*A
test h=B e=block*B;
means A / lsd lines e=block*A;
```

```
means B / lsd lines e=block*B;
```

Use the main effect error terms to test the main effect means.

Split-split-plot - A is the main plot, B is the sub-plot and C is the sub-sub-plot

```
proc glm;
class block A B C;
model response = block A block*A B A*B block*A*B C A*C B*C;
test h=A e=block*A;
test h=B A*B e=block*A*B;
means A / lsd lines e=block*A;
means B A*B / lsd lines e=block*A*B;
```

Use the main and sub-plot effect error terms to test the main and sub-plot effect means.

Repeated measures – a randomized complete block measured over dates

```
proc glm;
class block A B date;
model response = block A B A*B block*A*B date;
test h=A B A*B e=block*A*B;
means A B / lsd lines e=block*A*B;
means date / lsd lines;
```

A B and A*B should be tested with block*A*B because if the plots had not been measured over dates, block*A*B would have been residual error.

Notes:

The error terms listed in the test and means statements correspond to **error(A)**, **error(B)**, etc., that are commonly listed in ANOVA tables in classic statistical texts.

The examples use Fisher's LSD mean separation procedure in the **means** statement. Other mean separation procedures, such as, **tukey**, **duncan**, or **waller**, would work just as well.

Least squares means are also available. This statement produces tests that are similar to the split-plot means statement:

```
lsmeans A / pdiff lines e=block*A;
```

Use the **adjust=** option to get adjustments to the default pairwise t-tests. For example:

```
lsmeans A / pdiff adjust=tukey lines e=block*A;
```