An Introduction to the Analysis of Repeated Measures for Continuous Response Data using PROC GLM and PROC MIXED

Robert M. Hamer, UMDNJ R. W. Johnson Medical School
Pippa M. Simpson, University of Arkansas for Medical Sciences

Abstract

Many methods have been used for the analysis of repeated measures data. The classical approach is to treat the experimental units in a repeated measures study as blocks in a blocked design. Multivariate approaches make fewer assumptions than the classical approach, but in general are less powerful when the assumptions that the classical approach requires are met. The newest approach uses mixed models, which may not require as stringent assumptions as the classical approach, and under some circumstances may be more powerful and flexible than the multivariate approach. In this paper, we define repeated measures designs, and discuss their analysis with both PROC GLM and PROC MIXED.

Definitions

A repeated measures design is one in which at least one of the factors consists of repeated measurements on the same subjects or experimental units, under different conditions. Such a factor is commonly called a "within-subjects," factor. A "between-subjects" factor is one in which each level of the factor contains different experimental units. In this paper, we will concentrate on experimental designs with both within-subjects and between-subjects factors.

In a repeated factor, the repeated measurements are not simply replicates of each other, but there is some sort of qualitative or quantitative relationship among the levels of that factor. A repeated measures experiment differs from one in which the multiple measurements are simple replicates; e.g., subjects have blood drawn one occasion, and the sample is divided into several replicate subsamples, each of which is measured.

A brief discussion of repeated measures designs may be found in Spector, 1987 (SUGI 12 Proceedings, page 1174).} (See also O'Brien, R.G., and Kaiser, M.K., MANOVA Method for Analyzing Repeated Measures Designs, Psychological Bulletin,1985, 97, 316-415.)

We discuss univariate, multivariate, and mixed model approaches to repeated measures.

Our illustrations will be drawn from our medical research contexts, usually clinical trials. A clinical trial is an experiment in which human beings are subjects. Typical clinical trials consist of several groups (treatments), and the subjects in each group are measured repeatedly over time (occasions). For example, in a study to test a new medication for psychosis, we might have two groups (placebo and drug), and measure the subjects on six occasions (weeks -2, 0, 2, 4, 6, 8). (We have numbered the occasions in this manner because in many clinical trials, we screen subjects (week -2), and have to wait some time (here, 2 weeks) for various test results to become available, take baseline measures (week 0), randomize to group (here, placebo or drug), and then measure repeatedly for the remainder of the study (here, in an 8 week study, we are measuring every 2 weeks, that is, on weeks 2, 4, 6, and 8). As an aside, in a clinical trial, there are often missing data, which makes the analysis using standard repeated measures methods problematic. We will not discuss that here, but assume that we have complete data. In this example, treatment is our between-subjects factor, and occasion is our within-subjects factor. The reason we can't analyze this study as if it were a simple two-way factorial (with both factors treated as between-subjects factors) is that that analysis assumes all observations are independent of each other, after eliminating variation due to the factors. It should be obvious that observations taken repeatedly on the same subjects are not independent of each other.

Simple Repeated Measures

The simplest case of repeated measures is something you learned in your first statistics class: the paired t test. The way you analyzed it was by forming the difference between the paired values, and performing a t-test on those difference scores, testing the null hypothesis that

the population mean of those difference scores was zero.

Assume you have a single group of subjects (a sample from a single population), and you have measured them both before and after some treatment. You can use a data step to take the difference between the two measurements:

```
data simple;
   input subject pre post;
   diff=post-pre;
   cards;
[…data…]
;
```

The SAS code to perform a t test on these difference scores is:

```
proc means mean std stderr t prt
     data=simple;
   var diff;
   run; cesarean
```

(You could also use PROC UNIVARIATE to obtain the same test.)

You could obtain the same test in GLM with an intercept-only model:

```
proc glm data=simplest;
   model diff=;
   title 'ttest on difference
scores';
   run;
```

Although this is an F-test, rather than a t-test, the two are functions of each other, and the p-values are the same.

Finally, you could fit a traditional blocked model, using the subjects as blocks, and the repeated measures as dependent variable. This requires that the data be reshaped such that each observation consists of a subject code, a variable indicating whether the observed value was observed at time 0 or time 1, and the dependent variable value itself:

```
proc glm data=simplong;
   class id time;
```

```
   model y = id time;
   title 'Univariate approach one
within factor design';
   run;
```

In this analysis, subject id is treated as a blocking factor.

We could even analyse these data using PROC MIXED, simply treating it as a fixed effects blocked design:

```
proc mixed data=simplong;
   class id time;
   model y = id time;
   titlel 'fit uni approach with
proc mixed';
   run;
```

Notice that we have no REPEATED statement nor RANDOM statement in our invocation of PROC MIXED. The MODEL statement in MIXED is used to specify FIXED effects only, and is thus fitting exactly the same model as is GLM. (Further, we ignore the issue of whether a repeated measures analysis is most appropriate for these data, as opposed to an analysis of covariance, for example.)

Analyzing repeated measures designs treating subjects as blocks is necessary to meet the independence of observations assumption central to the usual fixed-effects ANOVA that GLM uses. The usual ANOVA assumes that all observations are independent of all other observations (within groups), and if multiple observations come from the same subject, that is obviously not true. That is why we cannot annualize repeated measures designs using the same methods we would use for factorial designs, treating the within-subjects factors as if they were between-subjects factors. We must somehow fit a model that extracts this dependence from the errors, and moves it into the model, thus rendering the errors either independent, or something enough like independent that we can handle it properly. The blocked ANOVA does this.

It turns out that it is not strictly necessary for all the measurements within a subject to be independent of each other, in order for ANOVA to produce appropriate estimates and tests, but

that they meet a less restrictive condition called sphericity. Spherically roughly corresponds to a statement that all correlations among differences among the occasions are a constant. With only two levels to the within-subjects factor (two time points), the sphericity assumption is met by definition, and we don't have to worry about it. If there are more than two time points, that assumption may be an issue. Thus, in the above example, with only two time points, sphericity was not an issue.

If there are more than two time points, sphericity may be an issue. There are several methods of handling that: one may (a) perform a test in which the null hypothesis is that the sphericity assumption is met, and if it is not, adjust for the extent to which it is not met, (b) perform the analysis via a multivariate approach to repeated measures, which doesn't make the sphericity assumption. The multivariate approach is less powerful when the sphericity assumption is met. Finally (c), one may use mixed models, which have the capacity to fit models with a variety of error covariance matrices, which may include several patterns of relationships among the occasions.

More Complex Models – One Between-Subjects and one Within-Subjects Factor

As stated earlier, we use a clinical trial for our example, with two groups, placebo and drug (the between-subjects factor with two levels), and three occasions (the within-subjects factor with three levels). We use a within-subjects factor with three levels (occasions). We chose three occasions because had we fewer we would not need to concern ourselves with the sphericity assumption, and more than three occasions would not teach us anything we can't learn with three occasions.

To analyse these data using the traditional univariate approach to repeated measures, with subjects treated as blocks, we could use PROC GLM:

```
proc glm data=nextlong;
   class group time id;
   model yy = group id(group) time
group*time;
   test h=group e=id(group);
```

```
   title 'Univariate Approach w/o
Correction';
   run;
```

In this analysis, we have one observation for each time point within each subject. Since group is between-subjects, we need a between-subjects mean square to use in the denominator of the F-statistics to test the null hypothesis about it. We construct this F test with the test statement. The id(group) in the model is subject nested within group, which is the within-group variation. That is the appropriate yardstick to use to examine the between-group variation. The usual residual error term is the within-subjects variation, which is the appropriate yardstick to use to examine the time and group*time effects.

This analysis implicitly assumes sphericity, but does not examine it. If we wish to perform an analysis that does not assume sphericity, and we wish to examine it more closely, we may use the REPEATED statement in GLM:

```
proc glm data=nextwide;
   class group;
   model y0-y2 = group;
   repeated time 3 / printe;
   title 'univariate and mult
approach with repeated
statement';
   run;
```

The REPEATED statement triggers several analyses: It produces the univariate approach to repeated measures that we discussed above; it produces the multivariate approach to repeated measures, which does not require the sphericity assumption. It tests (if we use the PRINTE option on the REPEATED statement) the sphericity assumption. Further, it produces two adjustments to the univariate approach, to allow us to use a univariate approach but adjust for the extent to which the sphericity assumption is violated. Finally, it produces a multivariate approach to repeated measures, in which the repeated observations are simply treated as multiple dependent variables, with contrasts among them used to produce the time effect and the group*time interaction. Since dependent variables in a multivariate ANOVA are allowed to have any pattern of correlation among them,

we don't have to worry whether the sphericity assumption is met.

Finally, PROC MIXED actually divides the model into two portions – one for the fixed effects (in this case, think of that as the between-subjects factor), and the other for the random effects (think of that as the within-subjects factor). The REPEATED statement in MIXED is a different statement than the REPEATED statement in GLM, with a different syntax. It is used in MIXED to specify the variable within-which the repeated measures occur. The between-subjects effects and the within-subjects effects are each allowed to have their own covariance matrices, which means that we can fit a wide range of models to a wide range of data, without necessarily assuming that the errors for all observations are independent of the error terms for all other observations.

To analyze these data with PROC MIXED, we would use the following code:

```
proc mixed data=nextlong;
   class group id time;
   model yy = group time
group*time;
   repeated / subject=group(id) r;
   title 'MIXED with one between
one within REPEATED stmt';
   run;
```

Note that the MODEL statement contains the effects in which we are interested, the group effect, the time effect, and the group by time interaction. The REPEATED statement is used to tell MIXED that there is some effect that occurs within subject within group. MIXED examines the values of each variable for each subject within each group, and realizes that the variable time occurs multiple times. It thus realizes that time (and hence the group*time interaction) are within-subject factors.

Finally, another advantage of MIXED is the ability to fit models in which the error between subjects are allowed to be correlated. You can fit various patterned covariance matrices, but one of particular interest to repeated measures designs is a compound symmetric covariance matrix. This is very close to the sphericity condition we discussed earlier. Specifying

TYPE=CS in the REPEATED statement forces MIXED to fit the data to a model with a compound symmetric covariance matrix:

```
proc mixed data=nextlong;
   class group id time;
   model yy = group time
group*time;
   repeated / type=cs
subject=group(id) r;
   title 'MIXED with one between
one within REPEATED stmt and
TYPE=CS';
   run;
```

If we are willing to assume compound symmetry, it makes sense to model it. MIXED is capable of fitting a wide variety of covariance structures.

In summary, we have used GLM and MIXED with several types of repeated measures designs, and compared and contrasted them. We have discussed some particular advantages to the use of the MIXED procedure.

Contact Information:

Robert M. Hamer, Ph.D. hamer@umdnj.edu
Department of Psychiatry
UMDNJ-RWJMS
675 Hoes Lane
Piscataway, NJ 08854

Pippa M. Simpson, Ph.D.
Arkansas Children's Hospital / UAMS