

1 Introduction

In statistics, life becomes more complicated when you collect **repeated measures** (or **longitudinal**) data, where subjects are observed/measured more than once. For example, this could occur in an experiment where you compare values of an outcome variable before and after a treatment or intervention, or you may be doing a study where you want to look at changes over time in one or more outcome variables.

This leaflet outlines a few possible strategies for the analysis of such data when the outcome (dependent) variable is quantitative. It does NOT explain how to carry out these methods but suitable references are suggested. The first section covers the relatively simple case where subjects are only measured twice; the next section covers situations where you have more than two measurements per subject.

2 Two measurements per subject

2.1 One group of subjects

This is the simple case where you have only one group of subjects and have measured them on two occasions only — for example, in a clinical trial this could be before and after treatment. The analysis of such data is straightforward — you would use a paired t-test (or the non-parametric equivalent if the assumptions for the paired t-test are not met).

2.2 Two or more groups of subjects

There are three options here:

1. Ignore the baseline measurements and simply compare the final measurements using the usual tests used for non-repeated data — e.g. an unpaired t-test or oneway ANOVA, depending on the number of groups being compared. If you wanted to take account of other variables, multiple regression could be used. Obviously this takes no account of the baseline measurements. This is only really a feasible option (although not necessarily recommended) if there is no reason to believe that the groups differed at baseline, such as in a study where subjects are randomly assigned to groups.
2. Calculate the change for each subject and compare the average change across groups using the same methods as just described. Unfortunately, this method tends to over-compensate for any baseline differences.
3. Use a method called **analysis of covariance** (or **ANCOVA**). This is a way of comparing the final values across groups but taking account of the baseline values. Essentially it is done by

carrying out a multiple regression, putting in the final values as the dependent variable and the baseline values, as well as the variable used to define the groups, as the independent variables. If applicable, other variables can also be taken account of by adding them into the multiple regression. Further details are given in the reference by Vickers and Altman. ANCOVA is also described in many books on regression (see, for example, the book by Edwards).

3 More than two measurements per subject

3.1 One group of subjects

In this case you would probably use repeated measures ANOVA (or, if the assumptions for this are not met, the non-parametric equivalent, Friedmans's test).

3.2 Two or more groups of subjects

The temptation with such data may be to compare subjects at each time point separately, perhaps with a series of unpaired t-tests. This is **not appropriate**. There are several analysis options available:

1. Use a mixed measures ANOVA (the variable defining the groups would be the between subjects factor). Note that in SPSS this is done in the same way as repeated measures ANOVA, except you specify a between subjects factor as well as defining your within subjects variable.
2. Use a suitable summary statistic for each individual (e.g. the mean) as the outcome variable and analyse the data using standard tests such as the unpaired t-test, oneway ANOVA, etc. The choice of summary measure depends on the purpose of the study — for example you may want to know whether one group reached their maximum value before another group, whether the 'overall' response was higher/lower, whether the rate of change was different. . . . The paper by Matthews et al. gives further details (although this paper is focused on medical research, the principles apply equally in other applications). The advantage of this method is that it is more easily understandable than more complex approaches.
3. Multi-level modelling is another possible approach to analysing repeated measures data. This cannot be done in all statistical packages — possibilities include S-Plus and a free package called R. Alternatively, there is a specialist multi-level modelling package called MLWin.

4 References

1. Vickers AJ and Altman DG (2001), **Analysing controlled trials with baseline and follow up measurements**, British Medical Journal, Vol. 323: 1123-1124.
2. Edwards AL (1985), **Multiple regression and the analysis of variance and covariance**, Freeman.
3. Matthews JNS, Altman DG, Campbell MJ, Royston P (1990), **Analysis of serial measurements in medical research**, British Medical Journal, Vol. 300: 230-235.
4. Pinheiro JC and Bates DM (2000), **Mixed-effects models in S and S-PLUS**, Springer.
5. Goldstein H (2003), **Multilevel statistical models**, Arnold.