Spreadsheets for Analysis of Controlled Trials, Crossovers and Time Series

Will G Hopkins

Institute of Sport Exercise & Active Living, Victoria University, Melbourne, Australia. Email. Reviewer: Alan M Batterham, Health and Social Care Institute, Teesside University, Middlesbrough, UK.

Spreadsheets at this site for the analysis of controlled trials, post-only crossovers and pre-post crossovers have been modified to allow adjustment for, and estimation of, the effect of two covariates (predictor variables). The user of the spreadsheet now has to specify a "custom" effect by inserting weighting factors to make any combination of each subject's repeated measurements. There is provision for estimating the custom effect as a trend in some or all of the repeated measurements representing a time series. Adjustment for a single covariate is obtained as previously, using the FORECAST function to fit a simple linear relationship between the covariate and the custom effect and thereby to predict the value of the custom effect at the mean or some other chosen value of the covariate. Adjustment for two covariates is achieved by predicting the value of the custom effect at chosen values of each predictor using the regression coefficients provided by the LINEST function. The standard error of the predicted value, which is required for inferential statistics, is not provided by LINEST but was derived from the standard errors of the regression coefficients by accounting for the correlation between the predictors. As before, individual responses to treatments are estimated by comparing variances of custom effects in experimental and reference (control) groups or treatments; in the case of post-only crossovers, two treatments are assumed to be repeats of a control or reference treatment, and these are now assigned to a separate spreadsheet to allow comparison of the variance of their change scores with the variance of change scores between reference and active treatments. The spreadsheets provide the usual analyses of the raw and log-transformed dependent variable, along with non-clinical and clinical magnitude-based inferences for smallest important effects defined by raw, percent, factor and/or standardized differences in the dependent variable. Values of the trial(s) or treatment(s) chosen to provide the standard deviation for standardizing are automatically selected as the values of one of the covariates, but these values can be replaced with those of another covariate. Accuracy of the computations in the spreadsheets was checked by analysing simulated datasets with the spreadsheets and with mixed modeling in the University Edition of the Statistical Analysis System (SAS Studio). KEYWORDS: design, inference, intervention, log-transformation, magnitude-based inference, mixed modeling, repeated measures.

Reprint pdf · Reprint docx
Spreadsheets: controlled trial, post-only crossover, pre-post crossover

Other resources. Adjustment for a covariate is explained in an article accompanying the spreadsheets that first included one covariate. Individual responses and a measure of reliability are explained in an article accompanying one of the first version of the controlled-trial and crossover spreadsheets. (Note that only log-transformation available now.) Deciding which kind of controlled trial to use is explained in

another article.

Update April 2019. Compatibility limits and magnitude-based decisions (MBD) for standardized effects now take into account uncertainty in the standardizing SD using formulae for the standard error and degrees of freedom of the standardized effect developed in the article on standardized effects in the 2019 issue. The limits and decisions are trustworthy for sample

sizes of the standardizing SD of 10 or more (degrees of freedom ≥9). A formula for the degrees of freedom of the standardizing SD in the pre-post crossover spreadsheet (in which the SD is the square root of the mean of the two treatment variances) was intuited from a simulation. The compatibility limits and MBD for standardized individual differences and responses have not been updated, on the reasonable assumption that the uncertainty in the SD representing individual responses is much greater than the contribution of the uncertainty in the standardizing.

Updates March 2018. The controlled-trial and pre-post crossover spreadsheets no longer have instructions on post-only analyses or simple comparison of means, because the approach with these spreadsheets is either irrelevant or too complex. Use instead the spreadsheet compare two group means, even though that spreadsheet as yet allows adjustment for only one covariate. I have also provided more analyses and comments in relevant cells for SDs representing individual responses. The comment in the cell Number of independent inferences has been updated to remove what I thought was an efficient method to constrain error rates.

Update June 2017. I have improved the selection of tests for baseline/reference and for standardizing. There is also now no real distinction between pre- and post-tests, to facilitate insertion of extra trials (of which five are now provided) and to allow any trial(s) to be selected for baseline or standardizing.

Update April 2017. Spreadsheets for plots of change scores are now available in each of the three workbooks. Such plots are useful to visualize trends in a time series when you have missing data, because missing data can add "noise" to the plots of the original scores (shown in Sheet1 of each workbook). Use plots of change scores to guide the inferences you want to make in Sheet1, but be aware that adjustment for baseline (X2) and any other covariate (X1) will change the mean change scores in Sheet1 but not in the plots.

IMPORTANT! If you have lots of repeated measurements and missing values, I strongly advise you to clear any log-transformed values of missing data, save the spreadsheet as a template, then save copies for each of the custom effects that you want to investigate. You can

then clear individual values of the custom effects and/or delete whole rows with the missing values without compromising the analyses of other custom effects that do not involve missing values.

The spreadsheets for analysis of controlled trials, crossovers and time series at the Sportscience site are based on reducing repeated measurement on each subject to a single effect (usually a change score). The resulting analyses are as effective as mixed modeling for mean effects and individual responses. When I added a covariate to the spreadsheets 11 years ago (Hopkins, 2006a) to allow adjustment for a modifying subject characteristic, I acknowledged in the article that "extending the analysis to two or more covariates is simple in theory, but it is practically impossible in Excel because of the bizarre awkwardness of the LINEST function (which performs the necessary multiple linear regression)." I have now solved this problem by getting the user of the spreadsheet to specify only one effect at a time, using a row of weighting factors to combine the repeated measurements into a single "custom" effect.

The only other challenge in updating the spreadsheets was to derive the custom effect's standard error, which is needed for confidence limits and magnitude-based inferences. This challenge is easiest to understand with the crossover spreadsheet. Inclusion of covariates with a crossover is simply a matter of fitting a multiple linear regression, in which the covariates predict the individual values of the custom effect. The value of the effect predicted at the means or some other chosen values of the covariates is the adjusted effect–that is, the effect adjusted to the mean values of the covariates–and hence the need for the standard error of a predicted value in a multiple linear regression. LINEST provides standard errors only for the regression coefficients, not for any linear combination of the regression coefficients. After a fruitless search on the Internet, I worked out the appropriate formula from first principles. The formula takes into account the correlation between the two predictors to reduce (for a positive correlation) the standard error given by the sum of the squares of the standard errors for the contributions of each predictor. You can find the formula in the cell beneath the label "SE pred" in each spreadsheet. For the controlled trial and pre-post crossover spreadsheets, the

standard errors of the adjusted custom effects in the two groups are combined as in the previous versions of the spreadsheets: using the Satterthwaite (1946) approximation for degrees of freedom (controlled trials), and using the correlation between the custom effects (pre-post crossover).

Accuracy of the computations in the spreadsheets was checked by analysing the simulated datasets shown there with mixed modeling in the University Edition of the Statistical Analysis System (SAS Studio). The programs can be found in the suite of materials for the mixed-model workshop (Hopkins, 2016).

The old spreadsheets can be accessed from the links in the 2006 article. The new spreadsheets have the following new or improved features…

- **Missing data not permitted when adjusting for both covariates**. The analyses for a single covariate will work as before with missing values, but LINEST does not allow blanks or non-numeric values. A warning message with instructions on how to deal with it appears when a custom effect includes a missing value. If you want to adjust for both covariates, you have to right-click-delete the entire row. Generating the error messages required use of the SUMIFS function, which is not available in the 97-2003 (.xls) version of Excel. The new spreadsheets are therefore saved as 2007 versions (.xlsx).

- **Specification of a "custom" effect with weighting factors**. See above. This approach makes it easier for you to specify averages of repeated measurements, once you understand that averaging n measurements requires weighting factors of 1/n (or -1/n). An error message appears if the sum of the weights does not equal exactly 0 and if the sum of the positive values does not equal exactly 1.

- **Provision for data representing an overt time series**. The values for the time of each measurement can be inserted and a custom effect specified by fitting lines to the data using the functions SLOPE (to determine overall linearized change) or FORECAST (to extrapolate a baseline and compare with observed values during or after an intervention). It is also possible to fit non-linear models (e.g., an exponential model for values reaching a plateau) using the Solver in Excel, but you will have to install the Data Analysis pack first. If you

can't see it at the far right of the Data menu, follow instructions for installing it at this link. Unfortunately you will have to run the Solver for each subject, so you might need to do it in a separate spreadsheet, then transfer the appropriate parameters from the non-linear model to the Custom effect column. Figure out in advance whether you should fit the non-linear model to the raw or the log-transformed data.

- **A line diagram and a time-series scatterplot of the means and SDs of each repeated measurement**. For log-transformed data, these graphs show the back-transformed values, and the Y axis is displayed as logarithmic to base 2. Excel does log plots poorly, so there is advice next to the plots about improving the appearance either by pasting into Powerpoint or by using a different graphing package.

- **Adjustment for each predictor separately and together**. The effects all sit in adjacent columns, so the reader can easily see how the values change when one, other, or both of the covariates are included. In general you should include a covariate if you think it could have a substantial effect, and you should keep it in the analysis even if it doesn't have a substantial or clear effect. Of course, with the current spreadsheet you can end up with only two covariates in the analysis, so you might have to keep one in (for example, the baseline or reference value) and report on each of the others when they are included separately. If the baseline or other covariate does not reduce the width of the confidence interval of the custom effect, removing it from the analysis is acceptable (but report its trivial effect, clear or unclear).

- **The appropriate baseline value of the dependent variable as a default predictor**. *Appropriate* here is the value of the trial or treatment or mean of the trials or treatments chosen to provide the standard deviation for standardizing. It is shown as the X2 covariate and its log transform. Use of the SUMIFS function was also required to display the values of the chosen trial(s) or treatment(s). You can replace the formulae in the raw-data X2 cells with the values of another covariate, in which case the values are not log transformed.

- **Improved advice on binary predictor variables**. A comment in the X1 covariate heading sug-

gests that you can code the variable as 0 and 1 to represent, for example, females and males. Comments in the cells **Adjusted to predictor** and **Effect of delta predictor** explain how to get the mean of females and males and their difference. It is better to analyze subgroups separately to allow for different errors and different effects of the other covariate, then compare and possibly combine the effects for the two (or more) groups using the combine/compare effects spreadsheet (Hopkins, 2006b).

- Improved instructions about mediator (mechanisms) analyses. Comments in several cells direct you to investigate a potential mediator by including its change scores as the X1 predictor. The change scores have to match those of the custom effect, of course. The custom effect adjusted to X1 = 0 is then the custom effect *independent of* the predictor, and setting **Effect of delta predictor** to the mean change of the predictor provides the custom effect *explained by* the predictor.

- Adjustment for an order effect in a crossover. Such adjustment is possible in the post-only spreadsheet only for a crossover of two treatments. For example, if the custom effect is TrtC-TrtA, for each subject assign X1 the value 1 if TrtC was the first treatment or 2 if it was the second treatment. Adjustment for the order effect is then achieved by adjusting to a value of X1 = 1.5 (which will not be the mean, if the order of treatments was not balanced). The order effect itself (second minus first treatment) is obtained by estimating the effect of delta X1 = 1. A similar approach applies to the pre-post crossover. These instructions are included in comments in several cells. Adjustment for, and estimation of, the order of more than two treatments requires mixed modeling with data in "long" format.

- New scatterplots. Two scatterplots, one for each predictor are the same kind as in the previous spreadsheets: individual values of the custom effect (e.g., a change score) plot-

ted against the value of the predictor without adjustment for the other predictor. There are two new scatterplots: one for each predictor, showing individual values of the custom effect adjusted to the chosen value of the other predictor. The adjustment is done with the model using both predictors.

- Individual responses to a treatment. These are estimated as before from the difference between the variances of control and experimental groups or treatments in the controlled trial and pre-post crossover spreadsheets. A new approach was needed for the post-only crossover: two treatments assumed to be repeats of a control or reference treatment are now assigned in a *Reference* copy of the spreadsheet. Any change to the number of rows in the main spreadsheet must be performed in the same sequence in this Reference spreadsheet.

- Improved cosmetics. It is now easier to see which cells need your data and which cells are results. I have also provided more extensive comments.

It is inevitable that the new spreadsheets have bugs, hopefully only minor formatting errors or ambiguous instructions. Please get back to me by email if you encounter problems. Also get back to me with any other comments. The new approach with LINEST lends itself to adding another covariate, which I will do if these spreadsheets are received with enthusiasm.

Hopkins WG (2006a). Spreadsheets for analysis of controlled trials, with adjustment for a subject characteristic. Sportscience 10, 46-50

Hopkins WG (2006b). A spreadsheet for combining outcomes from several subject groups. Sportscience 10, 51-53

Hopkins WG (2016). SAS (and R) for mixed models. Sportscience 20, iii

Satterthwaite FW (1946). An approximate distribution of estimates of variance components. Biometrics Bulletin 2, 110-114