StatNews #13

Differences Between ANOVA and Regression
November 1996

Two techniques for comparing two or more population means to each other are analysis of variance (ANOVA), and regression with properly constructed indicator variables. In both techniques, F statistics may be used to test hypotheses about relationships between explanatory and response variables. Each technique results in the same overall fit to the data, but the hypotheses tested by the computer-generated output are not necessarily the same. This newsletter will use an example that arose recently in consulting to illustrate how the hypotheses tested by the two techniques may differ.

A researcher wants to estimate the effect of two explanatory factors, minority status and sex, on test scores. Each factor has two categories, and is represented by a variable coded 1 for the first category and 0 for the second category. The data are collected in an observational study. As is typical in observational studies, the number of subjects in each cell (defined by the levels of minority status and sex) is not the same. The adjusted sums of squares and F statistics, also known as partial, unique, or type III, should be used in this case. The researcher is interested in testing the effects of minority status and sex on test scores, and the interaction between minority status and sex.

Using software for ANOVA or general linear models, the researcher estimates a two-way factorial ANOVA, which includes the main effects of minority status and sex, and their interaction. Later, the researcher estimates the comparable model, treating it as a regression. He creates an interaction term which is the product of the two dummy variables, and includes the dummy variables for minority status, sex, and the interaction directly in the model as covariates.

Both the ANOVA and regression techniques produce significance tests for minority status, sex, and the interaction. The ANOVA technique produces adjusted F statistics, and depending on the software, the regression technique produces adjusted F or t statistics, or both. However, the tests for the effects of minority status and sex using the two different techniques are not the same. Why do the tests differ?

In the ANOVA technique, the F statistic for minority status tests the hypothesis that there is a main effect of minority status on test score, averaging over the two levels of sex. Similarly, the F statistic for the sex tests the hypothesis that there is a main effect of sex on test score, averaging over the two levels of minority status. This technique tests the effect of each factor after controlling for the other, but not controlling for the interaction. This is the same hypothesis tested by a two-way ANOVA with no interaction. Note that although the two models, with and without the interaction, test the same hypothesis, they produce different results.

In the regression technique, the F statistic for minority status tests the hypothesis that there is an effect of minority status when sex equals zero. Similarly, the F-statistic for sex tests the hypothesis that there is an effect of sex when minority status equals zero. Note that these are not the main effects of minority status and sex. This technique tests the effect of minority status after controlling for sex and for the interaction, and tests the effect of sex after controlling for minority status and the interaction. Note that the regression technique without the interaction would test the same hypothesis that is tested by the ANOVA technique.

The different hypotheses tested by the ANOVA and regression techniques result from the different ways in which the information about minority status and sex is specified. This example illustrates that care is needed in the specification and interpretation of seemingly simple models.

A related document on using contrasts to test hypotheses in ANOVA is available from the Statistical Consulting office.

Author: Cara Olsen

(This newsletter was distributed by the Cornell Statistical Consulting Unit. Please forward it to any interested colleagues, students, and research staff. Anyone not receiving this newsletter who would like to be added to the mailing list for future newsletters should contact us at cscu@cornell.edu. Information about the Cornell Statistical Consulting Unit and copies of previous newsletters can be obtained at http://www.cscu.cornell.edu).