

Testing and Interpreting Interactions in Regression – In a Nutshell

The principles given here always apply when interpreting the coefficients in a multiple regression analysis containing interactions. However, given these principles, the meaning of the coefficients for categorical variables varies according to the method used to code the categorical variables.

The method assumed here is dummy coding, whereby each category except one is represented by a dummy (or indicator) variable which has a value of one for the category being represented and a value of zero for all other categories. The category for which there is no dummy variable consequently has a value of zero for all the dummy variables and is known as the reference category.

It is also assumed, for convenience, that the indicator variables are entered into the procedure used for the analysis in such a way that the procedure doesn't do its own coding, but leaves the variables exactly as they are coded. For example, if the GLM procedure in SPSS is used, it is assumed here that the indicator variables are entered as covariates rather than fixed factors (after 'with' rather than 'by' in syntax). If they are entered as fixed factors or after 'by' in GLM, the procedure always makes the highest-numbered category the reference category, which can be a bit difficult to get your head around. For example, if our variable gender is coded 0 for females and 1 for males (so that females are the reference category), GLM reverses this so that males are now the reference category (0) and females are represented by 1.

Note that the comments here apply to the regression coefficients shown in the parameter estimates table, not to the results in the ANOVA table. Exactly the same principles apply to the interpretation of the results shown in the ANOVA table, but programs like SPSS, if we allow them to do the coding of categorical variables for us, typically don't use dummy coding for the results shown in ANOVA tables. This is a topic in itself.

As the emphasis is on interpreting interactions, no reference is made in the following to interpreting the coefficient for the constant. However, a note at the end briefly describes the effects that the strategies used for interpreting interactions have on the constant.

Two Way Interactions

In the regression equation for the model

$$y = A + B + A*B \quad (\text{where } A * B \text{ is the product of } A \text{ and } B, \text{ which is a test of their interaction})$$

the regression coefficient for A shows the effect of A when B is zero and the coefficient for B shows the effect of B when A is zero. (The coefficient for A*B shows how the effect of A changes with a one-unit increase in B, but we won't be concentrating on that here.) This rule holds whether the interaction is significant or not: its mere presence changes the interpretation of the coefficients for A and B from unconditional (when there is no interaction term included) to conditional.

Categorical Variables

Imagine that A and B are single dummy (0,1) variables, and that A represents gender (0=females, 1=males) and B represents condition (0=control, 1=experimental). Then, the interaction shows whether the effect of condition is different for males and females (or, that the difference between males and females on y is different for the two conditions).

Given the rule given earlier, the coefficient for A shows the difference in y between males and females for the control condition, because it is coded zero. Likewise, the coefficient for B shows the difference between the control and experimental conditions for female subjects.

Once this is understood, a whole new world of possibilities opens up. To find out whether male and females scores differ for the experimental condition, we can run the analysis again with the codes for condition reversed (0=experimental, 1=control), so that now the coefficient for gender shows the difference between males and females in the experimental condition. Likewise, we can see whether the treatment (experimental versus control) had an effect for males by reversing the coding for gender, so that 0=males and 1=females. This is referred to as testing the simple effects of gender and condition.

One Categorical and One Numeric Variable

Now, imagine that in

$y = A + B + A*B$ (where $A * B$ is the product of A and B, which is a test of their interaction)

A is still gender, dummy coded as before, but B a continuous variable, let's say age in years.

Exactly the same rules apply. Let's start with the coefficient for B. It shows the effect of age (by effect I mean the slope of the regression line relating age and y) for females. The test of significance of the effect for age shows whether the slope of the line departs significantly from zero and, of course, the sign of the coefficient shows whether the relationship is positive or negative. In order to examine the relationship between age and y for males, we can reverse the coding for gender, so that now the coefficient for age is that for males. This is called testing the simple slopes.

What about the difference for males and females? Exactly the same principle applies. The coefficient for A shows the difference between males and females when age is zero. But hang on – if we have a sample of people aged from 18 up, what does this mean? Well, it's a perfectly valid result as far as the model is concerned, but the coefficient is pretty meaningless given that the sample contains no one of zero age (and very few samples are likely to). So, this where centring comes in. If we subtract the mean of age for the sample from each subject's age, the mean of age is now zero and, when we run the analysis again, the coefficient for gender now shows the difference between males and females at the mean of age for the sample, a much more meaningful value. (The exciting thing, of course, is that we don't have to stop at the

mean. Say the mean age of the sample is 35, but we have a goodly number of subjects aged from 18 to 30; it would be legitimate, instead of subtracting the mean age, to subtract (say) 25, and find out whether the model suggests that there is a significant difference between males and females aged 25.)

Two Numeric Variables

Now imagine that in

$$y = A + B + A*B \quad (\text{where } A * B \text{ is the product of } A \text{ and } B, \text{ which is a test of their interaction})$$

both A and B are numeric. Let's say B is age, as above, but now A is IQ. If y is a test score of some sort, the question might be whether the relationship between age and y differs according to IQ (or whether the relationship between IQ and y differs according to age).

The coefficient for A shows the relationship between IQ and y when age is zero, and the coefficient for B shows the relationship between age and y when IQ is zero. (Of course, the coefficient for A * B answers the research question, but we're concentrating on how including an interaction changes the meaning of the coefficients for the variables involved in the interaction.)

Once again, centring will produce more meaningful values for the coefficients. If we centre age and IQ at their respective means, the coefficient for A shows the effect (slope) of IQ at the mean of age and the coefficient for B shows the slope for age at the average IQ of the sample.

Three-Way Interactions

The same principles apply when we move from two-way to higher-level interactions. Here is an example of a model with a three-way interaction and all two-way interactions:

$$y = A + B + C + A*B + A*C + B*C + A*B*C$$

Now, as well as considering the effects of the inclusion of an interaction on the interpretation of coefficients for individual variables, we can consider the effects of including higher-order interactions on the interpretation of the coefficients for lower-order interactions. But, it all makes perfect sense.

Two Way Interactions

The rules are:

When the interaction A*B*C is included:

The coefficient for A*B shows the interaction between A and B when C is zero,
The coefficient for A*C shows the interaction between A and C when B is zero, and
The coefficient for B*C shows the interaction between B and C when A is zero.

The same sorts of things described above apply here: We can manipulate the values of variables to carry out specific tests. For example, if C is gender, where 0=female and 1=male, the original analysis will show whether the interaction of A and B is significant for females; if we reverse the coding for gender, the result for A*B will show whether the interaction A*B is significant for males. If C is a numeric variable, the analysis will show whether the interaction A*B is significant when that variable has a value of zero. The usefulness of this result depends on the meaning of zero on C. If C is age, it is unlikely that the result will be useful. However, if C is centred at the mean, we can see whether the A*B interaction is significant at the mean age.

Main Effects

The rules are:

When the interaction A*B*C and all two-way interactions are included:

The coefficient for A shows the effect of A when both B and C are zero,
The coefficient for B shows the effect of B when both A and C are zero, and
The coefficient for C shows the effect of C when both A and B are zero.

Therefore, the simple effects or simple slopes for each variable can be tested by the manipulating the values of the other two variables (see the example below).

Four-Way and Higher Interactions

The above principles extend directly to any order of interaction.

An Example

Say we have three IVs, gender (0=females, 1=male), age (mean 20) and IQ (mean 100), and an DV, creativity (the mean is irrelevant for our purposes)

The original analysis:

```
glm create with gender age iq/  
print=parameters/  
design=gender age iq gender*age gender*iq age*iq gender*age*iq.
```

Selected Examples

1. The coefficient for age*iq shows whether there is an interaction between age and iq for females. Is there such an interaction for males?

```
temporary.  
recode gender (0=1)(1=0).  
glm create with gender age iq/  
print=parameters/  
design=gender age iq gender*age gender*iq age*iq gender*age*iq.
```

2. In the original analysis, the coefficient for gender*iq shows whether the relationship between IQ and creativity differs for males and females aged zero. What is the interaction between gender and IQ at the mean age of the sample?

temporary.

compute age = age - 20.

glm create with gender age iq/

print=parameters/

*design=gender age iq gender*age gender*iq age*iq gender*age*iq.*

3. In the original analysis, the coefficient for age shows the relationship between age and creativity for females with zero IQ. What is the relationship between age and creativity for males with an average IQ?

temporary.

compute iq = iq - 100.

recode gender (0=1)(1=0).

glm create with gender age iq/

print=parameters/

*design=gender age iq gender*age gender*iq age*iq gender*age*iq.*

Two final points

1. *Don't get misled and worked up about the apparently very different results obtained for a variable with and without an interaction*

Sometimes the coefficient for, or the significance of, a variable involved in an interaction changes dramatically when an interaction term is included in an analysis, especially if it contains numeric variables. The most likely reason for this is that the coefficient for that variable is now showing something very different from what it was showing when there was no interaction term. Without the interaction term, the coefficient shows the relationship between the variable and the dependent variable averaged over all levels of the other variables. When the interaction is included (whether it is significant or not), the coefficient for the variable shows the effect of that variable when the other variable involved in the interaction is zero. This is called a conditional effect, and can be very different from the unconditional effect obtained when there is no product term included in the analysis. Centring the numeric variable(s) at the mean will often produce coefficients which are much more similar to the unconditional coefficients.

2. *The constant is affected by the coding and centring of variables in a regression analysis*

The constant in a regression equation shows the value of the dependent variable when all the independent variables are zero. Consequently, the constant may change dramatically when numeric variables are centred at the mean. The change is usually to a more sensible value (i.e., more likely to be within the range of the values of the dependent variable actually observed) than is obtained with the uncentred version of an independent variable which does not include zero in its range. The constant may

also change noticeably when a dummy code is reversed. For example, when a variable coded zero for females and one for males is reversed, the constant goes from showing the mean for females to showing the mean for males.

Alan Taylor
20th June 2007