Interaction Contrasts/Partial Interactions

When our data shows a significant interaction, we want to be able to understand, in as much detail as possible, what that interaction means. In some cases, doing this is a relatively simple matter, such as when we have a 2 x 2 design. Once we begin to add factors and levels to our design, though, interpreting interactions becomes somewhat more difficult. Interaction contrasts and partial interactions are statistical tests that you can use to give you a better handle on what your interaction means.

Definitions:

According to Keppel, an interaction contrast is performed when we test for differences in some contrast on factor A as a function of a contrast on factor B. Interaction contrasts have one degree of freedom. For example, we can test whether a specific contrast on A (2, -1, -1) is the same at B2 and B3:

	2	-1	-1
0			
1			
-1			

If you cross multiply these coefficients, you can see how they allow you to test whether the contrast on A is the same at B2 and B3. Note that you can test this contrast on A against a variety of contrasts on B. You can see if the contrast is the same at B1 and B3, for example, or if it is the same at B1 as it is at B2 and B3 combined. By setting your coefficients appropriately, you can do a variety of interaction contrasts that allow you to see where your interaction is coming from.

In contrast to this, a partial interaction is a test that pits a contrast on one factor against an intact factor, such as testing a contrast on A against an intact factor B. For instance, we might want to ask whether the (2, -1, -1) contrast on A is the same at all levels of factor B. To do this, we represent the contrast on A, and set up a set of coefficients to represent B:

	2	-1	-1
0			
1			
-1			

	2	-1	-1
2			
-1			
-1			

Here, we have used the coefficients (0, 1, -1) and (2, -1, -1) to represent B (i.e., they completely partition factor B). Any set of coefficients that represent a complete set of orthogonal contrasts on B (including trend coefficients) would have worked just as well. By summing these two interaction contrasts, we can get the A(contrast) x B partial interaction. The df for partial interactions will be equal to the df for the contrast on A (i.e., 1) times the df of the intact factor (i.e., B, where the df = 2).

So, partial interactions let you see if a contrast on one factor is the same at all levels of another factor. Interaction contrasts let you be more specific: is the contrast on one factor the same as a function of some contrast on the other factor?

Below are some data from a 3-factor design. They will allow you to perform simple interaction effects, simple main effects, interaction contrasts and partial interactions. To get more practice, you can re-arrange the data set into various configurations, omitting certain factors, or certain levels of a particular factor, as you wish. I will have the ANOVA table for the whole 3-factor experiment for you to check your answers against, but I will not necessarily have the answers for the particular interaction contrasts, partial interactions, etc., that you may have tested. Remember what tests partition what factors on the ANOVA table, and you will be able to see if your answers are right or not. If you have any questions, come see me.

	A1			A2			A3		
	B1	B2	B3	B1	B2	B3	B1	B2	B3
C1	9	6	9	5	9	1	3	3	3
	9	7	9	2	7	2	1	1	2
	8	8	9	3	7	2	2	1	4
C2	5	3	5	3	0	8	0	7	4
	4	4	5	3	1	9	0	7	4
	3	4	4	2	1	7	1	6	2
C3	1	4	8	1	3	3	2	5	3
	0	2	8	5	2	2	1	5	3
	2	5	7	3	5	4	1	7	4