Simple Effects Test Following a Significant Interaction

Simple effects tests are follow-up tests when the interaction is significant. They explore the nature of the interaction by examining the difference between groups within one level of one of the independent variables. For example, if there was a significant interaction between violence and training, a simple effects test could compare the difference between violence and no violence conditions for those people who did not receive eyewitness training. Note that there is generally no reason to conduct a simple effects test when the interaction is nonsignificant.

To conduct a simple effects test following a significant interaction, I use the MANOVA command in SPSS (the GLM syntax command could also be used).[1] MANOVA, which stands for multivariate analysis of variance, is only available through syntax. We are not conducting a multivariate analysis of variance here, because there is only one dependent variable involved (the statistician's definition of a multivariate test is that there are multiple dependent variables involved). However, this older SPSS command was used to conduct a number of different kinds of analyses prior to the addition of windows menus and the GLM procedure and is simple to use for this.

Syntax[2]

Below, I provide the syntax for testing the simple effect of violence within the no training group.

	1 No training	2 Training	
1 No violence	4.0	8.0	6.0
2 Violence	2.0	2.0	2.0
	3.0	5.0	4.0

For the violence factor 1=no violence, 2=violence. For the training factor, 1=no training and 2=training. To test a simple effect, the following syntax commands are used:

```
MEANS VARS=memory by violence by training.

MANOVA memory BY violence(1,2) training(1,2)
  /ERROR=WITHIN
   /DESIGN=training, violence WITHIN training(1).
```
[3]

The `MEANS` command simply generates means for all the cells and the marginal means. They can be obtained with an `OMEANS` subcommand in MANOVA, but I like the presentation with the `MEANS` command better.

The MANOVA line lists the dependent variable first and then the two factors. After each factor, the upper and lower values for that independent variable are listed in parentheses. The second subcommand, `/ERROR=WITHIN`, specifies the error term to use and is always the same in any problem. The first `/DESIGN` subcommand generates the full ANOVA (which you may or may not always want). The second `/DESIGN` command, `/DESIGN=violence WITHIN training(1)`, generates the simple effects test by specifying that the two levels of violence should be compared within the first level of training. The `(1)` in this statement refers to ***the first level of the training variable rather than the group coded 1***. If training had been coded 0=no training and 1=training, then the `(1)` in this statement would instruct SPSS to test the simple effect within the no training group (i.e., the group coded 0).

Note that there is a power (and likely accuracy) advantage of this type of analysis over separate standard *t*-tests, because *t*-tests use only half of the subjects to compute the error term and significance is only based on half the *df*. Using simple effects tests (like planned contrasts) will use the within-cell variation for all the cases in the data set and generally will result in a smaller and more reliable error term, thus leading to higher power.

---

[1] The following GLM syntax produces both simple effects within each level of training:

```
GLM memory BY violence training
   /emmeans = tables(violence*training) compare(violence).
```

[2] This example is based on syntax provided by Page, C., Braver, S.L., & MacKinnon, D.P. (2003). *Levine's guide to SPSS for analysis of variance* (2nd Edition). Mahway, NJ: Erlbaum. GLM could also be used.

[3] Specification of the main effect for the variable after WITHIN keyword is needed for unbalanced *n* case. If design is balanced, then the DESIGN command can just specify the simple effect without the main effect of the within factor (although there is no harm in specifying the main effect in the balanced case).

Means

Report

MEMORY memory for assailant

VIOLENCE violent	TRAINING received	Mean	N	Std. Deviation
1.00 no violence	1.00 no training	4.0000	5	.70711
	2.00 training	8.0000	5	1.22474
	Total	6.0000	10	2.30940
2.00 violence	1.00 no training	2.0000	5	1.22474
	2.00 training	2.0000	5	1.22474
	Total	2.0000	10	1.15470
Total	1.00 no training	3.0000	10	1.41421
	2.00 training	5.0000	10	3.36650
	Total	4.0000	20	2.71448

Manova

```
    * * * * * * * * * * * * * * * * * A n a l y s i s   o f   V a r i a n c e -- Design   1 * * * * * * * * * *
Tests of Significance for memory using UNIQUE sums of squares
Source of Variation          SS      DF        MS         F  Sig of F

WITHIN CELLS              20.00      16      1.25
TRAINING                  20.00       1     20.00     16.00      .001
VIOLENCE WITHIN TRAI      10.00       1     10.00      8.00      .012
NING(1)
- - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
```

R

The most convenient method I have found for testing simple effects in R is to use yet another R package, the *phia* package. Like the GLM method, the `testInteractions` function from the `phia` package produces simple effects comparisons across the groups for one independent variable within each of the levels of the other independent variable, and you may not want to use all of those comparisons. The `fixed` statement refers to the factor for which we want to make comparisons within a particular level (here, we are looking within one of the levels of training) and the `across` statement is to specify the factor levels we are comparing across (here, we are comparing across levels of violence). `adjustment="none"` statement is used to indicate no post hoc adjustment to the p-values. `mymodel` refers to the results from the `car` package `Anova` procedure demonstrated in the factorial handout, so that analysis must be specified first.

```
> library(phia)
> testInteractions(mymodel, fixed="training", across="violence", adjustment="none")

F Test:
P-value adjustment method: none
          Value      SE Df Sum of Sq  F       Pr(>F)
1             2  0.7071  1        10  8      0.01211 *
2             6  0.7071  1        90 72 0.0000002562 ***
Residuals       16.0000 20
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
```

Write-up

(This material would be added to the end of the factorial ANOVA write-up example, see the "Factorial ANOVA Example" handout).

A simple effects test examined the whether the violence manipulation had significant effects for those who did not receive any training, F(1,16) = 8.00, p = .01. Within the no training group, those who saw the violent incident had poorer memory (M = 2.00) than those who did not see the violent incident (M = 4.00).

I'm not sure this was the most logical simple effects test, and more than one might have been desired in this example. Eta-squared seems to be rarely if ever reported for these tests, probably because software packages do not generate them (although in GLM, you can add the subcommand `/print=etasq`. An eta-squared could be computed by hand as well, simply by dividing the sum of squares by the total sum of squares (between plus within), $\eta^2 = 10/(10 + 20) = .33$.