An Introduction to Factorial Analysis of Variance

Before reading further, point your browser to Hyperstat and click on Chapter 12: Introduction to Between Subjects ANOVA. Read Chapter 1 (Preliminaries), and Chapter 2. In Chapter 2, When David Lane talks about MSB (mean square between) and MSE (mean square error), note that other authors use different, but similar, terms. Two synonyms for MSE that you should be aware of are *residual mean square* and *mean square within*. These all refer to the same source of variance.

You can skip Subsection 1 (Two Estimates of Variance) of Section 2 (ANOVA with 1 between-subjects factor) if you want to—it is a little more technical than other sections in the text. You can scan Pages 3-5 in Subsection 3, Partitioning the Sums of Squares, but be sure to read Pages 6 & 7, where an ANOVA table is presented and explained.

In Section 3 (Introduction to Tests Supplementing a One-factor Between-Subjects ANOVA), read Subsection 1 (Introduction) and Subsection 2 (All pairwise comparisons among means), parts 1-4 and 8. Also, read Subsection 3 (Comparing Means with a Control) and all of Subsection 4 (Specific comparisons among means).

Finally, read Subsection 6 (Reporting results). If you are adventurous, you could try the Exercises given in Subsection 7.

Now that you have studied the one-way ANOVA, used when there is only one independent variable (IV) and one dependent variable (DV), we will turn to a consideration of factorial ANOVA, where there is more than one IV. Research designs with more than one IV are often more interesting than those with only one IV. When a research design includes two IVs and one DV, the design is called a two-way factorial design, or simply, a two-way design. When there are three IVs and one DV, the design is a three-way factorial design (or three-way design). The IVs in a multi-way design are called "factors;" hence the name, "factorial design."

In a factorial analysis, the levels of each factor are important in describing the design. For instance, in a two-way design, where we have, say, Factors A and B, suppose Factor A is sex and Factor B is "Greek Society Membership." The DV can be anything (for instance, ounces of beer consumed on a Friday evening). In this example, Factor A would have two levels: male and female. Factor B would also have two levels: member of a Greek society or nonmember. In this example, the design would be described as a 2 x 2 factorial design. Supposed we had another experiment where Factor A had three levels and Factor B had 4 levels. This design would be described as a 3 x 4 factorial design. Similarly, if we had three factors (three IVs), A, B, and C, and A had 3 levels, B, two levels, and C, 4 levels, the design would be described as a 3 x 2 x 4 factorial design.

Consider the 2 x2 design illustrated in the figure below. Subjects in each of four groups are given one of four possible combinations of barbiturates and alcohol. The DV is some measure of physical impairment. Here, Factor A is "dose of alcohol," Factor B, "dose of barbiturate." After conducting the experiment, we obtain the group means shown in the table.

Mean Level of Impairment

Barbiturate (Factor A)	Alcohol (Factor B) none	one	marginal
none	00	10	05
one	20	30	25
marginal	10	20	15

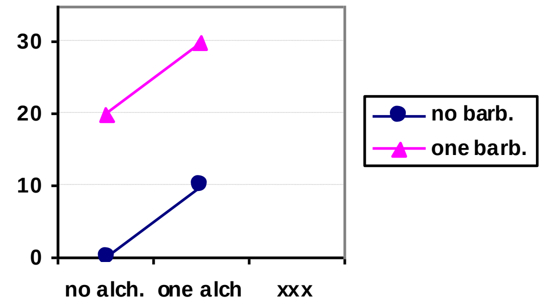

The 2 x 2 = 4 group means (0, 10, 20, 30) are called *cell* means. The cell means are averaged to obtain *marginal* means, which reflect the effect of one factor ignoring the other factor. For example, for factor A, ignoring B, participants who drank no alcohol averaged 10 on the impairment scale, those who did drink averaged 20. From such marginal means, we can compute the *main effect* of a factor, its effect ignoring the other factor. For factor B, that main effect is (20 - 10) = 10 (participants who drank alcohol were, on average, 10 units more impaired than those who did not. For factor A, the main effect is (25 - 5) = 20, the barbiturate tablet produced 20 units of impairment, on average.

A *simple main effect* is the effect of one factor at a specified level of the other factor. For example, the simple main effect of alcohol for participants who took no barbiturate is (10 - 0) = 10. For participants who did take a barbiturate, taking alcohol also made them (30 - 20) = 10 units more impaired. In this case, the simple main effect of B at level 1 of A is the same as it is at level 2 of A. The same is true of the simple main effects of A—they do not change across levels of B. When this is the case, i.e., when the simple main effects of one factor do not change across levels of the other factor, the two factors are said to be *additive*. In such cases the combined effects of A and B will equal the simple sum of the separate effects of A and B—for our example, taking only the drink makes an individual 10 units impaired, taking only the barbiturate makes an individual 20 units impaired, and taking both makes an individual (10 + 20) = 30 units impaired. The combination of A and B is, in this case, *additive*.

Look at the plot of the interaction next to the table above. One line is drawn for each simple main effect. The lower line represents the simple main effect of alcohol for subjects who had no barbiturate; the upper line represents the simple main effect of alcohol for subjects who had one barbiturate. The main effect of alcohol is evident by the fact that both lines have a *slope*. The main effect of barbiturate is evident by the separation of the lines in the vertical dimension.

Suppose that we also conduct our experiment on another set of participants and obtain the means given in the table on the next page. Note that only one cell mean, that of the participants who washed down their sleeping pill with a alcohol (a very foolish thing to do!) has been changed. Sober folks taking only the drink still get (10 - 0) = 10 units of impairment and sober folks taking the pill still get (20 - 0) = 20 units of impairment, but now when A and B are combined, we do not get (10 + 20) = 30, instead we get (10 + 20 + 10) = 40. Where did that extra 10 units of impairment come from? It came from the *interaction* of A and B in this *nonadditive* combination. An interaction exists when the effect of one

Mean Level of Impairment

	Alcohol		
Barbiturate	none	one	marginal
none	00	10	05
one	20	40	30
marginal	10	25	17.5

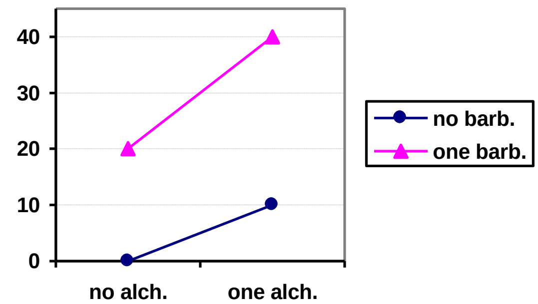

factor changes across levels of another factor, that is, when the simple main effects of one factor vary across levels of another. For this example, the simple main effect of drinking alcohol is (10 - 0) = 10 units of impairment if when a pill is not taken, but it is twice that, (40 - 20) = 20, if a pill has been taken. The alcohol has a greater effect if a barbiturate has already been consumed than if it has not. Likewise, the simple main effect of taking a barbiturate pill, if already impaired, is (40 - 10) = 30 units more impairment—more than the simple main effect of taking a pill if not already impaired, (20 - 0) = 20.

The interaction between alcohol and barbiturates, discussed here, is a called a *monotonic interaction*, one in which the simple main effects vary in magnitude but not in direction across levels of the other variable. Alcohol increases your impairment whether or not a barbiturate pill is taken; it just does it more so if you have taken a pill, also. Likewise, the pill increases your impairment whether or not you have taken alcohol; it just does more so if you have been drinking. With such a monotonic interaction, one can still interpret main effects—in the example, drinking alcohol or taking a pill increases impairment. In the plot, the fact that the two lines are not parallel indicates that there is an interaction. The fact that the direction of the slope is the same (positive) for both lines indicates that the interaction is monotonic.

Sometimes the direction of the simple main effects changes across levels of the other variable. In such a case the interaction may be described as a *nonmonotonic interaction*. For example, consider the following means from a third experiment:

Mean Level of Impairment

	Alcohol		
Barbiturate	none	one	marginal
none	00	20	10
one	30	10	20
marginal	15	15	15

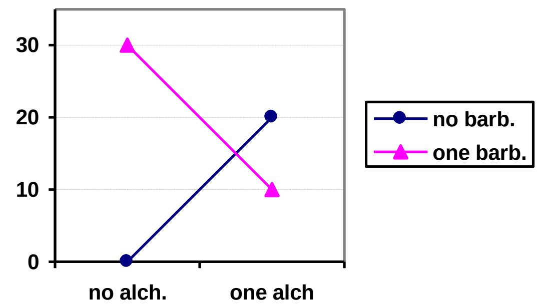

Alcohol has absolutely no main effect here, (15 - 15) = 0, and the barbiturate does, but the interesting effect is the strange interaction. For those who took no barbiturate the simple main effect of alcohol was to make them (20 - 0) = 20 units more impaired, but for those who had taken a pill beforehand the simple main effect of the alcohol was to make them (10 - 30) = 20 units less impaired. Likewise, the effect of a pill for those who did not drink alcohol, was (30 - 0) = +30, but for those who did drink alcohol, it was (10 - 20) = -10.

The presence of such a nonmonotonic interaction may make it unreasonable to interpret the main effects. For example, asked what the effect of alcohol is, we cannot honestly answer, "It makes them more impaired." The answer has to be qualified, or, "It depends. If they haven't been abusing barbiturates, it makes them more impaired, but if they have been abusing barbiturates it makes them less impaired."

The plot makes it clear that the interaction is nonmonotonic -- the direction of the slope for the one line is positive, for the other it is negative.

In a three-way factorial design (three IVs) one can evaluate three main effects (A, B, and C), three two-way interactions (A x B, A x C, and B x C) and one three-way interaction (A x B x C). A three-way interaction exists when the simple two-way interactions (the interaction between two factors at each level of a third factor) differ from one another. For the contrived data presented earlier, consider the third factor, C, to be age of participant, with level one being early adult(21-35), level two being young adult (36-50)*,"* and level three being older adult (51 and over). We have already seen that the A x B interactions differ across levels of C, so there is indication of a triple interaction in our 2 x 2 x 3 design.

You should later be able to generalize the concepts of main effects, simple effects, and interactions beyond the three-factor example being presented here, but do not be surprised if you have trouble understanding higher-order interactions, such as four-way interactions—most people do—three dimensions is generally as many as a most of us can simultaneously handle!

Chapters 12 and 13 in Salkind's book (the one with Excel) treats the topics of ANOVA. If you have this book, you should read those chapters. In Chapter 12, Salkind covers the one-way ANOVA. Read the introductory material and see if you can repeat the computations for the example given toward the end of the section, *Computing the F Test Statistic*. Try using Excel to do the computations. Then read the next section, *How do I Interpret….* Then, read the section, *Using the Amazing Analysis Toopak...*, and follow his steps in doing the example he gives.

Chapter 13 expands the discussion to a two-factor ANOVA design. Again, read all the introductory material, and the sections , *The Main Event…*, and *Even More Interesting….* Then, follow his steps and reproduce the analysis he gives in the section, *Computing the ANOVA F Statistic Using….* This may appear difficult, but it really is not.