

Using ANOVA to Examine Data from Groups and Dyads

Jamie DeCoster

Department of Psychology
University of Alabama
348 Gordon Palmer Hall
Box 870348
Tuscaloosa, AL 35487-0348

Phone: (205) 348-4431
Fax: (205) 348-8648

April 12, 2002

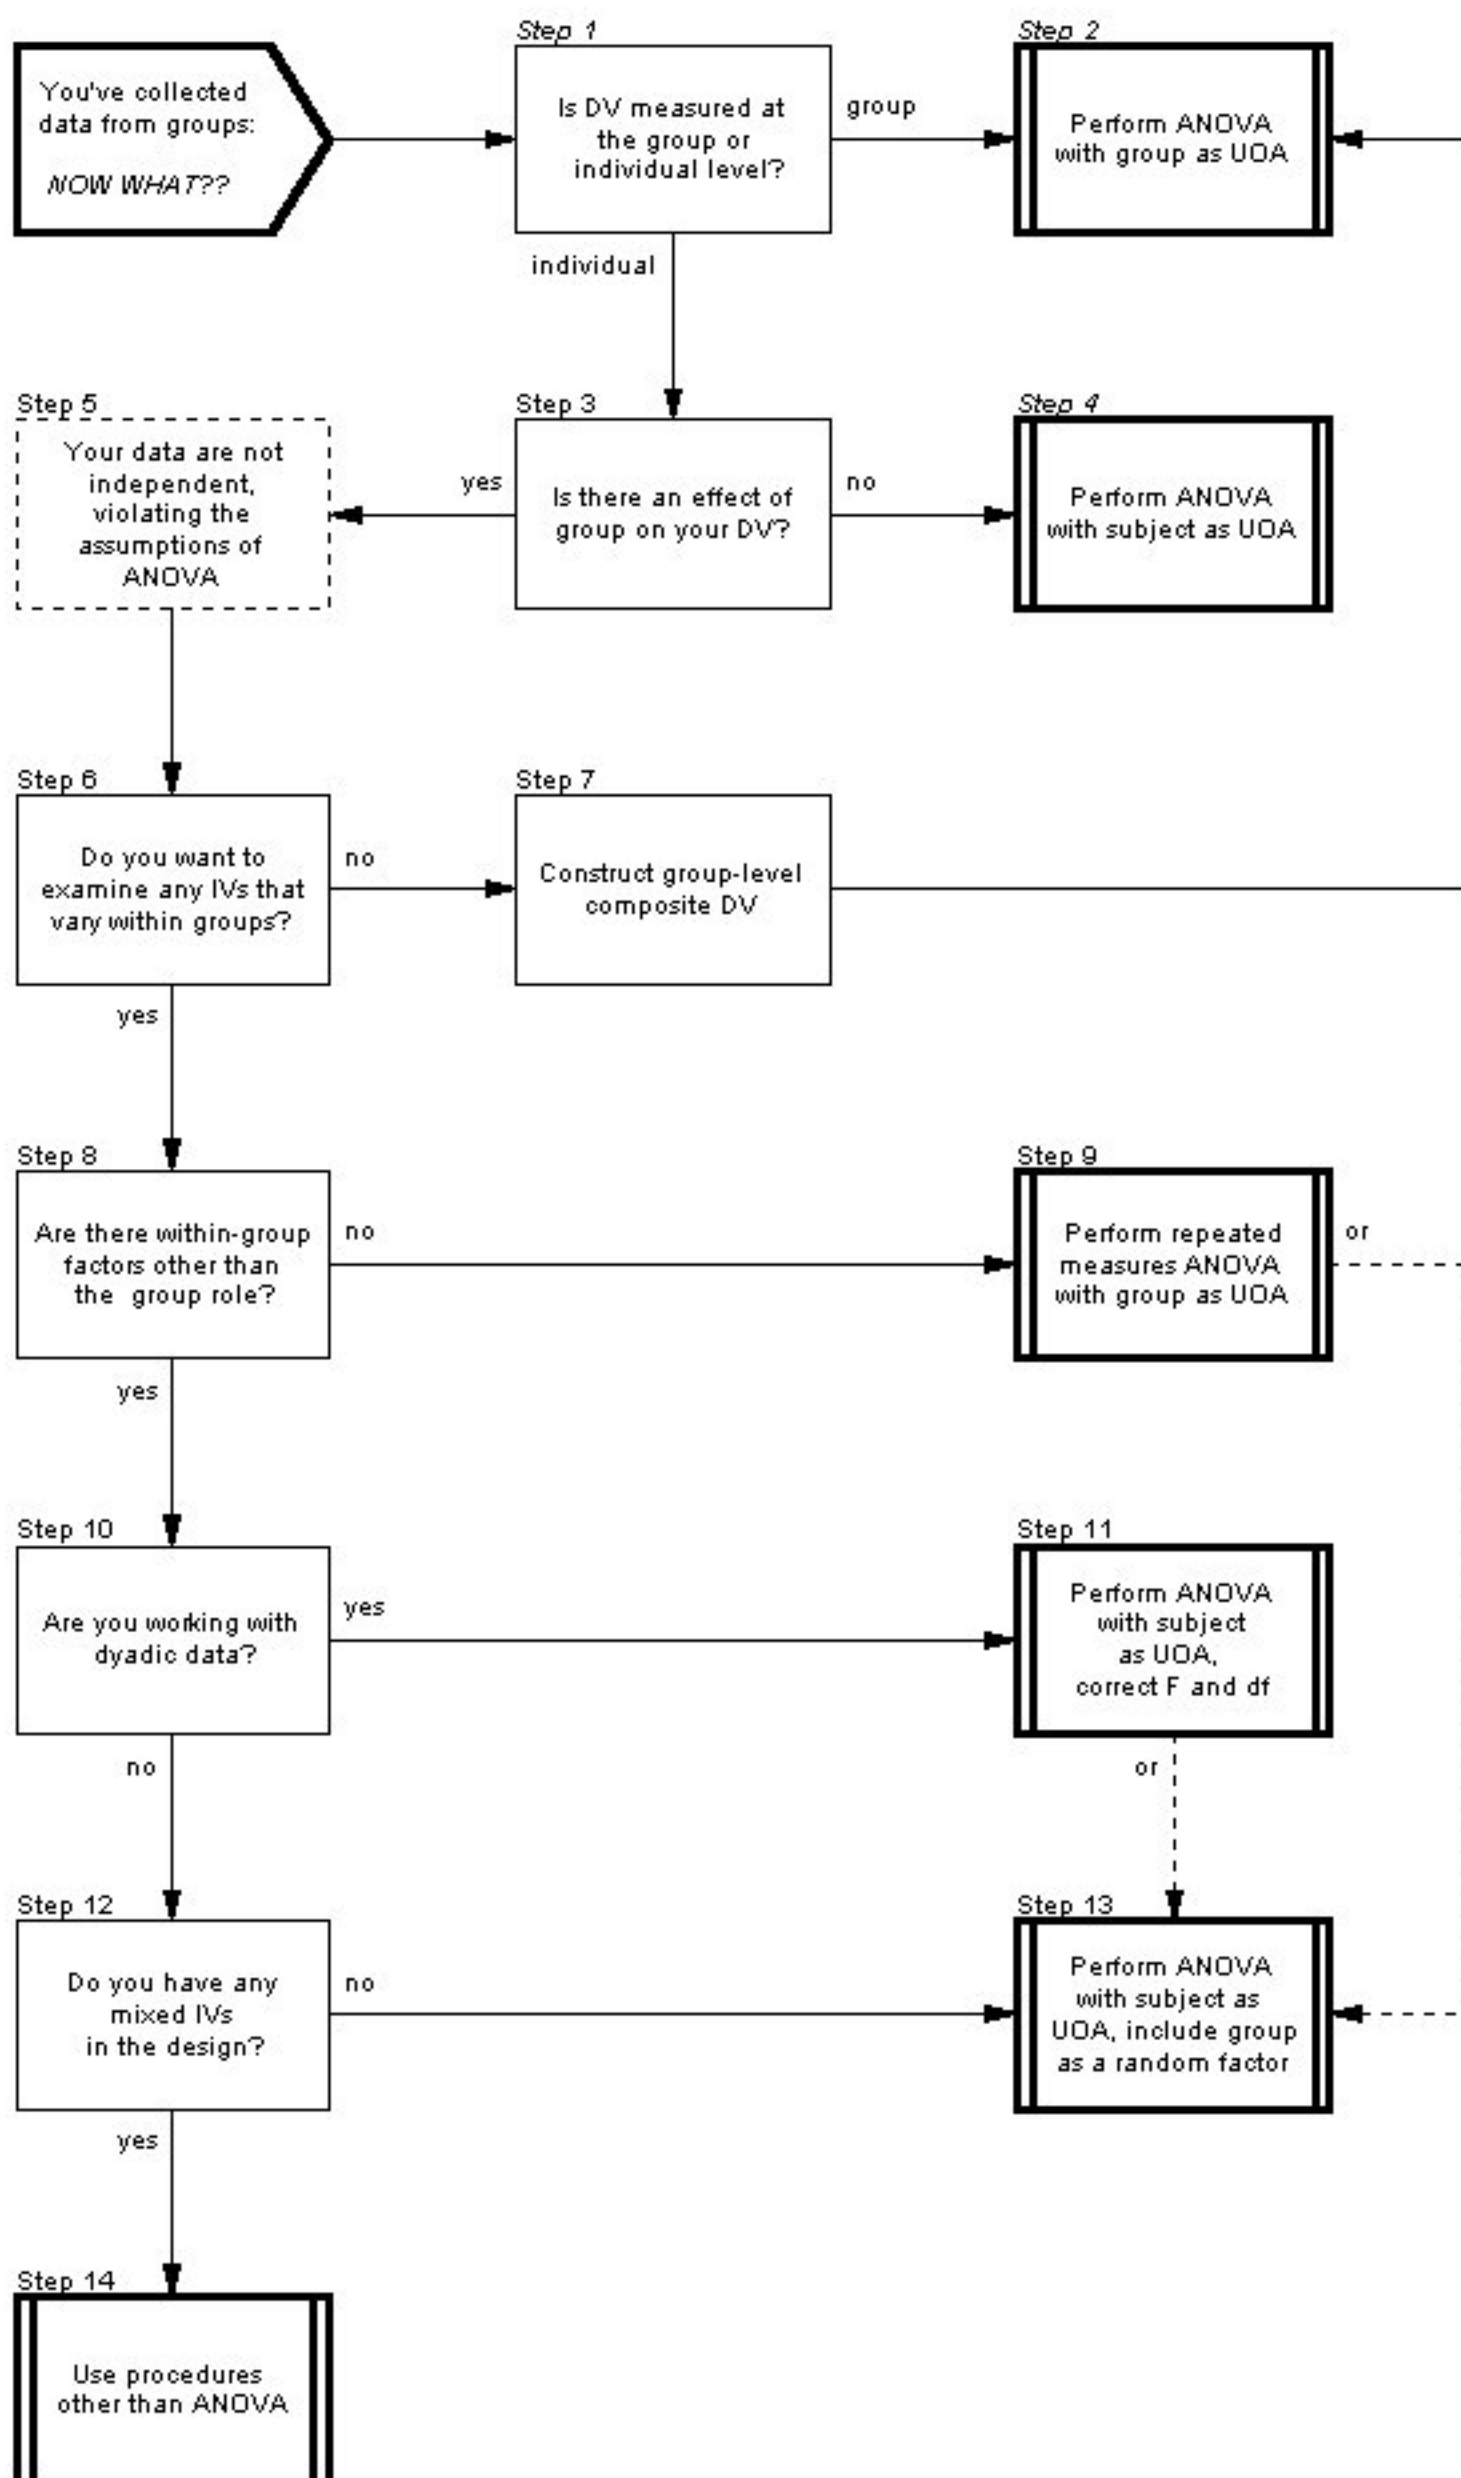
These notes were prepared with the support of a grant from the Dutch Science Foundation. If you wish to cite the contents of this document, the APA reference for them would be

DeCoster, J. (2002). *Using ANOVA to Examine Data from Groups and Dyads*. Retrieved <month, day, and year you downloaded this file> from <http://www.stat-help.com/notes.html>

For future versions of these notes or for help with data analysis visit
<http://www.stat-help.com>

ALL RIGHTS TO THIS DOCUMENT ARE RESERVED.

Flowchart for Analyzing Group Data



Introduction

- Many times researchers will perform experiments where participants must interact with each other in groups. Standard statistical procedures, however, assume that the individual data points from a study are independent. For analyses to be accurate they must properly account for the influence of the groups on the response measures.
- These notes are designed to provide step-by-step instructions on how to use analysis of variance (ANOVA) to examine data from interacting groups or dyads. The flowchart at the beginning of this document illustrates the steps researchers should take to determine the appropriate procedures that should be applied to their data. The following pages provide detailed information about how to perform each step.
- The design of the flowchart is fairly simple. Each box represents either decisions to be made or actions you should take. Starting at Step 1, follow the path of the arrows through the chart, taking each step you encounter along the way. When there are two arrows leading away from your current step you should follow the one whose label matches the characteristics of your data. Once you reach a bold box with double left and right edges you have determined what type of analysis you should perform. You should check the chart separately for each test that you want to conduct, considering the variables that are included in that particular analysis. If you perform multiple analyses on the same set of data it is quite possible that at least some of them will require different methods.
- Throughout the notes we will use the term “group” to refer to two or more interacting participants. Our use of the word “group” therefore includes what people usually refer to as “dyads.” In general, the same procedures are used whether a group has two, three, or more interacting participants. We will use the term “dyad” when specifically referring to a group of exactly two participants.
- ANOVA tests the ability of a set of predictor variables to explain the variability in a single response variable. The predictor variables are generally referred to as the *independent variables* (IVs) while the response is called the *dependent variable* (DV). These notes will not go into detail regarding how you should conduct your ANOVA - our main purpose is to help you determine what type of ANOVA you should perform.
- The object “type” that defines your cases or observations is called your *unit of analysis* (UOA). In our discussion we will be referring to data sets with the participant as the UOA and those with the group as the UOA. A data set with the participant as the UOA has a number of rows equal to the number of participants, and each row contains information about the conditions of and measurements made on a particular participant. Similarly, a data set with the group as the UOA has a number of rows equal to the number of groups, and each row contains information about the conditions of and measurements made on an entire group.

You can construct both types of data set from the results of an experiment on groups. Which one you should use depends on the type of analysis that you want to perform. Details on transforming data sets from one UOA to another can be found in our notes on Transforming and Restructuring Data (DeCoster, 2001).

- Groups can have two different types of members: *distinguishable* members and *exchangeable* members. If each member of the group fills a particular role then the members are said to be distinguishable. For members to be distinguishable the roles must be defined in such a way that every member of the group fits into one and only one of the possible roles. The roles must always be the same within each group, and each group must have someone filling each role. If the roles vary from group to group then your members are not distinguishable. An example of a group with distinguishable members would be a heterosexual dating couple. In each couple there will always be one male member and one female member.

If there are no theoretically important factors that can be used to define a set of roles that are consistent from group to group, then the members of your group are said to be exchangeable. In exchangeable groups your members are all theoretically equivalent and will generally be considered together during

analysis. An example of a group with exchangeable members would be a five-person workgroup where the workers all have the same rank and status. In this group there are no theoretically important criteria that can be used to distinguish one member from another. Sometimes you might be able to distinguish your group members using a theoretically irrelevant variable. If you don't expect that the role will have an impact on participants' performances then you should consider your members to be exchangeable.

- To determine what type of analysis you will need it is important to know how the other variables in your design relate to your group factor. If every group is only in a single condition of a manipulation, or if a single measure is taken from an entire group, the variable is said to vary *between groups*. Members of the same group always have the same value on the variable, so you will only observe between-group variability. For example, you might have a study where half of your groups are given 10 minutes to perform a task while the other half are given 30 minutes. You have two different conditions, but each group only sees one of them. When a factor varies between groups we often say that groups are *nested under the factor*. Conventionally, to indicate that one factor is nested under another we will put the higher level factor in parentheses after the factor which is nested under it, such as **group(group-level manipulation)**.

If a single group experiences more than a single level of a manipulation or measurements are made on individual members instead of on the group as a whole, the variable is said to vary *within groups*. In this case it is possible for members of the same group to have different values on the variable, so you will observe both within-group and between-group variability. If every group has an observation at every level of the within-group factor then we say that groups are *crossed with the factor*. Conventionally, to indicate that two factors are crossed with each other we display them together with an "X" in between them, such as **individual-level manipulation X group**. For example, you might conduct a study examining problem solving in dyads where one member of the dyad receives complete information while the other receives incomplete information. In this case the type of information received would be crossed with dyad, since you would have a measurement in both levels of the factor from every dyad.

You can also have a factor that varies within groups but where you don't always have observations at every level of the factor for each group. In this case the relationship between group and the factor is said to be *mixed*. For example, you might conduct a study where you randomly pair research participants to form interacting dyads. In this study participant gender would be a mixed factor since some dyads will contain members of both genders while others would only contain members of a single gender. The relationship between the factor and group is neither fully between or fully within, so we refer to this as a *mixed design*. Mixed variables can be spotted because they don't have a consistently crossed or consistently nested relationship with group. Group is always nested under any group-level factors, so you need only look at factors that concern individual-level characteristics. If you have an individual-level factor that has the same value for all members of a group then it is nested under group.

- As a general recommendation, you should always try to minimize the group-to-group variability that does not result from your manipulations. The power of your statistical analyses decreases as the amount of between-group variability increases.

Step 1: Determine the level of measurement for the DV

- The first thing you need to do is determine whether the DV is measured at the group level or at the individual level. Measurements made at the group level describe the characteristics or performance of the group as a whole, while those at the individual describe the characteristics and performance of group members.

This distinction is similar to, but not the same as, the difference between variables that vary within groups and those that vary between groups. If the DV varies between groups it will always be measured at the group level, and if it varies within groups it will *usually* be measured at the individual level. However, it is possible to have multiple measurements on a group that are not associated with the group members, such as the performance of the group at several different time periods. In this latter

case the DV is measured at the group level because the measurements concern the performance of the group as a whole, rather than the performance of individual group members.

- If the DV is measured at the group level you will need to analyze the data using group as the unit of analysis: **Go to Step 2**. If the DV is measured at the individual level you will need to consider whether there is reason to suspect within-group dependence: **Go to Step 3**.

Step 2: Perform an ANOVA with group as the UOA

- To analyze your data with group as the UOA you must make a data set where each row represents a group and each column represents a variable measuring some characteristic of the groups. A particular cell in the data set contains the value of the column variable for the group represented by the row.
- If your data set was originally organized using individuals as the UOA you may need to perform some transformations to obtain a data set with group as the UOA.
 - You can directly assign the level of a group-level measurement or manipulation to a variable in the data set. The value of this variable should be the same for all members of the same group, so the value for the group data set may be obtained from any participant in the group.
 - Measurements made on exchangeable group members must be somehow combined to form a single composite variable. Since exchangeable members are not uniquely identifiable you must always consider them together. Your composite variable will usually be an average of the scores, but sometimes you might want to use a more complicated function, such as the minimum or maximum value.
 - Measurements made on distinguishable group members may be included in the data set with the value of each member stored in a separate variable. You may also choose to combine them to form a single composite variable as with exchangeable members.
- Once you have your data set you may perform your ANOVA using standard procedures. In fact, you can perform any statistical test you would usually use on independent data, such as correlation and regression. The only difference is that your conclusions tell you about the relationship of your variables within and between groups instead of within and between individuals.

If you have measurements made on your group under multiple conditions or at multiple times you can analyze it using multivariate statistics. You would perform these analyses just like you would on individual-level data.

Step 3: Determine if there is within-group dependence

- If your participants were allowed to have any significant interactions with each other then there is a chance that those participating together might influence each others' responses. To test this we look for an influence of group on the DV.

Just because you have multiple participants in an experimental session doesn't necessarily mean that you need to test for a group effect. Conventionally we are allowed to assume that there will be no influence of group if the experimental environment is the same for all participants and the procedure prevents any meaningful interactions between participants.

- Myers (1979) and other researchers recommend that you use the relatively liberal critical value of $\alpha = .20$ when checking for within-group dependence because of the low power of these tests. Even if your tests for dependence produce nonsignificant results you might consider controlling for the effect of group anyway if there is a logical reason to suspect that observations from the same group might be related.

- If your group contains exchangeable members then you can determine the effect of group by calculating the *intraclass correlation* (ICC). The ICC can be interpreted as the typical correlation found between the responses of members of the same group. A positive ICC indicates that when one member of the group has a high score then other group members are also likely to have high scores. A negative ICC indicates that when one member of the group has a high score then other members are likely to have low scores. An ICC near zero indicates that there is no linear relationship between the scores of members of the same group. The absolute value of the ICC is the proportion of variability in the responses that can be explained by group membership.
- The ICC can be calculated using both SPSS and SAS. In both cases you will need a data set with group as the UOA. The data set should have variables representing the group number, the conditions of any between-group manipulations, as well as a variable containing each group member’s score on the DV. In SPSS versions 8.0 and higher, you can obtain the ICC by selecting **Analyze** → **Scale** → **Reliability**. At this menu you should click the “Statistics” button and then check the “Intraclass correlation coefficient” checkbox. Using the dialog box you should tell SPSS to calculate the ICC using a one-way random effects model. Nichols (1998) has produced a web page explaining the differences between the various types of ICCs available at <http://www.ats.ucla.edu/stat/spss/library/whichicc.htm>
While SAS does not provide any procedure that can calculate the ICC, the SAS Institute (2000) has provided a SAS macro that will do so at <http://ewe3.sas.com/techsup/download/stat/intracc.html>
If you use this macro you will want to use ICC(1,1) to test for an effect of group.
- If you need to calculate the ICC by hand, Griffin and Gonzales (1995) provide a simple method of calculating the ICC when your group is a dyad. To use this method you must first create a special data set. This data set should have two columns, one to represent the score for each member. The trick is that you will create two cases in the special data set for each dyad in your study. In the first case the score from the first member (determined arbitrarily) goes into the first column while the score from the second individual goes into the second column. In the second case the score from the first member goes into the *second* column while the score from the second member goes into the *first* column. For example, let us say you collected the data present in Table 1.

Table 1: ICC Example Actual Data

Dyad	Score 1	Score 2
1	7	5
2	4	3
3	2	3

For this method you would need to create the special data set shown in Table 2.

This can be done quite easily using a spreadsheet. All you need to do is copy the Score 1 column and paste it onto the bottom of Score 2, and copy the Score 2 column (excluding the added values) and paste it onto the bottom of Score 1.

Once you create the special data set you can determine the ICC simply by taking the standard Pearson correlation between the two columns holding the responses. However, the p-value associated with this correlation will be incorrect because it assumed you had twice as many dyads as you actually did. You will therefore need to calculate the p-value by hand. Using Fisher’s r-to-Z transformation (Fisher, 1928) you can calculate a test statistic Z^* , where

$$Z^* = \frac{Z_r}{s\{Z_r\}} = \frac{\frac{1}{2} \ln \left[\frac{1+r}{1-r} \right]}{\frac{1}{\sqrt{n-3}}} = \ln \left[\frac{1+r}{1-r} \right] \left(\frac{\sqrt{n-3}}{2} \right). \quad (1)$$

Table 2: ICC Example Special Data Set

Dyad	Score 1	Score 2
1	7	5
2	4	3
3	2	3
1	5	7
2	3	4
3	3	2

In this formula, r is your ICC and n is the number of dyads (rather than the number of cases in the data set, which is $2n$). The p-value for Z^* in the standard normal distribution will be the p-value for your ICC. This formula actually works for any type of correlation, not just the ICC.

- If your group has more than two members, Kashy and Kenny (2000) present a way to determine the ICC from the contents of an ANOVA table where you analyze the individual DV scores by group. To perform this analysis you need a data set with participant as the UOA including variables representing the group number and the DV. The generic form of the table is presented in Table 3.

Table 3: ANOVA table used to calculate ICC

Source of variation	SS	df	MS	F
Group	SS_G	$g - 1$	$MS_G = \frac{SS_G}{g-1}$	$\frac{MS_G}{MS_E}$
Error	SS_E	$g(n - 1)$	$MS_E = \frac{SS_E}{g(n-1)}$	

In the table g refers to the number of groups and n refers to the number of participants in each group. Be sure to code group as a categorical variable - if the effect of group only has 1 df in your table then you have accidentally coded it as numeric. If the F statistic for the effect of group is ≥ 1 then the ICC will be positive and can be calculated using the formula

$$\hat{r} = \frac{MS_G - MS_E}{MS_G + (n - 1)MS_E}. \quad (2)$$

In this case the F test for the group effect in the ANOVA table can be used as a test of the intraclass correlation.

If the F statistic for the effect of group is less than one then the ICC will be negative and can be calculated using the formula

$$\hat{r} = -\frac{MS_G + (n - 1)MS_E}{MS_G - MS_E}. \quad (3)$$

In this case the F for group is *not* a test of the ICC. You will need to determine the p-value for your correlation directly by finding the p-value associated with the Z^* statistic calculated using formula 1.

- If you have manipulated factors in the design then they will be confounded with your group factor, so you will need to remove their influence from your estimate of the ICC. The manipulated factors can either inflate or decrease the ICC, depending on what type of manipulation it is.

- If you have a within-groups manipulation then differences associated with the manipulation increases the variability within your groups, artificially decreasing the ICC. You can compute a version of the ICC that controls for the manipulation using information from the ANOVA table where you analyze the individual DV scores by the manipulation, group, and the manipulation X group interaction. The generic form of the table is presented in Table 4.

Table 4: ANOVA table used to calculate ICC when there is a within-group manipulation

Source of variation	SS	df	MS	F
Manipulation	SS_M	$m - 1$	$MS_M = \frac{SS_M}{m-1}$	$\frac{MS_M}{MS_{G(M)}}$
Group	SS_G	$g - 1$	$MS_G = \frac{SS_G}{g-1}$	$\frac{MS_G}{MS_E}$
Manipulation X Group	SS_{MG}	$(m - 1)(g - 1)$	$\frac{SS_{MG}}{(m-1)(g-1)}$	$\frac{MS_{MG}}{MS_E}$
Error	SS_E	$mg(n - 1)$	$MS_E = \frac{SS_E}{mg(n-1)}$	

In the table, m refers to the number of levels in the manipulated factor, g refers to the total number of groups, and n refers to the number of participants in each group. If the F statistic for group is ≥ 1 then the ICC corrected for the manipulation can be obtained using Formula 2, and the F test for group in the table can be used as a test of this correlation. Otherwise you will need to use calculate the ICC using formula 3 and determine its p-value from the Z^* statistic calculated using formula 1.

- If you have a between-groups manipulation then differences associated with the manipulation increases the variability between your groups, artificially inflating the ICC. You can compute a version of the ICC that controls for the manipulation using information from the ANOVA table where you analyze the individual DV scores by the manipulation and group. The generic form of the table is presented in Table 5. Note that in this analysis, group is nested under the manipulated factor.

Table 5: ANOVA table used to calculate ICC when there is a between-group manipulation

Source of variation	SS	df	MS	F
Manipulation	SS_M	$m - 1$	$MS_M = \frac{SS_M}{m-1}$	$\frac{MS_M}{MS_{G(M)}}$
Group(Manipulation)	$SS_{G(M)}$	$m(g - 1)$	$MS_{G(M)} = \frac{SS_{G(M)}}{m(g-1)}$	$\frac{MS_{G(M)}}{MS_E}$
Error	SS_E	$mg(n - 1)$	$MS_E = \frac{SS_E}{mg(n-1)}$	

In the table, m refers to the number of levels in the manipulated factor, g refers to the number of groups in each level of the factor (rather than the total number of groups), and n refers to the number of participants in each group. If the F statistic for group is ≥ 1 then the ICC corrected for the manipulation can be obtained using formula 2, and the F test for group in the table can be used as a test of this correlation. Otherwise you will need to use calculate the ICC using formula 3 and determine its p-value from the Z^* statistic calculated using formula 1.

If you don't have the software to perform an ANOVA with a nested factor then you can create the nested ANOVA table by combining the results from a crossed ANOVA table. To do this you must first create a data set with participant as the unit of analysis, and you must have variables for the manipulation, group, and your DV. The group variable must be coded so that you use the same numbers within each level of your between-group manipulation. So if you had 10 groups in each level of a 4-level manipulation, you should have the groups numbered 1-10 within each level. This would mean that you would have four labeled as group 1, four labeled as group 2, and so forth. You then should perform a standard ANOVA predicting your response from the manipulation, the group, and the manipulation X group interaction.

You can calculate the results of the nested ANOVA table you want from the results of the crossed ANOVA table. Abstractly what need to do is take the values for **manipulation** and **error** for the nested table directly from the crossed table, and pool together the effects of **group** and **group X manipulation** to obtain **group(manipulation)**. Practically, this means that you must add together the sum of squares and the df for **group** and **group X manipulation** in the crossed ANOVA table to obtain the sum of squares and df for **group(manipulation)** in the nested table. The mean squares and F tests can all be computed manually according to the formulas in Table 5 once you obtain the sums of squares and df.

The procedure is summarized in table 6.

Table 6: Constructing a nested ANOVA table from a crossed ANOVA table

Nested factor	Nested table entry	Computed from crossed table values as
Manipulation	SS_M df	SS_M $m - 1$
Group(Manipulation)	$SS_{G(M)}$ df	$SS_G + SS_{MG}$ $(g - 1) + (m - 1)(g - 1) = m(g - 1)$
Error	SS_E df	SS_E $mg(n - 1)$

where m refers to the number of levels in the manipulated factor, g refers to the total number of groups, and n refers to the number of participants in each group.

- If the members of your group are distinguishable you can test for within-group dependence by determining whether the responses of group members can be successfully used to predict each other. If your group only has two members you can calculate the Pearson correlation between their responses, and use this as a test for the effect of group.

If your group has more than two members you can use the ICC to estimate the group effect, if you can assume that the effect of group is the same for all of your group members. If you expect the effect of group to vary depending on the role of the participant you will instead need to examine a set of multiple regression models. To determine if there is any dependence you will need to test a total of (number of group members - 1) independent models. For example, if your group had three distinguishable members A, B, and C, you would need to examine the ability of A and B to predict C, and the ability of A and C to predict B. You do not need to test the final model (whether B and

C can predict A) because the relationships present in that model had already been tested in the other two. In each of the models you want to consider whether the model as a whole can predict a significant proportion of the variance in the target variable. To do this you can perform an F Test for Regression Relation (Neter, Kutner, Nachtsheim, & Wasserman, 1996, pp. 240-241), calculating the test statistic

$$F^* = \frac{MS_R}{MS_E}. \quad (4)$$

This statistic follows an F distribution with $(p - 1)$ numerator and $(n - p)$ denominator degrees of freedom, where p is the number of predictor variables in the model and n is the total number of participants.

When performing these tests you should apply a Bonferroni correction to your critical value to maintain a constant overall experimental error rate. To apply the correction you simply divide your critical value α by the total number of models you are testing. If the p-value associated with any of the tests is below this revised value then you must conclude that there is significant dependence between measurements made on members of the same group.

- If there is no evidence of within-group dependence you can ignore the effect of group: **Go to Step 4**. If you have evidence of dependence then you will need to take the effect of group into account in your analyses: **Go to Step 5**.

Step 4: Perform an ANOVA with participant as the UOA

- If there is no evidence of within-group dependence you can ignore the effect of group in your analysis. In this case you would treat the observations from the same group as independent and perform a standard ANOVA using participant as the unit of analysis.

Step 5: Effects of nonindependence

- Before going into detail about how to correct for the effect of group in your analysis it is worthwhile to spend a little time discussing the effect within-group dependence can have on your statistics if it is not taken into account.
- ANOVA, regression, and many other commonly used statistics are based on the *General Linear Model* (GLM). The GLM allows you to generate an equation predicting the value of a DV from one or more IVs. It also allows you to test the ability of individual predictors to account for variability in the DV. The validity of these tests, however, rests on a set of assumptions made by the GLM. These assumptions are that
 1. The error terms are normally distributed.
 2. The distribution of the error terms is the same at all of the levels of the IVs.
 3. The error terms are independent.

If your analysis violates these assumptions then the results that you receive will not be accurate. It is worth noting that all of these assumptions are made on the error terms from the model, rather than on the DVs themselves. We therefore don't need to worry about the influence of a variable if we account for it in our predictive model. This makes perfect sense, since any time an IV can account for a significant amount of variance in a DV it will change distribution of the DV and introduce dependence in the observations (those with similar values on the IV can be expected to have similar values on the DV).

- Statisticians have demonstrated that the GLM is relatively robust to violations of the assumptions of normality and homogeneity of variance (Boneau, 1960). However, the model is not robust to violations of the assumption of independence (Kenny & Judd, 1986). Dependence will create substantial biases in your results if there is a tendency for observations in the same condition to be more or less dependent than those in different conditions.

- Unaccounted within-group dependence will bias tests of both between-group and within-group factors, although the influence differs. Positive correlations among group members will artificially inflate statistics testing between-group factors while reducing the statistics testing within-group factors. Negative correlations among group members will artificially reduce statistics testing between-group factors while inflating statistics testing within-group factors. Dependence will also bias the statistics of mixed factors, though the exact effect will depend both on the correlation between the factor and group as well as the correlations among group members. For a more mathematical treatment of the effect of nonindependence see Kenny and Judd (1986).
- Now that you are aware of why you need to account for the effect of group in your analysis you must next determine how this should be done: **Go to Step 6.**

Step 6: Determine if you want to examine IVs that vary within groups

- At this step you need to decide if you are actually interested in examining individual-level differences on your response variable. Sometimes you may be interested in the effect of a group-level manipulation on an individual-level measurement. Other times you may collect information about individuals but may only be interested in their influence as a group-level aggregate.

It should also be noted at this step that it is typically easier to analyze your data if you only consider IVs that vary between groups. Once you want to examine within-group variability you will need to account for the fact that your participants are nested under groups, which leads to more complex analyses.

- If you only want to examine the effect of IVs that vary between groups then you will want to combine the individual response measures to form a group-level composite DV: **Go to Step 7.** If you want to examine any IVs that vary within groups you should **Go to Step 8.**

Step 7: Construct a group-level composite DV

- If you only want to examine IVs that are manipulated between groups then you will be able to perform your analyses using a data set using the group as the UOA. However, in order to do that you will need to have a DV that is measured at the group level.
- You will likely create your group-level DV from a mathematical function of the individual responses. Most commonly this is done by averaging the responses together, though it is also possible to use more complex functions such as the minimum, maximum, or median.

If your group members are distinguishable you may also consider only setting the group score equal to the response of a particular member, or calculate the group score only using a subset of the members. It is not appropriate to do this with exchangeable members since there is no way to justify using one particular member's response instead of that from another member.

- Once you have created the group-level DV you will then need to perform an ANOVA on a data set with group as the UOA: **Go to Step 2.**

Step 8: Determine if the members in your groups are distinguishable

- If your group has distinguishable members you may have the option of performing a repeated-measures ANOVA at the group level: **Go to Step 9.** If your group has exchangeable members **Go to Step 11.**

Step 9: Determine if group role is your only within-group IV

- Possibly the most common within-group factor you will see is the role that a particular member takes in the group. When this is the only within-group factor in your study then you can analyze your experiment using a repeated measures ANOVA. You can include additional between-group factors if you like, and you can test for interactions between these and your group role factor. However, this analysis will not allow you to include factors based on other participant-level measurements.
- If your group contains distinguishable members and the only within-group factor you wish to investigate is the group role then you can perform a repeated measures ANOVA with the responses of the different members as your repeated factor: **Go to Step 10**. If you wish to examine any other within-group factors (crossed or mixed) you will need to perform a standard ANOVA and account for the fact that your participants are nested under groups: **Go to Step 11**.

Step 10: Perform a repeated measures ANOVA with group as the UOA

- We know from basic statistics that if we make several different measurements on the same participant under different conditions then we can look for an effect of the conditions using a repeated measures analysis. This procedure can also be applied to groups. Lets consider a data set where group is the UOA. If we recorded several different measurements on the same group under different conditions, we should be able to look for an effect of those conditions using repeated measures. In this framework we can consider measuring different individuals from the same group to be repeated measurements of the group itself. If we code the responses we measure on different members as different variables in the group-level data set then we can look for an effect of role using repeated measures analysis.
- To perform a repeated measures ANOVA you will first need to create a data set that uses group as the UOA. For each between-group IV in the design you should have a variable in the data set recording information about the level of that factor for each of the groups. Additionally, you will need a number of variables equal to the number of group members to hold information about the DV. Each variable will correspond to a particular role, and will hold the measurement made on the participant in each group that was assigned that role.

For example, let us say that a researcher had groups of three people working on a decision-making task. In the study the three people received different types of information regarding the decision. One received information in favor of decision A, one received information in favor of decision B, and the third received both sets of information. After the discussion the researchers measured the private opinions of each of the three members. In the data set you would need to have a total of three variables to record the DV. The first variable would hold the opinion of the person who received information from set A, the second would hold the opinion of the person who received information from set B, and the third would hold the opinion of the person who received information from both sets. If there were any between-group factors in the design each of those would be included in this data set as a single variable indicating the level of the factor the group was in.

- Once you have created your data set you can perform your repeated measures analysis. You should specify the scores from your different group members as the different levels of your repeated factor. You can then add any between-group factors to your model. This analysis can be performed in SPSS using the **Analyze** → **General Linear Model** → **GLM - Repeated Measures** menu option. It can also be performed in SAS using the **repeated** subcommand under **proc glm**.
- Like any repeated measures analysis, you have the option of performing your analysis as a standard ANOVA where you treat your repeated measurements as multiple observations that are nested under your group factor. If you wish to do this then you can skip this step and instead **Go to Step 14**. However, this is typically leads to much more difficult analyses so you are usually best off working with the data as repeated measures when it is an option.

Step 11: Determine if you are working with dyadic data

- If your group only has two members then you have the option of applying a set of corrections to your statistics instead of accounting for the influence of group within the ANOVA itself. If you have dyadic data and you want to use these corrections you should **Go to Step 12**. Otherwise you should **Go to Step 13**.

Step 12: Correct the statistics from an ANOVA ignoring the effect of group

- If you are working with groups containing only two members then you can first perform an ANOVA ignoring the effect of group and then later correct the test statistics for the influence of dependence (Kenny, 1995). This is generally easier than directly accounting for the influence of group in your analyses.
- For this analysis you will first need to create a data set with the participant as the unit of analysis. Your data set should have variables for each of the within-group and between-group IVs, as well as a variable for your DV.
- Once you have your data set you should simply conduct a standard ANOVA, ignoring the fact that the participants were in dyads. However, both the F statistics and the denominator degrees of freedom from this analysis will be need corrections for them to be valid. These corrections require you to first calculate the ICC for your DV (see Step 3 for details). After correction you should look up the final p-value in a standard F table using the revised F and degrees of freedom.
- For between-dyad factors and interactions composed strictly of between-dyad factors, you should calculate a revised F statistic and df using the formulas

$$F' = F \left(\frac{1 - \frac{2r_y}{n-2}}{1 + r_y} \right) \quad (5)$$

and

$$df' = \frac{(n - 2 - 2r_y)^2}{n(1 - r_y^2) - 2(1 + r_y)^2}, \quad (6)$$

where F is the F statistic reported by the ANOVA, r_y is the ICC, and n is the number of participants (twice the number of dyads) in your study.

- For within-dyad factors and interactions including between-dyad and within-dyad factors, you should calculate a revised F statistic and df using the formulas

$$F' = \frac{F}{1 - r_y} \quad (7)$$

and

$$df' = \frac{n - 2}{1 + r_y^2} \quad (8)$$

where again F is the F statistic reported by the ANOVA, r_y is the ICC, and n is the number of participants in your study.

- For mixed factors and interactions including at least one mixed factor, the correction can be performed but is more complicated. The specifics, as well as a sample computer program that will provide the correction, can be found in Appendix B of Kenny (1995).
- These corrections are an alternative to performing an ANOVA including dyad as a nested factor in the design. If you would rather obtain your statistics directly you can **Go to Step 14** and instead perform an ANOVA including the effect of dyad in the analysis.

Step 13: Determine if you have any mixed IVs in your design

- Standard ANOVA procedures are only designed to handle designs that included either crossed or nested variables. If you have variables in your design have mixed relationships with each other you will need to use special procedures to analyze the data.
- You will observe mixed relationships most commonly when the levels of a factor are determined by individual difference measures. Since you don't have control over the characteristics of your participants if they are randomly assigned to your groups, the exact distribution of this variable will vary from group to group, so the factor will not have a consistently crossed or nested relationship with group.
- If you your design only includes factors that are either crossed or nested under group then you can perform a standard ANOVA with participant as the UOA, including group as a factor in the design: **Go to Step 14**. If you have factors that have a mixed relationship with group you will not be able to analyze your data using ANOVA: **Go to Step 15**.

Step 14: Perform an ANOVA with participant as the UOA, including group as a factor in your design

- For this analysis you must create a data set using the participant as the UOA, containing a variable for the group number, a variable for each of the within-group and between-group IVs, as well as a variable for your DV. If your group contains distinguishable members and you wish to test for an effect of group role you should simply code this as a within-group IV.
- Once you have prepared this data set then you need to determine your ANOVA model. The inclusion of group in your design will make this slightly more difficult. First off, group must be declared as being nested under your between-group factors. Secondly, group itself must be treated as a random factor in the design.

After specifying the design for your analysis you may also need to construct an estimated mean squares (EMS) table so that you know the appropriate error terms for your F tests. Sometimes your statistical software will choose the appropriate error term automatically, but sometimes you will need to specify it. Most notably, SAS's **proc glm** by default tests all of your factors against the overall error term, which would be inappropriate for your between-group factors. For information on how to create an EMS table and how to use it to determine the appropriate error term see Montgomery (1997, pp. 480-485).

- This analysis will be complex enough that you will definitely want to use specialized statistical software such as SAS or SPSS. In SAS you will need to use either **proc glm** or **proc mixed**. In SPSS you will need to use the **glm** procedure within syntax. You cannot perform an ANOVA with nested factors using the SPSS windows interface.

Step 15: Use procedures other than ANOVA

- If you are working with groups of three or more members you cannot use ANOVA to analyze the influence of mixed variables on your DV. In this circumstance you may need to perform Hierarchical Linear Modeling (Raudenbush & Bryk, 2002). This is a procedure designed to analyze data including groups and subgroups, and allows much more flexibility in design than ANOVA. In addition to allowing the inclusion of mixed factors, it also allows your groups to each have different numbers of members. You can download a free version of the program HLM, the most commonly used software for hierarchical linear modeling, at the website

<http://www.ssicentral.com/hlm/hlm.htm>

References

- Boneau, C. A. (1960). The effects of violations of assumptions underlying the t test. *Psychological Bulletin*, 57, 49-64.
- DeCoster, J. (2001). Transforming and restructuring data. Retrieved from the web March 25, 2002. <http://www.stat-help.com/notes.html>
- Fisher, R. A. (1928). *Statistical Methods for Research Workers* (2nd ed.). London: Oliver & Boyd.
- Griffin, D., & Gonzales, R. (1995). Correlational analysis of dyad-level data in the exchangeable case. *Psychological Bulletin*, 118, 430-439.
- Kashy, D. A., & Kenny, D. A. (2000). The analysis of data from dyads and groups. In H. T. Reis & C. M. Judd (Eds.), *Handbook of Research Methods in Social and Personality Psychology* (pp. 451-477). London: Cambridge University Press.
- Kenny, D. A. (1995). The effect of nonindependence on significance testing in dyadic research. *Personal Relationships*, 2, 67-75.
- Kenny, D. A., & Judd, C. M. (1986). Consequences of violating the independence assumption in analysis of variance. *Psychological Bulletin*, 99, 422-431.
- Montgomery, D. C. (1997). *Design and Analysis of Experiments*. New York: Wiley and Sons.
- Myers, J. L. (1979). *Fundamentals of Experimental Design* (3rd ed.). Boston: Allen & Bacon.
- Neter, J., Kutner, M. H., Nachtsheim, C. J., & Wasserman, W. (1996). *Applied Linear Statistical Models* (4th ed.). Chicago: Irwin.
- Nichols, D. P. (1998). Choosing an intraclass correlation coefficient. Retrieved from the web March 17, 2002. <http://www.ats.ucla.edu/stat/spss/library/whichicc.htm>
- Raudenbush, S. W., & Bryk, A. S. (2002). *Hierarchical Linear Models* (2nd ed.). Thousand Oaks: Sage.
- SAS Institute. (2000). Compute six intraclass correlation measures. Retrieved from the web March 26, 2002. <http://ewe3.sas.com/techsup/download/stat/intracc.html>