The Center for Research and Outreach
REACH
Leveraging Evidence and Science into Practice

UNDERSTANDING ANOVA

A HOW-TO GUIDE

USDA

UNIVERSITY OF MINNESOTA
Driven to Discover℠

Developed in collaboration with the Department of Defense's Office of Family Policy, the National Institute of Food and Agriculture, and the U.S. Department of Agriculture under Award No. 2009-48667-05833

Understanding ANOVA: A How-To Guide

Do children of different ethnicities attend youth programs at different rates? Do these differences depend on the gender of the youth? Could an intervention aimed at increasing youth participation in programs be effective? Does the effectiveness of the intervention depend on the gender of the youth? These questions and more can be answered (at least in part) by using Analysis of Variance (ANOVA). In this How-To Guide we will discuss the uses of ANOVA to answer such questions where differences are explored between three or more groups of individuals. This guide is intended for individuals who are unfamiliar with ANOVA and serves as a basic guide for the contexts in which ANOVA is typically used.

Analysis of Variance, more commonly referred to as ANOVA, is similar to *t*-tests (see the *Understanding t-Tests: A How-To Guide* for more information on *t*-tests). *t*-tests are used when you want to compare two means, ANOVA is used when you want to compare more than two means. We begin with asking the question: "Why can't we just do multiple *t*-tests between three or four groups?"

Let's say we have three groups and we want to compare the mean of each group to each other.

- Group 1 compared to Group 2
- Group 2 compared to Group 3
- Group 1 compared to Group 3

One option might be to conduct multiple *t*-tests, in this case three separate tests. The problem with conducting these *t*-tests has to do with the rate of error. For each *t*-test conducted there would be a certain degree of error associated with each test. In other words, each test has a probability of being wrong. The probability of being wrong is typically small, but when multiple tests are conducted the chances that at least one of the tests is wrong increases. To solve this problem of increasing error, we use ANOVA to make comparisons of means between multiple groups instead of lots of *t*-tests.

ANOVA, like a *t*-test, tells us if the means of different groups are the same or different, and the *p*-value associated with ANOVA tells us if in fact the groups are different, if that difference is *statistically significant* (see the How-To Guide on *Statistical Language* for more information on *p*-values and significance).

There are a few types of ANOVA including:

- One-Way ANOVA
- Factorial ANOVA
- Repeated Measures ANOVA

Below we discuss each of these different types of ANOVA and the contexts in which they might be used.

One-Way ANOVA

Use

The One-Way ANOVA is used when comparing the means from three or more groups. Specifically, a One-Way ANOVA is used when there is a single independent variable that has three or more categories. The One-Way ANOVA tells us if the three groups differ from one another on a dependent variable.

Example

Imagine a study in which researchers implemented two separate interventions, one designed to improve skill building among youth (Intervention A) and another designed to increase supportive relationships between instructors and youth (Intervention B). In addition, the researchers also employed a *control group*[1] in their study.

The researchers want to know if either of these two interventions improved instruction quality. In this study there are three groups, participants who received Intervention A, those that received Intervention B, and the control group. Below is a graph of what the means for each group might look like. Researchers interested in differences between these three groups in terms of instruction quality would implement a One-Way ANOVA. The One-Way ANOVA provides information about if there were statistically significant instruction quality differences between these three groups

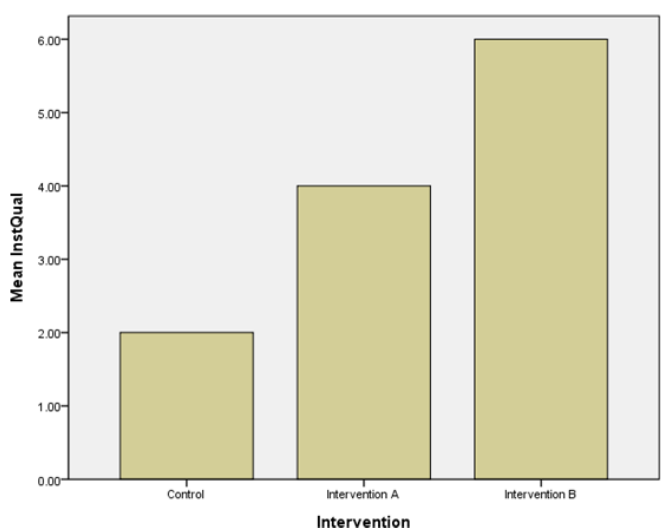

Interpretation

The result of a One-Way ANOVA indicates that there are differences between the three means. However, ANOVA on its own does not provide information about where these differences actually are. In this example there could be a difference in instructional quality between the control and Intervention A, between the control and Intervention B, and/or between Intervention A and Intervention B. To get at these differences additional analyses must be conducted. See the section below on *Post-Hoc Analyses and Planned Contrasts* for more information.

[1] A *control group* is used in experimental researcher designs. In such designs, typically one group of individuals is given some sort of treatment or intervention while the other group is not (or is given something similar but innocuous). The control group is the group that is not given the treatment or intervention. In well-designed experiments, difference between these groups would reflect the effect of the treatment or intervention.

Factorial ANOVA

Use

Unlike a One-Way ANOVA, a Factorial ANOVA is used when there is **more than one** independent variable. In the previous example there was only one independent variable with three levels (Intervention A, Intervention B, control group). Now suppose that a researcher also wanted to know if there were additional group differences between boys and girls in the youth program. When this second independent variable is added to the analysis, a Factorial ANOVA must be used.

Example

Factorial ANOVA's typically use a mathematical notation to indicate the kind of Factorial ANOVA being conducted. In the present example 3 x 2 Factorial ANOVA is being conducted. Below is an example of how to break down this notation:

3			x	2	
Intervention A	Intervention B	Control Group	By	Boy	Girl

In this example there are two independent variables, one with three levels (Intervention A, Intervention B, control group) and one with two levels (boy, girl). Thus, the analysis would be a 3 x 2 Factorial ANOVA.

The results of a Factorial ANOVA consist of several parts. First are called *main effects*. Main effects tell you if there is a difference between groups for each of the independent variables. For example, a main effect of intervention type (intervention A, intervention B, control group) would indicate that there is a significant difference between these three groups. That is, somewhere among these three means, at least two of them are significantly different from one another. Like the One-Way ANOVA, a Factorial ANOVA does not provide information on where these differences are and additional analyses are required, these additional analyses are discussed below (*Post-Hoc Analyses and Planned Contrasts*)

A *main effect* of gender would indicate that there is a difference between the means of boys and girls.

In addition to the main effects, Factorial ANOVA also often provides information about *interaction effects* as well. Interaction effects provide information about whether an observed group difference in one independent variable varies as a function of another independent variable.

An example will clarify:

Suppose that a researcher found a *main effect* for intervention type. This tells us that there is a significant difference in instruction quality between these three groups. If there were a significant *interaction* between intervention type and gender, this would mean that the differences found between intervention types were only there for one gender and not the other. Below is an example of such an interaction.

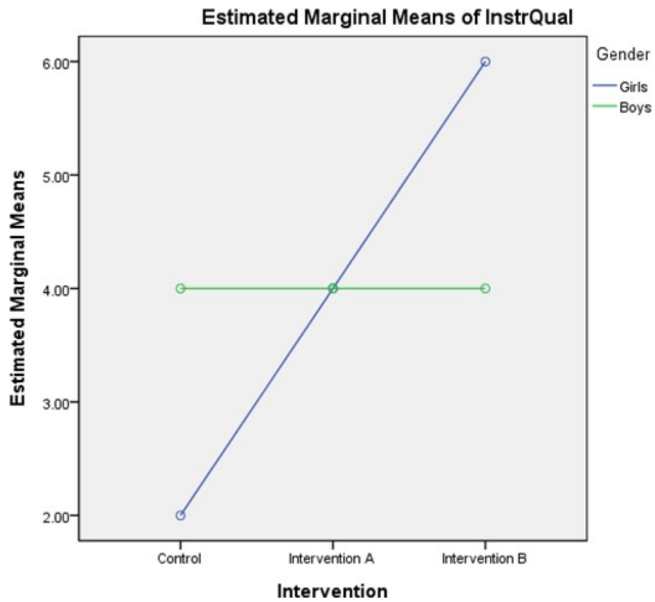

The graph shows that there are no differences for boys (green line) between the control, Intervention A, or Intervention B. However, there are clear differences for girls (blue line) between these three groups. The implication of an *interaction* is that differences between groups are **dependent** on membership in another group (in this case, being a boy or girl).

Interpretation

In the above example, because there is an interaction between intervention type and gender, mean differences in instruction quality among intervention types only exist for girls and not boys.

Repeated Measures ANOVA

Use

As discussed in the How-To Guide on *t*-tests, a dependent samples *t*-test is used when the scores between two groups are somehow dependent on each other. One example of such a dependency is when the same people are given the same measure over time to see whether there is change in that measure. Often this may take the form of pre- and post-test scores (see the How-To Guide on *t-tests* for a refresher on statistical dependency). The Repeated Measures ANOVA takes the dependent samples *t*-test one step further and allows the research to ask the question "Does the difference between the pre-test and post-test means differ as a function of group membership?"

Example

Imagine that pre-test and post-test data were collected regarding instruction quality among 100 youth program instructors. Between the pre- and post-tests an intervention was implemented aimed at improving instruction quality. If you were only interested in the difference between pre- and post-test means, a dependent samples t-test would be sufficient. However, suppose that a researcher wanted to know whether the difference in instruction quality between pre- and post-test was different between male and female instructors; in this case a Repeated Measures ANOVA would be used.

In this example, there would be two independent variables: (1) testing period (pre vs. post) and gender (male vs. female). In the language of a Repeated Measures ANOVA the testing period independent variable would be called a *within-subjects factor*[2] while the gender independent variable would be called a *between-subjects factor*[3].

Like the Factorial ANOVA, the Repeated Measures ANOVA has both main effects and an interaction. One main effect in this example would be for testing period. A significant main effect for testing period would indicate that the means for instruction quality between pre- and post-test were meaningfully different from one another.

Interpretation

In this example a significant main effect of gender would indicate that there was a meaningful difference in instructional quality between male and female instructors. In this case the analysis would test for an interaction between testing period and gender. This would provide information regarding if the difference observed in instruction quality between pre- and post-test differed depending on if the instructor were male or female. A significant interaction would indicate that the difference observed between pre- and post-test instruction quality may be present for female instructors and not male instructors. In the figure below we can see that the means for men (green line) are the same between pre-test (1) and post-test (2). However, there was a big difference between pre- and post-test means for women (blue line). From this analysis one might conclude that the intervention was effective for female instructors but not male instructors.

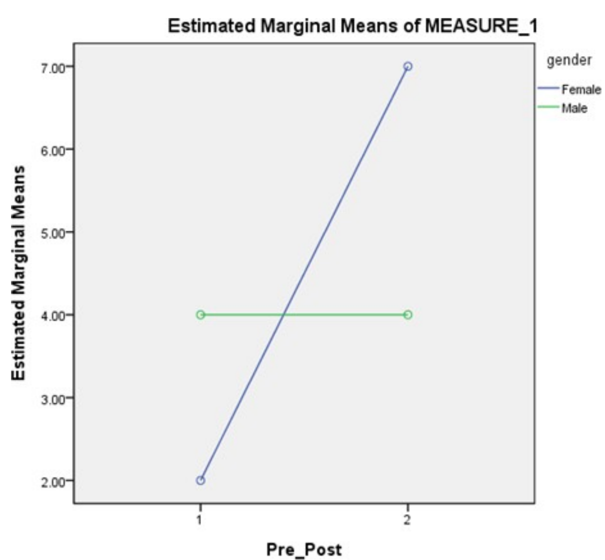

[2] A *within-subjects factor* is an independent variable that varies within the participants of the study. This means that the variance of the variable is centered on the changes observed within each study participant, such as between a pre-test and a post-test

[3] A *between-subjects factor* is an independent variable that varies only between participants but does not change within participants. Gender can be treated as a between-subjects factor as it is static within an individual but varies between individuals.

Post-Hoc Analyses and Planned Contrasts

The results of an ANOVA only tell you that there is a difference between means of different groups; it does not provide information about where these differences lie. In the Factorial ANOVA described above, an ANOVA would tell us that there is a difference among the means between intervention A, intervention B, and the control group. However, it does not tell us which means differ from each other. Is the difference between intervention A and intervention B? Is it between the control group and intervention A? Is it between all three means?

In order to answer these questions additional analyses have to be conducted.

There are two types of follow-up analyses that might be conducted with an ANOVA

- post-hoc analyses
- planned contrasts

Post-hoc analyses are used after the fact, that is these analyses are used when the researcher does not have specific predictions about where the differences between means lie. Planned contrasts are exactly that: planned. These analyses are used when the researcher has specific predictions about differences between the means.

Summary

In this How-To Guide we have gone over the basics of Analysis of Variance. We hope that we have provided useful information about ANOVA, the different types, and information about how and why ANOVA is used in research. In addition to this guide on ANOVA, our website has several other guides that may be of interest. These include guides on t-tests, correlation, regression, and others. We encourage readers to examine these guides as well to become critical consumer of youth program quality research.

Developed by
The Arizona Center for Research and Outreach (AZ REACH)

Lynne M. Borden, Principal Investigator
Gabriel Schlomer, Ph.D. (Primary Author)

Updated March 2021

The Center for Research and Outreach
REACH
Leveraging Evidence and Science into Practice

The Center for Research and Outreach (REACH Lab)
University of Minnesota
Email: REACHLab@umn.edu
Website: https://reachfamilies.umn.edu