An introduction to the one-way ANOVA

March 6, 2020 by Rebecca Bevans Date updated: July 17, 2020

ANOVA, which stands for Analysis of Variance, is a statistical test that is part of the General Linear Model.

A one-way ANOVA uses one independent variable, while a two-way ANOVA uses two independent variables.

When to use a one-way ANOVA
Use a one-way ANOVA when you have collected data about

  one categorical (DISCRETE) independent variable and
  one CONTINUOUS dependent variable.

ANOVA allows you to examine whether the dependent variable is affected by (depends on) the categories of the independent variable.

The null hypothesis (H_0) of ANOVA is that the dependent variable is NOT affected by (does not depend on) the categories of the independent variable.

The alternate hypothesis (H_a) is that the dependent variable IS affected by (does depend on) the categories of the independent variable.

One-way ANOVA Examples

Example 1:

As a crop researcher, you want to examine whether your crops' growth rates are affected by (i.e., depend on) the brand of fertilizer you use (Brand A, Brand B, and Brand C).

Your crops' growth rates are your CONTINUOUS dependent variable.

The brands of fertilizer (Brand A, Brand B, and Brand C) are your DISCRETE independent variable.

ANOVA allows you to examine whether the dependent variable (crop growth) is affected by (depends on) the categories of the independent variable (the three brands of fertilizer).

Example 2:

As a student in a popular college town, you want to examine whether apartments' rental prices are affected by (i.e., depend on) the apartments' location (near campus, west of campus, downtown).

The apartments' rental prices are your CONTINUOUS dependent variable.

The apartment's locations (near campus, west of campus, downtown) are your DISCRETE independent variable.

ANOVA allows you to examine whether the dependent variable (apartments' rental price) is affected by (depends on) the categories of the independent variable (apartments' location).

How does an ANOVA F-test work?

ANOVA calculates an *F*-test to assess the null hypothesis. The *F*-test compares the amount of explained variance to unexplained variance.

The F-test is actually the following ratio:

$$\frac{\text{variance in the dependent variable that's explained by the independent variable}}{\text{variance in the dependent variable that's unexplained by the independent variable}}$$

The more of the dependent variable's variance that can be explained by the discrete categories (of the independent variable), the higher the ratio will be.

A higher ratio (a larger *F*-test statistic) means that more of the dependent variable's variance can be explained by the independent variable's categories.

A lower ratio (a smaller *F*-test statistic) means that less of the dependent variable's variance can be explained by the independent variable's categories.

NOTE: This document has been modified.