

How to Perform a One-Way ANOVA to Compare Group Means

Authored by
stats writer

January 22, 2026

RECOMMENDED CITATION

stats writer (2026). *How to Perform a One-Way ANOVA to Compare Group Means*.
PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=126990>


The One-Way ANOVA (Analysis of Variance) represents a fundamental statistical method specifically designed for comparing the means of three or more distinct, independent groups. As a parametric test, its primary objective is to rigorously determine whether statistically significant differences exist between these group means, using the crucial ratio of variation observed between the groups versus the variation existing within the groups themselves. This powerful technique is indispensable across various fields of research and experimental design, allowing analysts to assess the influence of a single categorical independent variable on a continuous dependent variable. Understanding and properly executing the One-Way ANOVA is critical for identifying meaningful patterns and validating hypotheses regarding group differences, ultimately providing a single, efficient test statistic (the F-ratio) and corresponding p-value to summarize the findings.

Defining the One-Way Analysis of Variance

The One-Way ANOVA is a rigorous statistical test employed when the researcher aims to ascertain if the means of three or more independent populations or treatment conditions are significantly different from one another concerning a specific variable of interest. This test provides an elegant alternative to conducting multiple pairwise t-tests, which would inflate the Type I error rate (the risk of falsely rejecting the null hypothesis). Instead of comparing Groups 1 and 2, then 1 and 3, and then 2 and 3 separately, ANOVA performs a single omnibus test to determine if any significant variation exists across the entire set of groups.

For the results of the ANOVA to be reliable, the dependent variable--the variable of interest upon which the comparison is made--must meet several inherent criteria. Crucially, this variable must be continuous, meaning it can take on any value within a given range (like height or score), and it should ideally be normally distributed within each of the groups being compared. Furthermore, the variability or spread of the data must be relatively similar across all groups. The groups themselves must be defined as independent; that is, the selection of participants or data points in one group cannot influence or be related to the selection in any other group.

The underlying mechanism of the One-Way ANOVA centers around partitioning the total variance observed in the data into two primary components: the variance explained by the differences between the group means (the 'between-group' variance) and the variance attributable to individual differences or measurement error within each group (the 'within-group' or 'error' variance). The output, the F-statistic (or F-ratio), is calculated by dividing the between-group variance by the within-group variance. A larger F-ratio suggests that the differences between the group means are large relative to the random variability within the groups, thus indicating a likely significant effect of the categorical independent variable.



The One-Way ANOVA is also commonly referred to by several other names, including the One-Way ANOVA F-Test, or simply Analysis of Variance (ANOVA).

Critical Assumptions for Valid One-Way ANOVA Results

Every statistical method relies on a set of underlying assumptions about the data. If these assumptions are violated, the results derived from the statistical test may be inaccurate, misleading, or unreliable. The One-Way ANOVA, being a parametric test, is particularly sensitive to these requirements. Researchers must rigorously check these assumptions prior to interpreting the F-statistic and p-value to ensure the validity of their conclusions. Failing to meet these prerequisites often necessitates employing alternative non-parametric tests.

The core assumptions necessary for applying the One-Way ANOVA successfully are detailed below. These assumptions dictate the necessary properties of the dependent variable, the method of data collection, and the relationship between the variances of the groups being compared.

The essential assumptions for conducting a One-Way ANOVA include:

The dependent variable must be Continuous.

The scores within each group must be Normally Distributed.
The samples must be Randomly Sampled and Independent.
There must be an Adequate Sample Size (Enough Data).
The variance across groups must be Similar (Homogeneity of Variance).

We will now delve into the practical implications and detailed explanations for each of these assumptions, providing guidance on how to assess their fulfillment in real-world data analysis.

The Dependent Variable Must Be Continuous

The first and most critical assumption relates directly to the nature of the variable being measured across the groups. The dependent variable, which is the outcome measure you are investigating (and seeking differences in across the 3 or more groups), must be a continuous variable. A continuous variable is defined as one that can theoretically take on any value within a given range, exhibiting fine degrees of difference and measurement precision.

Examples of variables that satisfy the continuous requirement include physical measures such as age, weight, and height, psychological measures such as standardized test scores or aggregated survey scores (e.g., pain ratings on a 0-100 scale), and economic metrics like yearly salary. These variables are measured on either interval or ratio scales, allowing for meaningful calculation of means and standard deviations, which are the fundamental building blocks of the ANOVA calculation.


If your variable of interest is instead measured using discrete, ordinal, or categorical data (e.g., ranking categories, nominal grouping variables), the core mathematical operations of the ANOVA will be invalid. In such cases, non-parametric alternatives, which do not rely on the calculation of means, should be considered to ensure appropriate statistical inference.

Normality of the Dependent Variable


A fundamental assumption of the One-Way ANOVA is that the dependent variable must be approximately normally distributed within each of the independent groups. The concept of normality implies that when the data points for a specific group are plotted, they should resemble a bell curve, with most scores clustering around the mean and fewer scores appearing at the extremes. This characteristic ensures that the means calculated are accurate representatives of the central tendency of the data and that the F-ratio calculation remains robust.

While ANOVA is considered relatively robust against minor departures from normality, particularly with larger sample sizes, severe skewness or heavy tails can compromise the reliability of the p-value calculation. Therefore, researchers often employ statistical tests, such as the Kolmogorov-Smirnov test or the Shapiro-Wilk test, or visual tools like histograms and Q-Q plots, to formally

assess whether the distribution of the dependent variable within each group deviates significantly from the idealized bell shape.



A normal distribution.
It is bell shaped with most of the data in the middle



A skewed distribution.
It is leaning left or right with most of the data on the edge

If the variable of interest exhibits substantial and confirmed non-normality, particularly in smaller samples, a non-parametric alternative like the Kruskal-Wallis One-Way ANOVA should be used to provide a more conservative and accurate assessment of group differences.

Independence and Random Sampling

The independence assumption requires that the data points within and across the groups are unrelated. Crucially, the data points for each group in the analysis must have originated from a simple random sample of the population it represents. This means that every potential participant or observation in the target population has an equal chance of being included in the sample, and the assignment to treatment groups is generally random (if conducting an experiment).

For instance, if comparing the academic performance of students across three different teaching methodologies, the students within each methodology group must have been selected

independently of the students in the other groups. This strict adherence to random selection and assignment minimizes the threat of systematic error, or bias, which could otherwise lead to inaccurate estimates of the population parameters and fundamentally flawed conclusions regarding group differences.

If the data collection process fails to incorporate random sampling, the generalizability of the results is severely constrained, often limiting conclusions only to the specific sample studied. Furthermore, if the design involves paired or repeated measures--meaning the same subjects are measured multiple times across different conditions--the independence assumption is violated. In such scenarios, the correlation between measures must be accounted for using a specialized statistical approach.

If you are working with data where the samples are paired or related (e.g., three or more measurements taken from the identical subjects), the appropriate analysis shifts to a One-Way Repeated Measures ANOVA, which is designed to handle dependent samples.

Adequate Sample Size and Statistical Power

While there is no definitive, universal threshold for sample size, the requirement for having "Enough Data" dictates that the sample size (n) within each independent group must be sufficient to provide adequate statistical power and ensure the robustness of the ANOVA test. A common rule of thumb suggests having more than five observations in each group, although many statisticians advocate for larger sample sizes, especially if the normality assumption is weakly met.

The necessary sample size is intrinsically linked to the expected effect size--the magnitude of the true difference hypothesized across the groups. If the researcher expects a large, obvious difference in means (a large effect size), a relatively smaller sample might suffice to detect this difference. Conversely, if the anticipated difference across groups is subtle (a small effect size), a much larger sample size is imperative to achieve sufficient statistical power, which is the probability of correctly rejecting a false null hypothesis.


Therefore, researchers should ideally conduct a power analysis prior to data collection. This preparatory step helps determine the minimum sample size required to detect the smallest effect of interest with a desired level of confidence, ensuring that resources are used efficiently and that the resulting statistical test has a genuine opportunity to detect significant differences if they truly exist.

Homogeneity of Variances


The final key assumption is the homogeneity of variance, sometimes referred to as homoscedasticity. This assumption requires that the variances (the degree of spread or dispersion) of the dependent variable scores are approximately equal across all the independent groups being

compared. The ANOVA calculation inherently assumes that the within-group variance--the denominator of the F-ratio--is a pooled estimate representative of the variability in all populations.

To formally examine this assumption, statistical tests like Levene's Test or the Brown-Forsythe test are typically employed. Visually, researchers can examine box plots or scatterplots of the data, observing whether the spread of data points appears reasonably similar across the categorical groups, as illustrated in the figure below. When the assumption of homogeneity of variance is severely violated, the calculated F-statistic becomes unreliable because the pooled estimate of error variance is inaccurate, potentially leading to inflated Type I or Type II error rates.



Group 2 is much more spread out than Group 1.



Group 1 and Group 2 have a similar spread

If your analysis indicates that the groups possess substantially different spreads or unequal variances (heteroscedasticity), a modification to the standard ANOVA, known as the Welch's ANOVA variant, should be utilized, as it does not rely on the assumption of equal population variances.

Determining the Appropriate Use of One-Way ANOVA

Selecting the correct statistical procedure is paramount to sound research practice. The One-Way ANOVA is specifically tailored for a distinct set of research questions and data structures. It is the appropriate choice when the researcher's goal is explicitly focused on group comparisons based on means and when the underlying data characteristics align perfectly with the test's parametric requirements. Understanding the five core criteria for application ensures that the test yields meaningful and statistically valid inferences.

You should proceed with a One-Way ANOVA when the following conditions are simultaneously met:

The research question seeks to identify if multiple groups exhibit a significant difference in their central tendencies.

The outcome measure (dependent variable) is continuous.

There are exactly **three or more independent groups** being compared.

The data comes from **independent samples** (no repeated measures).

The dependent variable is approximately normally distributed within each group.

A common point of confusion arises when researchers try to use ANOVA for non-comparison goals. It is essential to remember that this test is fundamentally designed for difference-based inquiries, distinct from analyses that examine the strength of association between variables (correlation) or those designed for predicting an outcome based on predictor variables (regression).

Addressing Data Type Restrictions

The requirement for Continuous Data is non-negotiable for the standard One-Way ANOVA. The variable of interest must be measurable on a scale where arbitrary precision is possible, allowing for valid arithmetic operations such as calculating the mean. Examples include measurements like body temperature, reaction time, or the exact number of items completed in a timed task.

It is critical to distinguish continuous data from other data types that require alternative statistical methods. Data that are nominal (e.g., gender, eye color), ordinal (e.g., finishing place in a competition, consumer satisfaction rankings), or binary (e.g., success/failure, presence/absence of a condition) cannot be analyzed appropriately using a standard ANOVA because the mean is not a meaningful descriptor for these variable types. Using ANOVA on inappropriate data scales risks drawing erroneous conclusions.

The Mandatory Group Count

The inherent design of the One-Way ANOVA mandates that it be used to compare the means of exactly three or more groups on the chosen variable of interest. If the research design involves only two groups, the statistical landscape simplifies significantly, making a different test more appropriate and statistically powerful.

If your research design involves only two independent groups, the statistically correct procedure is the Independent Samples T-Test. Conversely, if you have only one sample group and wish to compare its mean to a pre-established or hypothesized population value, you would employ a Single Sample T-Test instead.

Ensuring Sample Independence

The requirement for Independent Samples is foundational to the One-Way ANOVA model. Independence implies that the observations drawn from one group are entirely unconnected to the observations drawn from any other group. For example, comparing the test scores of three separate cohorts of students taught by three different instructors satisfies this independence requirement, provided the students were randomly assigned to their cohorts.

Any scenario where the same individual contributes data to multiple groups (such as measuring performance before, during, and after an intervention) constitutes paired or repeated measures data. In these cases, the assumption of independence is violated, and the researcher must use specialized methods, such as the One-Way Repeated Measures ANOVA, which incorporates the correlation structure of the dependent data.

Reviewing Normality Requirements

As previously detailed, the Normality of the dependent variable within each group is a key prerequisite. This means the distribution of scores should be roughly bell-shaped. For practical application, normality is often assessed using visual techniques (histograms) or formal goodness-of-fit tests. Two widely recognized statistical tests for formally verifying whether a dataset deviates significantly from a normal distribution are the Kolmogorov-Smirnov test and the Shapiro-Wilk test. While these tests are valuable diagnostic tools, it is crucial to interpret their results alongside visual evidence, especially in very large or very small datasets.

A Practical Application of the One-Way ANOVA

Consider a controlled clinical trial designed to compare the efficacy of two new medical treatments against a standard control condition for a common infectious disease. The research question seeks to determine if the different interventions result in different recovery times. We define the three

independent groups and the dependent variable as follows:

Group 1: Patients who received Medical Treatment #1.

Group 2: Patients who received Medical Treatment #2.

Group 3: Patients who received a neutral placebo or the standard control condition.

Variable of interest: Time measured in days required for the patient to fully recover from the disease (a continuous variable).

This scenario perfectly aligns with the requirements for a One-Way ANOVA: we have three independent groups and a single, continuous variable of interest. After confirming that the recovery time data within each group meets the assumptions of normality and homogeneity of variance, we proceed to execute the analysis.

The statistical test begins with the formulation of the null hypothesis (H_0). The null hypothesis represents the expectation if the treatments have no effect: that the mean recovery times for all three groups are statistically identical ($\mu_1 = \mu_2 = \mu_3$). Conversely, the alternative hypothesis (H_a) posits that at least one group mean is different from the others. The goal of the ANOVA is to gather sufficient evidence to potentially reject the null hypothesis, thereby concluding that the treatments did indeed influence recovery time.

Upon running the One-Way ANOVA, two core values are produced: the F-statistic and the p-value. The F-statistic quantifies the ratio of between-group variance to within-group variance, effectively measuring how much the group means deviate from each other relative to the natural variability within the samples. A high F-statistic suggests that the observed differences between treatment groups are unlikely to be due merely to random chance.

The p-value is the probability of observing the current data (or data more extreme) if the null hypothesis were true--that is, assuming the treatments genuinely had no differential effect. Conventionally, if the p-value is less than or equal to a predetermined significance level (alpha, typically 0.05), the result is deemed statistically significant. A significant result allows the researcher to reject the null hypothesis and conclude that there is a statistically significant difference in recovery time among the groups. However, it is crucial to note that ANOVA only tells us that a difference exists somewhere among the groups; further post-hoc tests are required to pinpoint exactly which pairs of groups are significantly different from one another.