

How to Perform Levene's Test in SAS

Authored by
stats writer

November 19, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Perform Levene's Test in SAS*. PSYCHOLOGICAL SCALES.
Retrieved from <https://scales.arabpsychology.com/?p=97027>

In the realm of quantitative research and statistical analysis, certain parametric tests, such as Analysis of Variance (ANOVA), rely on critical underlying assumptions to ensure the validity and reliability of their results. One of the most fundamental of these prerequisites is the assumption of homogeneity of variance, often referred to as homoscedasticity. This assumption posits that the variance of the dependent variable must be approximately equal across all levels or groups defined by the independent variable. When this condition is violated--a state known as heteroscedasticity--the resulting F-statistics and p-values derived from standard ANOVA procedures may be inaccurate, potentially leading to erroneous conclusions regarding group differences.

Failing to verify this assumption can significantly undermine the rigor of a study. For instance, if one group exhibits much larger variance than others, it suggests that the precision of measurement or the underlying spread of data differs substantially across conditions. Traditional tests often pool these variances to estimate the error term; if the variances are unequal, this pooled estimate becomes unreliable, leading to inflated or deflated Type I or Type II error rates depending on the sample sizes and variance ratios. Therefore, before interpreting the results of any mean comparison test, it is incumbent upon the statistician to formally assess whether the assumption of equal variance holds true across the groups under investigation.

The practical application of statistical methods in fields ranging from biological sciences to market research necessitates robust methods for testing this critical assumption. While visual inspection using box plots or scatter plots can provide initial insights, a formal statistical test is required to provide objective evidence. This is where Levene's test enters the analytical workflow. As a widely recognized procedure, Levene's test provides a direct method for evaluating whether the dispersion of data points is statistically equivalent across various population subgroups, making it an essential tool for data validation prior to conducting more complex inferential analyses in environments like SAS.

Introducing Levene's Test: The Formal Procedure

The Levene's test is a powerful and robust statistical technique specifically designed to assess the equality of variances for a variable calculated for two or more independent groups. Unlike the Bartlett's test, which is more sensitive to non-normality in the data distributions, Levene's test is less susceptible to deviations from normality, making it a preferred choice for many real-world datasets that may not perfectly adhere to Gaussian distribution assumptions. The test operates by applying a one-way ANOVA to the absolute differences between the observed data points and the mean or median of their respective groups, effectively transforming the variance question into a mean comparison problem.

The fundamental logic behind Levene's test relies on setting up competing hypotheses that define the scope of the statistical inference. These hypotheses guide the interpretation of the resulting

test statistic and its associated p-value. Understanding these definitions is paramount, as the conclusion drawn from the test determines whether the original ANOVA (or similar test) should proceed using standard methods, or if alternative, robust procedures must be employed, such as Welch's ANOVA, which does not assume equal variances.

The hypotheses for Levene's test are structured as follows:

Null hypothesis (H₀): The population variances among all the groups being compared are statistically equal. This is the hypothesis we wish to retain if the evidence against it is weak.

Alternative hypothesis (H_A): At least one group's population variance is significantly different from the others. If the results strongly support rejecting H₀, we conclude that the assumption of homoscedasticity is violated.

Interpretation hinges on the comparison of the test's p-value against a pre-selected significance level (alpha, typically 0.05). If the calculated p-value is less than or equal to the significance level, we reject the Null hypothesis and conclude that we have enough evidence to state that the variance among the groups is not equal. Conversely, if the p-value is greater than alpha, we fail to reject the null hypothesis, thereby assuming that the variance is homogeneous across the groups, allowing us to proceed confidently with parametric analyses like standard ANOVA.

Preparing Data for Analysis in SAS

To demonstrate the implementation of Levene's Test, we will utilize the statistical power of SAS. Before running any statistical procedure, the data must be properly defined and structured within a SAS dataset. For this example, we consider a hypothetical experiment designed to evaluate the effectiveness of three different fertilizers (A, B, and C) on plant growth. The dataset captures the type of fertilizer used and the resulting plant height, measured in inches, for a total of 18 individual plants. Proper data definition is the crucial first step, ensuring that the categorical independent variable (fertilizer) and the continuous response variable (growth) are correctly recognized by the SAS system.

The structure requires the use of the DATA step in SAS, followed by the INPUT and DATALINES statements, which allows us to embed the raw data directly into the program for immediate processing. We must define the fertilizer variable as a character variable using the dollar sign (\$) notation, as it represents distinct categorical groups, while the growth variable remains numerical. This structured approach ensures that subsequent procedures, such as PROC GLM or PROC UNIVARIATE, can correctly identify the grouping factors necessary for variance comparisons.

The following code block illustrates the creation and initial viewing of this experimental dataset, which contains the fundamental measurements required for assessing both the mean differences and the variance equality between the three fertilizer treatment groups. We use the PROC PRINT

statement to confirm that the data has been read into the SAS environment accurately and is ready for the inferential statistical procedures. This preliminary step validates the integrity of the input data before proceeding to the actual Levene's test computation, setting the stage for a robust analysis.

SAS Code for Dataset Creation

```
/*create dataset*/  
data my_data;  
input fertilizer $ growth;  
datalines;  
A 29  
A 23  
A 20  
A 21  
A 33  
A 30  
B 19  
B 19  
B 17  
B 24  
B 25  
B 29  
C 21  
C 22  
C 30  
C 25  
C 24  
C 33  
;  
run;
```

```
/*view dataset*/  
proc print data=my_data;
```

Obs	fertilizer	growth
1	A	29
2	A	23
3	A	20
4	A	21
5	A	33
6	A	30
7	B	19
8	B	19
9	B	17
10	B	24
11	B	25
12	B	29
13	C	21
14	C	22
15	C	30
16	C	25
17	C	24
18	C	33

As shown in the data view, we have 18 observations categorized into three groups (Fertilizer A, B, and C). Our primary objective is twofold: first, to determine if there is a statistically significant difference in the mean growth attributable to the three fertilizers using a one-way ANOVA; and second, and more importantly for this tutorial, to formally test the required assumption of homogeneity of variance using Levene's test before interpreting the ANOVA results. This integrated approach is common practice in applied statistics, ensuring that the premises of the statistical model are met.

Executing Levene's Test Using PROC GLM in SAS

In SAS, Levene's test can be requested within several procedures, including PROC UNIVARIATE (for exploratory data analysis) and PROC GLM (General Linear Model) or PROC ANOVA (for hypothesis testing). When performing a formal ANOVA, PROC GLM is often the preferred choice due to its flexibility and robustness in handling different experimental designs. To integrate Levene's test directly into the ANOVA output, we utilize the specific HOVTEST option within the MEANS statement of the PROC GLM procedure. This command instructs SAS to perform the Homogeneity of Variance test simultaneously with the mean comparisons.

The crucial element is the syntax: **hovtest=levене(type=abs)**. The **hovtest** argument signals the request for a homogeneity test, while specifying **levене** selects the specific Levene's method. The optional but highly recommended argument, **type=abs**, dictates that the calculation of the test statistic should be based on the absolute deviations from the group median rather than the group mean. Although Levene originally proposed using the mean, using the median (which the **type=abs** option approximates) is generally considered more robust against non-normal data distributions, aligning the SAS output with the common practice in other statistical packages like R.

The following SAS code executes the one-way ANOVA using PROC GLM, defining **fertilizer** as the classification variable (CLASS statement) and **growth** as the dependent variable (MODEL statement). We then request the mean comparison for the fertilizer groups and explicitly append the Levene's test results to the output using the MEANS statement. This combined approach streamlines the analysis, providing both the primary test results and the assumption checks in one run.

SAS Code for PROC GLM and Levene's Test

```
/*perform one-way ANOVA along with Levene's test*/  
proc glm data = my_data;  
class fertilizer;  
model growth = fertilizer;  
means fertilizer / hovtest=levене(type=abs);  
run;
```

Interpreting the ANOVA Results

Upon execution of the PROC GLM code, the SAS output is generated in two primary sections. The first section provides the standard output for the one-way ANOVA, which addresses the primary research question: whether the mean plant growth differs significantly among the three fertilizer groups. This section includes the traditional ANOVA table, detailing the Sums of Squares, Degrees of Freedom, Mean Squares, the calculated F-statistic, and the corresponding p-value.

The first table in the output shows the results of the one-way ANOVA:

The GLM Procedure

Dependent Variable: growth

Source	DF	Sum of Squares	Mean Square	F Value	Pr > F
Model	2	56.3333333	28.1666667	1.17	0.3358
Error	15	359.6666667	23.9777778		
Corrected Total	17	416.0000000			

R-Square	Coeff Var	Root MSE	growth Mean
0.135417	19.85153	4.896711	24.66667

Source	DF	Type I SS	Mean Square	F Value	Pr > F
fertilizer	2	56.3333333	28.1666667	1.17	0.3358

Source	DF	Type III SS	Mean Square	F Value	Pr > F
fertilizer	2	56.3333333	28.1666667	1.17	0.3358

When examining the row corresponding to the **fertilizer** effect, we observe that the associated p-value in the ANOVA table is **0.3358**. In the context of hypothesis testing, we compare this value to our predetermined significance level, typically $\alpha = 0.05$. Since 0.3358 is substantially greater than 0.05, we do not have sufficient statistical evidence to reject the Null hypothesis of equal means. This leads us to the preliminary conclusion that there is no statistically significant difference in the mean plant growth between the three fertilizers. However, this conclusion is only fully reliable if the assumption of homogeneity of variance is met, which is confirmed by the subsequent Levene's test output.

Interpreting the Levene's Test Output

The second critical component of the PROC GLM output, generated specifically due to the **hovtest=levене(type=abs)** specification, is the table dedicated to the Homogeneity of Variance test results. This table presents the calculated F-statistic for Levene's test, along with its degrees of freedom and, most importantly, the corresponding p-value, which is essential for determining whether the variance across the groups is equal. This test directly addresses the core assumption required for the validity of the preceding ANOVA results.

We can see the output of Levene's test in the second table in the output:

The GLM Procedure

Levene's Test for Homogeneity of growth Variance ANOVA of Absolute Deviations from Group Means					
Source	DF	Sum of Squares	Mean Square	F Value	Pr > F
fertilizer	2	2.9753	1.4877	0.40	0.6745
Error	15	55.2037	3.6802		

From the results shown in this specialized table, focusing on the row where the Test Statistic is calculated using the absolute differences from the median (Type=ABS), we find that the p-value of Levene's test is reported as **0.6745**. Applying the standard decision rule, we compare this value to the significance level $\alpha = 0.05$. Since 0.6745 is significantly larger than 0.05, we **fail to reject the Null hypothesis (H0)**, which asserts that the variances are equal. Consequently, we conclude that there is no statistically significant evidence to suggest that the variances of plant growth differ across the three fertilizer groups.

In practical terms, successfully passing Levene's test confirms the assumption of homogeneity of variance. This validation ensures that the standard F-test performed in the ANOVA is appropriate and that the conclusion regarding the means--that there is no significant difference in plant growth--is reliable and not biased by unequal variance distributions. Had the test failed ($p < 0.05$), we would have been obligated to use a robust alternative like the GLM procedure with the 'unequal' variance option, or employ a non-parametric test.

Advanced Considerations: The 'TYPE=ABS' Argument

While the core function of Levene's test remains constant--testing variance equality--the exact mathematical formulation can vary depending on which measure of central tendency is used for calculating the deviations. SAS offers several options, but specifying **type=abs** within the **levене()** function is highly recommended for standard practice. This particular specification calculates the absolute deviations of the data points from their respective group medians, rather than the means.

The technical rationale for using the median is tied to the robustness of the test. When data distributions are heavily skewed or contain outliers, the group mean is sensitive and can distort the measure of central tendency, leading to an inflated Type I error rate (falsely rejecting the Null hypothesis of equal variances). The median, being resistant to extreme values, provides a more stable estimate, thereby making the resulting Levene's test more reliable under conditions of non-normality. This median-based approach is often termed the Brown-Forsythe test, although SAS groups it under the broader Levene's test framework when using **type=abs**.

By explicitly using the **type=abs** argument, the analyst ensures consistency with many other modern statistical software packages (such as R and SPSS) that default to or strongly recommend the median-based approach for variance testing due to its enhanced robustness. This choice is vital for producing defensible statistical results, particularly when analyzing biological or social science data where perfect normality is rarely achieved. The robust nature of the median-based Levene's test contributes significantly to the overall trustworthiness of the subsequent inferential statistics.

Further Resources for SAS Statistical Procedures

The following tutorials explain how to perform other common statistical tests in SAS, building upon the foundational knowledge of assumption testing demonstrated here:

Testing for Normality: Procedures like PROC UNIVARIATE also offer tests (Shapiro-Wilk, Kolmogorov-Smirnov) that assess whether data follows a normal distribution, another key assumption for parametric tests.

Robust ANOVA Methods: Exploring procedures for handling assumption violations, such as applying the **WELCH** option in PROC GLM when variance homogeneity is rejected.

Non-Parametric Alternatives: Learning when to use tests that do not rely on variance equality or normality, such as the Kruskal-Wallis test, which can be executed using PROC NPAR1WAY.