

Learn How to Perform Levene's Test for Equality of Variances in R

Authored by
Mohammed loot

November 9, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Learn How to Perform Levene's Test for Equality of Variances in R*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=14147>

In quantitative research and [statistical hypothesis testing](#), the reliability of our conclusions depends heavily on whether the underlying assumptions of the statistical models are met. One of the most fundamental requirements for parametric analyses, such as the [Analysis of Variance \(ANOVA\)](#) or the standard independent samples t-test, is the assumption of **homogeneity of variance**, often referred to as homoscedasticity.

This critical assumption demands that the spread, or [variance](#), of the dependent variable must be roughly equivalent across all defined groups or conditions being compared. When this condition is satisfied, we can proceed with confidence, knowing that the test statistics are calculated correctly. However, a failure to meet this requirement--a state known as [heteroscedasticity](#)--can severely compromise the validity of the primary statistical test, leading to erroneous interpretations of treatment effects or group differences. Consequently, formally testing for equal variances is not merely a procedural step but a necessary safeguard for rigorous inferential statistics.

The Critical Role of Equal Variances in Parametric Testing

The premise that population variances are equal across different groups is integral to the mathematical foundation of numerous widely used statistical models. For instance, both the independent samples t-test and [ANOVA](#) employ a technique called pooled variance estimation. This involves calculating a single, combined estimate of the population variance by merging data from all comparison groups. This pooling procedure is statistically sound and justifiable only if the variances are indeed similar across the underlying populations.

Should the group variances be substantially unequal, the pooled estimate becomes inherently biased. This bias primarily affects the calculation of the standard error, which is the denominator in the F-ratio or t-statistic formula. When the standard error is distorted, the resulting test statistic and, crucially, the associated [p-value](#) are inaccurate. In scenarios involving severe [heteroscedasticity](#), the test might become overly conservative (failing to detect a genuine effect) or, conversely, overly liberal (falsely indicating an effect where none exists). Validating the assumption of **homogeneity of variance** is thus paramount for maintaining the integrity and trustworthiness of any inferential analysis.

Introducing Levene's Test: Theory and Hypotheses

To formally and robustly assess whether the variances among two or more groups are equal, researchers commonly utilize [Levene's Test](#). This method is highly favored over alternatives such as Bartlett's test because it demonstrates greater resilience and robustness against violations of the often-related assumption of normality. This characteristic makes [Levene's Test](#) an exceptionally practical choice in real-world data analysis where perfect normality is rare.

The core mechanism of Levene's Test involves performing an [ANOVA](#) on the absolute residuals.

Specifically, it tests the differences between the absolute deviations of each observation from its group mean or, more robustly, its group median. If these differences in absolute residuals are statistically significant, it implies that the groups exhibit different levels of variability, suggesting the presence of [heteroscedasticity](#).

The formal statistical hypotheses guiding the interpretation of Levene's Test are defined precisely:

Null Hypothesis (H0): The population [variance](#) is statistically equal across all groups defined in the study (i.e., the assumption of Homogeneity of Variance holds).

Alternative Hypothesis (HA): The population [variance](#) is *not* equal across all groups (i.e., [Heteroscedasticity](#) exists).

Interpretation hinges on the resultant [p-value](#): if this value is less than the pre-specified significance level (alpha, typically set at 0.05), we reject the [Null hypothesis](#). Rejecting H0 means concluding that there is sufficient statistical evidence to assert that the group variances are unequal, thereby violating the critical assumption for standard parametric tests.

Setting Up the Analysis Environment in R

The statistical programming language [R](#) offers an efficient framework for conducting [Levene's Test](#). Although the specific function, `leveneTest()`, is not included in R's base distribution, it is readily accessible within the widely used **car package** (Companion to Applied Regression). This package is essential for many diagnostic and regression-related procedures.

Before any execution, analysts must first ensure that the [car package](#) is installed on their system and subsequently loaded into the current [R](#) session using the `library()` command. Once loaded, the `leveneTest()` function employs a straightforward formula interface, mirroring the syntax common to R's linear modeling functions:

```
leveneTest(response variable ~ group variable, data = data)
```

In this structure, the `response variable` represents the continuous measurement variable whose variance is being examined (the dependent variable), and the `group variable` is the categorical factor that defines the different populations or comparison groups. The mandatory `data` argument specifies the name of the [data frame](#) object that holds both of these variables.

Executing the Test: Example Data Preparation

To demonstrate the practical application of [Levene's Test](#) within [R](#), let us construct a realistic scenario involving the effectiveness of three different weight loss programs: A, B, and C. Our objective is to determine if the variability (or consistency) in the amount of weight lost is statistically

similar across all three programs. The following reproducible R code generates the necessary example [data frame](#):

```
#make this example reproducible
```

```
set.seed(0)
```

```
#create data frame
```

```
data <- data.frame(program = rep(c("A", "B", "C"), each = 30),
```

```
weight_loss = c(runif(30, 0, 3),
```

```
runif(30, 0, 5),
```

```
runif(30, 1, 7)))
```

```
#view first six rows of data frame
```

```
head(data)
```

```
# program weight_loss
```

```
#1 A 2.6900916
```

```
#2 A 0.7965260
```

```
#3 A 1.1163717
```

```
#4 A 1.7185601
```

```
#5 A 2.7246234
```

```
#6 A 0.6050458
```

The resulting synthetic dataset comprises 90 observations, with participants equally distributed among the three weight loss programs. Our current task is focused solely on assessing whether the spread of the `weight_loss` variable differs significantly depending on which `program` a participant was assigned. We now proceed to load the requisite [car package](#) and execute the `leveneTest()` function, maintaining our standard significance level (α) at 0.05 for decision-making:

```
#load car package
```

```
library(car)
```

```
#conduct Levene's Test for equality of variances
```

```
leveneTest(weight_loss ~ program, data = data)
```

```
#Levene's Test for Homogeneity of Variance (center = median)
```

```
# Df F value Pr(>F)
```

```
#group 2 4.1716 0.01862 *
```

```
# 87
```

```
#---
```

```
#Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Interpreting Statistical Outcomes and Drawing Conclusions

The structured output generated by the `leveneTest()` function closely resembles a simplified [ANOVA](#) summary table. Key metrics provided include the Degrees of Freedom (Df), the calculated F statistic (F value), and the associated probability value (Pr(>F)). For the purpose of hypothesis testing regarding the equality of variances, the [p-value](#) is the most critical component.

In the results displayed above, the resulting [p-value](#) (Pr(>F)) is calculated as **0.01862**. This figure must be evaluated against our predetermined significance threshold, $\alpha = 0.05$. Since 0.01862 is clearly less than 0.05, we possess sufficient evidence to statistically reject the [Null hypothesis](#) (H₀) of equal variances.

The definitive conclusion is that the [variance](#) in weight loss achieved by participants is significantly different across the three weight loss programs (A, B, and C). This significant finding confirms a violation of the assumption of **homogeneity of variance**. If the subsequent analytic step were a standard one-way ANOVA, this violation would necessitate the implementation of corrective measures, such as employing a robust alternative like Welch's ANOVA, which is specifically designed to handle unequal variances without requiring complex data transformation.


Visualizing Heterogeneity: The Power of Boxplots

While the numerical outcome of [Levene's Test](#) provides the definitive statistical judgment, visualizing the data distributions offers invaluable context. Visualization can help researchers identify which specific groups contribute most significantly to the variance inequality and enhance the overall clarity of the conclusion. Boxplots are an exceptionally effective graphical tool for visually comparing the central tendency and, more importantly here, the spread (variability) of continuous data across distinct categorical groups.

We can generate a clear and informative boxplot in [R](#) using the following base plotting command, which maps the continuous outcome variable (`weight_loss`) against the categorical factor (`program`):

```
boxplot(weight_loss ~ program,  
data = data,  
main = "Weight Loss Distribution by Program",  
xlab = "Program",  
ylab = "Weight Loss",  
col = "steelblue",
```

`border = "black")`



The resulting visualization strongly reinforces the statistical decision. By closely examining the vertical stretch of the boxplots--which reflects the interquartile range and serves as a robust indicator of data spread--it becomes visually evident that the distribution of weight loss for Program C is considerably wider and more dispersed than the distributions for Programs A and B. This patent difference in spread confirms the [heteroscedasticity](#) detected by [Levene's Test](#), solidifying the necessary rejection of the [Null hypothesis](#) and guiding the choice toward an appropriate robust follow-up test.

Additional Resources for Variance Testing

Mastering the ability to test and address the assumption of variance equality is absolutely foundational to performing rigorous quantitative analysis. While this guide focused on the [R](#) environment and the powerful functionality of the [car package](#), researchers often work across multiple software platforms. The following resources offer guidance on implementing Levene's Test in other common statistical software environments:

Detailed tutorial on performing Levene's Test in SPSS.

Comprehensive guide to conducting variance tests using Python libraries, specifically focusing on the SciPy package.

Official documentation regarding homogeneity of [variance](#) checks in enterprise statistical software packages such as SAS or Stata.