

Brown–Forsythe Test in R: Step-by-Step Example?

Authored by
stats writer

December 16, 2025

RECOMMENDED CITATION

stats writer (2025). *Brown–Forsythe Test in R: Step-by-Step Example?*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=107573>

The Brown-Forsythe Test in R is a robust statistical procedure essential for assessing the equality of variances across multiple independent groups. This test is particularly valuable in situations where traditional assumptions, such as normality, may not be strictly met, offering a more reliable assessment of homogeneity of variance (homoscedasticity).

Understanding the distribution of variances is a critical prerequisite for many parametric tests, most notably the One-Way Analysis of Variance (ANOVA). If the assumption of equal variances is violated, the reliability of the ANOVA results can be compromised. Therefore, the Brown-Forsythe test serves as an important diagnostic tool in the data analysis pipeline.

This comprehensive tutorial provides a detailed, step-by-step guide on how to implement and interpret the Brown-Forsythe Test using the R programming environment. We will utilize specific R functions and packages to analyze a practical dataset, ensuring clarity and applicability for statistical practitioners and researchers.

Understanding the Role of the Brown-Forsythe Test

The One-Way ANOVA is the standard statistical procedure employed to determine if there is a statistically significant difference between the means of three or more independent groups. While incredibly powerful, the ANOVA relies on several foundational assumptions regarding the nature of the data distributions.

Crucially, one of the primary assumptions for valid ANOVA results is the assumption of homogeneity of variance (or homoscedasticity). This mandates that the underlying populations from which the samples are drawn must possess equal variances. When this assumption is violated--a condition known as heteroscedasticity--the F-test statistic produced by the ANOVA may become unreliable, potentially leading to inaccurate statistical conclusions.

The **Brown-Forsythe test** is one of the most robust methods designed specifically to assess this assumption. Unlike the Levene's test, the Brown-Forsythe modification utilizes the absolute deviations from the median rather than the mean, making it less sensitive to departures from normality and extreme outliers. The test structure is built around a pair of competing hypotheses:

H₀ (Null Hypothesis): The variances among the populations are equal (Homoscedasticity holds).

H_A (Alternative Hypothesis): The variances among the populations are not equal (Heteroscedasticity is present).

To make a decision, we evaluate the resulting p-value against a predetermined significance level (commonly denoted as α , such as $\alpha = 0.05$). If the p-value is less than α , we must reject the null hypothesis, thereby concluding that the variances are significantly unequal across the groups. This result necessitates moving to alternative statistical approaches, which we

discuss later in this article.

Step 1: Setting Up the Data in R

To demonstrate the Brown-Forsythe Test, let us consider a practical research scenario. A team wants to investigate whether three distinct fitness programs (Program A, B, and C) yield significantly different outcomes in terms of weight loss over a one-month period.

The experimental design involved recruiting 90 participants and randomly assigning 30 individuals to each of the three programs. The dependent variable measured is the amount of weight loss (in pounds or kilograms) recorded for each person after the defined study duration.

In R, we simulate this dataset to create a reproducible example. The dataset, named `data`, contains two variables: `program` (the categorical grouping variable) and `weight_loss` (the continuous outcome variable). We use the `set.seed(0)` function to ensure that the randomized data generation process can be replicated precisely.

The following R code snippet outlines the creation of this dataframe and displays the initial six rows for verification:

```
#make this example reproducible
```

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

```
#create data frame
```

```
data <- data.frame(program = as.factor(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
```

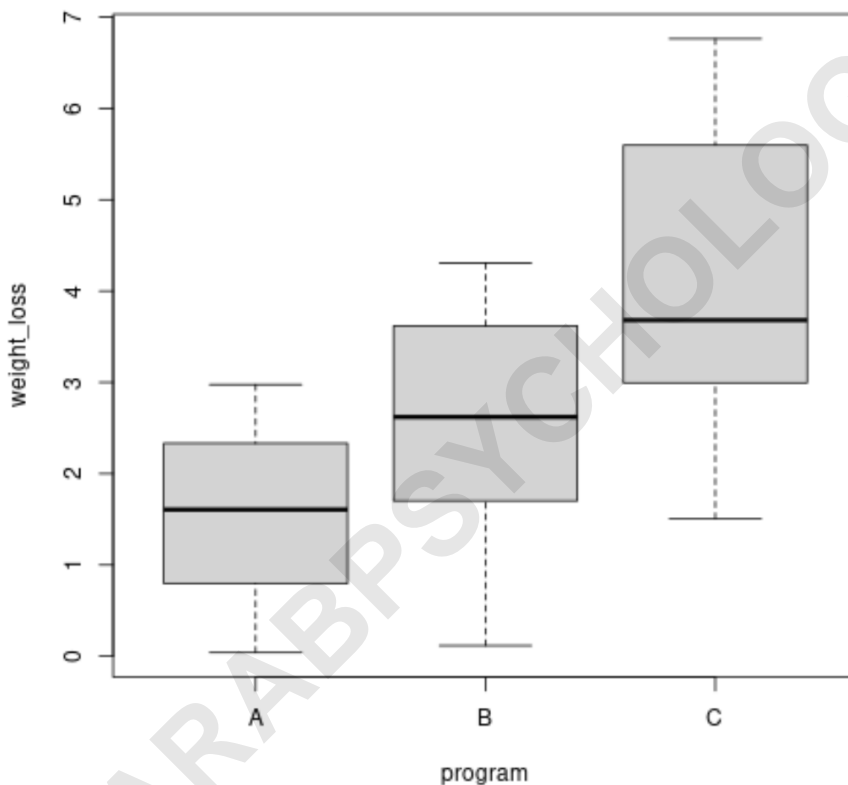
```
#6 A 0.6050458
```

Step 2: Preliminary Data Summarization and Visualization

Prior to executing the formal Brown-Forsythe Test, it is always recommended practice to perform exploratory data analysis. Visualizing the data provides initial insights into the distribution and potential differences in spread (heteroscedasticity) among the groups. A boxplot is particularly effective for this purpose, illustrating the median, quartiles, and range for each program.

We generate a comparative boxplot using R's base graphics functionality, specifying the formula `weight_loss ~ program`:

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



Observing the boxplot visually suggests that the spread of weight loss increases from Program A to Program C, indicating a potential violation of the homogeneity assumption. To quantify this observation, we calculate the exact sample variance for each group. This requires loading the `dplyr` package for efficient grouping and summarization operations in R.

```
#load dplyr package
```

```
library(dplyr)
```

```
#calculate variance of weight loss by group
```

```
data %>%
group_by(program) %>%
summarize(var=var(weight_loss))

# A tibble: 3 x 2
  program var
1 A 0.819
2 B 1.53
3 C 2.46
```

The calculated variances confirm the visual assessment: Program A has the smallest variance (0.819), while Program C exhibits the largest variance (2.46). This numerical disparity strongly suggests that the equality of variances assumption required for a standard One-Way ANOVA may be breached, thus justifying the formal use of the Brown-Forsythe Test.

Step 3: Execution and Interpretation of the Brown-Forsythe Test

To formally test the null hypothesis of equal variances, we utilize the R statistical environment. The Brown-Forsythe test is conveniently implemented using the `bf.test()` function, which is available within the highly useful `onewaytests` package. Ensure this package is installed and loaded before proceeding with the analysis.

The `bf.test()` function takes the same formula syntax common in R statistical modeling: `dependent_variable ~ independent_variable`, where `weight_loss` is the outcome and `program` defines the groups. The function returns a detailed output summarizing the test statistics and the conclusion based on the default significance level ($\alpha = 0.05$).

```
#load onewaytests package
```

```
library(onewaytests)
```

```
#perform Brown-Forsythe test
```

```
bf.test(weight_loss ~ program, data = data)
```

```
Brown-Forsythe Test (alpha = 0.05)
```

```
-----
data : weight_loss and program
```

```
statistic : 30.83304
```

```
num df : 2
```

```
denom df : 74.0272
```

p.value : 1.816529e-10

Result : Difference is statistically significant.

Upon examining the test output, the resulting p-value is reported as \$1.816529 times 10^{-10} \$. This value is extremely small, being substantially less than our conventional significance level of $\alpha = 0.05$ \$. Because the p-value falls below the threshold, we strongly reject the null hypothesis (H_0).

The conclusion is clear: the differences in variances among the three workout program groups are statistically significant. This confirms the initial visual and descriptive findings, indicating a serious violation of the homoscedasticity assumption required for a standard One-Way ANOVA.

Interpreting Results and Choosing the Next Statistical Path

The outcome of the Brown-Forsythe test dictates the subsequent steps in the mean comparison analysis. There are two primary scenarios based on whether the null hypothesis (H_0) is retained or rejected.

If the test yields a p-value greater than α (e.g., $p > 0.05$), we fail to reject the null hypothesis. This indicates that there is insufficient evidence to suggest that the population variances are unequal. In this favorable case, the assumption of homoscedasticity is met, and the researcher can confidently proceed with the standard One-Way ANOVA to compare group means.

However, as demonstrated in our example, rejecting the null hypothesis signifies a violation of the equal variances assumption. When facing statistically significant heteroscedasticity, the analyst must decide on the appropriate course of action from two main alternatives: applying corrections to the ANOVA or utilizing a non-parametric test.

Strategies for Handling Unequal Variances

Since the standard One-Way ANOVA relies heavily on the assumption of equal variances, its use when this assumption is violated requires careful consideration. Fortunately, the ANOVA is considered relatively **robust** to moderate violations, particularly if group sizes are equal (as in our example).

The first strategy is often to proceed with the One-Way ANOVA but confirm the severity of the variance inequality. Statistical theory suggests that ANOVA results remain reliable if the ratio of the largest group variance to the smallest group variance does not exceed 4. We can verify this ratio using the results from Step 2:

Smallest Variance (Program A): 0.819

Largest Variance (Program C): 2.46

Ratio: $2.46 / 0.819 = 3.003$

Since this calculated ratio (3.003) is less than the critical threshold of 4, we might conclude that the violation is moderate enough that we could proceed with a standard One-Way ANOVA, trusting its robustness. Alternatively, we could apply corrections (such as the Welch's ANOVA) designed specifically for unequal variances.

The second, more conservative strategy is required if the ratio of variances is greater than 4, or if the data are severely non-normal. In such cases, switching to a **Kruskal-Wallis Test** is highly advisable. This is considered the non-parametric analogue to the One-Way ANOVA, operating on ranks rather than raw data, making it immune to variance inequality and non-normality issues.

For those required to use the non-parametric route, resources detailing the implementation of the Kruskal-Wallis test in R are widely available for a comprehensive, step-by-step analysis.