

# How to Perform a Jarque-Bera Test in R to Check for Normal Distribution

Authored by  
**stats writer**

March 2, 2026

## RECOMMENDED CITATION

stats writer (2026). *How to Perform a Jarque-Bera Test in R to Check for Normal Distribution*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=133500>

The **Jarque-Bera test** represents a fundamental statistical procedure employed by data scientists and econometricians to assess the normality of a given dataset. In the realm of quantitative analysis, many statistical models, including **linear regression**, rely on the underlying assumption that residuals or data points follow a **normal distribution**. This goodness-of-fit test specifically examines whether the sample data exhibits **skewness** and **kurtosis** consistent with a Gaussian profile. By quantifying these characteristics, the test provides a rigorous mathematical basis for accepting or rejecting the assumption of normality, which is critical for the validity of subsequent parametric testing.

## Understanding the Fundamentals of the Jarque-Bera Test

At its core, the **Jarque-Bera test** is designed to measure the difference between the third and fourth moments of a sample and those of a true **normal distribution**. A perfectly normal distribution is characterized by a **skewness** of zero, indicating perfect symmetry, and a **kurtosis** of three, reflecting the "peakedness" or thickness of the distribution's tails. When a dataset deviates significantly from these benchmarks, the resulting test statistic increases, suggesting that the data may be biased or contain heavy tails that violate standard assumptions. This test is particularly prevalent in econometrics because financial returns and economic indicators often exhibit "fat tails," making a formal normality test indispensable for accurate risk assessment.

The execution of this test within the **R** environment is streamlined through specialized packages, allowing researchers to automate the calculation of complex mathematical formulas. Before performing the test, it is essential to understand that the **null hypothesis** ( $H_0$ ) posits that the data is normally distributed. Conversely, the alternative hypothesis ( $H_1$ ) suggests that the data does not follow a normal distribution. Interpreting the results requires a careful look at the **p-value** in relation to a predefined **significance level**, which is usually set at 0.05. A **p-value** lower than this threshold provides strong evidence to reject the null hypothesis.

Furthermore, the **Jarque-Bera test** is most reliable when applied to large sample sizes. Because the test statistic asymptotically follows a **chi-squared distribution** with two degrees of freedom, smaller samples may sometimes produce unreliable results or lack the statistical power to detect subtle deviations from normality. Therefore, practitioners often pair this test with visual inspections, such as Q-Q plots or histograms, to gain a comprehensive understanding of the data's shape and behavior before proceeding with advanced statistical modeling or hypothesis testing.

## The Mathematical Foundation of the Jarque-Bera Statistic

The mathematical architecture of the **Jarque-Bera test** is elegantly structured to capture the essence of a distribution's shape. The test statistic, often denoted as **JB**, is calculated using the sample's **skewness** ( $S$ ) and **kurtosis** ( $C$ ). The formula is generally defined as follows:  $JB = \frac{S^2}{6} + \frac{C-3}{2}$ . In

this equation,  $n$  represents the total number of observations in the sample, while  $k$  represents the number of regressors if the test is being applied to the residuals of a regression model. If the test is performed on a standalone dataset,  $k$  is typically assigned a value of 1. This formula effectively penalizes any deviation from a skewness of 0 and an excess kurtosis of 0 (since excess kurtosis is defined as  $C - 3$ ).

Because the **JB** statistic is a sum of squares, it is inherently non-negative. A value of zero would indicate that the sample data perfectly aligns with the expected skewness and kurtosis of a **normal distribution**. As the **JB** value moves further away from zero, the likelihood that the data originated from a normal population decreases. Under the **null hypothesis**, the test statistic follows a **chi-squared distribution** with two degrees of freedom, which allows for the calculation of the **p-value**. This mathematical rigor ensures that the test remains an objective standard for checking distributional assumptions in scientific research.

The components of the formula--**skewness** and **kurtosis**--describe different facets of the data. Skewness measures the asymmetry of the probability distribution about its mean; a positive skew indicates a longer tail on the right side, while a negative skew indicates a longer tail on the left. Kurtosis, on the other hand, measures the "tailedness" of the distribution. High kurtosis (leptokurtic) indicates that the data has outliers or a heavy tail, while low kurtosis (platykurtic) suggests a flatter distribution with fewer extreme values. By combining these two metrics into a single statistic, the Jarque-Bera test offers a holistic view of the distribution's deviation from Gaussian ideals.

## Preparing the R Environment for Normality Testing

To implement the **Jarque-Bera test** in the **R** programming environment, users must first ensure they have the appropriate tools installed. The most common function for this test is found within the **tseries package**, which is a comprehensive library for time series analysis and computational finance. Installing the package is a straightforward process using the standard *install.packages* command, and it only needs to be performed once. Once installed, the library must be loaded into the current **R** session using the *library()* or *require()* functions to make the testing capabilities accessible.

Data preparation is the next critical step. Before running the **jarque.bera.test** function, the data should be cleaned and formatted as a numeric vector or a single column within a data frame. Missing values (NAs) should be handled appropriately, as they can cause the function to return errors or skewed results. In many real-world scenarios, the data is imported from external sources like CSV files or SQL databases. Ensuring the data integrity at this stage is paramount for obtaining a reliable **p-value**. Once the environment is configured and the data is loaded, the user can invoke the function by passing the dataset as the primary argument.

The **tseries package** is favored by many because its implementation of the Jarque-Bera test is highly efficient and provides clear, concise output. Beyond just the **JB** statistic, the output includes the degrees of freedom and the associated **p-value**. This ease of use makes **R** an ideal platform for both students learning statistics and professional analysts performing complex data audits. By following these setup steps, researchers can ensure a smooth transition from data acquisition to rigorous statistical validation.

## Executing the Jarque-Bera Test with Normal Data

To demonstrate the practical application of the test, let us first examine a scenario where the data is known to follow a **normal distribution**. We can simulate this in **R** using the `rnorm()` function, which generates random variables from a Gaussian population. In the following example, we generate a sample of 100 observations. Because these values are drawn from a normal distribution, we expect the **Jarque-Bera test** to yield a low test statistic and a high **p-value**, leading us to fail to reject the **null hypothesis**.

```
#install (if not already installed) and load tseries package  
if(!require(tseries)){install.packages('tseries')}
```

```
#generate a list of 100 normally distributed random variables  
dataset <- rnorm(100)
```

```
#conduct Jarque-Bera test  
jarque.bera.test(dataset)
```

The code block above illustrates the standard workflow for conducting the test. By checking for the presence of the **tseries package** and loading it conditionally, the script remains robust across different user environments. The `dataset` variable holds the generated values, which are then analyzed by the `jarque.bera.test()` function. The resulting output provides the statistical evidence needed to conclude whether the sample possesses the characteristics of a normal population.

Upon running the script, **R** produces a summary of the test results. This summary is essential for documenting the analysis and verifying that the distributional assumptions of the study are met. In the case of normally distributed random variables, the **JB** statistic will typically be close to zero, reflecting the fact that the sample's **skewness** and **kurtosis** do not deviate significantly from the theoretical values of a normal curve.

## Analyzing Results and Interpreting the P-Value

Interpreting the output of the **Jarque-Bera test** requires an understanding of statistical inference. When the test is executed in **R**, the console displays a specific set of values that describe the

outcome. For instance, consider the output generated from our normally distributed dataset:

```
Jarque Bera Test

data: dataset
X-squared = 0.67446, df = 2, p-value = 0.7137
```

In this specific trial, the test statistic is calculated as 0.67446, with an associated **p-value** of 0.7137. Since this **p-value** is significantly higher than the standard **significance level** of 0.05, we fail to reject the **null hypothesis**. This result indicates that there is no statistically significant evidence to suggest that the data deviates from a **normal distribution**. This is exactly what we expected, given that the data was artificially generated using a normal random number generator.

It is important to remember that failing to reject the null hypothesis does not strictly "prove" that the data is normal; rather, it suggests that the data is consistent with a normal distribution based on the evidence provided by its **skewness** and **kurtosis**. In practical research, this result allows the analyst to proceed with other parametric tests, such as t-tests or ANOVA, which require normality as a prerequisite. If the test statistic had been much larger, it would have pushed the **p-value** below the threshold, signaling that the assumption of normality had been violated.

When reporting these results in an academic or professional setting, one should always include the test statistic, the degrees of freedom, and the **p-value**. Providing the full context of the test allows other researchers to verify the findings and understand the limitations of the data. Furthermore, comparing these results across different subsets of data can reveal underlying patterns or anomalies that might otherwise go unnoticed during a cursory data review.

## Testing Non-Normal Distributions: The Uniform Distribution Example

To truly appreciate the sensitivity of the **Jarque-Bera test**, it is helpful to apply it to a dataset that is intentionally non-normal. A **uniform distribution** serves as an excellent contrast because every value within a specific range has an equal probability of occurring. Unlike the **normal distribution**, which is bell-shaped with thin tails, a uniform distribution is flat and lacks the central tendency and tail characteristics that the Jarque-Bera test looks for. We can simulate this in **R** using the `runif()` function.

```
#install (if not already installed) and load tseries package  
if(!require(tseries)){install.packages('tseries')}
```

```
#generate a list of 100 uniformly distributed random variables
```

```
dataset <- runif(100)

#conduct Jarque-Bera test
jarque.bera.test(dataset)
```

In this second example, we generate 100 random variables that follow a uniform distribution. When we subject this dataset to the **Jarque-Bera test**, we expect the resulting test statistic to be significantly higher and the **p-value** to be much lower, likely falling below the 0.05 **significance level**. This outcome would lead us to reject the **null hypothesis**, correctly identifying that the data does not conform to a normal distribution.

The difference in output between the normal and uniform datasets highlights the utility of the test in identifying distributional deviations. By systematically testing different types of data, analysts can better understand how various data-generating processes affect the shape of their samples. This knowledge is vital when choosing the appropriate statistical methods for analysis, as using parametric tests on non-normal data can lead to inaccurate conclusions and flawed decision-making.

## Evaluating the Results of the Uniform Distribution Test

After executing the Jarque-Bera test on our uniformly distributed data, we can observe a distinct shift in the statistical output. The **p-value** should reflect the data's lack of normality, providing clear evidence that the **null hypothesis** is no longer tenable. Let us look at the actual output generated by the **R** console for this scenario:

```
Jarque Bera Test

data: dataset
X-squared = 8.0807, df = 2, p-value = 0.01759
```

As shown in the image, the test statistic has increased to 8.0807, and the **p-value** is now 0.01759. Because 0.01759 is less than our 0.05 **significance level**, we reject the null hypothesis. We now have sufficient statistical evidence to conclude that this dataset does not follow a **normal distribution**. This finding is entirely consistent with the fact that the data was generated from a uniform distribution, which possesses fundamentally different **skewness** and **kurtosis** properties than a normal distribution.

This contrast between the two examples--the normal distribution and the uniform distribution--demonstrates the reliability of the **Jarque-Bera test** in **R**. It effectively differentiates between

datasets that meet the assumptions of normality and those that do not. For practitioners, this means they can use the test as a "gatekeeper" to decide whether to use standard parametric tools or to pivot toward non-parametric alternatives that do not require distributional assumptions, such as the Wilcoxon rank-sum test or bootstrapping methods.

Ultimately, the ability to conduct and interpret the **Jarque-Bera test** is an essential skill for anyone working with data in **R**. Whether analyzing financial markets, biological trends, or social science surveys, verifying the distribution of your data is a critical step in ensuring that your statistical inferences are sound. By using the **tseries package** and carefully evaluating the **p-value**, you can maintain a high standard of analytical rigor in all your projects.

## Best Practices and Limitations of the Jarque-Bera Test

While the **Jarque-Bera test** is a powerful tool, it is important to be aware of its limitations and use it within the context of best statistical practices. One primary consideration is the sample size. The test relies on large-sample properties, meaning its reliability increases as the number of observations grows. For very small samples, the Jarque-Bera test may lack the sensitivity to detect non-normality, or it may provide misleading results. In such cases, alternative tests like the **Shapiro-Wilk test** are often recommended, as they are specifically optimized for smaller datasets.

Another best practice is to always supplement the **Jarque-Bera test** with visual data exploration. Graphical methods like Q-Q plots provide a visual representation of how closely the sample quantiles match the theoretical quantiles of a **normal distribution**. If the points on a Q-Q plot deviate significantly from the diagonal line, it provides a qualitative confirmation of the quantitative results provided by the Jarque-Bera test. Combining these two approaches--statistical testing and visual inspection--yields a much more robust assessment of the data's properties.

Finally, researchers should be mindful of the impact of outliers. Because the Jarque-Bera test uses **skewness** and **kurtosis**, which involve raising values to the third and fourth powers, extreme outliers can have a disproportionate effect on the test statistic. A single extreme value can sometimes cause the test to reject the **null hypothesis** even if the rest of the data is perfectly normal. Therefore, it is always wise to investigate the presence of outliers and determine if they represent genuine data points or errors in data collection before finalizing any conclusions about the distribution's normality.