

How to Perform a Jarque-Bera Test in Python

Authored by
stats writer

December 25, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Perform a Jarque-Bera Test in Python*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=108742>

In the field of statistics and data science, assessing whether a dataset adheres to a normal distribution is often a critical prerequisite for performing various parametric analyses. Many powerful statistical models, such as linear regression or t-tests, rely fundamentally on the assumption that the underlying data is normally distributed. Failure to confirm this assumption can lead to unreliable conclusions and invalid model results.

The Jarque-Bera test (JB test) stands out as a robust goodness-of-fit test specifically designed to quantify this departure from normality. Unlike visual inspection methods like Q-Q plots, the JB test provides a definitive numerical result--the test statistic and the associated p-value--allowing practitioners to make an objective decision about the data's distributional properties. This comprehensive guide will detail the theoretical foundations of the Jarque-Bera test and demonstrate its practical application using the powerful data analysis capabilities available in Python.

The Core Principle of the Jarque-Bera Test

The Jarque-Bera test, named after economists Carlos Jarque and Anil Bera, is a powerful statistical test used to determine if the sample data has the same skewness and kurtosis coefficients as a theoretical normal distribution. For any perfect normal distribution, the skewness is exactly zero, indicating symmetry around the mean, and the excess kurtosis (kurtosis minus 3) is also zero, indicating a specific peakedness and tail weight.

The JB test is a joint test of these two moments. It calculates a single test statistic based on the third and fourth standardized moments of the sample data. Since the test statistic is based on squared values, it is always a non-negative number. A value of the test statistic close to zero suggests that the sample data is close to the characteristics of a normal distribution. Conversely, a large test statistic provides strong evidence that the data significantly deviates from normality, primarily due to asymmetry or unusual peakedness.

The Role of Skewness and Kurtosis in Normality Assessment

Understanding the components that underpin the JB statistic is crucial for proper interpretation. The normal distribution is defined by two key properties: location (mean) and spread (standard deviation). However, the JB test focuses on the shape parameters--skewness and kurtosis--which measure deviations from perfect Gaussian shape.

Skewness: This measures the asymmetry of the probability distribution of a real-valued random variable about its mean. If skewness is positive, the tail on the right side of the distribution is longer or fatter. If it is negative, the tail on the left side is longer. A symmetrical distribution, like the normal distribution, has zero skewness.

Kurtosis: This measures the "tailedness" of the probability distribution. High kurtosis means more

extreme outliers are present (heavy tails), while low kurtosis means fewer outliers and lighter tails. The standard normal distribution has a kurtosis value of 3. The JB test often uses **Excess Kurtosis**, which is defined as Sample Kurtosis minus 3, making the ideal value for a normal distribution 0.

The Jarque-Bera test statistic is proportional to the sum of the squared skewness and the squared excess kurtosis, weighted by the sample size (N). Because it incorporates both measures simultaneously, it offers a more comprehensive assessment of distributional shape than checking skewness or kurtosis in isolation.

Formulating and Interpreting the Jarque-Bera Hypotheses

As with all formal hypothesis tests, the Jarque-Bera procedure requires the establishment of a null hypothesis (H_0) and an alternative hypothesis (H_A). These hypotheses directly govern how we interpret the resulting p-value and make a final determination regarding the data's normality.

Null Hypothesis (H_0): The data is normally distributed. This implies that both the skewness and the excess kurtosis of the sample are equal to zero.

Alternative Hypothesis (H_A): The data is not normally distributed. This means that the skewness and/or the excess kurtosis are significantly different from zero.

The decision rule is based on comparing the calculated p-value to a predetermined significance level (α), typically set at 0.05. If the p-value is less than α (e.g., $p < 0.05$), we reject the null hypothesis (H_0). Rejecting H_0 means there is sufficient evidence to conclude that the data does not follow a normal distribution. Conversely, if $p \geq 0.05$, we fail to reject H_0 , meaning we do not have enough evidence to claim a significant departure from normality.

Implementing the Jarque-Bera Test in Python: Scipy vs. Statsmodels

Python offers excellent tools for performing statistical analysis, primarily through the Scipy library and the Statsmodels library. Although the Jarque-Bera test is available in both packages, the implementation within Scipy's statistics module, `scipy.stats`, is widely used for quick normality assessment.

The `scipy.stats.jarque_bera()` function provides a streamlined way to execute the test. It requires the data array (or series) as its primary input. It is important to note the specific outputs returned by this function, as they differ slightly from other statistical packages.

The general syntax for using the Scipy function is:

```
stats.jarque_bera(x)
```

Where:

x: Represents the array of observations (e.g., a NumPy array or a Pandas Series) upon which the normality test is to be performed.

This function returns a tuple containing two values: the Jarque-Bera **test statistic** and the corresponding **p-value**. For analyses requiring more comprehensive output, such as the individual skewness and kurtosis metrics, the Statsmodels implementation may be preferred, but the core function remains the same.

Case Study 1: Testing Normally Distributed Data

To confirm the validity and expected behavior of the Jarque-Bera test, we first examine data that is theoretically known to be normally distributed. We utilize the `numpy.random.normal()` function to generate a large sample of 5,000 observations drawn from a standard normal distribution (mean 0, standard deviation 1). We expect the test to confirm that this dataset adheres to the normal distribution assumption.

```
import numpy as np
import scipy.stats as stats

#generate array of 5000 values that follow a standard normal distribution
np.random.seed(0)
data = np.random.normal(0, 1, 5000)

#perform Jarque-Bera test
stats.jarque_bera(data)

(statistic=1.2287, pvalue=0.54098)
```

The results confirm our expectation. The Jarque-Bera test statistic is calculated as **1.2287**, and the corresponding p-value is **0.54098**. Since the p-value (0.54098) is significantly greater than the standard significance level ($\alpha=0.05$), we **fail to reject the null hypothesis**. This failure to reject H_0 means that there is insufficient statistical evidence to conclude that this sample data exhibits skewness or kurtosis significantly different from a theoretical normal distribution. This outcome is consistent with the process used to generate the data.

Case Study 2: Detecting Non-Normal Distributions

Next, we test a dataset that we know violates the assumption of normality. We generate 5,000 observations using the `numpy.random.uniform()` function, creating data drawn from a uniform

distribution. While a uniform distribution is symmetric (skewness is 0), its kurtosis is significantly different from 3, resulting in negative excess kurtosis (it is flatter than the normal distribution).

```
import numpy as np
```

```
import scipy.stats as stats
```

```
#generate array of 5000 values that follow a uniform distribution
```

```
np.random.seed(0)
```

```
data = np.random.uniform(0, 1, 5000)
```

```
#perform Jarque-Bera test
```

```
stats.jarque_bera(data)
```

```
(statistic=300.1043, pvalue=0.0)
```

In this second scenario, the Jarque-Bera test returns a massive test statistic of **300.1043**, which is substantially further from zero than the previous example. Crucially, the p-value is calculated as **0.0** (or extremely close to zero). Since $p < 0.05$, we **reject the null hypothesis**. This statistically confirms that the sample data is not drawn from a normal distribution, aligning perfectly with the fact that we generated the sample using a uniform distribution.

The highly significant result is expected because the uniform distribution has an excess kurtosis that is far from zero, which drastically violates the normality assumption tested by the Jarque-Bera statistic. The large test statistic reflects this significant departure from the ideal shape of the normal distribution.

Contextualizing the Jarque-Bera Test: Limitations and Best Practices

While the Jarque-Bera test is highly effective for identifying departures from normality based on shape parameters, it is essential to understand its specific application context. The JB test is derived assuming the data is independent and identically distributed (i.i.d.), and its statistical properties rely on large sample sizes. Generally, the test is considered reliable only when the sample size (N) is substantial, often cited as $N > 2000$ observations.

For smaller datasets, other tests of normality are generally preferred. The **Shapiro-Wilk test**, for instance, is often recommended as the most powerful test for small to moderate sample sizes ($N < 50$). Using the Jarque-Bera test on small samples can lead to inaccurate p-values and a loss of statistical power, potentially resulting in a Type II error (failing to reject a false null hypothesis when it should have been rejected).

Therefore, the primary utility of the Jarque-Bera test lies in analyzing large datasets, particularly

those encountered in financial econometrics or large-scale observational studies, where initial confirmation of distributional assumptions is necessary before applying advanced parametric modeling techniques.

ARABPSYCHOLOGY.COM