How to Perform and Interpret the Ljung-Box Test for Time Series Data

Authored by
stats writer

March 7, 2026

RECOMMENDED CITATION

stats writer (2026). *How to Perform and Interpret the Ljung-Box Test for Time Series Data*. PSYCHOLOGICAL SCALES. Retrieved from https://scales.arabpsychology.com/?p=134300

The Ljung-Box test is a statistical test used to determine if a time series data is exhibiting significant autocorrelation. It is based on the Ljung-Box Q statistic, which measures the correlation between a series and its lagged values. This test is commonly used in time series analysis to check if the residuals of a statistical model are random and do not contain any significant patterns. A low p-value from the Ljung-Box test indicates that the data is not random and the model needs to be improved. For example, if we are analyzing the stock prices of a company over a period of time, the Ljung-Box test can be used to determine if there is any significant relationship between the stock prices on different days.

Ljung-Box Test: Definition + Example

The Ljung-Box test, named after statisticians Greta M. Ljung and George E.P. Box, is a statistical test that checks if autocorrelation exists in a time series.

The Ljung-Box test is used widely in econometrics and in other fields in which time series data is common.

The Basics of the Ljung-Box Test

Here are the basics of the Ljung-Box test:

Hypotheses

The Ljung-Box test uses the following hypotheses:

H0: The residuals are independently distributed.

HA: The residuals are not independently distributed; they exhibit serial correlation.

Ideally, we would like to fail to reject the null hypothesis. That is, we would like to see the p-value of the test be greater than 0.05 because this means the residuals for our time series model are independent, which is often an assumption we make when creating a model.

Test Statistic

The test statistic for the Ljung-Box test is as follows:

Q = n(n+2) Σpk2 / (n-k)

where:

n = sample size

Σ = a fancy symbol that means "sum" and is taken as the sum of 1 to *h*, where *h* is the number of lags being tested.

pk = sample autocorrelation at lag *k*

Rejection Region

We reject the null hypothesis and say that the residuals of the model are not independently distributed if Q >

X21-α, h

Example: How to Conduct a Ljung-Box Test in R

To conduct a Ljung-Box test in R for a given time series, we can use the Box.test() function, which uses the following notation:

Box.test(x, lag =1, type=c("Box-Pierce", "Ljung-Box"), fitdf = 0)

where:

x: A numeric vector or univariate time serieslag: Specified number of lagstype: Test to be performed; options include Box-Pierce and Ljung-Boxfitdf: bDegrees of freedom to be subtracted if x is a series of residuals

The following example illustrates how to perform the Ljung-Box test for an arbitrary vector of 100 values that follow a normal distribution with mean = 0 and variance = 1:

```
#make this example reproducible
set.seed(1)
```

```
#generate a list of 100 normally distributed random variables
data <- rnorm(100, 0, 1)

#conduct Ljung-Box test
Box.test(data, lag = 10, type = "Ljung")
```

This generates the following output:

```
Box-Ljung test

data: data
X-squared = 6.0721, df = 10, p-value = 0.8092
```

The test statistic of the test is Q = 6.0721 and the p-value of the test is 0.8092, which is much larger than 0.05. Thus, we fail to reject the null hypothesis of the test and conclude that the data values are independent.

Note that we used a lag value of 10 in this example, but you can choose any value that you would like to use for the lag, depending on your particular situation.

How to Perform a Ljung-Box Test in Python