

How to Easily Perform the Breusch-Godfrey Test for Autocorrelation in R

Authored by
stats writer

December 5, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Easily Perform the Breusch-Godfrey Test for Autocorrelation in R*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=106038>

The Breusch-Godfrey test is a fundamental statistical tool utilized primarily within the realm of time-series econometrics to rigorously assess the validity of a key assumption in linear regression modeling: the absence of serial correlation among the error terms. Unlike simpler diagnostics, this test offers flexibility, allowing analysts to check for autocorrelation at higher lag orders, providing a comprehensive evaluation of model fit and reliability. In the statistical computing environment R, this procedure is efficiently executed using the `lmtest::bgtest()` function. This function returns a comprehensive output object, including the necessary test statistic, the crucial p-value, and contextual information detailing the test parameters. It is imperative that the `lmtest` package is properly installed and loaded into the session before attempting to run the `bgtest()` function, ensuring access to this powerful diagnostic tool.

The Importance of Independent Residuals in Regression

One of the fundamental assumptions underpinning the validity and efficiency of Ordinary Least Squares (OLS) regression estimates is that the residuals (the differences between the observed and predicted values) are independent. This independence implies that the error observed at one point in time or space is not correlated with the error observed at any other point. When this assumption is violated, typically resulting in a condition known as serial correlation or autocorrelation, the standard errors calculated for the coefficient estimates become biased, often leading to misleadingly small p-values. Consequently, the statistical inferences drawn from the model might be unreliable, making hypothesis tests invalid and confidence intervals inaccurate.

For decades, practitioners relied heavily on the Durbin-Watson test to check for first-order autocorrelation (where the current residual is correlated only with the immediately preceding residual). While effective for this specific scenario, the Durbin-Watson test is restrictive. If the structure of the data suggests that autocorrelation might persist over several periods--for instance, if the error today relates to errors from two or three days ago--a more sophisticated test is required. The Breusch-Godfrey test, also known as the LM (Lagrange Multiplier) test for serial correlation, overcomes this limitation by enabling the specification of autocorrelation detection up to an arbitrary lag order, denoted as p .

The flexibility of the Breusch-Godfrey test makes it the preferred diagnostic tool when analyzing time series data where complex dependency structures are common. By allowing the user to specify the maximum lag order p , the test provides a robust assessment against general forms of autocorrelation, ensuring that the researcher does not overlook higher-order serial dependencies that could compromise the integrity of the econometric results. Understanding and addressing these dependencies is crucial for generating reliable forecasts and drawing sound causal conclusions from the regression output.

Understanding the Hypotheses and Test Statistic

The Breusch-Godfrey test operates by setting up formal hypotheses that define the presence or absence of serial correlation up to the specified order p . This framework ensures a clear, quantitative basis for decision-making regarding the model's structure. The test utilizes the following formal hypotheses:

H0 (null hypothesis): There is no autocorrelation among the residuals at any order less than or equal to p . This hypothesis asserts that the model is statistically sound concerning the independence of errors.

HA (alternative hypothesis): There exists autocorrelation at some order less than or equal to p . This indicates a failure of the independence assumption, suggesting that the model is misspecified, potentially omitting important lagged relationships.

The computational basis of the test involves regressing the residuals from the original regression model onto the original predictor variables and the lagged residuals up to order p . The resulting test statistic, typically referred to as the LM statistic, follows a Chi-Square distribution with p degrees of freedom under the assumption that the null hypothesis is true. This statistic quantifies the collective explanatory power of the lagged residuals on the current residual term.

The decision to reject or fail to reject the null hypothesis hinges on the p-value associated with the calculated test statistic. If this p-value is less than the predetermined significance level (commonly set at 0.05 or 5%) then we possess sufficient statistical evidence to reject the null hypothesis. Rejecting H0 leads to the conclusion that significant autocorrelation exists among the residuals at one or more orders up to p . Conversely, if the p-value is greater than the significance level, we fail to reject H0, suggesting that the model residuals are acceptably independent within the tested lag range.

Prerequisites: Setting Up R and the lmtest Package

To successfully execute the Breusch-Godfrey test, the appropriate tools must be available within the R environment. The primary functionality for this test is encapsulated within the lmtest package, a robust collection of linear model diagnostic functions specifically designed for R. Before proceeding with the analysis, the user must ensure that this package is installed and then explicitly loaded into the current R session. Installation is a one-time process typically achieved using the `install.packages("lmtest")` command, provided the user has an active internet connection and the necessary permissions.

Once installed, the package must be activated using the `library(lmtest)` command. This action makes all functions within the package, including `bgtest()`, accessible to the user. Neglecting this

step will result in an error when attempting to call the test function, as R will not be able to locate the definition of `bgtest()`. The `lmtest` package is an essential component of econometric analysis in R, offering not just the Breusch-Godfrey test but also other vital diagnostics like the Durbin-Watson test and tests for heteroskedasticity.

The general syntax for performing the Breusch-Godfrey test in R is highly intuitive, following the standard formula interface common to many R statistical functions: `bgtest(formula, order = p, data = df)`. Here, `formula` represents the original regression model specification (e.g., $y \sim x_1 + x_2$), `order = p` specifies the maximum lag order for which to test autocorrelation, and `data = df` points to the dataframe containing the variables used in the model. This clear syntax allows for quick and accurate application of the diagnostic procedure to the target model.

Practical Example: Setting Up the Data in R

To illustrate the application of the Breusch-Godfrey test, we will first construct a small, representative dataset. This dataset will mimic typical time-series or cross-sectional data, consisting of a response variable and two predictor variables. Although the dataset is artificially generated, it serves the necessary purpose of establishing the structure upon which the subsequent regression and diagnostic tests will be performed. Establishing the data structure clearly is the foundational step in any statistical analysis within R.

We begin by defining the variables `x1`, `x2`, and `y`, and then consolidating them into an R dataframe named `df`. This process ensures that the variables are correctly aligned and labeled for use in the linear modeling function, `lm()`, which is implicitly used by `bgtest()` to derive the residuals. The initial creation and display of the data are essential checks to verify data integrity before moving to complex modeling stages, confirming the data type and format are appropriate for numerical analysis.

#create dataset

```
df <- data.frame(x1=c(3, 4, 4, 5, 8, 9, 11, 13, 14, 16, 17, 20),
x2=c(7, 7, 8, 8, 12, 4, 5, 15, 9, 17, 19, 19),
y=c(24, 25, 25, 27, 29, 31, 34, 34, 39, 30, 40, 49))
```

```
#view first six rows of dataset
```

```
head(df)
```

```
x1 x2 y
1 3 7 24
2 4 7 25
3 4 8 25
4 5 8 27
```

```
5 8 12 29
```

```
6 9 4 31
```

The code segment above demonstrates the construction of the dataframe and provides a quick verification using the `head()` function, confirming the structure of the data: three numerical columns corresponding to one response variable (y) and two independent variables (x_1 and x_2). With this data frame established, we are now prepared to estimate the linear regression model that will serve as the basis for the Breusch-Godfrey diagnostic test. The test operates directly on the residuals derived from this specified model.

Executing the Breusch-Godfrey Test in R

Once the data is prepared and the `lmtest` package is successfully loaded, the execution of the diagnostic test is straightforward. We apply the `bgtest()` function, specifying the formula that defines the regression relationship ($y \sim x_1 + x_2$) and critically, the order of autocorrelation we wish to check. For this specific demonstration, we will set the order parameter to `p = 3`. This means the test will evaluate whether the current residual is significantly correlated with residuals lagged up to three periods prior (i.e., lag 1, lag 2, and lag 3).

Choosing the correct lag order p is often informed by theoretical considerations or visual inspection of the residuals' Autocorrelation Function (ACF) plot. If the data is quarterly, p might be set to 4 to capture seasonal effects; if the relationship is expected to decay quickly, a lower p might suffice. For this example, setting `order=3` provides a comprehensive check beyond the basic first-order test. The R commands below perform the necessary setup and execute the test against the specified linear model.

```
#load lmtest package
```

```
library(lmtest)
```

```
#perform Breusch-Godfrey test
```

```
bgtest(y ~ x1 + x2, order=3, data=df)
```

```
Breusch-Godfrey test for serial correlation of order up to 3
```

```
data: y ~ x1 + x2
```

```
LM test = 8.7031, df = 3, p-value = 0.03351
```

The output provides a clear summary of the diagnostic results. It explicitly confirms that the test examined serial correlation up to order 3. The key elements are the LM test statistic, which is the calculated value of the test, and the associated degrees of freedom (df), which equals the specified

order p . Most critically, the output provides the p-value, which dictates the statistical decision concerning the null hypothesis.

Interpreting the Breusch-Godfrey Results

The interpretation phase is where the statistical output translates into practical conclusions regarding the quality of the regression model. From the output generated in the previous step, we observed that the LM test statistic is $X^2 = 8.7031$, calculated with 3 degrees of freedom. The corresponding p-value is determined to be **0.03351**. This p-value is the probability of observing a test statistic as extreme as 8.7031, assuming the null hypothesis (no autocorrelation up to order 3) is true.

The standard threshold for the significance level, denoted as α , is typically 0.05. Since the calculated p-value of 0.03351 is demonstrably less than 0.05, we must adhere to the rule of statistical inference: we reject the null hypothesis (H_0). The rejection of H_0 signifies a crucial finding: we conclude that significant serial correlation exists among the model residuals at some lag order less than or equal to 3. This result suggests that the error terms are systematically related across time, violating a core assumption of OLS estimation, and necessitating corrective action.

A rejection of the null hypothesis in the Breusch-Godfrey test implies that the standard errors of the regression coefficients are likely underestimated, leading to inflated t-statistics and potentially erroneous conclusions about the statistical significance of the predictor variables. The presence of autocorrelation indicates model misspecification, often suggesting that relevant lagged variables, time trends, or other dynamic elements have been omitted from the functional form. Recognizing this issue is the first critical step toward building a more robust and statistically valid model.

Strategies for Addressing Autocorrelation

Detecting autocorrelation through the Breusch-Godfrey test is only half the battle; the subsequent and more challenging step is implementing corrective measures to mitigate its impact. If the null hypothesis of no autocorrelation is rejected, researchers have several established econometric techniques to address the serial correlation present in the residuals. The most appropriate strategy often depends on whether the correlation is positive, negative, or seasonal in nature.

Specific remedial actions include:

For **positive serial correlation**, which often signals a missing dynamic component, consider augmenting the model by adding lags of the dependent variable and/or independent variables to capture the time-dependent relationship.

For **negative serial correlation**, which can sometimes arise from over-correction or "overdifferencing," check the data preparation steps carefully to make sure that none of your variables are excessively transformed, potentially reverting to a less aggressive model specification.

For **seasonal correlation**, which occurs at fixed intervals (e.g., lag 4 for quarterly data), consider adding seasonal dummy variables to the model or incorporating seasonal lags to absorb the periodic effects.

Alternatively, utilize robust standard errors, such as Heteroskedasticity and Autocorrelation Consistent (HAC) estimators (e.g., Newey-West), which allow for valid inference even when autocorrelation persists, provided the underlying OLS coefficient estimates remain consistent.

Ultimately, the goal is to refine the regression model until the Breusch-Godfrey test indicates that the assumption of independent residuals is statistically acceptable. This iterative process of diagnosis and remediation is central to high-quality econometric analysis.