How to Perform a Durbin-Watson Test in R to Detect Autocorrelation

Authored by
stats writer

March 10, 2026

RECOMMENDED CITATION

stats writer (2026). *How to Perform a Durbin-Watson Test in R to Detect Autocorrelation*. PSYCHOLOGICAL SCALES. Retrieved from https://scales.arabpsychology.com/?p=135046

ARABPSYCHOLOGY.COM

Foundations of the Durbin-Watson Test in Statistical Modeling

The **Durbin-Watson test** serves as a fundamental diagnostic tool in the field of econometrics and statistics, specifically designed to detect the presence of **autocorrelation** within the residuals of a regression analysis. **Autocorrelation**, often referred to as serial correlation, occurs when the error terms in a model are not independent of one another, meaning the value of an error at one point in time is related to the value of an error at another point. This phenomenon is particularly prevalent in **time series analysis**, where data points are collected at successive time intervals, but it can also emerge in cross-sectional data if the observations possess a specific spatial or logical ordering.

By utilizing the **Durbin-Watson test**, researchers can evaluate whether the assumptions of **Ordinary Least Squares (OLS)** regression are being met. One of the core requirements for **OLS** to provide the Best Linear Unbiased Estimator (BLUE) is that the **residuals** must be uncorrelated. If this assumption is violated, the standard errors of the coefficients may be underestimated, leading to misleadingly high t-statistics and potentially incorrect conclusions regarding the significance of **independent variables**. Consequently, performing this test is a non-negotiable step for any rigorous **statistical** workflow involving linear models.

In the **R** programming environment, the test is highly accessible through various libraries, most notably the **car package** and the **lmtest package**. These tools allow users to calculate the Durbin-Watson statistic, which typically ranges from 0 to 4. A value near 2 suggests no **autocorrelation**, while values approaching 0 indicate positive correlation and values approaching 4 indicate negative correlation. Understanding these nuances is essential for ensuring the **validity** of your predictive or explanatory models.

The Critical Role of Residual Independence in Linear Regression

To appreciate the necessity of the **Durbin-Watson test**, one must first understand the broader context of **linear regression** assumptions. When we construct a model to explain the relationship between a **dependent variable** and one or more predictors, we assume that the **residuals** (the differences between observed and predicted values) represent random noise. This randomness implies that knowing the error for one observation provides no information about the error for the next. When this independence is lost, the model fails to capture some systematic pattern in the data, which is then "leaked" into the **residuals**.

When **autocorrelation** is present, it often signals that the model is misspecified. This could be due to the omission of an important **independent variable**, such as a trend or seasonal factor, or because the functional form of the relationship is non-linear. Ignoring serial correlation can lead to inefficient estimates, where the variance of the coefficient estimates is larger than it needs to be.

ARABPSYCHOLOGY.COM

More dangerously, it can result in biased estimates of the standard errors, which invalidates hypothesis tests and **confidence intervals**, potentially leading a researcher to claim a relationship exists when it does not.

The **Durbin-Watson test** specifically targets first-order **autocorrelation**, meaning it checks if an error at time *t* is correlated with the error at time *t-1*. While it does not detect higher-order correlations, it remains the most popular diagnostic for initial checks. By ensuring **residual independence**, the analyst can be more confident that the model's parameters are reliable and that the inferential statistics derived from the **linear regression** are mathematically sound.

Formulating the Hypotheses for the Durbin-Watson Procedure

Every **statistical test** begins with a clear formulation of hypotheses, and the **Durbin-Watson test** is no exception. Before running the code in **R**, it is vital to understand what the **p-value** actually represents. The test is structured around two competing claims regarding the nature of the **residuals** within the fitted model. These claims allow us to use **probability** to decide whether any observed correlation is likely due to chance or represents a genuine pattern.

The **null hypothesis** (H0) for this test states that there is no correlation among the **residuals**. In mathematical terms, this implies that the **autocorrelation** coefficient (rho) is equal to zero. If the **null hypothesis** holds true, we can assume that the errors are independent and that the standard regression results are valid. This is the desired outcome for most researchers, as it simplifies the interpretation of the model.

Conversely, the **alternative hypothesis** (HA) posits that the **residuals** are **autocorrelated**. Depending on the specific software implementation, this can be a two-sided test (rho is not equal to zero) or a one-sided test (rho is greater than zero for positive correlation, or less than zero for negative correlation). In the **R** output, a low **p-value** (typically below 0.05) provides evidence to reject the **null hypothesis**, suggesting that the model suffers from serial correlation issues that must be addressed.

Data Preparation and Regression Modeling in R

Before we can execute the **Durbin-Watson test**, we must have a **linear model** to evaluate. For the purposes of this tutorial, we will utilize the **mtcars dataset**, a classic dataset built into **R** that contains various performance characteristics for 32 automobiles. This dataset is ideal for demonstrating **regression** techniques because it contains clear numerical variables that are often related in a linear fashion.

Our objective is to model the **dependent variable** **mpg** (miles per gallon) using two **independent variables**: **disp** (displacement) and **wt** (weight). By fitting this **linear regression**, we aim to see

APABPSYCHOLOGY.COM

how well these engine and physical characteristics predict fuel efficiency. The first step involves loading the data and inspecting its structure to ensure there are no missing values or anomalies that could skew the **residuals**.

Once the data is ready, we use the `lm()` function to create our model object. This object contains all the necessary information about the coefficients, fitted values, and **residuals**. The **Durbin-Watson test** will eventually be applied directly to this model object to check for any patterns remaining in the unexplained variance. Below is the **R** code used to prepare the dataset and fit the initial **linear regression** model:

```
#load mtcars dataset
data(mtcars)

#view first six rows of dataset
head(mtcars)

mpg cyl disp hp drat wt qsec vs am gear carb
Mazda RX4 21.0 6 160 110 3.90 2.620 16.46 0 1 4 4
Mazda RX4 Wag 21.0 6 160 110 3.90 2.875 17.02 0 1 4 4
Datsun 710 22.8 4 108 93 3.85 2.320 18.61 1 1 4 1
Hornet 4 Drive 21.4 6 258 110 3.08 3.215 19.44 1 0 3 1
Hornet Sportabout 18.7 8 360 175 3.15 3.440 17.02 0 0 3 2
Valiant 18.1 6 225 105 2.76 3.460 20.22 1 0 3 1

#fit regression model
model <- lm(mpg ~ disp+wt, data=mtcars)
```

Implementing the Durbin-Watson Test via the car Package

With our regression model successfully fitted, we can now proceed to the diagnostic phase. While there are multiple ways to perform the **Durbin-Watson test** in **R**, the `durbinWatsonTest()` function from the **car package** is widely regarded as one of the most robust and user-friendly options. This function not only calculates the test statistic but also uses **bootstrapping** techniques to provide a **p-value**, which is essential for determining **statistical significance**.

To use this function, you must first ensure that the **car package** is installed and loaded into your **R** session. The function takes the model object as its primary argument. It then analyzes the **residuals** of that model to calculate the D-W statistic. This statistic is defined as the sum of squared differences between adjacent **residuals** divided by the **residual sum of squares**.

Running the test is straightforward. The output will provide the calculated **autocorrelation** (rho),

ARABPSYCHOLOGY.COM

the D-W statistic, and the **p-value**. It is important to note that the function also specifies the **alternative hypothesis** being tested--usually that rho is not equal to zero. Below is the syntax for loading the necessary library and executing the test on our previously created model:

#load car package
library(car)

#perform Durbin-Watson test
durbinWatsonTest(model)

Loading required package: carData
lag Autocorrelation D-W Statistic p-value
1 0.341622 1.276569 0.034
Alternative hypothesis: rho != 0

Analyzing the Statistical Output and p-values

Interpreting the results of the **Durbin-Watson test** requires a careful look at both the test statistic and the **p-value**. In our specific example using the **mtcars** data, the output shows a D-W Statistic of approximately **1.2766**. Since this value is considerably lower than 2, it suggests the presence of positive **autocorrelation**, where a positive residual for one car tends to be followed by another positive residual for the next observation in the dataset.

The most critical component of the output is the **p-value**, which is reported as **0.034**. In the context of **hypothesis testing**, we compare this value to a predetermined significance level, typically alpha = 0.05. Because 0.034 is less than 0.05, we have sufficient evidence to reject the **null hypothesis**. This leads us to the conclusion that the **residuals** in our regression model are indeed **autocorrelated**, and the assumption of independence has been violated.

While a **p-value** of 0.034 is significant, it is also useful to look at the "lag 1 **Autocorrelation**" estimate, which is 0.3416. This value quantifies the strength of the relationship between consecutive **residuals**. A value of 0.34 indicates a moderate positive correlation. When such a result is found, the researcher must decide whether the degree of correlation is "serious enough" to warrant corrective measures or if the model can still be used with caution, perhaps by employing robust standard errors.

Practical Solutions for Correcting Detected Autocorrelation

If the **Durbin-Watson test** indicates that **autocorrelation** is a problem, you have several strategies to improve your model's **validity**. The choice of solution depends largely on the nature of the correlation detected and the type of data being analyzed. Addressing these issues is vital for

ARABPSYCHOLOGY.COM

ensuring that your **regression analysis** remains credible and accurate.

For positive **serial correlation**, which is the most common type, you should consider adding **lags** of the **dependent variable** or the **independent variables** to the model. This allows the model to account for the temporal or sequential dependency directly.
For negative **serial correlation**, which is rarer but possible, you should investigate whether any of your variables are **overdifferenced**. Overdifferencing can introduce artificial patterns into the **residuals** that were not present in the original data.
For correlation that follows a specific pattern over time, such as **seasonal correlation**, consider adding **seasonal dummy variables**. These variables help capture cyclical fluctuations that are not explained by the other predictors in the model.

Beyond these steps, you might also consider using **Generalized Least Squares (GLS)** instead of **OLS**. **GLS** is specifically designed to handle situations where the **residuals** have a known correlation structure. Alternatively, calculating **Newey-West standard errors** can provide **autocorrelation**-consistent estimates, allowing you to keep your original coefficients while ensuring your hypothesis tests remain valid.

Exploring Advanced Diagnostic Alternatives in R

While the **Durbin-Watson test** is a powerful and widely used tool, it is not the only diagnostic available for checking **autocorrelation**. Advanced users often supplement the D-W test with the **Breusch-Godfrey test**, which is more flexible. Unlike the Durbin-Watson, the **Breusch-Godfrey test** can detect higher-order **autocorrelation** (lags greater than 1) and is valid even when the model includes lagged **dependent variables**.

In **R**, the `bgtest()` function from the **lmtest package** can be used to perform this more comprehensive check. Furthermore, visualizing the **residuals** is often just as important as the formal tests. Plotting the **autocorrelation** function (ACF) of the **residuals** using the `acf()` function provides a clear graphical representation of any remaining patterns in the data across various lags.

In conclusion, the **Durbin-Watson test** is an essential first step in validating a **linear regression** model. By identifying **autocorrelation** early, you can take the necessary steps to refine your model, whether through variable selection, data transformation, or the use of more advanced **statistical** estimators. Maintaining a rigorous diagnostic routine ensures that your findings are not just **statistically significant**, but also robust and reliable for real-world decision-making.

ARABPSYCHOLOGY.COM