

Breusch-Pagan Test in R: Detecting Heteroscedasticity in Regression Models

Authored by
Mohammed Iooti

November 8, 2025

RECOMMENDED CITATION

Mohammed Iooti (2025). *Breusch-Pagan Test in R: Detecting Heteroscedasticity in Regression Models*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=13380>

The [Breusch-Pagan Test](#) stands as an indispensable diagnostic instrument in modern quantitative research, especially within the field of [regression analysis](#). Its primary purpose is to formally detect the presence of [heteroscedasticity](#)--a serious violation of the core assumptions underpinning classical linear models. A foundational requirement for efficient Ordinary Least Squares (OLS) estimation is **homoscedasticity**, meaning the variance of the error terms (or residuals) must remain constant across all levels of the independent variables.

When the variance of the errors is not constant, the issue is termed **heteroscedasticity**. This violation does not bias the coefficient estimates themselves, but it severely compromises the reliability of statistical inference by producing biased [standard errors](#) and, consequently, unreliable t-statistics and F-statistics. If [heteroscedasticity](#) is ignored, the traditional OLS approach loses its desirable property of efficiency, rendering p-values and confidence intervals untrustworthy for hypothesis testing. This detailed guide provides expert instruction on executing the [Breusch-Pagan Test](#) within the R statistical environment, ensuring practitioners can validate their models effectively.

The Core Problem: Understanding Heteroscedasticity

In statistical modeling, particularly when analyzing complex datasets such as economic, financial, or large-scale sociological data, the assumption of homoscedasticity is often tenuous and easily violated. [Heteroscedasticity](#) fundamentally implies that the spread of the residuals around the regression line systematically changes as the predictor variables change. A classic example arises in income modeling: the variability (variance) of savings or discretionary spending habits (the error term) might be consistently low for individuals with low incomes but may expand dramatically for high-income individuals who possess a wider range of spending options. This systematic pattern fundamentally violates the underlying assumption required for reliable OLS [regression analysis](#).

The critical consequence of this violation lies in the realm of statistical inference. Although the calculated coefficient estimates (the Betas) remain consistent and unbiased even if heteroscedasticity is present, the calculated variances of these estimates--the foundation for computing [standard errors](#)--become inaccurate. If the standard errors are flawed, any subsequent calculations relying on them, including t-statistics, F-statistics, and confidence intervals, will also be flawed. This can easily lead researchers to commit Type I or Type II errors, resulting in the incorrect rejection or failure to reject the [null hypothesis](#), thereby undermining the validity of the study's conclusions.

Because subjective visual inspection of residual plots can often be misleading, the [Breusch-Pagan Test](#) was specifically designed to offer an objective, formal, and quantifiable method for assessing the constancy of the error variance. It establishes a rigorous test based on an auxiliary regression, which determines whether the variance of the errors is systematically related to the explanatory

variables. This objective evidence is essential for rigorous statistical reporting and ensuring model adequacy.

Mechanism of the Breusch-Pagan Test

The mechanics of the [Breusch-Pagan Test](#) revolve around examining the correlation between the squared residuals derived from the initial OLS model and the independent variables. The testing procedure involves three critical steps: First, the researcher fits the standard OLS regression model to the data. Second, the residuals from this primary model are extracted and squared (\hat{u}^2). Third, an auxiliary regression is performed where these squared residuals are regressed onto the original independent variables. The core logical premise is simple: if the assumption of homoscedasticity holds true (i.e., the variance is constant), the independent variables should collectively possess zero explanatory power over the magnitude of the squared error terms.

The resulting test statistic, often referred to as BP, is calculated directly from the R-squared value of this auxiliary regression and is scaled appropriately. Under the [null hypothesis](#), this BP statistic follows a Chi-squared distribution. For the Breusch-Pagan Test, the [null hypothesis](#) (H_0) asserts that the variances of the errors are constant (homoscedasticity is present). Conversely, the alternative hypothesis (H_a) posits that the error variances are not constant (signifying the presence of [heteroscedasticity](#)).

A significant test statistic, resulting in a low p-value (typically below the 0.05 significance level), provides strong evidence against the [null hypothesis](#) of homoscedasticity. Such a finding suggests that the independent variables are indeed systematically related to the variation in the squared residuals. Conversely, a small BP test statistic and a high p-value indicate a lack of systematic relationship between the predictors and the error variance, thereby confirming that the assumption of homoscedasticity is likely valid for the model. This clear decision framework is vital for diagnosing model adequacy before proceeding to interpret the primary regression coefficients.

Practical Example: Executing the Test in R

To solidify understanding, we will now walk through the execution of the [Breusch-Pagan Test](#) using R. This demonstration utilizes the familiar built-in R dataset, `mtcars`, which compiles specifications for 32 different automobiles. We will fit a simple multiple linear regression model and then apply the dedicated `bptest` function, which is conveniently located within the essential `lmtest` library, to formally check for non-constant variance. This standardized procedure is a requirement for validating linear models in robust statistical practice.

Step 1: Fitting the Linear Regression Model

Our initial requirement is to establish the specific linear model that necessitates diagnostic testing. For this example, we designate the fuel efficiency in miles per gallon (**mpg**) as the response variable. We select engine displacement (**disp**) and gross horsepower (**hp**) as our explanatory variables, seeking to model how these mechanical characteristics influence fuel efficiency. Once the model is successfully fitted using R's base **lm()** function, it is standard procedure to inspect the summary output to confirm basic fit details, including the coefficient estimates, their associated [standard errors](#), and overall model statistics.

#load the dataset

```
data(mtcars)
```

```
#fit a regression model
```

```
model <- lm(mpg~disp+hp, data=mtcars)
```

```
#view model summary
```

```
summary(model)
```

Coefficients:

```
Estimate Std. Error t value Pr(>|t|)
```

```
(Intercept) 30.735904 1.331566 23.083 < 2e-16 ***
```

```
disp -0.030346 0.007405 -4.098 0.000306 ***
```

```
hp -0.024840 0.013385 -1.856 0.073679 .
```

```
---
```

```
Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
Residual standard error: 3.127 on 29 degrees of freedom
```

```
Multiple R-squared: 0.7482, Adjusted R-squared: 0.7309
```

```
F-statistic: 43.09 on 2 and 29 DF, p-value: 2.062e-09
```

Step 2: Running the Breusch-Pagan Diagnostic

With the linear model successfully estimated, the subsequent and vital step is to formally execute the diagnostic test. This requires loading the **lmtest** library, which provides access to the powerful **bptest()** function. We apply this function directly to the fitted regression object (`model`). The R function automatically handles the complex internal steps--calculating the squared residuals, running the auxiliary regression, and computing the final BP test statistic and associated p-value--thereby offering a definitive, quantitative assessment of whether the homoscedasticity assumption has been maintained.

#load lmtest library

```
library(lmtest)
```

```
#perform Breusch-Pagan Test
```

```
bptest(model)
```

```
studentized Breusch-Pagan test
```

```
data: model
```

```
BP = 4.0861, df = 2, p-value = 0.1296
```

Interpreting the Statistical Results

The output generated by the **bptest()** function yields the necessary statistical metrics for informed decision-making. In the example provided, the calculated Breusch-Pagan test statistic (BP) is reported as **4.0861**, paired with 2 degrees of freedom (df). Most critically, the resulting p-value is calculated to be **0.1296**. The process of interpreting this result revolves around comparing the p-value against a chosen significance level (α), which is conventionally set at 0.05 (or 5%) for most scientific and social science applications.

It is crucial to recall that the [null hypothesis](#) (H_0) for this test assumes that the model exhibits homoscedasticity (constant error variance). The decision rule is unambiguous: if the computed p-value is less than the predetermined α (0.05), we must reject the null hypothesis, leading us to conclude that significant [heteroscedasticity](#) is present. Conversely, if the p-value is greater than or equal to α , we formally fail to reject the null hypothesis, indicating that there is insufficient statistical evidence to conclude that the error variance is non-constant.

Applying this rule to our `mtcars` regression example, the p-value of **0.1296** is distinctly greater than the 0.05 threshold. Consequently, we formally fail to reject the [null hypothesis](#). The firm conclusion is that we lack sufficient statistical evidence to assert the presence of [heteroscedasticity](#) in this linear [regression analysis](#). Because the assumption of homoscedasticity is upheld, the standard OLS estimates and their associated [standard errors](#) can be considered reliable, allowing the researcher to proceed confidently with the interpretation of the coefficients from the original model summary.

Remedial Strategies for Addressing Non-Constant Variance

Had the [Breusch-Pagan Test](#) led to the rejection of the null hypothesis (i.e., p-value < 0.05), confirming the presence of significant heteroscedasticity, immediate corrective measures would be mandatory. Ignoring this issue means that while the coefficient estimates themselves remain unbiased, the standard errors shown in the regression output would be invalid, leading to

fundamentally inaccurate hypothesis testing and potentially misleading conclusions. Statisticians have developed several robust techniques to mitigate or correct for heteroscedasticity, thereby preserving the integrity of inferential statistics.

These solutions generally fall into two categories: altering the structure of the variables to stabilize variance, or employing advanced estimation techniques that explicitly account for the non-constant variance structure. The selection of the appropriate method typically depends on the underlying data characteristics and the source of the detected heteroscedasticity. Below are the principal strategies employed to resolve this critical modeling challenge:

Transforming the Response Variable.

A fundamental approach is the application of a mathematical transformation to the dependent (response) variable in an attempt to stabilize the error variance across the range of predictors. The logarithmic transformation, where the natural log of the response variable is used in place of the original variable, is highly effective and widely used. Log transformation often compresses the scale of higher values--which are frequently responsible for the largest variance--thereby successfully reducing or eliminating heteroscedasticity. Other common transformations include using the square root or the inverse of the response variable. This method is generally preferred when the source of heteroscedasticity is intrinsic to the scale of the response variable itself.

Using Weighted Least Squares (WLS) Regression.

A more advanced estimation method is the utilization of [weighted regression](#), or **Weighted Least Squares (WLS)**. Unlike OLS, which assumes all data points are equally reliable, WLS assigns a specific weight to each observation based on the estimated variance of its error term. Data points associated with higher variance (the source of the heteroscedasticity) receive smaller weights, effectively diminishing the influence of their squared residuals during the parameter optimization process. This selective weighting ensures that observations with more precise measurements (smaller variances) exert a greater influence on the final parameter estimates. When the correct weights--which are inversely proportional to the error variance--are applied, WLS produces the Best Linear Unbiased Estimator (BLUE), thereby resolving the efficiency loss inherent to heteroscedasticity.

Employing Heteroscedasticity-Consistent Standard Errors (HCSE).

The most straightforward and often most robust solution in contemporary econometrics is the use of **Heteroscedasticity-Consistent Standard Errors**, commonly referred to as White [Standard Errors](#) or robust standard errors. This technique is unique in that it does not attempt to eliminate the heteroscedasticity in the residuals; instead, it provides a corrected formula for calculating the standard errors that remains statistically valid even when the variance is non-constant. By

employing robust standard errors, the researcher can maintain the original OLS coefficient estimates while guaranteeing that subsequent hypothesis tests and confidence intervals are statistically sound. Due to its simplicity and generality, most statistical software packages, including R, offer easy implementation of robust standard errors, making this the preferred method when the primary concern is the reliability of statistical inference.