

How to Perform the Goldfeld-Quandt Test in R?

Authored by
stats writer

December 16, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Perform the Goldfeld-Quandt Test in R?*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=107595>

The Goldfeld-Quandt Test is a foundational econometric procedure designed to verify a crucial assumption in statistical modeling: the homogeneity of variance. Specifically, it determines whether a fitted linear regression model exhibits heteroscedasticity (unequal variance of errors) or homoscedasticity (equal variance of errors). In the R programming environment, this test is executed using the `gqtest()` function, which is part of the robust lmtest package. The test requires a pre-fitted linear model and returns a calculated test statistic alongside a critical p-value. This p-value dictates the conclusion: if it is statistically significant, we reject the null hypothesis, concluding that the variance of the residuals is not constant, thereby confirming the presence of heteroscedasticity.

The **Goldfeld-Quandt test** is a powerful tool used to confirm the presence of heteroscedasticity in a regression model. Understanding this phenomenon is essential for maintaining the integrity of statistical inference.

Heteroscedasticity refers to the situation where the scatter of residuals is unequal at different levels of a predictor variable in a regression model. This means the variability of the errors is not uniform across the range of observations.

If heteroscedasticity is identified, it violates one of the core assumptions of the Classical Linear Model (CLM), which postulates that the residuals must be equally scattered (homoscedastic) at each level of the response variable. Violation of this assumption can lead to inefficient parameter estimates and unreliable standard errors, jeopardizing the validity of hypothesis testing.

This detailed guide provides a practical, step-by-step example of how to implement the Goldfeld-Quandt test in R to conclusively determine whether or not heteroscedasticity is present in a given linear regression model, ensuring your statistical conclusions are sound.

The Role of the Goldfeld-Quandt Test

The Goldfeld-Quandt test operates by dividing the dataset into two distinct segments, excluding a small portion of central observations. It then compares the variance of the residuals calculated from the two remaining subsets. This comparison is facilitated using an F-statistic, which essentially tests the ratio of the residual sum of squares (RSS) from the larger variance group to the RSS of the smaller variance group.

The primary purpose of identifying variance issues is to ensure the estimated coefficients are the Best Linear Unbiased Estimators (BLUE). When variance is unequal, the standard errors of the regression coefficients become biased and inconsistent. Consequently, the confidence intervals and t-tests derived from the model summary become untrustworthy, potentially leading to incorrect conclusions about the significance of the predictor variables.

Unlike graphical methods, which rely on subjective visual inspection of residual plots, the Goldfeld-Quandt test provides an objective, statistical measure (the p-value) for making a definitive judgment regarding the presence of non-constant error variance. This quantitative approach is crucial for rigorous academic and professional statistical analysis.

Understanding Model Assumptions and Violation

In a standard ordinary least squares (OLS) regression, the assumption of homoscedasticity is paramount. This assumption mathematically states that the expected value of the square of the error term, conditional on the independent variables, is constant: $E(u^2|X) = \sigma^2$. When this condition is met, the OLS estimator is both unbiased and efficient.

The presence of heteroscedasticity often occurs in cross-sectional data, particularly when modeling economic or financial phenomena where the scale of the dependent variable differs significantly across observations. For instance, in modeling income, the variation in spending habits might be much higher for high-income earners than for low-income earners, causing the error variance to increase as income increases.

While heteroscedasticity does not introduce bias into the coefficient estimates themselves (they remain unbiased), it fundamentally compromises their efficiency and, critically, invalidates the standard error calculations. Therefore, performing tests like the Goldfeld-Quandt test is a necessary diagnostic step before interpreting the final model results.

Preparing the Data and Environment in R

To execute this test in R, two prerequisites must be met: first, a linear model must be successfully estimated using the `lm()` function, and second, the necessary R package containing the test function must be loaded. The `gqtest()` function resides within the **lmtest** package, which provides a comprehensive suite of diagnostic tests for linear models.

Before proceeding, ensure the **lmtest** package is installed and loaded into your R session. If it is not installed, the command `install.packages("lmtest")` should be run first. Loading the library is essential for accessing the `gqtest()` function without specifying the namespace, simplifying the syntax for the user.

For this tutorial, we will utilize the publicly available **mtcars** dataset, which is integrated into the R distribution. This dataset contains 32 observations on 11 variables related to performance and design aspects of 32 automobiles, providing a convenient basis for our regression analysis and subsequent diagnostic testing.

Step 1: Constructing the Linear Regression Model

Our first procedural step involves building a linear regression model. We will use the built-in **mtcars** dataset in R, modeling miles per gallon (`mpg`) as the dependent variable and predicting it based on two independent variables: engine displacement (`disp`) and horsepower (`hp`). This process initializes the residuals upon which the Goldfeld-Quandt test will operate.

The `lm()` function is the standard implementation for fitting linear models in R. By assigning the resulting model object to a variable (here, `model`), we capture all the necessary statistical outputs and parameters required by subsequent diagnostic functions like `gqtest()`.

After fitting the model, it is good practice to review the summary to confirm the coefficients and overall model fit (R-squared, F-statistic), although these values are preliminary until we verify the underlying assumptions regarding the error terms.

```
#fit a regression model using mpg, displacement, and horsepower
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: Executing the Goldfeld-Quandt Test

Following the construction of the model, we utilize the **gqtest()** function from the **lmtest** package to perform the Goldfeld-Quandt test itself. This function critically evaluates the assumption of homoscedasticity based on the variance split across the predictor variables.

The generalized syntax for the function is outlined below, highlighting the necessary arguments required for proper execution and ensuring the test is correctly ordered based on the variables

suspected of driving the variance inequality.

gqtest(model, order.by, data, fraction)

The parameters hold specific importance for the test mechanism:

model: This is the fitted linear regression model object generated by the `lm()` command, containing the residuals to be analyzed.

order.by: Specifies the predictor variable(s) by which the observations should be sorted before splitting. This is vital, as the test assumes that variance increases or decreases systematically along the axis of this specified variable.

data: Refers to the original dataset used to fit the model, ensuring the sorting and splitting are conducted on the correct observations.

fraction: Represents the number of central observations to remove from the dataset. This removal is crucial for maximizing the discriminatory power of the test, as the central observations often mask the difference in variance between the two extremes.

The core mechanism of the Goldfeld-Quandt test involves removing a specified `fraction` of observations located precisely in the center of the dataset. The test then compares the residual sums of squares (RSS) derived from the first and last segments of the data. The objective is to see if the spread of residuals in the first segment differs significantly from the spread in the second segment, indicating a pattern of variance change.

A common practice suggests removing approximately 20% of the total observations to achieve optimal test power. Since the `mtcars` dataset contains 32 total observations, removing 7 central observations (approximately 21.8%) is a statistically sound choice. We order the observations by both predictor variables, `disp` and `hp`, as both may contribute to the potential change in variance.

#load lmtest library

```
library(lmtest)
```

```
#perform the Goldfeld Quandt test, removing 7 central observations
```

```
gqtest(model, order.by = ~disp+hp, data = mtcars, fraction = 7)
```

Goldfeld-Quandt test

data: model

GQ = 1.0316, df1 = 10, df2 = 9, p-value = 0.486

alternative hypothesis: variance increases from segment 1 to 2

Interpreting the Results: Statistical Significance

The output of the `gqtest()` function provides the necessary components for statistical inference. We focus specifically on the calculated Goldfeld-Quandt statistic (GQ) and the corresponding p-value.

For our example, the key results are:

The test statistic (GQ) is **1.0316**, which represents the ratio of the estimated error variances from the two split groups.

The corresponding p-value is **0.486**.

The Goldfeld-Quandt test operates under the standard framework of null and alternative hypotheses:

Null Hypothesis (H0): Homoscedasticity is present (Variance of errors is constant).

Alternative Hypothesis (HA): Heteroscedasticity is present (Variance of errors is non-constant).

The decision rule requires comparing the p-value to a predetermined significance level (α), typically set at 0.05. Since our calculated p-value of 0.486 is substantially greater than 0.05, we consequently fail to reject the null hypothesis. This means we do not possess sufficient statistical evidence to conclude that heteroscedasticity is present in the regression model based on the variance pattern of `disp` and `hp`. We can proceed with interpreting the original OLS model with confidence in the reliability of the standard errors.

Addressing Heteroscedasticity: Remedial Measures

If, contrary to the result above, you were to reject the null hypothesis of the Goldfeld-Quandt test, it signals that heteroscedasticity is indeed present in the data. In this scenario, the standard errors reported in the regression output table are unreliable, making the t-statistics and associated p-values misleading for hypothesis testing.

Fortunately, several common econometric techniques can be employed to mitigate or eliminate the consequences of non-constant variance, ensuring the reliability of your model's inferences.

There are two primary approaches for dealing with this issue:

1. Employing Heteroscedasticity-Consistent Standard Errors (HCSE).

The most straightforward solution is often the use of robust standard errors, such as White's standard errors (HC3 in R). These errors adjust the variance-covariance matrix of the coefficient estimates to account for the presence of heteroscedasticity without requiring transformation of the

variables or changing the estimation method. The OLS coefficients remain the same, but the standard errors become consistent and reliable for inference, even when the underlying variance assumption is violated.

2. Transforming the Response Variable.

An alternative method involves modifying the structure of the model by performing a transformation on the response variable, such as taking the natural logarithm (log) of the dependent variable. Log transformations often compress the scale of the variable, which can stabilize the variance of the residuals across the range of the predictor variables, thereby satisfying the homoscedasticity assumption.

3. Using Weighted Least Squares (WLS) Regression.

Weighted regression is a method that explicitly assigns a weight to each data point based on the estimated variance of its fitted value. The core principle is to give smaller weights to observations that exhibit higher variance, effectively minimizing their contribution to the squared residuals used in the estimation process. When the correct weights (inversely proportional to the estimated error variance) are determined and applied, weighted regression can effectively resolve the problem of heteroscedasticity and restore the OLS estimator to its BLUE status.

Conclusion on Diagnostic Testing

The successful application and interpretation of the Goldfeld-Quandt test are fundamental steps in validating the integrity of any linear regression analysis. By systematically checking for non-constant variance using this F-test based procedure, researchers ensure that their calculated standard errors and, consequently, their statistical inferences regarding variable significance are robust and reliable. Integrating such diagnostic testing into the standard workflow guarantees a higher quality of econometric modeling.