

# How to perform a Likelihood Ratio Test in R?

Authored by  
**stats writer**

December 17, 2025

## RECOMMENDED CITATION

stats writer (2025). *How to perform a Likelihood Ratio Test in R?*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=107667>

The **Likelihood Ratio Test (LRT)** is a powerful statistical tool utilized widely in econometrics and biostatistics to compare the goodness of fit between two competing statistical regression models. Specifically, the LRT determines whether a more complex model, often referred to as the **full model**, provides a significantly better explanation of the observed data compared to a simpler, constrained model, known as the **reduced model**. This methodology is particularly relevant when researchers are engaged in model selection, aiming to achieve the most parsimonious representation of the underlying relationships without sacrificing predictive power. By evaluating the ratio of the maximized likelihoods of the two models, the test provides a formal decision-making framework.

The core principle behind the Likelihood Ratio Test rests on comparing the maximized likelihood function for both models. The likelihood function itself measures how probable the observed data are given the estimated parameters of the model. Intuitively, a model that fits the data well will have a higher maximized likelihood. The ratio calculated is always between zero and one; however, for computational simplicity and theoretical convenience, the test statistic is typically derived using the logarithm of this ratio, which follows an asymptotic **Chi-squared distribution** under the null hypothesis. This distribution allows us to calculate the probability of observing the test statistic, thus yielding the necessary p-value for hypothesis testing.

It is absolutely critical that the models being compared satisfy the condition of **nestedness**. A model is considered nested within another if it can be derived by imposing constraints--typically setting certain predictor coefficients to zero--on the more complex model. If the models are not nested, the standard LRT framework cannot be validly applied, and alternative model comparison techniques, such as the **Akaike Information Criterion (AIC)** or **Bayesian Information Criterion (BIC)**, would be more appropriate. Ensuring proper structure is the first step toward a statistically sound analysis using the LRT, guaranteeing that the comparison reflects the incremental explanatory power of the additional variables.

## Understanding Nested Models in Regression Analysis

A pair of regression models are considered **nested** if the simpler model is a special, constrained case of the more complex model. This means that the set of predictor variables used in the reduced model must be a strict subset of the predictor variables utilized in the full model. The simplest way to visualize this concept is through a comparison of their respective equations, demonstrating how the reduced model is formed by effectively forcing one or more parameters in the full model to be zero. This constraint is the mathematical basis for the comparison performed by the Likelihood Ratio Test.

Consider a scenario involving a prediction task where four potential predictor variables ( $x_1$ ,  $x_2$ ,  $x_3$ , and  $x_4$ ) are available. The **full model** would incorporate all these variables to explain the

dependent variable  $Y$ , maximizing its potential explanatory power. This is represented mathematically as:

$$Y = \beta_0 + \beta_1x_1 + \beta_2x_2 + \beta_3x_3 + \beta_4x_4 + \varepsilon$$

A corresponding **reduced model**, which is nested within the full model, would intentionally omit specific predictor variables. For instance, if we hypothesize that  $x_3$  and  $x_4$  do not significantly contribute to the prediction of  $Y$ , we would construct a reduced model that only includes  $x_1$  and  $x_2$ . This is equivalent to setting the coefficients for  $x_3$  and  $x_4$  (i.e.,  $\beta_3$  and  $\beta_4$ ) to zero in the full model equation. The resulting nested model is:

$$Y = \beta_0 + \beta_1x_1 + \beta_2x_2 + \varepsilon$$

The primary objective of the LRT, in this context, is to formally test whether the simplification achieved by moving from the full model to the reduced model--that is, the constraint that  $\beta_3 = \beta_4 = 0$ --leads to a statistically significant decrease in the overall model fit. If the fit difference is negligible, we should select the reduced model due to its greater simplicity and interpretability, adhering to the principle of **parsimony**. If the difference is significant, however, the full model is preferred.

## Formulating the Null and Alternative Hypotheses

The purpose of the Likelihood Ratio Test is to provide a statistical measure for determining whether the inclusion of the extra parameters in the full model is justified. This comparison is strictly formalized through the definition of the null ( $H_0$ ) and alternative ( $H_A$ ) hypotheses. Understanding these hypotheses is crucial, as they dictate the interpretation of the calculated test statistic and the resulting p-value. The hypotheses structure directly addresses the question of whether the additional complexity of the full model is worth the effort.

The **Null Hypothesis ( $H_0$ )** posits that the constraints imposed to create the reduced model are valid. Specifically,  $H_0$  states that the predictor variables excluded from the reduced model have coefficients that are effectively zero, and thus, they contribute no significant explanatory power. Consequently, under  $H_0$ , the overall fit provided by the full model is statistically indistinguishable from the fit provided by the nested model. This leads to the conclusion that if we fail to reject  $H_0$ , we should favor the simpler, **nested model**.

**$H_0$ :** The full model and the nested model fit the data equally well. The additional parameters in the full model are unnecessary. Thus, you should **use the nested model**.

Conversely, the **Alternative Hypothesis ( $H_A$ )** claims that the extra parameters in the full model are indeed statistically significant and necessary for accurately modeling the data.  $H_A$  asserts that removing these variables (as done in the reduced model) leads to a substantial and measurable

degradation in the model's goodness of fit. Therefore, if we reject  $H_0$  in favor of  $H_A$ , we conclude that the increased complexity is justified, and the **full model** is the preferred choice for prediction or inference.

**HA:** The full model fits the data significantly better than the nested model. The additional parameters provide valuable explanatory power. Thus, you should **use the full model**.

## Preparing the R Environment for Model Comparison

To demonstrate the practical application of the Likelihood Ratio Test, we utilize R statistical software, a widely used platform for statistical computing. Prior to running the test, it is essential to load the necessary packages and prepare the data. The function required for performing the LRT on linear models in R is `lrtest()`, which is part of the **lmtest** package. We must first ensure this package is installed and loaded into the current R session.

For this tutorial, we will rely on the well-known **mtcars** dataset, which is built into R and contains information about various automobile characteristics. We will formulate two specific regression models designed to predict fuel efficiency (miles per gallon, mpg) based on different sets of predictors. The response variable mpg will be modeled using variables such as engine displacement (`disp`), carburetor count (`carb`), horsepower (`hp`), and cylinder count (`cyl`).

We define our two competing models, ensuring the condition of nestedness is met: The **Full Model** includes four predictors, while the **Reduced Model** includes only two. This configuration allows us to formally test whether the extra variables (horsepower and cylinders) significantly improve the model's ability to predict mpg.

**Full Model (MF):**  $\text{mpg} = \beta_0 + \beta_1\text{disp} + \beta_2\text{carb} + \beta_3\text{hp} + \beta_4\text{cyl}$

**Reduced Model (MR):**  $\text{mpg} = \beta_0 + \beta_1\text{disp} + \beta_2\text{carb}$

The test will specifically assess the joint hypothesis that  $\beta_3$  (for `hp`) and  $\beta_4$  (for `cyl`) are simultaneously equal to zero. If this hypothesis holds true, the reduced model is sufficient. We will use the `lrtest()` function from the **lmtest** package to perform the likelihood ratio test on these two models.

## Example 1: Performing the LRT in R using **lmtest**

Once the hypotheses are defined and the models are structured, we proceed to fit both models using the standard `lm()` function in R statistical software. It is important to note that the LRT assumes that the errors are independently and identically distributed, typically following a normal distribution. We assign the fitted models to distinct objects, `model_full` and `model_reduced`, allowing the `lrtest()` function to easily compare their statistical characteristics.

The following sequence of code illustrates the fitting process and the execution of the Likelihood Ratio Test using the `lrtest()` function, providing the full model first, followed by the nested model, as required by the function's syntax.

### **library(lmtest)**

```
#fit full model
```

```
model_full <- lm(mpg ~ disp + carb + hp + cyl, data = mtcars)
```

```
#fit reduced model
```

```
model_reduced <- lm(mpg ~ disp + carb, data = mtcars)
```

```
#perform likelihood ratio test for differences in models
```

```
lrtest(model_full, model_reduced)
```

Likelihood ratio test

Model 1: mpg ~ disp + carb + hp + cyl

Model 2: mpg ~ disp + carb

#Df LogLik Df Chisq Pr(>Chisq)

1 6 -77.558

2 4 -78.603 -2 2.0902 0.3517

The output provides key diagnostic statistics necessary for interpretation. The difference in degrees of freedom (Df) is 2, corresponding to the two variables (hp and cyl) that were dropped in the reduced model. The Chisq value is the calculated test statistic, which is then used to determine the associated p-value (Pr(>Chisq)).

### **Interpreting the Likelihood Ratio Test Results**

The final and most critical step in performing a Likelihood Ratio Test is interpreting the resulting p-value. The p-value, reported as Pr(>Chisq) in the R output, quantifies the probability of observing a difference in model fit as large as (or larger than) the one calculated, assuming the null hypothesis (H0) is true. We compare this p-value against a predetermined significance level,  $\alpha$  (alpha), which is typically set at 0.05.

In Example 1, the calculated p-value is **0.3517**. Since this value is considerably greater than the common significance level of  $\alpha = 0.05$ , we must **fail to reject the null hypothesis**. This failure to reject H0 provides strong evidence that the inclusion of the variables hp and cyl in the full model does not lead to a statistically significant improvement in the overall model fit compared to the nested model. In practical terms, the increase in complexity provided by the full model is not

justified by the minor gain in explanatory power.

Therefore, the statistically sound decision is to adopt the simpler, **nested model** ( $\text{mpg} \sim \text{disp} + \text{carb}$ ). By choosing the model with fewer predictors when the difference in fit is non-significant, we adhere to the crucial statistical principle of **parsimony**. This selection process guards against overfitting the data and often results in a model that is more stable, easier to interpret, and generalizes better to new, unseen data, which is a primary goal in applied regression analysis.

## Example 2: Sequential Testing and Model Refinement

The process of model selection often involves sequential comparisons, where we iteratively test increasingly simplified nested models against the current best-fitting model. Following the conclusion of Example 1, where the two-predictor model ( $\text{mpg} \sim \text{disp} + \text{carb}$ ) was deemed sufficient, we might now question whether even one of those remaining predictors can be dropped without losing significant explanatory power. We will now compare the two-predictor model against a further reduced model containing only a single predictor.

For this second comparison, our new **Full Model (MF')** is  $\text{mpg} \sim \text{disp} + \text{carb}$ , and the new **Reduced Model (MR')** is  $\text{mpg} \sim \text{disp}$ . This test specifically evaluates whether the variable `carb` (carburetor count) contributes significantly to the model fit, above and beyond the contribution of `disp` (displacement). The null hypothesis ( $H_0$ ) states that the coefficient for `carb` is zero.

We repeat the fitting and testing process in R statistical software:

### library(lmtest)

```
#fit full model (M_F' equivalent to M_R from Example 1)
model_full <- lm(mpg ~ disp + carb, data = mtcars)
```

```
#fit reduced model
model_reduced <- lm(mpg ~ disp, data = mtcars)
```

```
#perform likelihood ratio test for differences in models
lrtest(model_full, model_reduced)
```

Likelihood ratio test

Model 1:  $\text{mpg} \sim \text{disp} + \text{carb}$

Model 2:  $\text{mpg} \sim \text{disp}$

```
#Df LogLik Df Chisq Pr(>Chisq)
```

```
1 4 -78.603
```

```
2 3 -82.105 -1 7.0034 0.008136 **
```

---

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

Upon reviewing the output for this sequential test, we observe a dramatic change in the result. The calculated p-value ( $\Pr(>\text{Chisq})$ ) is **0.008136**. Since this is less than the conventional threshold of  $\alpha = 0.05$ , we **reject the null hypothesis**. This rejection implies that the two models do not fit the data equally well; specifically, the model containing two predictors (`disp` and `carb`) fits the data significantly better than the model containing only one predictor (`disp`).

The conclusion is that the variable `carb` provides non-redundant, statistically significant information necessary for predicting `mpg`, leading to a demonstrable improvement in the model's goodness of fit. Therefore, we retain the two-predictor model as our final, most appropriate model, as it achieves an excellent balance between parsimony and explanatory power. Based on this rigorous comparison using the Likelihood Ratio Test, the final selected model is represented by the equation:

$$\text{mpg} = \beta_0 + \beta_1\text{disp} + \beta_2\text{carb}$$

### Further Reading and Related Topics

To deepen your understanding of regression methods and model assessment in R, the following resources may be helpful:

[How to Perform Simple Linear Regression in R](#)

[How to Perform Multiple Linear Regression in R](#)

[How to Interpret Significance Codes in R](#)