

# Learn How to Extract P-Values from Linear Regression Models in R

Authored by  
**Mohammed looti**

November 16, 2025

## RECOMMENDED CITATION

Mohammed looti (2025). *Learn How to Extract P-Values from Linear Regression Models in R*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=2719>

This comprehensive guide details effective methods for extracting **p-values** from the **lm() function** in **R**, a crucial step in interpreting **statistical significance** within your **regression models**. Understanding how to precisely obtain these values is fundamental for accurate statistical reporting and robust decision-making in complex data analysis workflows.

The **lm() function** in **R** is the standard tool used to fit **linear regression models**, providing essential insights into the relationships between dependent and independent variables. While the built-in **summary() function** offers a complete overview of the model diagnostics, isolating specific **p-values** for either the overall model fit or individual **coefficients** often requires specialized, targeted extraction techniques. This article will present two primary, efficient methods to achieve this, complemented by a practical, step-by-step example to illustrate their direct application in a typical data scenario.

## Understanding P-Values and Hypothesis Testing in Regression

Before implementing the extraction methods, it is crucial to solidify your understanding of what **p-values** represent within the context of **regression analysis**. Fundamentally, a **p-value** is a quantitative measure of the evidence against the **null hypothesis**. It specifically quantifies the probability of observing results as extreme as, or more extreme than, the observed data, assuming that the **null hypothesis** is truly correct.

In the realm of **regression analysis**, we encounter two critical categories of **p-values**: those assessing the overall model fit and those determining the significance of individual predictor variables. The overall model **p-value**, derived from the **F-statistic**, evaluates whether the entire set of independent variables, when considered together, significantly accounts for the variation observed in the dependent variable. Conversely, the **p-values** associated with individual **coefficients** assess the unique statistical significance of each predictor's contribution to the model, controlling for the effects of all other variables present.

The decision rule for interpretation relies on comparing the **p-value** to a predetermined **significance level**, conventionally denoted as alpha (typically 0.05 or 0.01). If the calculated **p-value** is less than this alpha threshold, we possess sufficient evidence to reject the **null hypothesis**. For the overall model, rejecting the null means we conclude that the regression equation is statistically useful. For individual **coefficients**, a low **p-value** confirms that the corresponding predictor variable is a statistically reliable contributor to the predictive power of the model.

### Method 1: Programmatic Extraction of the Overall Model P-Value

The overall **p-value** for any **regression model** is intrinsically linked to the **F-statistic**. This specific test is designed to evaluate the joint significance of all predictor variables, meaning it tests the **null**

**hypothesis** that all slope **coefficients** are simultaneously equal to zero. If the model is not statistically significant, this null hypothesis cannot be rejected. Therefore, a very small **p-value** associated with the **F-statistic** is strong evidence that the model, as a whole, is statistically meaningful.

While the standard output generated by the **summary() function** conveniently displays this overall **p-value** at the bottom, scenarios requiring automated processing, batch analysis, or streamlined report generation necessitate extracting this value programmatically. The following custom **R function** provides a robust and efficient method to retrieve only the numeric overall **p-value** from an **lm() model** object, separating it from the extensive summary output.

**#define function to extract overall p-value of model**

```
overall_p <- function(my_model) {  
  f <- summary(my_model)$fstatistic  
  p <- pf(f,f,f,lower.tail=F)  
  attributes(p) <- NULL  
  return(p)  
}
```

```
#extract overall p-value of model  
overall_p(model)
```

This custom function, intelligently named **overall\_p**, accepts your fitted **lm() model** as its input argument. Internally, it first accesses the necessary components--the **F-statistic** value and its corresponding degrees of freedom--from the model's summary object. It then utilizes the **pf() function**, which is the cumulative distribution function for the F-distribution, to compute the exact **p-value**. Crucially, setting the argument **lower.tail=F** ensures that the calculation yields the upper tail probability, which is the correct methodology for hypothesis testing using the **F-statistic**. Finally, **attributes(p) <- NULL** is used to guarantee a clean, purely numeric output suitable for scripting.

## Method 2: Isolating P-Values for Individual Regression Coefficients

While the overall model **p-value** confirms the model's general utility, determining which specific predictors are driving the relationship requires examining the individual **coefficients**. Each **coefficient** in a **model** is tested using a **t-statistic**, and this test yields an associated **p-value**. This individual **p-value** specifically tests the **null hypothesis** that the true population coefficient for that predictor is zero, implying that the variable has no linear relationship with the outcome when controlling for all other variables in the model.

To extract these essential individual **p-values** directly from your fitted **lm() model** object in **R**, you must access the specific matrix within the model's summary output. The **summary() function** returns a comprehensive matrix of coefficients that includes the estimated values, their standard errors, the calculated **t-values**, and the final column containing the corresponding **p-values**.

The syntax for isolating this specific column is highly concise and efficient. Since the individual **p-values** are consistently stored in the fourth column of the **coefficients matrix** (following the Estimate, Std. Error, and t value columns), you can use matrix subsetting to pull the exact data needed for reporting or further analysis.

```
summary(model)$coefficients
```

## Practical Application: Demonstrating P-Value Extraction in R

To fully grasp the utility of the two methods described above, let us walk through a practical scenario involving the analysis of student performance metrics. We will begin by generating a sample **data frame** and subsequently fitting a **multiple linear regression model** in **R** utilizing the powerful **lm() function**.

```
#create data frame
```

```
df <- data.frame(rating=c(67, 75, 79, 85, 90, 96, 97),  
points=c(8, 12, 16, 15, 22, 28, 24),  
assists=c(4, 6, 6, 5, 3, 8, 7),  
rebounds=c(1, 4, 3, 3, 2, 6, 7))
```

```
#fit multiple linear regression model
```

```
model <- lm(rating ~ points + assists + rebounds, data=df)
```

Once the model is fitted, the essential first step is to generate a comprehensive statistical overview using the **summary() function**. This output is critical as it contains all the diagnostic measures necessary for interpretation, including the estimated **coefficients**, standard errors, **t-values**, the overall **F-statistic**, and, most importantly, all associated **p-values** for both the individual predictors and the model as a whole.

```
#view model summary
```

```
summary(model)
```

Call:

```
lm(formula = rating ~ points + assists + rebounds, data = df)
```

Residuals:

```
1 2 3 4 5 6 7
-1.5902 -1.7181 0.2413 4.8597 -1.0201 -0.6082 -0.1644
```

Coefficients:

```
Estimate Std. Error t value Pr(>|t|)
(Intercept) 66.4355 6.6932 9.926 0.00218 **
points 1.2152 0.2788 4.359 0.02232 *
assists -2.5968 1.6263 -1.597 0.20860
rebounds 2.8202 1.6118 1.750 0.17847
---
Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
Residual standard error: 3.193 on 3 degrees of freedom
Multiple R-squared: 0.9589, Adjusted R-squared: 0.9179
F-statistic: 23.35 on 3 and 3 DF, p-value: 0.01396
```

From the inspection of the full summary output, we can immediately observe that the overall **p-value** for this **regression model** is **0.01396**. This vital piece of information is located at the bottom of the summary, directly accompanying the **F-statistic**. Given that this **p-value** is significantly below the typical **significance level** of 0.05, we confidently conclude that the model, as a collective unit, is statistically significant, meaning that the predictors jointly explain a meaningful amount of variance in the student ratings.

To programmatically isolate and extract just this overall **p-value** using the custom function defined earlier, we execute the following code:

```
#define function to extract overall p-value of model
```

```
overall_p <- function(my_model) {
f <- summary(my_model)$fstatistic
p <- pf(f,f,f,lower.tail=F)
attributes(p) <- NULL
return(p)
}
```

```
#extract overall p-value of model
```

```
overall_p(model)
```

```
0.01395572
```

As expected, the function returns a precise numeric value of approximately **0.01396**, which

perfectly validates the figure presented in the full model summary. This demonstrates the accuracy and practical utility of using a custom function for streamlined extraction in large-scale scripting environments.

Finally, to extract the specific [p-values](#) for each individual [regression coefficient](#) within the model, we apply the direct matrix subsetting technique detailed in Method 2:

```
#extract p-values for individual regression coefficients in model  
summary(model)$coefficients
```

```
(Intercept) points assists rebounds  
0.002175313 0.022315418 0.208600183 0.178471275
```

This command efficiently yields a named vector containing the [p-values](#) for all components: **(Intercept)**, **points**, **assists**, and **rebounds**. A quick inspection reveals that the predictor variable 'points' ( $p=0.0223$ ) is statistically significant at the 0.05 [significance level](#), whereas 'assists' ( $p=0.2086$ ) and 'rebounds' ( $p=0.1785$ ) do not meet this threshold. These individual [p-values](#) are crucial for determining precisely which predictors have a reliable, non-zero association with the dependent variable within the constructed [model](#).

## Interpreting P-Values and Statistical Significance

Correctly interpreting [p-values](#) is arguably more important than the mechanical act of extracting them. The standard procedure involves comparing the calculated [p-value](#) against a predetermined [significance level](#) ( $\alpha$ ). If the [p-value](#) is below this threshold (e.g.,  $p < 0.05$ ), the result is deemed [statistically significant](#), providing strong evidence to reject the [null hypothesis](#) in favor of the alternative hypothesis.

However, it is vital to maintain perspective on the role of the [p-value](#). It does not quantify the magnitude or practical importance of an effect; it merely indicates the strength of the statistical evidence against the [null hypothesis](#). A result that is statistically significant does not automatically translate into practical significance in the real world. For instance, even a small effect size can achieve statistical significance if the sample size is extremely large.

For truly robust analysis, analysts must always consider the context of their research, examine the magnitude of the estimated [coefficients](#), and review other essential diagnostic measures alongside [p-values](#). These additional diagnostics might include measures like [R-squared](#) (to assess goodness of fit) and residual plots (to check model assumptions). A holistic interpretation ensures that statistical conclusions are both numerically sound and practically meaningful.

## Conclusion

Mastering the accurate extraction of [p-values](#) from [lm\(\) models](#) in [R](#) is a fundamental requirement for competent [statistical analysis](#). Whether your objective is to confirm the overall significance of your model or to isolate the individual contributions of specific predictors, the methods detailed in this guide--the custom function for the overall model and the matrix subsetting for individual coefficients--offer clear, efficient, and reliable solutions.

By effectively implementing these programmatic techniques, you can significantly streamline your data analysis workflows and ensure precise, automated reporting of your [hypothesis tests](#). This ability is particularly valuable when working with multiple models or integrating results into production reports.

Ultimately, remember that numerical results, such as [p-values](#), must always be synthesized with careful consideration of the underlying data structure, model assumptions, and the real-world implications of your findings. Adopting this balanced approach is the key to reaching trustworthy and scientifically sound statistical conclusions.

## Additional Resources for R Programming

The following tutorials explain how to perform other common statistical and programming tasks in [R](#):