

How do you perform a partial F-test?

Authored by
stats writer

December 17, 2025

RECOMMENDED CITATION

stats writer (2025). *How do you perform a partial F-test?*. PSYCHOLOGICAL SCALES.
Retrieved from <https://scales.arabpsychology.com/?p=107682>

Introduction to the Partial F-Test and its Purpose

The **Partial F-test** stands as a fundamental statistical procedure within the domain of multiple linear regression analysis. Its essential function is to quantitatively determine whether a specific group of predictor variables--when considered together--contributes a statistically significant amount of predictive power to a model, beyond what is already explained by other variables included in a simpler, constrained version of that model. This sophisticated comparison allows statisticians and data analysts to rigorously test specific hypotheses about the joint explanatory power of a subset of covariates.

The core utility of the Partial F-test lies in its ability to compare two models that are related through a process known as nesting. Specifically, it tests the statistical difference between a **full regression model**, which includes all potential predictors, and a **nested model** (or reduced model), which is derived by excluding a select group of those predictors. By evaluating whether the drop in model fit (the increase in error) resulting from the exclusion of these variables is statistically significant, the test provides robust evidence for informed model selection, helping to balance predictive accuracy with model parsimony.

This test is particularly valuable in situations where researchers suspect multicollinearity or wish to simplify a complex structure without sacrificing meaningful explanatory capability. If the Partial F-test indicates that the reduced model is not significantly worse than the full model, the principle of parsimony suggests selecting the reduced model. This simplifies interpretation, enhances generalization to new data, and avoids the unnecessary inclusion of variables that only add noise or complexity, thereby strengthening the foundation of the resulting regression model.

Understanding Full and Nested Regression Models

A clear understanding of the full and nested model paradigm is prerequisite to executing a Partial F-test. The **full model** represents the hypothesized true relationship, incorporating all predictor variables (x_1, x_2, \dots, x_k) thought to influence the response variable (Y). For instance, if we consider four predictors, the full model equation is comprehensive:

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

In contrast, a **nested model** is constructed by imposing restrictions on the full model, typically by setting a subset of the β coefficients equal to zero. This effectively removes the corresponding predictor variables. A model is considered nested if its set of predictors is a strict subset of the predictors in the full model. For example, to test the joint effect of x_3 and x_4 , the nested model would omit these terms:

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

The Partial F-test is designed precisely to evaluate the implications of these restrictions. It assesses the statistical cost--measured as the increase in prediction error--incurred when moving from the unrestricted (full) model to the restricted (nested) model. If the observed increase in error is trivial, it supports the statistical validity of the restrictions, meaning the coefficients of the excluded variables are likely zero in the population.

The Mathematical Formulation of the Partial F-Statistic

The Partial F-test operates by comparing the variance unexplained by the two models, utilizing the Residual Sum of Squares (RSS). The F-statistic is essentially a ratio of variances: the variance explained by the added predictors (normalized by their count) versus the overall error variance of the full model. This ratio determines the magnitude of improvement per added degree of freedom relative to the baseline noise.

The formula for the Partial F test-statistic is structured as follows:

$$F = \frac{(RSS_{\text{reduced}} - RSS_{\text{full}})/p}{(RSS_{\text{full}})/(n-k)}$$

This statistic follows an F-distribution, characterized by its degrees of freedom. The numerator, $(RSS_{\text{reduced}} - RSS_{\text{full}})$, quantifies the reduction in error achieved by the p additional predictors in the full model. Dividing this by p provides the mean square reduction due to the included variables. The denominator, $RSS_{\text{full}}/(n-k)$, calculates the Mean Square Error (MSE) of the full model, which serves as the best unbiased estimate of the underlying population error variance (σ^2).

A detailed breakdown of the components provides clarity on the necessary calculations:

RSS_{reduced} : The measure of unexplained variance for the model containing fewer predictors (the nested model).

RSS_{full} : The measure of unexplained variance for the comprehensive model (the full model).

p : The number of coefficients that are restricted to zero when moving from the full model to the reduced model; this is the number of predictors being jointly tested.

n : The total number of observations in the dataset.

k : The total number of parameters estimated in the **full model**, including the intercept (β_0).

Defining the Hypotheses and Degrees of Freedom

The Partial F-test is structured around a precise set of statistical hypotheses that define the nature of the inquiry into the parameters of the regression model. The hypotheses focus exclusively on the coefficients (β 's) associated with the p variables that distinguish the full model from the

nested model.

The formal structure is as follows:

H_0 : The null hypothesis asserts that the collective explanatory power of the p predictors removed is zero. This implies that all their corresponding population coefficients are simultaneously zero. If x_3 and x_4 are the removed variables, $H_0: \beta_3 = \beta_4 = 0$.

H_A : The alternative hypothesis claims that the p predictors collectively hold significant explanatory power. This means at least one of their corresponding population coefficients is non-zero, suggesting the full model provides a significantly better fit than the reduced model. If x_3 and x_4 are the removed variables, $H_A: \text{At least one of } \beta_3 \text{ or } \beta_4 \text{ is non-zero}$.

The critical aspect of performing the test involves selecting the correct degrees of freedom for the F-distribution. The distribution is defined by two degrees of freedom: df_1 (numerator) and df_2 (denominator). The numerator degrees of freedom, df_1 , equals p , the number of variables whose coefficients are hypothesized to be zero. The denominator degrees of freedom, df_2 , equals $n - k$, which is the residual degrees of freedom of the full model. Correctly specifying these parameters is vital for calculating an accurate p-value and ensuring reliable inference regarding the significance of the tested predictors.

Interpreting the Residual Sum of Squares (RSS) Dynamics

The Residual Sum of Squares (RSS) is the metric upon which the Partial F-test hinge. It represents the aggregate squared distance between the observed values and the values predicted by the model. A fundamental property of multiple linear regression is that adding any predictor variable to a model will always result in a reduction of the RSS, or, in the worst case, the RSS will remain unchanged. Therefore, RSS_{full} can never be greater than RSS_{reduced} .

The critical insight derived from RSS comparison is encapsulated in the quantity $RSS_{\text{reduced}} - RSS_{\text{full}}$. This difference represents the exact amount of variance reduction attributable solely to the inclusion of the p variables being tested. If these variables are truly irrelevant, this difference will be small, reflecting only random fluctuations. If the variables are highly relevant, the difference will be large, indicating that the full model significantly reduced the prediction error compared to the restricted model.

If the calculated F-statistic is large--meaning the numerator (improvement due to p variables) is disproportionately large compared to the denominator (baseline error variance)--the p-value will be small. This outcome leads to the rejection of the null hypothesis, confirming that the added predictors contribute significantly to the model fit. Conversely, a small F-statistic implies that the

observed reduction in RSS is minor and likely due to chance, failing to provide enough evidence to overturn the assumption that the excluded parameters are zero.

Practical Implementation using Statistical Software (R Example)

In most statistical analyses, the comparison between the full and nested model is conducted using an ANOVA procedure, which efficiently calculates the required F-statistic and p-value. Let us use the widely referenced **mtcars** dataset in R to demonstrate a concrete application of the Partial F-test methodology. We aim to determine if horsepower (hp) and the number of cylinders (cyl) are jointly necessary for predicting miles per gallon (mpg), assuming displacement ($disp$) and carburetors ($carb$) are already included.

We begin by defining and fitting the two models based on our hypothesis:

Full model: Predicts mpg using $disp$, $carb$, hp , and cyl .

Reduced model: Predicts mpg using only $disp$ and $carb$. (We are testing the joint significance of hp and cyl , so $p=2$.)

The following R code blocks demonstrate the fitting of these models and the subsequent comparative ANOVA, which outputs the result of the Partial F-test:

Step 1: Fit the full regression model

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

Step 2: Fit the nested (reduced) regression model

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

Step 3 & 4: Perform ANOVA to test for differences in models (Partial F-test)

```
anova(model_reduced, model_full)
```

Analysis of Variance Table

Model 1: $mpg \sim disp + carb$

Model 2: $mpg \sim disp + carb + hp + cyl$

	Res.Df	RSS	Df	Sum of Sq	F	Pr(>F)
1	29	254.82				
2	27	238.71	2	16.113	0.9113	0.414

Interpreting the Results and Drawing Conclusions

The output from the ANOVA table provides the decisive statistics required to finalize the Partial F-test. We can extract the crucial values directly from the last row of the table, which compares

Model 1 (reduced) to Model 2 (full).

The residual sum of squares for the reduced model (RSS_{reduced}) is **254.82**.

The residual sum of squares for the full model (RSS_{full}) is **238.71**.

The degrees of freedom difference (p) is **2**.

The calculated Partial F test-statistic is **0.9113**.

The corresponding p-value ($Pr(>F)$) is **0.414**.

To reach a conclusion, we compare the p-value to a predetermined level of significance, typically $\alpha = 0.05$. In this instance, the p-value of 0.414 is significantly greater than 0.05. According to the rules of hypothesis testing, we **fail to reject the null hypothesis (H_0)**.

The statistical conclusion is that there is insufficient evidence to suggest that the group of predictors (X_p and X_{yl}) jointly contributes significantly to the prediction of Y when X_{disp} and X_{carb} are already present in the model. Consequently, the simpler, nested multiple linear regression model (Model 1) is statistically preferred, as the added complexity of Model 2 does not yield a statistically significant improvement in explanatory power. This demonstrates the critical role of the Partial F-test in simplifying complex models without sacrificing statistical fidelity.