

How to Perform Fisher's Exact Test in R: A Step-by-Step Guide

Authored by
stats writer

December 31, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Perform Fisher's Exact Test in R: A Step-by-Step Guide*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=110135>

The Fisher's Exact Test is a powerful and precise statistical method designed for analyzing the association between two nominal or categorical variables. Unlike approximate tests, this method computes the exact probability of observing the data given the marginal totals, making it particularly reliable when dealing with **small sample sizes**. Its primary application is to assess whether the proportions of one variable differ significantly across the categories of the second variable.

In the world of statistical computing, conducting this test is straightforward, especially when utilizing the statistical programming language R. The built-in `fisher.test()` function streamlines the process. This function requires a contingency table (typically 2x2, though extensions exist) as input and yields essential outputs, including the test statistic, the critical p-value, and confidence intervals for the odds ratio. Interpreting these results is crucial for determining the presence and strength of the relationship between the two variables under investigation.

The core utility of Fisher's Exact Test lies in its ability to handle scenarios where other conventional tests fail. Specifically, it is used to determine whether or not there is a statistically significant association between two categorical variables, forming the foundation of many epidemiological and experimental analyses.

When to Use Fisher's Exact Test

Understanding the context for applying the Fisher's Exact Test is essential for robust statistical analysis. This test shines particularly brightly when dealing with data structured in a 2x2 contingency table, especially in situations where the assumptions required for the standard Pearson's Chi-Squared Test are violated. Specifically, the Chi-Squared Test relies on the expectation that cell counts are sufficiently large. A common rule of thumb suggests that if one or more of the expected cell frequencies in a 2x2 table falls below the threshold of 5, the Chi-Squared approximation becomes unreliable, increasing the risk of **Type I error**.

In such cases of small sample size or sparse data, Fisher's Exact Test provides the mathematically precise solution. It calculates the probability of observing the given distribution of frequencies (or a distribution even more extreme) under the assumption that the null hypothesis is true. This calculation uses the hypergeometric distribution, ensuring that the results remain reliable regardless of how small the cell counts are, provided the **marginal totals** are fixed.

Therefore, whenever your research involves assessing the relationship between two binary categorical variables, and the dataset is limited--resulting in low cell counts--Fisher's Exact Test is the recommended methodology to ensure the validity and accuracy of your statistical inference.

Formalizing the Hypotheses

Like all hypothesis testing procedures, the Fisher's Exact Test operates under a defined set of competing hypotheses that structure the statistical inquiry. These hypotheses focus strictly on the independence (or dependence) of the variables represented in the contingency table. It is crucial to define these statements clearly before conducting the test.

The following are the established null and alternative hypotheses used in the analysis:

H0: (Null Hypothesis) The two categorical variables are **independent**. There is no significant association between the rows and columns of the contingency table. Equivalently, the distribution of one variable is the same across the categories of the other variable.

HA: (Alternative Hypothesis) The two categorical variables are **not independent**. There is a statistically significant association or relationship between the row and column variables.

The goal of executing the test is to determine whether the resulting p-value is small enough (typically less than 0.05) to provide sufficient evidence to reject the Null Hypothesis (H0) in favor of the Alternative Hypothesis (HA).

Setting Up Data for Analysis in R

To execute the Fisher's Exact Test using R, the data must first be organized into a proper contingency table structure. While the test can technically handle tables larger than 2x2, it is most commonly applied to the 2x2 format where the cell counts are low. This structure ensures that both rows and columns represent two distinct categorical outcomes (e.g., success/failure, treatment/control).

For our practical demonstration, we will generate a synthetic 2x2 matrix containing small cell counts, reflecting a typical scenario where the Chi-Squared Test approximation would be inappropriate. This matrix represents the joint frequency distribution of two variables, such as 'Outcome (Success/Failure)' and 'Group (A/B)'.

We use the `matrix()` function in R to construct this data object. The input vector `c(2, 5, 9, 4)` dictates the cell counts, and `nrow = 2` ensures it is arranged into two rows, filling the matrix column-wise by default. The following code demonstrates the creation and immediate viewing of the dataset:

```
#create 2x2 dataset
data = matrix(c(2,5,9,4), nrow = 2)
```

```
#view dataset
```

```
data
#
# 2 9
# 5 4
```

In this resulting matrix, the cells contain the observed frequencies: 2 and 5 in the first column, and 9 and 4 in the second column. Note that the smallest count here is 2, confirming that this is an ideal dataset for applying the [Fisher's Exact Test](#).

Executing the Fisher's Exact Test

Once the data is correctly structured as a matrix object in R, executing the test is remarkably simple. The core function required is `fisher.test()`, which is part of R's base installation and does not require any external packages. We simply pass the contingency table object, `data`, directly into this function.

The default execution of `fisher.test()` performs a two-sided test and calculates the exact p-value. If a one-sided hypothesis were required (e.g., testing for a strictly greater or strictly less association), additional arguments like `alternative="greater"` or `alternative="less"` would be included. However, for general association testing, the default is sufficient:

fisher.test(data)

This single line of code executes the necessary calculations, based on factorials and probabilities derived from the hypergeometric distribution, yielding a comprehensive output detailing the results of the exact test.

Analyzing the Test Results

Running the `fisher.test(data)` command generates a detailed summary that allows us to draw statistical conclusions about the relationship between our two variables. The output typically includes the p-value, the confidence interval for the odds ratio, and the alternative hypothesis used.

The visual output from R, based on the synthetic data we created, provides the necessary metrics for interpretation:

Fisher's Exact Test for Count Data

```
data: data
p-value = 0.1597
alternative hypothesis: true odds ratio is not equal to 1
95 percent confidence interval:
 0.0130943 1.8397543
sample estimates:
odds ratio
 0.1957871
```

The Null Hypothesis (H0) for this test posits that the two variables are independent, which is mathematically equivalent to stating that the true population odds ratio relating the two variables is precisely equal to 1. If the variables are independent, the odds of a specific outcome in one group are the same as the odds of that outcome in the other group.

Interpreting the P-Value

The most critical component of the output for hypothesis testing is the **p-value**. This value represents the probability of observing our specific data (or data more extreme) if the null hypothesis of independence were actually true. We compare this p-value against a predefined significance level, typically denoted as alpha (α), which is most commonly set at **0.05**.

In the example provided by the output image, the resulting p-value is **0.1597**. Since this value (0.1597) is significantly greater than the standard alpha level of 0.05, we must conclude that we do not have sufficient statistical evidence to reject the null hypothesis (H0). This finding indicates that, based on our sample data, we cannot assert that there is a statistically significant association or difference between the two categorical variables.

In practical terms, this suggests that the observed pattern in the 2x2 table could reasonably have occurred by random chance, even if the true underlying relationship between the variables were one of **complete independence**. Thus, we fail to reject H0.

The Confidence Interval for the Odds Ratio

Beyond the p-value, the output of the `fisher.test()` function in R also provides the estimated odds ratio and its associated confidence interval. The odds ratio quantifies the strength and direction of the association between the two variables. An odds ratio of 1 indicates perfect independence (no association), a value greater than 1 indicates a positive association, and a value less than 1 indicates a negative association.

The output for our example provides the 95% confidence interval for the true population odds ratio:

95% Confidence Interval for Odds Ratio: **(0.0130943, 1.8397543)**

This interval suggests that we are 95% confident that the true population odds ratio falls between 0.013 and 1.84. Critically, we must observe whether the value 1 is contained within this interval. If the number 1 is included, it means that the possibility of the odds ratio being equal to 1 (i.e., the null hypothesis of independence) cannot be statistically ruled out at the specified alpha level (0.05).

Drawing the Final Conclusion

Since the calculated 95% confidence interval **(0.0130943, 1.8397543)** clearly encompasses the value 1, this result perfectly aligns with the conclusion derived from the high p-value (0.1597). Both metrics confirm that there is no statistically significant evidence of a relationship between the row and column variables in our sample dataset. Had the interval been entirely above 1 (e.g., 1.5 to 5.0) or entirely below 1 (e.g., 0.1 to 0.5), we would have rejected H_0 and concluded dependence.

In summary, the Fisher's Exact Test provides a robust and reliable method for testing independence in 2x2 contingency tables, particularly when cell frequencies are small. By understanding both the p-value and the confidence interval for the odds ratio, researchers can make precise inferences regarding the association between categorical factors.

Further Resources on Exact Tests

For those interested in delving deeper into the theoretical background and advanced applications of exact statistical procedures, the following tutorials and resources provide additional information about Fisher's Exact Test and related methodologies: