

Conduct Fisher's Exact Test in R

Authored by
Mohammed loot

November 9, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Conduct Fisher's Exact Test in R*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=14361>

Understanding Fisher's Exact Test: Context and Purpose

The **Fisher's Exact Test** is a powerful statistical tool utilized in the analysis of [categorical variables](#). Specifically, it is designed to determine whether a statistically significant non-random association exists between two different classifications. This test is foundational in fields such as biological research, social sciences, and epidemiology, where researchers frequently analyze data presented in contingency tables. Unlike approximations, Fisher's method provides an exact probability (p-value) of observing the data, or data more extreme, given the marginal totals are fixed.

In essence, the test evaluates if the distribution of observations across the categories of one variable is dependent upon the categories of the second variable. When dealing with experimental design, especially when outcomes are binary (e.g., success/failure, presence/absence), the ability of Fisher's test to handle small sample sizes makes it indispensable. Its rigor ensures that conclusions drawn regarding the association between variables are based on precise probabilities, mitigating risks associated with relying on asymptotic approximations, particularly when data are sparse. The test is fundamental in establishing whether observed differences in frequencies between groups are likely due to chance or represent a true underlying relationship.

The methodology behind the test leverages the principles of combinatorics to calculate the precise probability of obtaining the observed cell frequencies under the assumption of independence. This exact calculation is what grants the test its high reliability, especially when analyzing data derived from small studies or rare events where statistical power might be otherwise limited.

Why Use Fisher's Exact Test? Comparison to Chi-Square Assumptions

While the **Chi-Square Test of Independence** is the most commonly taught method for analyzing associations in contingency tables, it relies on certain underlying assumptions regarding expected cell frequencies. The [Chi-Square test](#) uses an approximation of the sampling distribution, which is valid only when sample sizes are sufficiently large. A critical rule of thumb for validating the Chi-Square approximation is that the expected frequency in every cell of the contingency table should be five or greater. Failure to meet this expectation can lead to a highly distorted p-value, skewing the results and potentially leading to incorrect statistical inference.

When this assumption of adequate expected cell counts is violated--often occurring in studies with limited sample sizes or when outcomes are very rare--the results of the standard Chi-Square test become unreliable. This is precisely where the [Fisher's Exact Test](#) steps in as a robust alternative. Fisher's Exact Test calculates the probability directly using the hypergeometric distribution, considering every possible configuration of the table that preserves the marginal totals. Because it does not rely on large sample approximations, it remains valid even when one or more cell counts in the 2x2 table are small (less than 5).

Consequently, researchers must first assess the suitability of their data structure before choosing a statistical procedure. If a 2x2 table exhibits low cell counts, transitioning immediately to Fisher's Exact Test ensures the integrity of the statistical analysis. This careful selection of the appropriate statistical methodology is paramount for drawing scientifically sound conclusions about the relationship between two **categorical variables**. Furthermore, the reliance on exact calculation means that Fisher's test is inherently more conservative than the Chi-Square test when sample sizes are small, providing a more cautious estimate of significance.

Formulating Hypotheses for Categorical Association

Statistical testing begins with the clear articulation of the null and alternative hypotheses. These hypotheses frame the question being asked of the data and guide the interpretation of the resulting [p-value](#). For Fisher's Exact Test, these hypotheses relate directly to the concept of independence between the two variables presented in the contingency table, often referred to as a test of association.

The formal structure of the hypotheses is defined as follows:

H0: (Null Hypothesis) The two categorical variables are [statistically independent](#). This means there is no association or relationship between the row variable and the column variable. If the **null hypothesis** is true, the distribution of one variable is unaffected by the state of the other, implying that any observed differences are simply due to random sampling variability.

HA: (Alternative Hypothesis) The two categorical variables are *not* independent. This is the hypothesis of association, suggesting that the distribution of one variable is significantly related to the distribution of the other. The goal of the test is usually to find sufficient evidence to reject the **null hypothesis** in favor of this alternative hypothesis, thereby confirming a relationship.

The subsequent analysis in [R](#) will generate a p-value which quantifies the probability of observing the current data (or data more extreme) assuming the null hypothesis is true. A small p-value (conventionally less than 0.05) provides strong evidence against the assumption of independence, leading to the rejection of H0. Conversely, a large p-value suggests that the observed data are consistent with the assumption of independence.

Step-by-Step Example: Implementing the Test in R

Performing the Fisher's Exact Test in [R](#) is straightforward, requiring only the creation of a 2x2 contingency table (matrix). The primary function needed is `fisher.test()`, which is part of R's base installation and does not require loading additional packages. To illustrate the process, let us construct a hypothetical dataset representing the outcomes of two groups--for instance, comparing treatment success rates.

First, we must structure the data correctly. The data must be input as a matrix where the rows typically represent one variable (e.g., Group A vs. Group B) and the columns represent the outcomes (e.g., Outcome 1 vs. Outcome 2). The following R code demonstrates how to create and view such a 2x2 matrix using the `matrix()` function. Note that in R, data is input column by column by default unless specified otherwise. Our example uses the specific cell counts of 2, 5, 9, and 4.

```
#create 2x2 dataset
```

```
data = matrix(c(2,5,9,4), nrow = 2)
```

```
#view dataset
```

```
data
```

```
# 2 9
```

```
# 5 4
```

Once the data object (named `data` in this case) is correctly defined, executing the test is a single command. The simplicity of the R command structure makes complex statistical analysis highly accessible. We simply pass the matrix object directly to the `fisher.test()` function, allowing R to handle the complex underlying combinatorial calculations instantly.

```
fisher.test(data)
```

This execution immediately provides the comprehensive statistical output required for interpretation, including the exact [p-value](#), the estimate of the [odds ratio](#), and the associated 95% [confidence interval](#). R handles both one-sided and two-sided tests by default, though specific arguments can be used to control the alternative hypothesis if required by the research design.

Interpreting the Output: P-Values and Independence

Upon running the `fisher.test(data)` command, R produces a detailed summary of the results. The central component of this output, which dictates the outcome of the hypothesis test, is the **p-value**. The output, visually represented below, contains all necessary metrics for decision-making regarding the independence of the variables:

This produces the following output:

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
```

In the context of the [Fisher's Exact Test](#), the **null hypothesis** posits that the two variables--represented by the rows and columns--are statistically independent. Our decision criterion relies on comparing the calculated p-value against a predetermined significance level, typically denoted as alpha (α), which is most commonly set at 0.05. If the p-value is less than or equal to α , we reject the null hypothesis, concluding that the association is statistically significant.

In the provided example output, the calculated p-value is **0.1597**. Since 0.1597 is greater than the standard alpha level of 0.05, we fail to reject the null hypothesis (H_0). This means that based on the observed data, there is insufficient statistical evidence to conclude that a significant association exists between the row variable and the column variable. The probability of observing this data pattern (or one more extreme) purely by chance, assuming no relationship exists, is relatively high (nearly 16%).

Thus, we conclude that the two columns (or groups) are likely independent of each other. It is important to remember that failing to reject the null hypothesis does not prove independence; it simply means that the data do not provide enough evidence to support the alternative hypothesis of dependence. The conclusion is that we cannot state there is any statistically significant difference or association between the two columns based on this sample, assuming a significance threshold of 0.05.

Detailed Analysis: The Odds Ratio and Confidence Intervals

Beyond the p-value, the output of the `fisher.test()` function provides crucial information regarding the magnitude and direction of the association, quantified through the [odds ratio](#) (OR) and its associated [confidence interval](#). The **odds ratio** is a fundamental measure used in analyzing 2x2 contingency tables, representing the ratio of the odds of an event occurring in one group compared to the odds of it occurring in another group. It is a critical metric for assessing the strength of the association.

The null hypothesis that the two variables are independent is statistically equivalent to stating that

the true population **odds ratio** is equal to 1. An odds ratio significantly greater than 1 suggests a positive association (increased odds in one group), while an odds ratio significantly less than 1 suggests a negative association (decreased odds in one group). In our example output, the estimated odds ratio is provided as a point estimate of the effect size observed in the sample data.

The output of the test also gives us a 95% **confidence interval** for the odds ratio, which provides a range of plausible values for the true population odds ratio. This interval is calculated to contain the true odds ratio 95% of the time if the experiment were hypothetically repeated many times.

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

To confirm the result derived from the p-value, we examine this confidence interval. If the number 1 (the value representing no association) is contained within the calculated interval, then we cannot conclude, at the 0.05 significance level, that the **odds ratio** is significantly different from 1. Since the interval (0.0130943, 1.8397543) clearly includes 1, this analysis strongly reinforces the conclusion reached by the p-value: there is no statistically significant evidence of dependence between the two variables. The consistency between the p-value (which is greater than 0.05) and the inclusion of 1 in the **confidence interval** is a vital check for the validity of the statistical inference, providing a measure of certainty regarding the lack of association.

Conclusion and Further Study

The **Fisher's Exact Test** provides a critical, non-parametric method for testing the independence of two categorical variables, particularly when dealing with small sample sizes or low expected cell frequencies where the traditional [Chi-Square test](#) would be inappropriate. The implementation in [R](#) is highly efficient, requiring only a simple matrix input and the `fisher.test()` function.

Interpreting the output requires careful consideration of both the **p-value** and the 95% confidence interval for the **odds ratio**. In our specific example, both metrics indicated a failure to reject the null hypothesis, confirming the lack of a statistically significant association between the variables. Mastery of this test is essential for accurate statistical reporting when analyzing sparse 2x2 contingency data, ensuring that conclusions are based on exact probabilities rather than asymptotic approximations. For researchers needing to analyze tables larger than 2x2 with small cell counts, specialized algorithms or permutation tests (like Monte Carlo simulations available in R) are often required, although the basic principles derived from Fisher's methodology remain relevant.

For those interested in delving deeper into the theoretical underpinnings or seeking applications beyond the 2x2 scenario, further exploration of statistical documentation and advanced R packages is highly recommended. The following resources offer additional information and context regarding this powerful statistical procedure:

The following tutorials provide additional information about Fisher's Exact Test: