

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

Authored by  
**stats writer**

March 2, 2026

## RECOMMENDED CITATION

stats writer (2026). *How to Perform Fisher's Exact Test in R: A Step-by-Step Guide*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=133460>

## Understanding the Foundations of Fisher's Exact Test

In the vast landscape of statistical methodologies, **Fisher's Exact Test** stands as a cornerstone for researchers dealing with categorical data analysis. Named after the legendary statistician **Sir Ronald A. Fisher**, this test provides a method to determine if there are non-random associations between two **categorical variables**. Unlike many other statistical tests that rely on large-sample approximations, this procedure calculates the exact probability of the observed data under the null hypothesis, making it an "exact" test. This precision is particularly valuable in specialized fields such as genomics, clinical trials, and rare disease research, where acquiring large datasets is often logistically or ethically challenging.

The conceptual framework of **Fisher's Exact Test** is rooted in the **hypergeometric distribution**. By assuming that the marginal totals of a **contingency table** are fixed, the test calculates the probability of obtaining the specific cell frequencies observed, as well as the probabilities of all more extreme configurations. This rigorous approach eliminates the need for the assumptions required by asymptotic tests. While it was originally popularized through the famous "Lady Tasting Tea" experiment, its modern application in **R** allows users to handle complex data structures with just a few lines of code, ensuring that even small-scale experiments yield reliable **statistical significance**.

One of the defining features of this test is its independence from the requirement of a normal distribution. In many scenarios, data collected from surveys or experimental observations do not follow a bell curve, particularly when dealing with binary outcomes like "success/failure" or "presence/absence." By focusing on the exact combinations of outcomes within a 2x2 or larger matrix, **Fisher's Exact Test** provides a robust alternative to parametric methods. As we delve deeper into the implementation within the **R** environment, it is essential to appreciate the mathematical elegance that allows this test to maintain its validity even when sample sizes are minimal.

Furthermore, the utility of **Fisher's Exact Test** extends beyond simple 2x2 tables, although it is most frequently applied in that context. In **R**, the algorithm can be scaled to larger tables, provided that the computational resources are available to calculate the permutations. This makes it a versatile tool for any data scientist or researcher who needs to validate the relationship between discrete variables without the risk of the errors associated with approximation methods. Understanding when and why to use this test is the first step toward conducting high-quality statistical analysis.

## When to Prefer Fisher's Exact Test Over Chi-Square

A common dilemma for researchers is deciding between **Fisher's Exact Test** and the **Chi-square**

**test** of independence. The primary factor influencing this decision is the sample size and the distribution of frequencies across the cells of the **contingency table**. The **Chi-square test** is an asymptotic test, meaning its accuracy improves as the sample size increases. However, when the expected frequencies in any cell of the table fall below five, the **Chi-square test** can provide misleading results, often overestimating significance and leading to a Type I error.

In contrast, **Fisher's Exact Test** does not rely on these large-sample approximations. It is the gold standard when dealing with small datasets where at least one cell has an expected count of less than five. While some practitioners apply the Yates' correction to the **Chi-square test** to compensate for small samples, **Fisher's Exact Test** is inherently more accurate because it computes the actual probability distribution. This makes it indispensable for clinical research involving rare side effects or pilot studies where data points are scarce and expensive to obtain.

Another critical distinction lies in the marginal totals. **Fisher's Exact Test** assumes that both the row and column totals of the **contingency table** are fixed in advance by the experimental design. This is common in "tasting" experiments or controlled laboratory settings. While this assumption is sometimes debated in purely observational studies, the test remains highly robust. For **R** users, the convenience of the `fisher.test()` function means that there is little reason to settle for an approximation when an exact calculation is computationally feasible.

Ultimately, the choice depends on the balance between computational intensity and the need for precision. For very large datasets with high cell counts, the **Chi-square test** is computationally efficient and provides nearly identical results to the exact test. However, as datasets become smaller or more sparse, the **Fisher's Exact Test** becomes the necessary choice to maintain scientific integrity. In the following sections, we will explore how to set up the **null hypothesis** and prepare your data to execute this analysis effectively.

## Defining the Null and Alternative Hypotheses

Before executing any code in **R**, it is vital to clearly define the **null hypothesis** ( $H_0$ ) and the **alternative hypothesis** ( $H_A$ ). In the context of **Fisher's Exact Test**, these hypotheses center on the concept of independence between variables. The **null hypothesis** typically posits that there is no association between the two **categorical variables** being studied. In other words, knowing the classification of an observation in one variable provides no information about its classification in the other.

Conversely, the **alternative hypothesis** suggests that a significant relationship or association does exist. In mathematical terms, for a 2x2 table, the **null hypothesis** states that the **odds ratio** is equal to one. If the **odds ratio** significantly deviates from one, it provides evidence in favor of the **alternative hypothesis**. This structured approach allows researchers to apply a rigorous framework to their observations, ensuring that any conclusions drawn are backed by **statistical**

## **significance.**

Fisher's Exact Test uses the following null and alternative hypotheses:

**H0: (null hypothesis)** The two variables are independent, implying no association exists between them.

**HA: (alternative hypothesis)** The two variables are *not* independent, indicating a significant association.

It is important to note that **Fisher's Exact Test** can be conducted as either a one-tailed or a two-tailed test. A two-tailed test is the default in **R** and is generally preferred as it tests for any difference in either direction. However, if a researcher has a strong theoretical reason to expect an association in a specific direction (e.g., a treatment only improving outcomes, not worsening them), a one-tailed test may be appropriate. Clearly stating these hypotheses beforehand prevents the pitfalls of "p-hacking" and ensures the transparency of the statistical analysis.

Once the hypotheses are established, the next step involves preparing the data for the **R** environment. This requires organizing observations into a **contingency table** format, which **R** can then process to calculate the **p-value**. By adhering to this formal process, researchers can confidently move from raw data to meaningful statistical inference, providing a solid foundation for their experimental findings.

## **Constructing the Contingency Table in R**

The first practical step in conducting **Fisher's Exact Test** in **R** is the creation of a **contingency table**. This table, often referred to as a cross-tabulation or matrix, summarizes the frequency distribution of the variables. In **R**, the most efficient way to generate this is by using the ``matrix()`` function. You must provide the data points in a vector format and specify the number of rows or columns to define the dimensions of the table.

For example, let's generate a 2x2 dataset to use as an example:

```
#create 2x2 dataset
```

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

```
#view dataset
```

```
data
```

```
# 2 9
```

```
# 5 4
```

In the code snippet above, we create a matrix with four values organized into two rows. The values

2, 5, 9, and 4 represent the counts for each combination of categories. It is crucial to ensure that the data is entered in the correct order, as **R** fills matrices by column by default. If your data is currently in a data frame format, you can also use the ``table()`` function to generate the **contingency table** automatically from your variables, which is highly efficient for larger datasets.

Beyond simple creation, naming the rows and columns of your matrix can significantly improve the readability of your output. Using the ``rownames()`` and ``colnames()`` functions allows you to label the categories (e.g., "Treatment," "Control," "Success," "Failure"). This practice not only helps in interpreting the results but also makes the code more accessible to collaborators. A well-organized table is the prerequisite for an accurate and interpretable **Fisher's Exact Test**.

Once your matrix is ready, you can perform a quick visual check by printing the object to the console. Ensuring that the row and column totals align with your raw data is a vital quality control step. In statistical programming, the integrity of the input data directly dictates the reliability of the output. With the **contingency table** properly formatted, we are now ready to apply the core statistical function to obtain our results.

## Executing the `fisher.test()` Function

With the data organized into a matrix, the execution of the test itself is remarkably straightforward. The **R** environment provides a built-in function titled ``fisher.test()`` which is specifically designed to handle these calculations. This function is highly optimized and can handle both 2x2 tables and larger **contingency table** formats using an extension of the original Fisher algorithm. By passing your matrix object as the primary argument, **R** will perform the necessary permutations to arrive at the **p-value**.

To conduct Fisher's Exact Test, we simply use the following code:

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

This single line of code initiates a complex mathematical process. For a 2x2 table, the function calculates the probability of the observed table using the **hypergeometric distribution** formula. It then identifies all other possible tables that could have been formed given the fixed marginal totals and sums the probabilities of those tables that are equally or less likely than the observed one. This summation results in the **p-value**, which is the primary metric used to evaluate **statistical significance**.

The ``fisher.test()`` function also offers several optional parameters that allow for greater flexibility. For instance, the ``alternative`` argument can be set to "two.sided", "greater", or "less" to specify the direction of the test. Additionally, the ``conf.level`` argument allows you to adjust the **confidence interval** for the **odds ratio**, with 0.95 being the standard default. Understanding these parameters

allows a researcher to tailor the test to the specific requirements of their experimental design.

Executing the function is only half the battle; the real value lies in the output it generates. **R** provides a comprehensive summary that includes the **p-value**, the estimated **odds ratio**, and the **confidence interval**. This output is presented in a clear, text-based format in the console, providing all the information needed for a formal research report. In the next section, we will break down an example of this output to understand what these numbers actually mean for your data.

## Interpreting the Statistical Output

After running the `fisher.test()` function, **R** will return a detailed summary of the results. This output is the most critical part of the process, as it tells you whether your findings are statistically meaningful. The most prominent feature of the result is the **p-value**. This value represents the probability of observing a result at least as extreme as the one obtained, assuming the **null hypothesis** is true. A common threshold for significance is 0.05, though this can vary depending on the field of study.

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 **Fisher's Exact Test**, the **null hypothesis** is that the two columns are independent (or equivalently, that the **odds ratio** is equal to 1). To determine if the two columns are independent, we can look at the **p-value** of the test. In this specific case, the **p-value** is **0.1597**. Because this value is greater than the typical alpha level of 0.05, we conclude that we do not have sufficient evidence to reject the **null hypothesis**. Therefore, we cannot state that there is a statistically significant association between the variables.

Another vital component of the output is the **odds ratio**. The **odds ratio** provides a measure of effect size, indicating the strength of the association. An **odds ratio** of 1 suggests no association. In our example, **R** also provides a 95% **confidence interval** for the **odds ratio**: **(0.0130943, 1.8397543)**. Since the number 1 is contained within this interval, it further confirms that our result is not significantly different from the null, assuming an alpha level of 0.05.

Interpreting these results requires a balanced view of both the **p-value** and the **confidence interval**. While the **p-value** tells you about the existence of an effect, the **odds ratio** and its interval tell you about the magnitude and precision of that effect. Together, they provide a comprehensive picture of the relationship between your **categorical variables**. In many scientific journals, reporting both the **p-value** and the **odds ratio** with its interval is considered best practice for transparency.

## Advanced Considerations and Larger Tables

While **Fisher's Exact Test** is most famous for 2x2 tables, it is capable of analyzing larger **contingency table** structures, such as 2x3 or 3x3 matrices. However, as the table size and the total sample size increase, the number of possible permutations grows exponentially. This can lead to significant computational demands. In **R**, if the calculation becomes too complex, the `fisher.test()` function may return an error regarding "workspace size."

To address this, the function includes a `workspace` argument that allows you to increase the memory allocated to the calculation. Furthermore, for very large tables where an exact calculation is impossible, **R** provides an option to use **Monte Carlo simulation** by setting the `simulate.p.value` argument to `TRUE`. This approach estimates the **p-value** by randomly sampling a large number of tables with the same marginal totals, providing a highly accurate approximation when an "exact" result is unreachable.

Another advanced consideration is the assumption of fixed marginals. In many experimental designs, only one set of marginals (either rows or columns) is truly fixed. For instance, in a clinical trial, the number of patients in the treatment and control groups is fixed, but the number of successes and failures is not. While **Fisher's Exact Test** is still widely used in these scenarios, some statisticians argue for alternative tests like Barnard's Test. However, for most practical purposes, the **Fisher's Exact Test** in **R** remains the standard due to its conservative nature and broad acceptance in the scientific community.

Finally, it is worth exploring the "hybrid" approach often used in modern software. **R** handles the internal logic seamlessly, but as a researcher, understanding these nuances ensures that you choose the right tools for your specific data structure. Whether you are dealing with a simple 2x2 table or a multi-dimensional matrix, **Fisher's Exact Test** provides a level of mathematical certainty that approximation-based tests simply cannot match. By mastering these advanced options, you can handle a wider array of data challenges with confidence.

## Summary and Best Practices

Conducting **Fisher's Exact Test** in **R** is a powerful skill for any data analyst. By following a structured workflow--importing data, organizing it into a **contingency table**, executing the

`fisher.test()` function, and carefully interpreting the results--you can derive meaningful insights from even the smallest datasets. This test bridges the gap between raw observation and **statistical significance**, providing a robust framework for making data-driven decisions.

To ensure the highest quality of analysis, always remember to check your cell counts before choosing between the **Chi-square test** and **Fisher's Exact Test**. If your sample size is small or your data is sparse, the exact test is almost always the better choice. Additionally, clearly define your **null hypothesis** and **alternative hypothesis** before running the test to maintain the integrity of your research. Proper labeling of your **R** objects will also make your code more reproducible and easier to share with others.

As you continue to explore the capabilities of **R**, you will find that **Fisher's Exact Test** is just one of many tools available for rigorous statistical analysis. Combining this test with other descriptive statistics and visualization techniques will provide a more holistic view of your data. The following tutorials and official documentation are excellent resources for further expanding your knowledge of categorical data analysis and the versatile world of **R** programming.