

How to Perform Fisher's Exact Test in Python

Authored by
stats writer

December 25, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Perform Fisher's Exact Test in Python*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=108746>

The Fisher's exact test is a rigorous statistical significance test utilized specifically for comparing two **categorical variables**. Its primary role is to determine whether these two variables are statistically independent. This powerful non-parametric test is particularly valuable when analyzing small sample sizes or sparse data, conditions under which asymptotic tests like the Chi-Square test may provide inaccurate results. In the Python ecosystem, this analysis is readily performed using the `scipy.stats.fisher_exact()` function, which accepts a contingency table as input and delivers two crucial metrics: the Odds Ratio and the precise p-value.

The Odds Ratio (1/5) quantifies the strength and direction of the association between the two variables, serving as a measure of effect size. Meanwhile, the p-value (1/5) indicates the probability that the observed results, or results more extreme, occurred purely due to random chance, assuming no association exists. By evaluating the p-value against a predetermined significance level (α), researchers can determine whether there is sufficient evidence to conclude that the two categorical variables are significantly associated, thereby rejecting the null hypothesis (1/5) of independence.

Why Fisher's Exact Test is Essential

The Fisher's exact test (1/5) is specifically designed to determine whether there is a significant association between two categorical variables, making it a cornerstone of non-parametric analysis. It is most often employed as the gold standard alternative to the Pearson's Chi-Square test for independence, especially in situations involving small datasets. While the Chi-Square test relies on large sample approximations and mandates that expected cell counts within the contingency table (1/5) be sufficiently large (typically greater than five), Fisher's test makes no such assumption, calculating probabilities exactly.

The necessity of using an exact test arises when data sparsity or extreme imbalance leads to small expected counts in one or more cells of a 2x2 table. If a standard asymptotic test is used under these conditions, the calculated test statistic will not accurately follow the theoretical Chi-Square distribution, leading to unreliable p-values and potentially erroneous conclusions. By focusing on the exact probability of observing the data given fixed marginal totals--utilizing the **hypergeometric distribution**--Fisher's test guarantees accuracy, regardless of the sample size or cell count magnitude.

This tutorial will guide you through the process of setting up and executing the Fisher's exact test (2/5) within a Python environment, leveraging the statistical capabilities provided by the SciPy library (1/5). We will walk through data preparation, function execution, and the proper interpretation of the resulting Odds Ratio and p-value.

Case Study: Gender and Political Preference Association

To demonstrate the practical application of this test, we examine a common scenario in social research: exploring whether or not gender is significantly associated with political party preference among a cohort of university students. The goal is to determine if the likelihood of identifying as Democrat versus Republican is dependent on the student's gender classification. Since we are testing a relationship between two binary categorical variables (Gender: Female/Male; Party: Democrat/Republican), and anticipating a relatively small sample size, Fisher's Exact Test is the appropriate method.

For this example, a random sample of 25 students was polled. The joint frequencies of gender and political affiliation were recorded, establishing our observed data. This data must be structured into a 2x2 contingency table (2/5), which is the standard input format for the statistical analysis. The arrangement of the table is crucial, as it defines the context for the resultant Odds Ratio (2/5).

The observed frequencies from the survey are presented below:

	Democrat	Republican
Female	8	4
Male	4	9

Step 1: Structuring the Data for Python

The initial step in executing the statistical test is transforming the raw table counts into a Python data structure suitable for the SciPy library (2/5). The `fisher_exact()` function expects the data as a nested list or a two-dimensional array, representing the 2x2 contingency table (3/5). We ensure that the order of rows and columns is maintained consistently (e.g., Rows = Female, Male; Columns = Democrat, Republican).

Based on the observed counts, we define the data structure as follows, where the first list represents the Female row counts, and the second list represents the Male row counts. This clean, programmatic representation of the data is essential for accurate computation.

```
data = ,  
]
```

It is important to emphasize that correct data structuring avoids potential misinterpretation of the test results, particularly the Odds Ratio (3/5). By defining the matrix explicitly, we are setting up the exact comparison for which the Fisher's test will calculate the probability of observing such an outcome under the assumption of independence.

Step 2: Executing Fisher's Exact Test using SciPy

Once the contingency data is prepared, the next phase involves invoking the statistical function from the [SciPy library](#) (3/5). We must first import the `scipy.stats` module to access the `fisher_exact` function. The function is highly optimized and requires minimal input, focusing mainly on the data matrix itself.

The general syntax for the function is: `stats.fisher_exact(table, alternative='two-sided')`. The `alternative` parameter allows for specifying the direction of the test. By default, `'two-sided'` is used, which tests against the alternative hypothesis that the variables are dependent in either direction (positive or negative association). If specific directional hypotheses are being tested, `'less'` or `'greater'` can be specified to perform a one-sided test, focusing the power of the analysis on detecting association in only one tail of the distribution.

The following code demonstrates the execution of the test on our specific dataset, returning the calculated results:

```
import scipy.stats as stats
```

```
print(stats.fisher_exact(data))
```

```
(4.5, 0.1152)
```

Interpreting the Null and Alternative Hypotheses

The output `(4.5, 0.1152)` provides the statistical evidence we need. The first value, 4.5, is the [Odds Ratio](#) (4/5), and the second value, 0.1152, is the [p-value](#) (2/5). Before concluding, we must formally state the statistical hypotheses that the test is designed to evaluate, which provides the framework for decision-making:

H0 (Null Hypothesis): The two variables are **independent**. There is no significant association between gender and political party preference.

H1 (Alternative Hypothesis): The two variables are **not independent**. There is a significant association between gender and political party preference.

The [p-value](#) (3/5) is the metric used to test the veracity of the [null hypothesis](#) (2/5). It tells us the probability of observing a [contingency table](#) (4/5) with such an extreme distribution, given that the marginal totals are fixed and the null hypothesis is true.

Drawing Statistical Conclusions

In our example, the calculated p-value (4/5) for the test is **0.1152**. To make a statistical decision, we compare this value to our predetermined significance level, α . Typically, α is set at 0.05. The decision rule is simple: If the p-value is less than α , we reject the null hypothesis; otherwise, we fail to reject it.

Since 0.1152 is not less than 0.05, we **do not reject the null hypothesis (3/5)**. This means that, based on the statistical evidence gathered from this sample, we do not have sufficient grounds to conclude that a significant association exists between gender and political party preference at this college. While the Odds Ratio (5/5) of 4.5 suggests a tendency for females to have 4.5 times the odds of being Democrats compared to males (relative to being Republicans), this observed magnitude of effect is not statistically significant given the sample size and distribution.

In summary, the statistical conclusion reached through the Fisher's exact test (3/5) indicates that the variation observed in the contingency table (5/5) is likely attributable to random sampling variability. Therefore, we conclude that, statistically speaking, gender and political party preference are independent variables in this context.