

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

Authored by  
**stats writer**

March 9, 2026

## RECOMMENDED CITATION

stats writer (2026). *How to Perform Fisher's Exact Test in Stata: A Step-by-Step Guide*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=134707>

## Introduction to Fisher's Exact Test in Categorical Data Analysis

In the realm of quantitative research, researchers frequently encounter scenarios where they must evaluate the relationship between two **categorical variables**. While the **Chi-squared test** is the most common tool for this purpose, it relies on large-sample approximations that may falter under specific conditions. **Fisher's Exact Test** serves as a powerful alternative, specifically designed to calculate the exact probability of the observed distribution of data in a **contingency table**. By utilizing the **hypergeometric distribution**, this test provides a precise **p-value** without necessitating the large-sample assumptions required by asymptotic tests. In the **Stata** environment, performing this analysis is both straightforward and computationally efficient, making it an essential skill for data scientists and statisticians working with sparse or limited datasets.

The primary utility of **Fisher's Exact Test** is found in its name: it is "exact" because the significance of the deviation from a **null hypothesis** can be calculated exactly, rather than relying on an approximation that becomes accurate only as the sample size grows toward infinity. This is particularly crucial in clinical trials, rare event analysis, or pilot studies where the total number of observations is small. Within **Stata**, the implementation of this test allows researchers to bypass the pitfalls of traditional methods, ensuring that the resulting statistical inferences are valid even when the data structure is inherently limited or unbalanced across different categories.

When analyzing the association between **categorical variables**, it is important to understand that the test focuses on the proportions of observations falling into different groups. For instance, if a researcher is examining the success rate of a new drug compared to a placebo across a small cohort, the resulting 2x2 table might contain "cells" with very low frequencies. In such cases, the mathematical foundations of the **Chi-squared test** break down, potentially leading to incorrect conclusions. **Fisher's Exact Test** solves this by considering all possible combinations of the table that could result in the same marginal totals, thereby determining how extreme the observed data truly is in the context of the **null hypothesis**.

## Theoretical Foundations and When to Use Exact Methods

Understanding the theoretical underpinnings of **Fisher's Exact Test** requires a look into the work of **Sir Ronald A. Fisher**, who popularized the method through his famous "Lady Tasting Tea" experiment. The test is grounded in the principle of permutation, where the marginal totals of the **contingency table** are treated as fixed. By assuming these totals are constant, the test calculates the probability of obtaining the observed cell frequencies--and any more extreme configurations--under the assumption of independence. This approach differs fundamentally from the **Chi-squared test**, which compares the observed frequencies to expected frequencies based on a continuous probability distribution.

The "Rule of Five" is a common heuristic used by statisticians to determine when to pivot from a standard **Chi-squared test** to an exact test. If any expected cell frequency in a 2x2 table is less than five, the approximation provided by the **Chi-squared distribution** may become unreliable. In these instances, **non-parametric statistics** or exact methods like **Fisher's Exact Test** become the gold standard. While modern computing allows **Stata** to perform these calculations on larger tables, the test is traditionally associated with 2x2 tables where the sample size is small, and the data is highly sensitive to individual observations.

Beyond small sample sizes, **Fisher's Exact Test** is also preferred when dealing with highly unbalanced data. If one category has significantly more observations than another, the resulting **contingency table** may have cells with zero or very low counts even if the overall sample size is relatively large. Because **Stata** uses highly optimized algorithms to compute factorials and permutations, researchers can apply this test to ensure **statistical significance** is measured accurately regardless of how the data is distributed across the categorical levels. This level of precision is vital for maintaining the integrity of the research findings and avoiding **Type I errors**.

## Implementing the Tabulate Command for Exact Analysis

To perform **Fisher's Exact Test** in **Stata**, the most common approach involves the `tabulate` command. This command is versatile, used primarily for creating frequency tables, but it includes several options for association testing. When you have two **categorical variables** in your active dataset, you can invoke the test by appending the `exact` option to the command. The basic syntax follows the structure: `tabulate variable1 variable2, exact`. This command tells **Stata** to generate the cross-tabulation of the two variables and then compute the exact **p-value** for the association between them.

One of the advantages of using the `tabulate` command is the ability to include additional descriptors that enrich the output. For example, adding `row`, `column`, or `cell` options allows the researcher to see the percentage distributions alongside the raw counts. While the `exact` option provides the **p-value**, seeing the percentages helps in interpreting the direction and magnitude of the association. This is particularly helpful when the result is statistically significant, as it allows for a more nuanced discussion of the relationship between the **categorical variables** in the final research report.

For researchers who do not have their data already loaded into a dataset but have the summary counts from a **contingency table**, **Stata** offers the `tabi` command, or "tabulate immediate." This is an incredibly efficient way to run **Fisher's Exact Test** without creating a `.dta` file. The syntax requires the user to input the cell counts row by row, separated by a backslash. For a 2x2 table, the command would look like `tabi a b c d, exact`. This flexibility makes **Stata** a preferred tool for quickly verifying findings from published literature or preliminary field data.

## A Practical Case Study: Gender and Political Preference

To illustrate the practical application of this test, let us consider a study exploring whether gender is associated with political party preference at a university. In this scenario, the researcher is working with a total sample of 25 students--a sample size that is clearly too small for a standard **Chi-squared test**. The objective is to determine if the distribution of Democrats and Republicans differs significantly between male and female students. The raw data for this study is organized into the following **contingency table**:

### Table 1: Student Political Preference by Gender

**Male Students:** 4 Democrats and 9 Republicans

**Female Students:** 8 Democrats and 4 Republicans

This table demonstrates a clear disparity in counts, with more males leaning Republican and more females leaning Democrat. However, with only 25 total participants, we must determine if this observation is a result of random chance or if it reflects a true **statistical significance**. Because several of the cell counts are low (specifically the 4 male Democrats and 4 female Republicans), **Fisher's Exact Test** is the appropriate choice to evaluate the **null hypothesis** that gender and political preference are independent variables.

In **Stata**, we can process this data using the immediate command. By entering the counts from left to right, row by row, we construct the command that will provide our results. The logical flow of the data entry ensures that the software understands the relationship between the rows (Gender) and columns (Political Party). This specific case study serves as a perfect template for any researcher needing to analyze a binary relationship with a limited number of observations.

## Executing the Command and Interpreting Stata Output

To execute the analysis for our case study, we use the following command in the **Stata** command window:

```
tabi 4 9 8 4, exact
```

This command instructs the software to create a 2x2 table with the first row containing 4 and 9, and the second row containing 8 and 4. The `exact` option triggers the calculation of the **Fisher's Exact Test p-value**. Upon pressing enter, **Stata** generates the following output, which includes the frequency table and the calculated exact probabilities:

```
. tabi 4 9 \ 8 4
```

row	col		Total
	1	2	
1	4	9	13
2	8	4	12
Total	12	13	25

```

Fisher's exact = 0.115
1-sided Fisher's exact = 0.081

```

When reviewing the output, the researcher should focus on several key areas. First, the **Output table** confirms that the data was entered correctly, reflecting the 4, 9, 8, and 4 distribution. Below the table, **Stata** provides two distinct **p-values**. The "Fisher's exact" value represents the **two-sided test**, which is the standard choice when the researcher does not have a pre-defined direction for the relationship. In this case, the two-sided **p-value** is 0.115. Additionally, the software provides a "1-sided Fisher's exact" value of 0.081, which would only be used if there were a specific directional **alternative hypothesis** established before the data collection.

The interpretation of these values is governed by the chosen **significance level**, typically denoted as alpha ( $\alpha = 0.05$ ). Since our two-sided **p-value** of 0.115 is greater than the threshold of 0.05, we fail to reject the **null hypothesis**. This means that based on this specific sample of 25 students, there is not enough evidence to conclude that gender and political party preference are significantly associated. While the raw numbers show a trend, the small sample size prevents us from ruling out the possibility that this distribution occurred by chance.

## Advanced Options and One-Sided vs. Two-Sided Tests

In statistical testing, the choice between a **one-sided** and a **two-sided test** is critical. A **two-sided test** evaluates whether there is any difference between the groups, regardless of direction. This is generally considered the more conservative and robust approach in scientific research. **Fisher's Exact Test** in **Stata** defaults to providing both, but the two-sided value is the one most commonly reported in academic journals unless a strong theoretical justification for a one-sided approach exists.

A **one-sided test** is appropriate only when the researcher is testing for a difference in a specific direction--for example, testing only if females are more likely to be Democrats than males, rather than testing if there is any difference at all. Because the one-sided **p-value** (0.081 in our example) is smaller than the two-sided value, it is easier to achieve **statistical significance**. However, using a one-sided test solely to reach a "significant" result is considered poor practice and can lead to

inflated **Type I error** rates. Researchers must decide on the tails of the test during the design phase of the study, not after looking at the results.

Furthermore, **Stata** allows for the analysis of tables larger than 2x2 using the `exact` option. While the math becomes significantly more complex as the number of rows and columns increases, the software's algorithms can handle larger **contingency tables** to provide an "exact" significance level. This is particularly useful in multi-level **categorical variables** where cell counts remain sparse. For very large tables where an exact **p-value** might take too long to compute, **Stata** also offers a **Monte Carlo** option (`tabulate var1 var2, exact(mcc)`), which estimates the exact **p-value** through repeated sampling, providing a balance between precision and computational speed.

## Concluding Thoughts on Statistical Inference in Stata

The ability to perform **Fisher's Exact Test** is an invaluable asset for any researcher working with **categorical variables** and small datasets. By following the steps outlined in this guide--from understanding the **null hypothesis** of independence to executing the `tabi` command--users can ensure that their statistical conclusions are based on rigorous, exact calculations rather than potentially flawed approximations. While our example regarding gender and political preference did not yield **statistical significance**, the process demonstrated how **Stata** provides the clarity needed to make informed decisions about the data.

It is important to remember that failing to reject the **null hypothesis** does not necessarily mean that no relationship exists; rather, it suggests that the current sample does not provide enough evidence to confidently identify one. In such cases, the researcher might consider a larger sample size in future studies to increase **statistical power**. Regardless of the outcome, **Fisher's Exact Test** remains a cornerstone of **non-parametric statistics**, offering a level of reliability that is unmatched by other methods in the context of sparse **contingency tables**.

In summary, the transition from raw data to a finished statistical report is made seamless by the powerful features of **Stata**. By mastering the `tabulate` and `tabi` commands with the `exact` option, you can handle a wide variety of research questions across disciplines such as sociology, medicine, and political science. Always ensure that you are reporting the correct **p-value** and contextualizing your findings within the broader theoretical framework of your study. With these tools at your disposal, your analysis of **categorical variables** will be both technically sound and professionally presented.